
Best Practice Guide Intel Xeon Phi v0.1

Michaela Barth, KTH Sweden

Mikko Byckling, CSC Finland

Nevena Ilieva, NCSA Bulgaria

Sami Saarinen, CSC Finland

Michael Schliephake, KTH Sweden

Volker Weinberg (Editor), LRZ Germany <weinberg@lrz.de>

31-03-2013



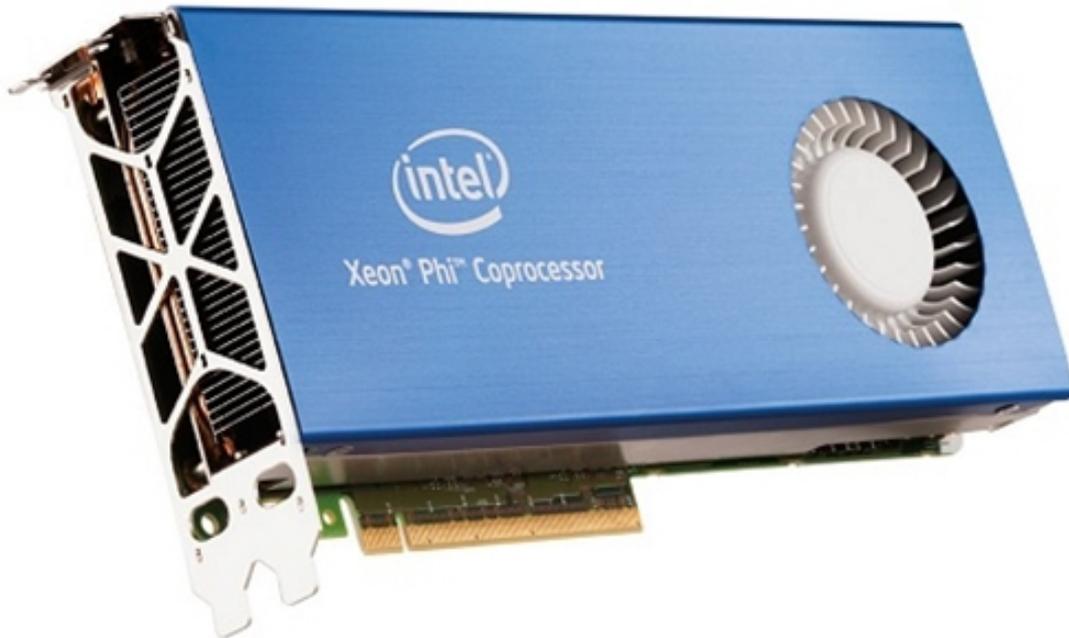
Table of Contents

1. Introduction	4
2. MIC architecture, system overview	4
2.1. The Intel MIC architecture	4
2.1.1. Intel Xeon Phi coprocessor architecture overview	4
2.1.2. The cache hierarchy	6
2.2. Network configuration & system access	7
3. Native compilation	11
4. Intel compiler's offload pragmas	12
4.1. Simple example	12
4.2. Obtaining informations about the offloading	13
4.3. Syntax of pragmas	14
4.4. Recommendations	15
4.4.1. Explicit worksharing	15
4.4.2. Persistent data on the coprocessor	15
4.4.3. Optimizing offloaded code	16
4.5. Intel Cilk Plus parallel extensions	17
5. OpenMP and hybrid	17
5.1. OpenMP	17
5.1.1. Programming models and offload	17
5.1.2. Threading and affinity	17
5.1.3. Loop scheduling	18
5.1.4. Scalability improvement	18
5.2. Hybrid OpenMP/MPI	18
5.2.1. Programming models	18
5.2.2. Threading of the MPI ranks	18
6. MPI	19
6.1. Setting up the MPI environment	19
6.2. MPI programming models	19
6.2.1. Coprocessor-only model	20
6.2.2. Symmetric model	20
6.2.3. Host-only model	20
6.3. Simplifying launching of MPI jobs	21
7. Intel MKL (Math Kernel Library)	21
7.1. MKL usage modes	21
7.1.1. Automatic Offload (AO)	21
7.1.2. Compiler Assisted Offload (CAO)	22
7.1.3. Native Execution	23
7.2. Example code	23
7.3. Intel Math Kernel Library Link Line Advisor	23
8. TBB: Intel Threading Building Blocks	23
8.1. Advantages of TBB	24
8.2. Using TBB natively	25
8.3. Offloading TBB	27
9. Debugging	27
9.1. Native debugging with gdb	27
9.2. Remote debugging with gdb	27
10. Tuning	28
10.1. Advanced OpenMP	28
10.1.1. OpenMP thread affinity	28
10.1.2. Example: Thread affinity	29
10.1.3. Multiple parallel regions and barriers	29
10.1.4. Example: Multiple parallel regions and barriers	30
10.1.5. False sharing	31
10.1.6. Example: False sharing	31
10.1.7. Memory limitations	32

10.1.8. Example: Memory limitations	32
10.1.9. Nested parallelism	33
10.1.10. Example: Nested parallelism	33
10.1.11. Load balancing	34
10.1.12. Example: Load balancing	34
11. Performance analysis tools	35
Further documentation	35

1. Introduction

Figure 1. Intel Xeon Phi coprocessor



This best practice guide provides information about Intel's MIC architecture and programming models for the Intel Xeon Phi coprocessor in order to enable programmers to achieve good performance of their applications. The guide covers a wide range of topics from the description of the hardware of the Intel Xeon Phi coprocessor through information about the basic programming models as well as information about porting programs up to tools and strategies how to analyze and improve the performance of applications.

This is a very preliminary draft of the Best Practice Guide which is still under development. Several informations from publicly available Intel documents and especially recent webinars [9] have been incorporated into the current draft. It further contains contributions from CSC, KTH, LRZ, and NCSA. The final version will include more experience gained within PRACE with the Intel Xeon Phi coprocessor and will be available in the end of December 2013.

Recently the first book about programming the Intel Xeon Phi coprocessor [1] has been published. We also recommend a book about structured parallel programming [2]. Useful online documentation about the Intel Xeon Phi coprocessor can be found in Intel's developer zone for Xeon Phi Programming [4] and the Intel Many Integrated Core Architecture User Forum [5]. To get things going quickly have a look on the Intel Xeon Phi Coprocessor Developer's Quick Start Guide [13] and also on the paper [20].

2. MIC architecture, system overview

2.1. The Intel MIC architecture

2.1.1. Intel Xeon Phi coprocessor architecture overview

The Intel Xeon Phi coprocessor consists of up to 61 cores connected by a high performance on-die bidirectional interconnect. The coprocessor runs a full service Linux operating system and supports all important Intel development tools, like C/C++ and Fortran compiler, MPI and OpenMP, high performance libraries like MKL, debugger and tracing tools like Intel VTune Amplifier XE. It is connected to an Intel Xeon processor - the "host" - via the PCI Express (PICE) bus. The implementation of a virtualized TCP/IP stack allows to access the coprocessor like

a network node. Summarized information about the hardware architecture can be found in [14]. In the following we cite the most important properties of the MIC architecture from the System Software Developers Guide [15], which includes many details about the MIC architecture:

- Core**
 - the processor core (scalar unit) is an in-order architecture (based on the Intel Pentium processor family)
 - fetches and decodes instructions from *four hardware threads*
 - supports a 32-bit and 64-bit execution environment, along with Intel Initial Many Core Instructions
 - does not support any previous Intel SIMD extensions like MME, SSE, SSE2, SSE3, SSE4.1 SSE4.2 or AVX instructions
 - new vector instructions provided by the Intel Xeon Phi coprocessor instruction set utilize a dedicated 512-bit wide vector floating-point unit (VPU) that is provided for each core
 - high performance support for reciprocal, square root, power and exponent operations, scatter/gather and streaming store capabilities to achieve higher effective memory bandwidth
 - can execute 2 instructions per cycle, one on the U-pipe and one on the V-pipe (not all instruction types can be executed by the V-pipe, e.g. vector instructions can only be executed on the U-pipe)
 - contains the L1 Icache and Dcache
 - each core is connected to a ring interconnect via the Core Ring Interface (CRI)
- Vector Processing Unit (VPU)**
 - the VPU includes the EMU (Extended Math Unit) and executes 16 single-precision floating point, 16 32bit integer operations or 8 double-precision floating point operations per cycle. Each operation can be a fused multiply-add, giving 32 single-precision or 16 double-precision floating-point operations per cycle.
 - contains the vector register file: 32 512-bit wide registers per thread context, each register can hold 16 singles or 8 doubles
 - most vector instructions have a 4-clock latency with a 1 clock throughput
- Core Ring Interface (CRI)**
 - hosts the L2 cache and the tag directory (TD)
 - connects each core to an Intel Xeon Phi coprocessor Ring Stop (RS), which connects to the interprocessor core network
- Ring**
 - includes component interfaces, ring stops, ring turns, addressing and flow control
 - a Xeon Phi coprocessor has 2 of these rings, one travelling each direction
- SBOX**
 - Gen2 PCI Express client logic
 - system interface to the host CPU or PCI Express switch
 - DMA engine
- GBOX**
 - coprocessor memory controller

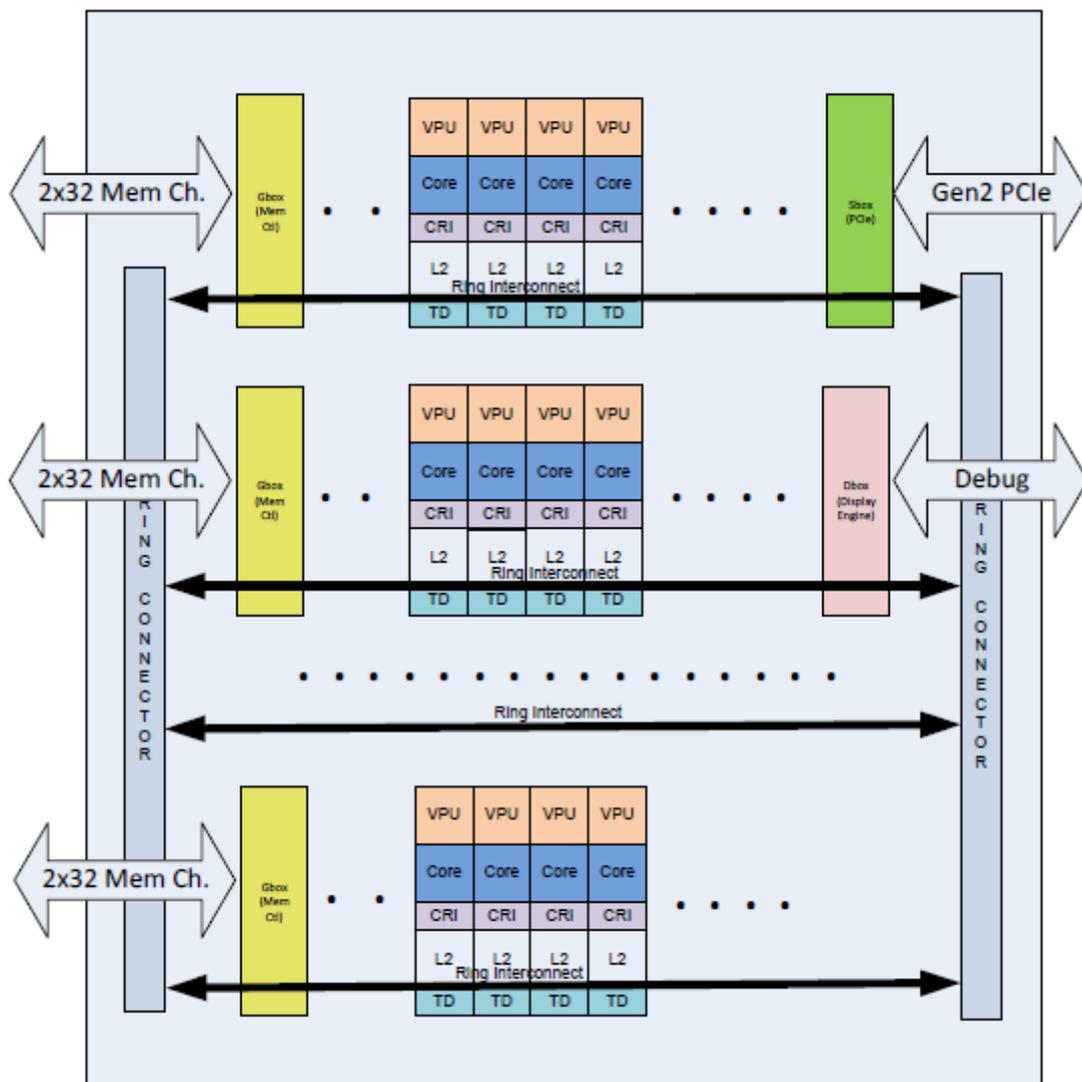
- consists of the FBOX (interface to the ring interconnect), the MBOX (request scheduler) and the PBOX (physical layer that interfaces with the GDDR devices).
- There are 8 memory controllers supporting up to 16 GDDR5 channels. With a transfer speed of up to 5.5 GT/s a theoretical aggregated bandwidth of 352 GB/s is provided.

Performance Monitoring Unit (PMU)

- allows data to be collected from all units in the architecture
- does not implement some advanced features found in mainline IA cores (e.g. precise event-based sampling, etc.)

The following picture (from [15]) illustrates the building blocks of the architecture.

Figure 2. MIC architecture overview



2.1.2. The cache hierarchy

Details about the L1 and L2 cache can be found in the System Software Developers Guide [15]. We only cite the most important features here.

The L1 cache has a 32 KB L1 instruction cache and 32 KB L1 data cache. Associativity is 8-way, with a cache line-size of 64 byte. It has a load-to-use latency of 1 cycle, which means that an integer value loaded from the L1 cache can be used in the next clock by an integer instruction. (Vector instructions have different latencies than integer instructions.)

The L2 cache is a unified cache which is inclusive of the L1 data and instruction caches. Each core contributes 512 KB of L2 to the total global shared L2 cache storage. If no cores share any data or code, then the effective total L2 size of the chip is up to 31 MB. On the other hand, if every core shares exactly the same code and data in perfect synchronization, then the effective total L2 size of the chip is only 512 KB. The actual size of the workload-perceived L2 storage is a function of the degree of code and data sharing among cores and thread.

Like for the L1 cache, associativity is 8-way, with a cache line-size of 64 byte. The raw latency is 11 clock cycles. It has a streaming hardware prefetcher and supports ECC correction.

The main properties of the L1 and L2 caches are summarized in the following table (from [15]):

Parameter	L1	L2
Coherence	MESI	MESI
Size	32 KB + 32 KB	512 KB
Associativity	8-way	8-way
Line Size	64 bytes	64 bytes
Banks	8	8
Access Time	1 cycle	11 cycles
Policy	pseudo LRU	pseudo LRU
Duty Cycle	1 per clock	1 per clock
Ports	Read or Write	Read or Write

2.2. Network configuration & system access

Details about the system startup and the network configuration can be found in [16] and in the documentation coming with MPSS [8].

To start the Intel Manycore Platform Software Stack (Intel MPSS) and initialize the Xeon Phi coprocessor the following command has to be executed as root or during host system start-up:

```
weinberg@knf1:~> sudo service mpss start
```

During start-up details are logged to `/var/log/messages`.

If MPSS with OFED support is needed, further the following commands have to be executed as root:

```
weinberg@knf1:~> sudo service openibd start
weinberg@knf1:~> sudo service opensmd start
weinberg@knf1:~> sudo service ofed-mic start
```

Per default IP addresses 172.31.1.254 , 172.31.2.254 , 172.31.3.254 etc. are then assigned to the attached Intel Xeon Phi coprocessors. The IP addresses of the attached coprocessors can be listed via the traditional `ifconfig` Linux program.

```
weinberg@knf1:~> /sbin/ifconfig
...
```

```
mic0      Link encap:Ethernet  HWaddr 86:32:20:F3:1A:4A
          inet addr:172.31.1.254  Bcast:172.31.1.255  Mask:255.255.255.0
          inet6 addr: fe80::8432:20ff:fef3:1a4a/64 Scope:Link
          UP BROADCAST RUNNING MTU:64512 Metric:1
          RX packets:46407 errors:0 dropped:0 overruns:0 frame:0
          TX packets:42100 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:8186629 (7.8 Mb)  TX bytes:139125152 (132.6 Mb)

mic1      Link encap:Ethernet  HWaddr FA:7A:0D:1E:65:B0
          inet addr:172.31.2.254  Bcast:172.31.2.255  Mask:255.255.255.0
          inet6 addr: fe80::f87a:dff:fe1e:65b0/64 Scope:Link
          UP BROADCAST RUNNING MTU:64512 Metric:1
          RX packets:5097 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4240 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:522169 (509.9 Kb)  TX bytes:124671342 (118.8 Mb)
```

Further information can be obtained by running the `micinfo` program on the host. To get also PCIe related details the command has to be run with root privileges. Here is an example output for a B0 stepping Knights Corner prototype:

```
weinberg@knfl:~> sudo /opt/intel/mic/bin/micinfo
MicInfo Utility Log
```

Created Tue Mar 12 15:00:32 2013

System Info

```
Host OS           : Linux
OS Version        : 3.0.13-0.27-default
Driver Version    : 4346-16
MPSS Version      : 2.1.4346-16
Host Physical Memory : 66056 MB
CPU Family        : GenuineIntel  Family 6  Model 45  Stepping 5
CPU Speed         : 2594.169
Threads per Core  : 2
```

Device No: 0, Device Name: Intel(R) Xeon Phi(TM) coprocessor

Version

```
Flash Version     : 2.1.01.0375
UOS Version       : 2.6.34.11-g65c0cd9
Device Serial Number : ADKC22600276
```

Board

```
Vendor ID         : 8086
Device ID         : 225c
SubSystem ID      : 2500
MIC Processor Stepping ID : 1
PCIe Width        : x16
PCIe Speed        : 5 GT/s
```

```
PCIe Max payload size      : 256 bytes
PCIe Max read req size    : 4096 bytes
MIC Processor Model       : 0x01
MIC Processor Model Ext   : 0x00
MIC Processor Type       : 0x00
MIC Processor Family     : 0x0b
MIC Processor Family Ext  : 0x00
MIC Silicon Stepping     : B0
Board SKU                 : ES2-P1750
ECC Mode                 : Enabled
SMC HW Revision          : Product 300W Active CS
```

Core

```
Total No of Active Cores: 61
Voltage                 : 972000 uV
Frequency               : 1090909 kHz
```

Thermal

```
Fan Speed Control      : On
SMC Firmware Version   : 1.6.3983
FSC Strap              : 14 MHz
Fan RPM                : 2700
Fan PWM                : 50
Die Temp               : 75 C
```

GDDR

```
GDDR Vendor            : Elpida
GDDR Version           : 0x1
GDDR Density           : 2048 Mb
GDDR Size              : 7936 MB
GDDR Technology        : GDDR5
GDDR Speed             : 5.500000 GT/s
GDDR Frequency         : 2750000 kHz
GDDR Voltage           : 1000000 uV
```

Device No: 1, Device Name: Intel(R) Xeon Phi(TM) coprocessor

...

Users can log in directly onto the Xeon Phi coprocessor via ssh.

```
weinberg@knf1:~> ssh mic0
[weinberg@knf1-mic0 weinberg]$ hostname
knf1-mic0
[weinberg@knf1-mic0 weinberg]$ cat /etc/issue
Intel MIC Platform Software Stack release 2.1
Kernel 2.6.34.11-g65c0cd9 on an k1om
```

Per default the home directory on the coprocessor is `/home/username` .

Since the access to the coprocessor is ssh-key based users have to generate a private/public key pair via `ssh-keygen` before accessing the coprocessor for the first time.

After the keys have been generated, the following commands have to be executed as root to populate the filesystem image for the coprocessor on the host (`/opt/intel/mic/filesystem/mic0/home`) with the new keys. Since the coprocessor has to be restarted to copy the new image to the coprocessor, the following commands have to be used (preferably only by the system administrator) with care.

```
weinberg@knf1:~> sudo service mpss stop
weinberg@knf1:~> sudo micctrl --resetconfig
weinberg@knf1:~> sudo service mpss start
```

Since a Linux kernel is running on the coprocessor, further information about the cores, memory etc. can be obtained from the virtual Linux /proc or /sys filesystems:

```
[weinberg@knf1-mic0 weinberg]$ tail -n26 /proc/cpuinfo
```

```
processor          : 243
vendor_id         : GenuineIntel
cpu family       : 11
model            : 1
model name       : 0b/01
stepping        : 1
cpu MHz         : 1090.908
cache size      : 512 KB
physical id     : 0
siblings        : 244
core id         : 60
cpu cores       : 61
apicid          : 243
initial apicid  : 243
fpu             : yes
fpu_exception   : yes
cpuid level     : 4
wp              : yes
flags           : fpu vme de pse tsc msr pae mce cx8 apic mtrr mca pat fxsr
                  ht syscall lm lahf_lm
bogomips        : 2190.14
clflush size    : 64
cache_alignment : 64
address sizes   : 40 bits physical, 48 bits virtual
power management:
```

```
[weinberg@knf1-mic0 weinberg]$ head -5 /proc/meminfo
```

```
MemTotal:       7876656 kB
MemFree:        7449188 kB
Buffers:         0 kB
Cached:         199688 kB
SwapCached:     0 kB
```

To run MKL, OpenMP or MPI based programs on the coprocessor, some libraries (exact path may differ depending on the version) need to be copied to the coprocessor. Root privileges are necessary for the destination directories given in the following example:

```
scp /opt/intel/composerxe/mkl/lib/mic/libmkl_intel_lp64.so root@mic0:/lib64/
scp /opt/intel/composerxe/mkl/lib/mic/libmkl_intel_thread.so root@mic0:/lib64/
scp /opt/intel/composerxe/mkl/lib/mic/libmkl_core.so root@mic0:/lib64/

scp /opt/intel/composerxe/lib/mic/libiomp5.so root@mic0:/lib64/
scp /opt/intel/composerxe/lib/mic/libimf.so root@mic0:/lib64/
scp /opt/intel/composerxe/lib/mic/libsvml.so root@mic0:/lib64/
```

```
scp /opt/intel/composerxe/lib/mic/libirng.so root@mic0:/lib64/
scp /opt/intel/composerxe/lib/mic/libintl.so.5 root@mic0:/lib64/

scp /opt/intel/mpi-rt/4.1.0/mic/lib/libmpi_mt.so.4 root@mic0:/lib64/
scp /opt/intel/mpi-rt/4.1.0/mic/lib/libmpigf.so.4 root@mic0:/lib64/
scp /opt/intel/impi/4.1.0.024/mic/lib/libmpi.so.4 root@mic0:/lib64/

scp /opt/intel/impi/4.1.0.024/mic/bin/mpiexec.hydra root@mic0:/bin
scp /opt/intel/impi/4.1.0.024/mic/bin/pmi_proxy root@mic0:/bin
```

3. Native compilation

The simplest model of running applications on the Intel Xeon Phi coprocessor is native mode. Detailed information about building a native application for Intel Xeon Phi coprocessors can be found in [22].

In native mode an application is compiled on the host using the compiler switch `-mmic` to generate code for the MIC architecture. The binary can then be copied to the coprocessor and has to be started there.

```
weinberg@knf1:~/c> . /opt/intel/composerxe/bin/compilervars.sh intel64
weinberg@knf1:~/c> icc -O3 -mmic program.c -o program
weinberg@knf1:~/c> scp program mic0:
program                               100%  10KB  10.2KB/s   00:00
weinberg@knf1:~/c> ssh mic0 ~/program
hello, world
```

To achieve good performance one should mind the following items. More details about the techniques mentioned will be added in a future version of this guide.

- Data should be **aligned to 64 Bytes (512 Bits)** for the MIC architecture, in contrast to 32 Bytes (256 Bits) for AVX and 16 Bytes (128 Bits) for SSE.
- Due to the large SIMD width of 64 Bytes **vectorization is even more important for the MIC architecture than for Intel Xeon!** The MIC architecture offers new instructions like gather/scatter, fused multiply-add, masked vector instructions etc. which allow more loops to be parallelized on the coprocessor than on an Intel Xeon based host.
- Use pragmas like `#pragma ivdep`, `#pragma vector always`, `#pragma vector aligned`, `#pragma simd` etc. to achieve autovectorization. Autovectorization is enabled at default optimization level `-O2`. Requirements for vectorizable loops can be found in [31].
- Let the compiler generate vectorization reports using the compiler option `-vecreport2` to see if loops were vectorized for MIC (Message `"*MIC* Loop was vectorized"` etc). The options `-opt-report-phase hlo` (High Level Optimizer Report) or `-opt-report-phase ipo_inl` (Inlining report) may also be useful.
- Explicit vector programming is also possible via Intel Cilk Plus language extensions (C/C++ array notation, vector elemental functions, ...) or the new SIMD constructs from OpenMP 4.0 RC1.
- Vector elemental functions can be declared by using `__attributes__((vector))`. The compiler then generates a vectorized version of a scalar function which can be called from a vectorized loop.
- One can use intrinsics to have full control over the vector registers and the instruction set. Include `<immintrin.h>` for using intrinsics.
- Hardware prefetching from the L2 cache is enabled per default. In addition, software prefetching is on by default at compiler optimization level `-O2` and above. Since Intel Xeon Phi is an in-order architecture, care about

prefetching is more important than on out-of-order architectures. The compiler prefetching can be influenced by setting the compiler switch `-opt-prefetch=n`. Manual prefetching can be done by using intrinsics (`_mm_prefetch()`) or pragmas (`#pragma prefetch var`).

4. Intel compiler's offload pragmas

One can simply add OpenMP-like pragmas to C/C++ or Fortran code to mark regions of code that should be offloaded to the Intel Xeon Phi Coprocessor and be run there. This approach is quite similar to the accelerator pragmas introduced by the PGI compiler, CAPS HMPP or OpenACC to offload code to GPGPUs. When the Intel compiler encounters an offload pragma, it generates code for both the coprocessor and the host. Code to transfer the data to the coprocessor is automatically created by the compiler, however the programmer can influence the data transfer by adding data clauses to the offload pragma. Details can be found under "Offload Using a Pragma" in the Intel compiler documentation [26].

4.1. Simple example

In the following we show a simple example how to offload a matrix-matrix computation to the coprocessor.

```
main(){

    double *a, *b, *c;
    int i,j,k, ok, n=100;

    // allocated memory on the heap aligned to 64 byte boundary
    ok = posix_memalign((void**)&a, 64, n*n*sizeof(double));
    ok = posix_memalign((void**)&b, 64, n*n*sizeof(double));
    ok = posix_memalign((void**)&c, 64, n*n*sizeof(double));

    // initialize matrices
    ...
    //offload code
#pragma offload target(mic) in(a,b:length(n*n)) inout(c:length(n*n))
    {
        //parallelize via OpenMP on MIC
#pragma omp parallel for
        for( i = 0; i < n; i++ ) {
            for( k = 0; k < n; k++ ) {
#pragma vector aligned
#pragma ivdep
                for( j = 0; j < n; j++ ) {
                    //c[i][j] = c[i][j] + a[i][k]*b[k][j];
                    c[i*n+j] = c[i*n+j] + a[i*n+k]*b[k*n+j];
                }
            }
        }
    }
}
```

This example (with quite bad performance) shows how to offload the matrix computation to the coprocessor using the `#pragma offload target(mic)`. One could also specify the specific coprocessor `num` in a system with multiple coprocessors by using `#pragma offload target(mic:num)`.

Since the matrices have been dynamically allocated using `posix_memalign()`, their sizes must be specified via the `length()` clause. Using `in`, `out` and `inout` one can specify which data has to be copied in which direction. It is recommended that for Intel Xeon Phi data is 64-byte aligned. `#pragma vector aligned` tells the compiler that all array data accessed in the loop is properly aligned. `#pragma ivdep` discards any data dependencies assumed by the compiler.

Offloading is enabled per default for the Intel compiler. Use `-no-offload` to disable the generation of offload code.

4.2. Obtaining informations about the offloading

Using the compiler option `-vec-report2` one can see which loops have been vectorized on the host and the MIC coprocessor:

```
weinberg@knf1:~/c> icc -vec-report2 -openmp offload.c
offload.c(57): (col. 2) remark: loop was not vectorized: vectorization
                    possible but seems inefficient.
...
offload.c(57): (col. 2) remark: *MIC* LOOP WAS VECTORIZED.
offload.c(54): (col. 7) remark: *MIC* loop was not vectorized: not inner loop.
offload.c(53): (col. 5) remark: *MIC* loop was not vectorized: not inner loop.
```

By setting the environment variable `OFFLOAD_REPORT` one can obtain information about performance and data transfers at runtime:

```
weinberg@knf1:~/c> export OFFLOAD_REPORT=2
weinberg@knf1:~/c> ./a.out
[Offload] [MIC 0] [File]           offload2.c
[Offload] [MIC 0] [Line]          50
[Offload] [MIC 0] [CPU Time]      12.853562 (seconds)
[Offload] [MIC 0] [CPU->MIC Data] 9830416 (bytes)
[Offload] [MIC 0] [MIC Time]     12.208636 (seconds)
[Offload] [MIC 0] [MIC->CPU Data] 3276816 (bytes)
```

If a function is called within the offloaded code block, this function has to be declared with `__attribute__((target(mic)))`.

For example one could put the matrix-matrix multiplication of the previous example into a subroutine and call that routine within an offloaded block region:

```
__attribute__((target(mic))) void mxm( int n, double * restrict a,
                                     double * restrict b, double *restrict c ){
    int i,j,k;
    for( i = 0; i < n; i++ ) {
        ...
    }
}

main(){
    ...

#pragma offload target(mic) in(a,b:length(n*n)) inout(c:length(n*n))
{
    mxm(n,a,b,c);
}
```

```
}
}
}
```

Mind the C99 restrict keyword that specifies that the vectors do not overlap. (Compile with -std=c99)

4.3. Syntax of pragmas

The following offload pragmas are available (from [9]):

Pragma	Syntax	Semantic
C/C++		
Offload pragma	<code>#pragma offload <clauses> <statement></code>	Allow next statement to execute on coprocessor or host CPU
Variable/function offload properties	<code>_attribute__ ((target(mic)))</code>	Compile function for, or allocate variable on, both host CPU and coprocessor
Entire blocks of data/code defs	<code>#pragma offload_attribute(push, target(mic))</code> ... <code>#pragma offload_attribute(pop)</code>	Mark entire files or large blocks of code to compile for both host CPU and coprocessor
Fortran		
Offload directive	<code>!dir\$ omp offload <clauses> <statement></code>	Execute OpenMP parallel block on coprocessor
Variable/function offload properties	<code>!dir\$ attributes offload:<mic> :: <ret-name> OR <var1,var2,...></code>	Compile function or variable for CPU and coprocessor
Entire code blocks	<code>!dir\$ offload begin <clauses></code> ... <code>!dir\$ end offload</code>	Mark entire files or large blocks of code to compile for both host CPU and coprocessor

The following clauses can be used to control data transfers:

Clause	Syntax	Semantic
Multiple coprocessors	<code>target(mic[:unit])</code>	Select specific coprocessors
Inputs	<code>in(var-list modifiers)</code>	Copy from host to coprocessor
Outputs	<code>out(var-list modifiers)</code>	Copy from coprocessor to host
Inputs & outputs	<code>inout(var-list modifiers)</code>	Copy host to coprocessor and back when offload completes
Non-copied data	<code>nocopy(var-list modifiers)</code>	Data is local to target

The following (optional) modifiers are specified:

Modifier	Syntax	Semantic
Specify copy length	<code>length(N)</code>	Copy N elements of pointer's type

Modifier	Syntax	Semantic
Coprocessor memory allocation	<code>alloc_if (bool)</code>	Allocate coprocessor space on this offload (default: TRUE)
Coprocessor memory release	<code>free_if (bool)</code>	Free coprocessor space at the end of this offload (default: TRUE)
Control target data alignment	<code>align (N bytes)</code>	Specify minimum memory alignment on coprocessor
Array partial allocation & variable relocation	<code>alloc (array-slice) into (var-expr)</code>	Enables partial array allocation and data copy into other vars & ranges

4.4. Recommendations

4.4.1. Explicit worksharing

To explicitly share work between the coprocessor and the host one can use OpenMP sections to manually distribute the work. In the following example both the host and the coprocessor will run a matrix-matrix multiplication in parallel.

```
#pragma omp parallel
{
#pragma omp sections
{
#pragma omp section
{
//section running on the coprocessor
#pragma offload target(mic) in(a,b:length(n*n)) inout(c:length(n*n))
{
    mxm(n,a,b,c);
}
}
#pragma omp section
{
//section running on the host
    mxm(n,d,e,f);
}
}
}
```

4.4.2. Persistent data on the coprocessor

The main bottleneck of accelerator based programming are data transfers over the slow PCIe bus from the host to the accelerator and vice versa. To increase the performance one should minimize data transfers as much as possible and keep the data on the coprocessor between computations using the same data.

Defining the following macros

```
#define ALLOC    alloc_if(1)
#define FREE     free_if(1)
#define RETAIN   free_if(0)
```

```
#define REUSE    alloc_if(0)
```

one can simply use the following notation:

- to allocate data and keep it for the next offload

```
#pragma offload target(mic) in (p:length(1) ALLOC RETAIN)
```

- to reuse the data and still keep it on the coprocessor

```
#pragma offload target(mic) in (p:length(1) REUSE RETAIN)
```

- to reuse the data again and free the memory. (FREE is the default, and does not need to be explicitly specified)

```
#pragma offload target(mic) in (p:length(1) REUSE FREE)
```

More information can be found in the section "Managing Memory Allocation for Pointer Variables" under "Offload Using a Pragma" in the compiler documentation [26].

4.4.3. Optimizing offloaded code

The implementation of the matrix-matrix multiplication given in Section 4.1 can be optimized by defining appropriate ROWCHUNK and COLCHUNK chunk sizes, rewriting the code with 6 nested loops (using OpenMP collapse for the 2 outermost loops) and some manual loop unrolling (thanks to A. Heinecke for input for this section).

```
#define ROWCHUNK 96
#define COLCHUNK 96

#pragma omp parallel for collapse(2) private(i,j,k)
    for(i = 0; i < n; i+=ROWCHUNK ) {
        for(j = 0; j < n; j+=ROWCHUNK ) {
            for(k = 0; k < n; k+=COLCHUNK ) {
                for (ii = i; ii < i+ROWCHUNK; ii+=6) {
                    for (kk = k; kk < k+COLCHUNK; kk++ ) {
#pragma ivdep
#pragma vector aligned
                        for ( jj = j; jj < j+ROWCHUNK; jj++){
                            c[(ii*n)+jj] += a[(ii*n)+kk]*b[kk*n+jj];
                            c[((ii+1)*n)+jj] += a[((ii+1)*n)+kk]*b[kk*n+jj];
                            c[((ii+2)*n)+jj] += a[((ii+2)*n)+kk]*b[kk*n+jj];
                            c[((ii+3)*n)+jj] += a[((ii+3)*n)+kk]*b[kk*n+jj];
                            c[((ii+4)*n)+jj] += a[((ii+4)*n)+kk]*b[kk*n+jj];
                            c[((ii+5)*n)+jj] += a[((ii+5)*n)+kk]*b[kk*n+jj];
                        }
                    }
                }
            }
        }
    }
```

Using intrinsics with manual data prefetching and register blocking can still considerably increase the performance. Generally speaking, the programmer should try to get a suitable vectorization and write cache and register efficient code, i.e. values stored in registers should be reused as often as possible in order to avoid cache and memory access. The tuning techniques for native implementations discussed in Section 3 also apply for offloaded code, of course. Informations about task pinning and finding the optimal thread number are given in Section 5.1 and Section 10.1.

4.5. Intel Cilk Plus parallel extensions

More complex data structures can be handled by Virtual Shared Memory. In this case the same virtual address space is used on both the host and the coprocessor, enabling a seamless sharing of data. Virtual shared data is specified using the `_Cilk_shared` allocation specifier. This model is integrated in Intel Cilk Plus parallel extensions and is only available in C/C++. There are also Cilk functions to specify offloading of functions and `_Cilk_for` loops. More information on Intel Cilk Plus can be found online under [10] and will be included in a future version of this guide.

5. OpenMP and hybrid

5.1. OpenMP

5.1.1. Programming models and offload

OpenMP parallelization on an Intel Xeon + Xeon Phi coprocessor machine can be applied in four different programming models that can be realized with different compiler options: native OpenMP on the Xeon host; serial Xeon host with OpenMP offload (see Section 4); OpenMP on the Xeon host with OpenMP offload (see Section 4.4.1) and native OpenMP on the Xeon Phi coprocessor.

OpenMP threads on the host CPU and on the Xeon Phi coprocessor do not interfere with each other and when an offload/pragma section of the code is encountered, it is offloaded as a unit and uses a number of threads based on the available resources on the Xeon Phi coprocessor. Within this construct apply usual semantics of shared and private data.

Offload to the Xeon Phi coprocessor can be done at any time by multiple host CPUs until the filling of the available resources. If there are no free threads, the task meant to be offloaded may be done on the host. For offload schemes the maximal amount of threads that can be used on the Xeon Phi coprocessor is 4 times the total number of cores minus one, because one core is reserved for the OS and its services.

5.1.2. Threading and affinity

The most important considerations for OpenMP threading and affinity are the total number of threads that should be utilized and the scheme for binding threads to processor cores.

The Xeon Phi coprocessor supports 4 threads per core. Unlike some CPU-intensive HPC applications that are run on Xeon architecture, which do not benefit from hyperthreading, applications run on Xeon Phi coprocessors do and using more than one thread per core is recommended. When running applications natively on the Xeon Phi coprocessor the full amount of threads can be used.

The default settings are as follows:

	OMP_NUM_THREADS
OpenMP on host without HT	1 x ncore-host
OpenMP on host with HT	2 x ncore-host
OpenMP on Xeon Phi in native mode	4 x ncore-phi
OpenMP on Xeon Phi in offload mode	4 x (ncore-phi - 1)

If OpenMP regions exist on the host and on the part of the code offloaded to the Xeon Phi, two separate OpenMP runtimes exist. Environment variables for controlling OpenMP behavior are to be set for both runtimes, for example the `KMP_AFFINITY` variable which can be used to assign a particular thread to a particular physical node. For Intel Xeon Phi it can be done like this:

```
export MIC_ENV_PREFIX=MIC
```

```
# specify affinity for all cards
export MIC_KMP_AFFINITY=...
#specify number of threads for all cards
export MIC_OMP_NUM_THREADS=120
# specify the number of threads for card #2
export MIC_2_OMP_NUM_THREADS=200
# specify number of threads and affinity for card #3
export MIC_3_ENV="OMP_NUM_THREADS=60 | KMP_AFFINITY=balanced"
```

One can also use special API calls to set the environment for the coprocessor only, e.g.

```
omp_set_num_threads_target()
omp_set_nested_target()
```

More details can be found in Section 10.1.1.

5.1.3. Loop scheduling

OpenMP accepts four different kinds of loop scheduling - static, dynamic, guided and auto. In this way the amount of iterations done by different threads can be controlled. The `schedule` clause can be used to set the loop scheduling at compile time. Another way to control this feature is to specify `schedule(runtime)` in your code and select the loop scheduling at runtime through setting the `OMP_SCHEDULE` environment variable. More details can be found in Section 10.1.11.

5.1.4. Scalability improvement

If the amount of work that should be done by each thread is non-trivial and consists of nested for-loops, one might use the `collapse()` directive to specify how many for-loops are associated with the OpenMP loop construct. This often improves scalability of OpenMP applications (see Section 4.4.3).

Another way to improve scalability is to reduce barrier synchronization overheads by using the `nowait` directive. The effect of it is that the threads will not synchronize after they have completed their individual pieces of work. This approach is applicable combined with static loop scheduling because all threads will execute the same amount of iterations in each loop.

5.2. Hybrid OpenMP/MPI

5.2.1. Programming models

For hybrid OpenMP/MPI programming there are two major approaches: an MPI offload approach, where MPI ranks reside on the host CPU and work is offloaded to the Xeon Phi coprocessor and a symmetric approach in which MPI ranks reside both on the CPU and on the Xeon Phi. An MPI program can be structured using either model.

When assigning MPI ranks, one should take into account that there is a data transfer overhead over the PCIe, so minimizing the communication from and to the Phi is a good idea. Another consideration is that there is limited amount of memory on the coprocessor which favors the shared memory parallelism ideology.

5.2.2. Threading of the MPI ranks

For hybrid OpenMP/MPI applications use the thread safe version of the Intel MPI Library by using the `-mt_mpi` compiler driver option. A desired process pinning scheme can be set with the `I_MPI_PIN_DOMAIN` environment variable. It is recommended to use the following setting:

```
$ export I_MPI_PIN_DOMAIN=omp
```

By using this, one sets the process pinning domain size to be `OMP_NUM_THREADS`. In this way, every MPI process is able to create `$OMP_NUM_THREADS` number of threads that will run within the corresponding domain. If this variable is not set, each process will create a number of threads per MPI process equal to the number of cores, because it will be treated as a separate domain.

Further, to pin OpenMP threads within a particular domain, one could use the `KMP_AFFINITY` environment variable.

6. MPI

Details about using the Intel MPI library on Xeon Phi coprocessor systems can be found in [17].

6.1. Setting up the MPI environment

The following commands have to be executed to set up the MPI environment:

```
# copy MPI libraries and binaries to the card (as root)
# only copying really necessary files saves memory

scp /opt/intel/impi/4.1.0.024/mic/lib/* mic0:/lib
scp /opt/intel/impi/4.1.0.024/mic/bin/* mic0:/bin

# setup Intel compiler variables
. /opt/intel/composerxe/bin/compilervars.sh intel64

# setup Intel MPI variables
. /opt/intel/impi/4.1.0.024/bin64/mpivars.sh
```

The following network fabrics are available for the Intel Xeon Phi coprocessor:

Fabric Name	Description
shm	Shared-memory
tcp	TCP/IP-capable network fabrics, such as Ethernet and InfiniBand (through IPoIB)
ofa	OFA-capable network fabric including InfiniBand (through OFED verbs)
dapl	DAPL-capable network fabrics, such as InfiniBand, iWarp, Dolphin, and XPMEM (through DAPL)

The Intel MPI library tries to automatically use the best available network fabric detected (usually shm for intra-node communication and InfiniBand (dapl, ofa) for inter-node communication).

The default can be changed by setting the `I_MPI_FABRICS` environment variable to `I_MPI_FABRICS=<fabric>` or `I_MPI_FABRICS=<intra-node fabric>:<inter-nodes fabric>`. The availability is checked in the following order: shm:dapl, shm:ofa, shm:tcp.

6.2. MPI programming models

Intel MPI for the Xeon Phi coprocessors offers various MPI programming models:

- Symmetric model** The MPI ranks reside on both the host and the coprocessor. Most general MPI case.
- Coprocessor-only model** All MPI ranks reside only on the coprocessors.

Host-only model

All MPI ranks reside on the host. The coprocessors can be used by using offload pragmas. (Using MPI calls inside offloaded code is not supported.)

6.2.1. Coprocessor-only model

To build and run an application in coprocessor-only mode, the following commands have to be executed:

```
# compile the program for the coprocessor (-mmic)
mpiicc -mmic -o test.MIC test.c

#copy the executable to the coprocessor
scp test.MIC mic0:/tmp

#set the I_MPI_MIC variable
export I_MPI_MIC=1

#launch MPI jobs on the coprocessor mic0 from the host
#(alternatively one can login to the coprocessor and run mpirun there)
mpirun -host mic0 -n 2 /tmp/test.MIC
```

6.2.2. Symmetric model

To build and run an application in symmetric mode, the following commands have to be executed:

```
# compile the program for the coprocessor (-mmic)
mpiicc -mmic -o test.MIC test.c

# compile the program for the host
mpiicc -mmic -o test test.c

#copy the executable to the coprocessor
scp test.MIC mic0:/tmp/test.MIC

#set the I_MPI_MIC variable
export I_MPI_MIC=1

#launch MPI jobs on the host knf1 and on the coprocessor mic0
mpirun -host knf1 -n 1 ./test : -n 1 -host mic0 /tmp/test.MIC
```

6.2.3. Host-only model

To build and run an application in host-only mode, the following commands have to be executed:

```
# compile the program for the host,
# mind that offloading is enabled per default
mpiicc -o test test.c

# launch MPI jobs on the host knf1, the MPI process will offload code
# for acceleration
mpirun -host knf1 -n 1 ./test
```

6.3. Simplifying launching of MPI jobs

Instead of specifying the hosts and coprocessors via `-n hostname` one can also put the names into a hostfile and launch the jobs via

```
mpirun -f hostfile -n 4 ./test
```

Mind that the executable must have the same name on the hosts and the coprocessors in this case.

If one sets

```
export I_MPI_POSTFIX=.mic
```

the `.mix` postfix is automatically added to the executable name by `mpirun`, so in the case of the example above `test` is launched on the host and `test.mic` on the coprocessors. It is also possible to specify a prefix using

```
export I_MPI_PREFIX=./MIC/
```

In this case `./MIC/test` will be launched on the coprocessor. This is specially useful if the host and the coprocessors share the same NFS filesystem.

7. Intel MKL (Math Kernel Library)

The Intel Xeon Phi coprocessor is supported since MKL 11.0. Details on using MKL with Intel Xeon Phi coprocessors can be found in [23], [24] and [25]. Also the MKL developer zone [6] contains useful information. All functions can be used on the Xeon Phi, however the optimization level for wider 512-bit SIMD instructions differs.

As of Intel MKL 11.0 Update 2 the following functions are highly optimized for the Intel Xeon Phi coprocessor:

- BLAS Level 3, and much of Level 1 & 2
- Sparse BLAS: ?CSRMV, ?CSRMM
- Some important LAPACK routines (LU, QR, Cholesky)
- Fast Fourier Transformations
- Vector Math Library
- Random number generators in the Vector Statistical Library

Intel plans to optimize a wider range of functions in future MKL releases.

7.1. MKL usage modes

The following 3 usage models of MKL are available for the Xeon Phi:

1. Automatic Offload
2. Compiler Assisted Offload
3. Native Execution

7.1.1. Automatic Offload (AO)

In the case of automatic offload the user does not have to change the code at all. For automatic offload enabled functions the runtime may automatically download data to the Xeon Phi coprocessor and execute (all or part of) the computations there. The data transfer and the execution management is completely automatic and transparent for the user.

As of Intel MKL 11.0.2 only the following functions are enabled for automatic offload:

- Level-3 BLAS functions
 - ?GEMM (for $m, n > 2048$, $k > 256$)
 - ?TRSM (for $M, N > 3072$)
 - ?TRMM (for $M, N > 3072$)
 - ?SYMM (for $M, N > 2048$)
- LAPACK functions
 - LU ($M, N > 8192$)
 - QR
 - Cholesky

In the above list also the matrix sizes for which MKL decides to offload the computation are given in brackets.

To enable automatic offload either the function `mkl_mic_enable()` has to be called within the source code or the environment variable `MKL_MIC_ENABLE=1` has to be set. If no Xeon Phi coprocessor is detected the application runs on the host without penalty.

To build a program for automatic offload, the same way of building code as on the Xeon host is used:

```
icc -O3 -mkl file.c -o file
```

By default, the MKL library decides when to offload and also tries to determine the optimal work division between the host and the targets (MKL can take advantage of multiple coprocessors). In case of the BLAS routines the user can specify the work division between the host and the coprocessor by calling the routine

```
mkl_mic_set_Workdivision(MKL_TARGET_MIC, 0, 0.5)
```

or by setting the environment variable

```
MKL_MIC_0_WORKDIVISION=0.5
```

Both examples specify to offload 50% of computation only to the 1st card (card #0).

7.1.2. Compiler Assisted Offload (CAO)

In this mode of MKL the offloading is explicitly controlled by compiler pragmas or directives. In contrast to the automatic offload mode, all MKL function can be offloaded in CAO-mode.

A big advantage of this mode is that it allows for data persistence on the device.

For Intel compilers it is possible to use AO and CAO in the same program, however the work division must be explicitly set for AO in this case. Otherwise, all MKL AO calls are executed on the host.

MKL functions are offloaded in the same way as any other offloaded function (see section Section 4). An example for offloading MKL's `sgemm` routine looks as follows:

```
#pragma offload target(mic) \  
  in(transa, transb, N, alpha, beta) \  
  in(A:length(N*N)) \  
  out(C:length(N*N))
```

```
in(B:length(N*N)) \  
in(C:length(N*N)) \  
out(C:length(N*N) alloc_if(0)) {  
  
    sgemm(&transa, &transb, &N, &N, &N, &alpha, A, &N, B, &N, &beta, C, &N);  
  
}
```

To build a program for compiler assisted offload, the following command is recommended by Intel:

```
icc -O3 -openmp -mkl \  
-offload-option,mic,ld, "-L$MKLROOT/lib/mic -Wl,\  
--start-group -lmkl_intel_lp64 -lmkl_intel_thread \  
-lmkl_core -Wl,--end-group" file.c -o file
```

To avoid using the OS core, it is recommended to use the following environment setting (in case of a 61-core coprocessor):

```
MIC_KMP_AFFINITY=explicit,granularity=fine,proclist=[1-240:1]
```

Setting larger pages by the environment setting `MIC_USE_2MB_BUFFERS=16K` usually increases performance. It is also recommended to exploit data persistence with CAO.

7.1.3. Native Execution

In this mode of MKL the Intel Xeon Phi coprocessor is used as an independent compute node.

To build a program for native mode, the following compiler settings should be used:

```
icc -O3 -mkl -mmic file.c -o file
```

The binary must then be manually copied to the coprocessor via ssh and directly started on the coprocessor.

7.2. Example code

Example code can be found under `$MKLROOT/examples/mic_ao` and `$MKLROOT/examples/mic_offload`.

7.3. Intel Math Kernel Library Link Line Advisor

To determine the appropriate link line for MKL the Intel Math Kernel Library Link Line Advisor available under [7] has been extended to include support for the Intel Xeon Phi specific options.

8. TBB: Intel Threading Building Blocks

The Intel TBB library is a template based runtime library for C++ code using threads that allows us to fully utilize the scaling capabilities within our code by increasing the number of threads and supporting task oriented load balancing.

Intel TBB is open source and available on many different platforms with most operating systems and processors. It is already popular in the C++ community. You should be able to use it with any compiler supporting ISO C++. So this is one of the advantages of Intel TBB when you intend to keep your code as easily portable as possible.

Typically as a rule of thumb an application must scale well past one hundred threads on Intel Xeon processors to profit from the possible higher parallel performance offered with e.g. the Intel Xeon Phi coprocessor. To check if

the scaling would profit from utilising the highly parallel capabilities of the MIC architecture, you should start to create a simple performance graph with a varying number of threads (from one up to the number of cores).

From a programming standpoint we treat the coprocessor as a 64-bit x86 SMP-on-a-chip with an high-speed bi-directional ring interconnect, (up to) four hardware threads per core and 512-bit SIMD instructions. With the available number of cores we have easily 200 hardware threads at hand on a single coprocessor. The multi-threading on each core is primarily used to hide latencies that come implicitly with an in-order microarchitecture. Unlike hyper-threading these hardware threads cannot be switched off and should never be ignored. Generally it should be impossible for a single thread per core to approach the memory or floating point capability limit. Highly tuned codesnippets may reach saturation already at two threads, but in general a minimum of three or four active threads per cores will be needed. This is one of the reasons why the number of threads per core should be parameterized as well as the number of cores. The other reason is of course to be future compatible.

TBB offers programming methods that support creating this many threads in a program. In the easiest way the one main production loop is transformed by adding a single directive or pragma enabling the code for many threads. The chunk size used is chosen automatically.

The new Intel Cilk Plus which offers support for a simpler set of tasking capabilities fully interoperates with Intel TBB. Apart from that Intel Cilk Plus also supports vectorization. So shared memory programmers have Intel TBB and Intel Cilk Plus to assist them with built-in tasking models. Intel Cilk Plus extends Intel TBB to offer C programmers a solution as well as help with vectorization in C and C++ programs.

Intel TBB itself does not offer any explicit vectorization support. However it does not interfere with any vectorization solution either.

In relevance to the Intel Xeon Phi coprocessor TBB is just one available runtime-based parallel programming model alongside OpenMP, Intel Cilk Plus and pthreads that are also already available on the host system. Any code running natively on the coprocessor can put them to use just like it would on the host with the only difference being the larger number of threads.

8.1. Advantages of TBB

There exists a variety of approaches to parallel programming, but there are several advantages to using Intel TBB when writing scalable applications:

- TBB relies on generic programming: Writing the best possible algorithms with the fewest possible constraints enables to deliver high performance algorithms which can be applied in a broader context. Other more traditional libraries specify interfaces in terms of particular types or base classes. Intel TBB specifies the requirements on the types instead and in this way keeps the algorithms themselves generic and easily adaptable to different data representations.
- It is easy to start: You don't have to be a threading expert to leverage multi-core performance with the help of TBB. Normally you can successfully thread some programs just by adding a single directive or pragma to the main production loop.
- It obeys to logical parallelism: Since with TBB you specify tasks instead of threads, you automatically produce more portable code which emphasizes scalable, data parallel programming. You are not bound to platform-dependent threading primitives; most threading packages require you to directly code on low-level constructs close to the hardware. Direct programming on raw thread level is error-prone, tedious and typically hard work since it forces you to efficiently map logical tasks into threads and it is not always leading to the desired results. With the higher level of data-parallel programming on the other hand, where you have multiple threads working on different parts of a collection, performance continues to increase as you add more cores since for a larger number of processors the collections are just divided into smaller chunks. This is a great feature when it comes to portability.
- TBB is compatible with other programming models: Since the library is not designed to address all kinds of threading problems, it can coexist seamlessly with other threading packages.
- The template-based approach allows Intel TBB to make no excuses for performance. Other general-purpose threading packages tend to be low-level tools that are still far from the actual solution, while at the same time

supporting many different kinds of threading. In TBB every template solves a computationally intensive problem in a generic, simple way instead.

All of these advantages make TBB popular and easily portable while at the same time facilitating data parallel programming.

Further advanced concepts in TBB that are not MIC specific can be found in the Intel TBB User Guide or in the Reference Manual, both available under [12].

8.2. Using TBB natively

Code that runs natively on the Intel Xeon Phi coprocessor can apply the TBB parallel programming model just as they would on the host, with no unusual complications beyond the larger number of threads.

In order to initialize your compiler environment variables needed to set up TBB correctly, typically the `/opt/intel/composerxe/tbb/bin/tbbvars.csh` or `tbbvars.sh` script with `intel64` as the argument is called by the `/opt/intel/composerxe/bin/compilervars.csh` or `compilervars.sh` script with `intel64` as argument. (e.g. `source /opt/intel/composerxe/bin/compilervars.sh intel64`)

Normally there is no need to call the `tbbvars` script directly and it is not advisable either since the `compilervars` script also calls other subscripts taking i.e. care of the debugger or Intel MKL and running the subscripts out of order might result in unpredictable behavior.

A minimal C++ TBB example looks as follows:

```
#include "tbb/task_scheduler_init.h"
#include "tbb/parallel_for.h"
#include "tbb/blocked_range.h"

using namespace tbb;

int main() {

    task_scheduler_init init;

    return 0;
}
```

The `using` directive imports the namespace `tbb` where all of the library's classes and functions are found. The namespace is explicit in the first mention of a component, but implicit afterwards. So with the `using namespace` statement present you can use the library component identifiers without having to write out the namespace prefix `tbb` before each of them.

The task scheduler is initialized by instantiating a `task_scheduler_init` object in the main function. The definition for the `task_scheduler_init` class is included from the corresponding header file. Actually any thread using one of the provided TBB template algorithms must have such an initialized `task_scheduler_init` object. The default constructor for the `task_scheduler_init` object informs the task scheduler that the thread is participating in task execution, and the destructor informs the scheduler that the thread no longer needs the scheduler. With the newer versions of Intel TBB as used in a MIC environment the task scheduler is automatically initialized, so there is no need to explicitly initialize it if you don't need to have control over when the task scheduler is constructed or destroyed. When initializing it you also have the further possibility to tell the task scheduler explicitly how many worker threads there are to be used and what their stack size would be.

In the simplest form scalable parallelism can be achieved by parallelizing a loop of iterations that can each run independently from each other.

The `parallel_for` template function replaces a serial loop where it is safe to process each element concurrently.

A typical example would be to apply a function `Foo` on all elements of an array over the iterations space of type `size_t` going from 0 to `n-1`:

```
void SerialApplyFoo( float a[], size_t n ) {
for( size_t i=0; i!=n; ++i )
Foo(a[i]);
}
```

becomes

```
void ParallelApplyFoo( float a[], size_t n) {
parallel_for(size_t(0), n, [=](size_t i) {Foo(a[i]);});
}
```

This is the TBB short form of a `parallel_for` over a loop based on a one-dimensional iteration space consisting of a consecutive range of integers (which is one of the most common cases). The expression `parallel_for(first,last,step,f)` is synonymous to `for(auto i=first; i!=last; i +=step) f(i)` except that each `f(i)` can be evaluated in parallel if resources permit. The omitted step parameter is optional. The short form implicitly uses automatic chunking.

The long form would be:

```
void ParallelApplyFoo( float* a, size_t n ) {
parallel_for( blocked_range<size_t>(0,n),
[=](const blocked_range<size_t>& r) {
for(size_t i=r.begin(); i!=r.end(); ++i)
Foo(a[i]);
});
}
```

Here the key feature of the TBB library is more clearly revealed. The template function `tbb::parallel_for` breaks the iteration space into chunks, and runs each chunk on a separate thread. The first parameter of template function call `parallel_for` is a `blocked_range` object that describes the entire iteration space from 0 to `n-1`. The `parallel_for` divides the iteration space into subspaces for each of the over 200 hardware threads. `blocked_range<T>` is a template class provided by the TBB library describing a one-dimensional iteration space over type `T`. The `parallel_for` class works just as well with other kinds of iteration spaces. The library provides `blocked_range2d` for two-dimensional spaces. There exists also the possibility to define own spaces. The general constructor of the `blocked_range` template class is `blocked_range<T>(begin,end,grainsize)`. The `T` specifies the value type. `begin` represents the lower bound of the half-open range interval `[begin,end)` representing the iteration space. `end` represents the excluded upper bound of this range. The `grainsize` is the approximate number of elements per sub-range. The default `grainsize` is 1.

A parallel loop construct introduces overhead cost for every chunk of work that it schedules. The MIC adapted Intel TBB library chooses chunk sizes automatically, depending upon load balancing needs. The heuristic normally works well with the default `grainsize`. It attempts to limit overhead cost while still providing ample opportunities for load balancing. For most use cases automatic chunking is the recommended choice. There might be situations though where controlling the chunk size more precisely might yield better performance.

When compiling programs that employ TBB constructs, be sure to link in the Intel TBB shared library with `-ltbb`. If you don't undefined references will occur.

```
icc -mmic -ltbb foo.cpp
```

Afterwards you can use `scp` to upload the binary and any shared libraries required by your application to the coprocessor. On the coprocessor you can then export the library path and run the application.

8.3. Offloading TBB

The Intel TBB header files are not available on the Intel MIC target environment by default (the same is also true for Intel Cilk Plus). To make them available on the coprocessor the header files have to be wrapped with `#pragma offload` directives as demonstrated in the example below:

```
#pragma offload_attribute (push,target(mic))
#include "tbb/task_scheduler_init.h"
#include "tbb/parallel_for.h"
#include "tbb/blocked_range.h"
#pragma offload_attribute (pop)
```

Functions called from within the offloaded construct and global data required on the Intel Xeon Phi coprocessor should be appended by the special function attribute `__attribute__((target(mic)))`.

Codes using Intel TBB with an offload should be compiled with `-tbb` flag instead of `-ltbb`.

9. Debugging

Information about debugging on Intel Xeon Phi coprocessors can be found in [18].

The GNU debugger (gdb) has been enabled by Intel to support the Intel Xeon Phi coprocessor. The debugger is now part of the recent MPSS release and does not have to be downloaded separately any more.

There are 2 different modes of debugging supported: native debugging on the coprocessor or remote cross-debugging on the host.

9.1. Native debugging with gdb

- Run gdb on the coprocessor

```
ssh -t mic0 /usr/bin/gdb
```

- One can then attach to a running application with process ID pid via

```
(gdb) attach pid
```

- or alternatively start an application from within gdb via

```
(gdb) file /path/to/application
(gdb) start
```

9.2. Remote debugging with gdb

- Run the special gdb version with Xeon Phi support on the host

```
/usr/linux-k1om-4.7/bin/x86_64-k1om-linux-gdb
```

- Start the gdbserver on the coprocessor by typing on the host gdb

```
(gdb) target extended-remote| ssh -T mic0 gdbserver -multi -
```

- Attach to a remotely running application with the remote process ID pid

```
(gdb) file /local/path/to/application
(gdb) attach pid
```

- It is also possible to run an application directly from the host gdb

```
(gdb) file /local/path/to/application  
(gdb) set remote exec-file /remote/path/to/application
```

10. Tuning

Information on performance tuning from Intel can be found in [28] and [29].

A single Xeon Phi core is slower than a Xeon core due to lower clock frequency, smaller caches and lack of sophisticated features such as out-of-order execution and branch prediction. To fully exploit the processing power of a Xeon Phi, parallelism is needed. As with any modern CPU in an HPC system, with Xeon Phi parallelism exists on three different levels: simd, thread and node.

In the following sections, we focus on the different aspects of parallelism one by one. After a selection of a suitably parallel algorithm, the first step is to maximize the performance of a code on a single core. This should be followed by parallelization with threads in shared memory. Final step is then to apply distributed memory parallelism over several compute nodes.

Performance figures will be included in a future version of this guide (currently NDA-restricted).

10.1. Advanced OpenMP

The easiest and (arguably) the most productive way to exploit threading parallelism with an Intel Xeon Phi is to use OpenMP. The threading constructs are equivalent for both offload and native models. We expect the basic concepts and syntax of OpenMP to be known. For OpenMP, see for instance the book [3], the OpenMP forum [11] and references therein.

The high level of parallelism available on an Intel Xeon Phi available is very likely to reveal any performance problems related to threading previously been unnoticed in the code. In the following, we introduce a few of the most common OpenMP performance problems and suggest some ways to correct them. We begin by considering thread to core affinity.

10.1.1. OpenMP thread affinity

On a modern HPC system, each node contains a shared-memory environment with several processor sockets, each socket containing several physical cores which, in turn, are divided into several logical cores. We refer to this as node topology.

Each memory bank usually resides closer to some of the cores in the topology and therefore access to data laying in a memory bank attached to another socket is generally more expensive. Such a non-uniform memory access (NUMA) can create performance issues if threads are allowed to migrate from one logical core to another during their execution.

In order to extract maximum performance, consider binding OpenMP threads to logical and physical cores across different sockets on one compute node. The layout of this binding in respect to the node topology has performance implications depending on the computational task and is referred as thread affinity.

We now briefly show how to set thread affinity using Intel compilers and OpenMP-library. For a complete description, see "Thread affinity interface" in the Intel compiler manual.

The thread affinity interface of the Intel runtime library can be controlled by using the `KMP_AFFINITY` environment variable or by using a proprietary Intel API. We now focus on the former and note that a standardized method for setting affinity is expected to be available with OpenMP 4.0 during early 2013.

```
KMP_AFFINITY=[modifier,...]<type>[,<permute>][,<offset>]]
```

modifier	default =noverbose, respect, granularity=core
	granularity=<{fine, thread, core}>, norespect, noverbose, nowarnings,
	proclist={<proc-list>}, respect, verbose, warnings.

type	default =none balanced, compact, disabled, explicit, none, scatter
permute	default =0. >=0, integer. Not valid with type=explicit,none,disabled.
offset	default =0. >=0, integer. Not valid with type=explicit,none,disabled

In most cases it is sufficient only to specify the affinity and granularity. The most important affinity types supported by Intel Xeon Phi are

balanced	Thread affinity balanced is a mixture of scatter and compact affinities. Threads from <1> to <n_p> will be spread across the topology as evenly as possible in the granularity context, where <n_p> denotes the number of physical cores. For thread <k> from threads <n_p+1> to <n> will be assigned as close as possible to thread <k+1>.
compact	Thread <k+1> will be assigned as close as possible to thread <k> in the granularity context according to which the threads are placed.
none	Threads are not bound to any contexts. Use of affinity none is not recommended in general.
scatter	Threads from <1> to <n> will be spread across the topology as evenly as possible in the granularity context according to which the threads are placed.

The most important granularity types supported by Intel Xeon Phi are

core	Threads are bound to a single core, but allowed to float within the context of a physical core.
fine/thread	Threads are bound to a single context, i.e., a logical core.

10.1.2. Example: Thread affinity

We now consider the effect of thread affinity to matrix-matrix multiply. Let $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times n}$ and $C=AB \in \mathbb{R}^{n \times n}$ with $n=1000$. We implement the operation in Fortran90 without blocking by using jki loop-ordering and OpenMP as follows

```
!$OMP PARALLEL DO DEFAULT(NONE) &
!$OMP SHARED(A,B,C) &
!$OMP PRIVATE(i,j,k) &
DO j=1,n
  DO k=1,n
    DO i=1,n
      C(i,j)=C(i,j)+A(i,k)*B(k,j)
    END DO
  END DO
END DO
!$OMP END PARALLEL DO
```

10.1.3. Multiple parallel regions and barriers

Whenever an OpenMP parallel region is encountered, a team of threads is formed and launched to execute the computations. Whenever the parallel region is ended, threads are joined and the computation proceeds with a

single thread. Between different parallel regions it is up to the OpenMP implementation to decide whether the threads are shut down or left in an idle state.

Intel OpenMP -library leaves the threads in a running state for a predefined amount of time before setting them to sleep. The time is defined by `KMP_BLOCKTIME` and `KMP_LIBRARY` environment variables. The default is 200ms. For more details, see Sections "Intel Environment Variables Extensions" and "Execution modes" in the Intel compiler manual.

Repeatedly forming and disbanding thread-teams and setting idle threads to sleep has some overhead associated with it. Another common source of threading overhead in OpenMP computations are implicit or explicit barriers. Recall that many OpenMP constructs have an implicit barrier attached to the end of the construct. Then, especially if the amount of work done inside an OpenMP construct is relatively small, thread synchronization with several threads may be a source of significant overhead. If the computations are independent, the implicit barrier at the end of OpenMP constructs can be removed with the optional `NOWAIT` parameter.

10.1.4. Example: Multiple parallel regions and barriers

We now consider the effect of multiple parallel regions and barriers to performance. Let v in \mathbb{R}^{set^n} , with $n=1000$. Let $f(x)$ denote a function, defined as $f(x)=x+1$.

We implement an OpenMP loop to apply $f(x)$ successively to a given vector v several times and count the number of applications of $f(x)$ explicitly. We consider three different implementations. In the first one, OpenMP parallel region is re-initialized for each successive application of $f(x)$. The second one initializes the parallel region once, but contains two implicit barriers from OpenMP constructs. In the third implementation the parallel region is initialized once and one barrier is used to synchronize the repetitions.

Implementation 1: parallel region re-initialized repeatedly.

```
DO rep=1, repeats
  !$OMP PARALLEL DO DEFAULT(NONE) NUM_THREADS(threads) &
  !$OMP SHARED(vec1, rep, n) &
  !$OMP PRIVATE(i)
  DO i=1, n
    vec1(i)=vec1(i)+1D0
  END DO
  !$OMP END PARALLEL DO

  ops = ops + 1
END DO
```

Implementation 2: parallel region initialized once, two implicit barriers from OpenMP constructs.

```
!$OMP PARALLEL DEFAULT(NONE) NUM_THREADS(threads) &
!$OMP SHARED(vec1, repeats, ops, n) &
!$OMP PRIVATE(i, rep)
DO rep=1, repeats
  !$OMP DO
  DO i=1, n
    vec1(i)=vec1(i)+1D0
  END DO
  !$OMP END DO

  !$OMP SINGLE
  ops = ops + 1
  !$OMP END SINGLE
END DO
```

```
!$OMP END PARALLEL
```

Implementation 3: parallel region initialized once, one explicit barrier from OpenMP construct.

```
!$OMP PARALLEL DEFAULT(NONE) NUM_THREADS(threads) &  
!$OMP SHARED(vec1, repeats, ops, n) &  
!$OMP PRIVATE(i, rep)  
DO rep=1,repeats  
  !$OMP DO  
  DO i=1,n  
    vec1(i)=vec1(i)+1D0  
  END DO  
  !$OMP END DO NOWAIT  
  
  !$OMP SINGLE  
  ops = ops + 1  
  !$OMP END SINGLE NOWAIT  
  
  ! Synchronize threads once per round  
  !$OMP BARRIER  
  END DO  
!$OMP END PARALLEL
```

10.1.5. False sharing

On a multiprocessor shared-memory system, each core has some local cache, which must be kept coherent the among the cores in the system. Processor cache is organized into several cache lines, each of which map to some part of the main memory. On an Intel Xeon Phi, cache line size is 64 bytes. For reference, see Intel Xeon Phi system software developers guide [15].

If more than one core accesses the same data in the main memory, a cache line is shared. Whenever a shared cache line is updated, to maintain coherency an update is forced to the caches of all the cores accessing the cache line.

False sharing occurs when several cores access and update different variables which happen to reside on a single shared cache line. The resulting updates to maintain cache coherency may cause a significant performance degradation. We note that the processors may not be actually sharing any data, it is sufficient that the data resides on a same cache line.

Given a code with performance problems, false sharing may be hard to localize. Tools such as Intel VTune Performance Analyzer can be extremely useful in pointing out the places where false sharing takes place. When writing code, false sharing can be avoided by carefully considering write access to shared variables. If a variable is updated often, it may be worthwhile to use a private variable in stead of a shared one and do a reduction at the end of the work sharing loop.

10.1.6. Example: False sharing

We now consider a simple example where false sharing occurs. Let $v \in \mathbb{R}^n$, with $n=1E+08$. Let $f(x)$ denote a function which counts the number of entries v_j of v for which $v_j < 0$ holds.

We implement $f(x)$ with OpenMP in two different ways. In the first implementation, each thread counts the number of negative entries it has found in v to a globally shared array. To avoid race conditions, each thread uses its own entry in the shared array, uniquely determined by thread id. When a thread has finished its portion of vector, a global counter is atomically incremented. The second implementation is practically equivalent to the first one, except that each thread has its own private array for counting the data.

Implementation 1: False sharing of array counter.

```
!$OMP PARALLEL DEFAULT(NONE) NUM_THREADS(threads) &
!$OMP SHARED(vec, count, counter, n) &
!$OMP PRIVATE(i,TID)

TID=1
!$ TID=omp_get_thread_num()+1

!$OMP DO
DO i=1,n
  IF (vec(i)<0) counter(TID)=counter(TID)+1
END DO
!$OMP END DO

!OMP ATOMIC
count = counter(TID)+count

!$OMP END PARALLEL
```

Implementation 2: Private array used to avoid false sharing.

```
!$OMP PARALLEL DEFAULT(NONE) NUM_THREADS(threads) &
!$OMP SHARED(vec, count, n) &
!$OMP PRIVATE(i, counter, TID)

TID=1
!$ TID=omp_get_thread_num()+1

!$OMP DO
DO i=1,n
  IF (vec(i)<0) counter(TID)=counter(TID)+1
END DO
!$OMP END DO

!OMP ATOMIC
count = counter(TID)+count

!$OMP END PARALLEL
```

We note that a better implementation for this particular problem will be given in the next section.

10.1.7. Memory limitations

Available memory per core on an Intel Xeon Phi is very limited. When an application is run using all the available threads, approximately 30Mb of memory is available per core in the case when all memory references are private. Excessive memory allocation per thread is therefore highly discouraged. Care should be also taken when assigning private variables in order to avoid unnecessary data duplication among threads.

10.1.8. Example: Memory limitations

We now return to the example given in the previous section. In the example, we prevented threads from doing false sharing by modifying the definition of the vector containing the counters. What is important to note is that in doing so, each thread now implicitly allocates a vector of length nthreads, i.e., the memory consumption is quadratic. A better alternative is to let each thread to store the local result in a temporary variable and use a reduction to count the number of elements smaller than zero.

Implementation 3: Temporary variable with reduction used to store local results.

```

count = 0
!$OMP PARALLEL DEFAULT(NONE) NUM_THREADS(threads) &
!$OMP SHARED(vec, n) &
!$OMP PRIVATE(i, Lcount, TID) &
!$OMP REDUCTION(+:count)

TID=1
!$ TID=omp_get_thread_num()+1

Lcount = 0
!$OMP DO
DO i=1,n
  IF (vec(i)<0) Lcount=Lcount+1
END DO
!$OMP END DO

count = Lcount+count
!$OMP END PARALLEL

```

10.1.9. Nested parallelism

Due to the limited amount of memory available, sometimes using all available threads on an Intel Xeon Phi to parallelize the outer loop of some computation is not possible. In some cases this may not be due to inefficient structure of the code, but because the data needed for computations per thread is too large. In this case, to take advantage of all the processing power of the coprocessor, an option is to use nested OpenMP parallelism.

When nested parallelism is enabled, any inner OpenMP parallel region which is enclosed within an outer parallel region will be executed with multiple threads. The performance impact of using nested parallelism is similar to performance impact of using multiple parallel regions and barriers.

Enabling OpenMP nested parallelism is done by setting environment variable `OMP_NESTED=TRUE` or with an API call to `omp_set_nested` -function. The number of nested threads within each OpenMP parallel region is done by setting the environment variable `OMP_NUM_THREADS=n_1, n_2, n_3, . . .`, where n_j refers to the number of threads on the j th level. The number or threads within each nesting level can be also set with an API call to `omp_set_num_threads` -function normally.

10.1.10. Example: Nested parallelism

Consider a case where several independent matrix-matrix multiplications have to be computed. Let $A \in \mathbb{R}^{m \times n}$ and let matrix $B_k \in \mathbb{R}^{n \times n}$ be defined as $B_k = A^T A$.

We let $m=200$, $n=80$ and $k=2000$ and study the effect of parallelizing the computation of B_k 's in three different ways. The first case is to use parallelize the loop over k with all available threads. A second case is to parallelize the computation over different k 's to physical Intel Xeon Phi cores and use nested parallelism with hardware threads in the computation of the matrix-matrix multiplications. The third case, which uses parallelization only over physical cores, is included for comparison. For this, we have the following implementation.

Implementation: possibly nested parallel loop for computing $A^T A$.

```

!$OMP PARALLEL DO DEFAULT(NONE) NUM_THREADS(nthreads) &
!$OMP SCHEDULE(STATIC) &
!$OMP SHARED(A, B, m, n, k) PRIVATE(elem)
DO i=1,k
  CALL DGEMM('T','N', n, n, m, 1D0, A, m, A, m, 0D0, B(1,1,k), n)

```

```
END DO
!$OMP END PARALLEL DO
```

10.1.11. Load balancing

In parallel processing, some processes may require more resources and be more time consuming than others. This results in a load imbalance between different processes or threads executing the computation.

In a regular distributed memory program, load balancing is generally tedious and usually requires some parts of the computation to be redistributed among the processors. In a shared memory program with a runtime, such as OpenMP, load balancing can be in some cases automatically handled by the runtime itself with little overhead.

OpenMP loop constructs support an additional SCHEDULE-clause. The syntax for this is

SCHEDULE(<kind>[,chunk_size]),

kind	default =STATIC. <STATIC,DYNAMIC,GUIDED,RUNTIME>
chunk_size	default =value depends on the schedule kind. >0, integer.

Different schedule kinds supported by OpenMP runtime on an Intel Xeon Phi are

STATIC	With the static scheduling policy, iteration indices are divided into chunks of chunk_size and distributed to threads in a round-robin fashion. If chunk_size is not defined, iteration indices are divided into chunks that are roughly equal in size.
DYNAMIC	With the dynamic scheduling policy, iteration indices are assigned to threads in chunks of chunk_size. Threads request and process new chunks with assigned iteration indices until the whole index range has been processed.
GUIDED	With the guided scheduling policy, iteration indices are assigned to threads in chunks of size chunk_size at minimum. In the beginning of the iteration, the chunk_size actually assigned to be processed is proportional to the number of unassigned iteration indices versus the number of available threads. The assigned chunk_size and can be larger than the minimum.
RUNTIME	The scheduling policy and chunk size will be decided at runtime based on the OMP_SCHEDULE environment variable.

10.1.12. Example: Load balancing

We now consider a case where OpenMP load balancing is beneficial. Let $v \setminus Rset^{1E+6}$ and a function $f(v)$, defined as $f(v_j)=v_j+1$, j times.

If parallelized with OpenMP over vector index j , the function f is load-imbalanced by construction. There are $(n+1)n/2$ operations in total, but the operations are distributed in a linearly increasing fashion with the index j . We implement the function in Fortran as follows.

Implementation: Load balancing with OpenMP scheduling policies.

```
!$OMP PARALLEL DO DEFAULT(NONE) NUM_THREADS(nthreads) &
!$OMP SCHEDULE(<STATIC|DYNAMIC|GUIDED>) &
!$OMP SHARED(vec, n) PRIVATE(elem)
```

```
DO i=1,n
  elem = 0
  DO j=1,i
    elem = elem + 1D0
  END DO
  vec(i)=elem
END DO
!$OMP END PARALLEL DO
```

11. Performance analysis tools

The following performance analysis tools have been enabled for the Intel Xeon Phi coprocessor:

- Intel trace analyzer and collector (ITAC)
- Intel VTune Amplifier XE

More information on performance analysis can be found in [27] [30]. Details will be included in a future version of this guide.

Further documentation

Books

- [1] *James Reinders, James Jeffers, Intel Xeon Phi Coprocessor High Performance Programming, Morgan Kaufman Publ Inc, 2013* <http://lotsofcores.com> [<http://lotsofcores.com>].
- [2] *Michael McCool, James Reinders, Arch Robison, Structured Parallel Programming: Patterns for Efficient Computation, Morgan Kaufman Publ Inc, 2013* <http://parallelbook.com> [<http://parallelbook.com>].
- [3] *Barbara Chapman, Gabriele Jost and Ruud van der Pas, Using OpenMP, MIT Press Cambridge, 2007*, <http://mitpress.mit.edu/books/using-openmp> [<http://mitpress.mit.edu/books/using-openmp>].

Forums, Download Sites, Webinars

- [4] *Intel Developer Zone: Intel Xeon Phi Coprocessor*, <http://software.intel.com/en-us/mic-developer> [<http://software.intel.com/en-us/mic-developer>].
- [5] *Intel Many Integrated Core Architecture User Forum*, <http://software.intel.com/en-us/forums/intel-many-integrated-core> [<http://software.intel.com/en-us/forums/intel-many-integrated-core>].
- [6] *Intel Developer Zone: Intel Math Kernel Library*, <http://software.intel.com/en-us/forums/intel-math-kernel-library> [<http://software.intel.com/en-us/forums/intel-math-kernel-library>].
- [7] *Intel Math Kernel Library Link Line Advisor*, <http://software.intel.com/sites/products/mkl/> [<http://software.intel.com/sites/products/mkl/>].
- [8] *Intel Manycore Platform Software Stack (MPSS)*, <http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss> [<http://software.intel.com/en-us/articles/intel-manycore-platform-software-stack-mpss>].
- [9] *Intel Xeon Processors & Intel Xeon Phi Coprocessors – Introduction to High Performance Applications Development for Multicore and Manycore – Live Webinar, 26.-27.2.2013, recorded* <http://>

software.intel.com/en-us/articles/intel-xeon-phi-training-m-core [<http://software.intel.com/en-us/articles/intel-xeon-phi-training-m-core>].

[10] *Intel Cilk Plus Home Page*, <http://cilkplus.org/> [<http://cilkplus.org/>].

[11] *OpenMP forum*, <http://openmp.org/wp/> [<http://openmp.org/wp/>].

[12] *Intel Threading Building Blocks Documentation Site*, <http://threadingbuildingblocks.org/documentation> [<http://threadingbuildingblocks.org/documentation>].

Manuals, Papers

[13] *Intel Xeon Phi Coprocessor Developer's Quick Start Guide*, <http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-developers-quick-start-guide> [<http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-developers-quick-start-guide>].

[14] *Intel Xeon Phi Coprocessor (codename Knights Corner)*, <http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-codename-knights-corner> [<http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-codename-knights-corner>].

[15] *Intel Xeon Phi Coprocessor System Software Developers Guide*, <https://secure-software.intel.com/sites/default/files/article/334766/intel-xeon-phi-systemsoftwaredevelopersguide.pdf> [<https://secure-software.intel.com/sites/default/files/article/334766/intel-xeon-phi-systemsoftwaredevelopersguide.pdf>].

[16] *System Administration for the Intel Xeon Phi Coprocessor*, <http://software.intel.com/sites/default/files/article/373934/system-administration-for-the-intel-xeon-phi-coprocessor.pdf> [<http://software.intel.com/sites/default/files/article/373934/system-administration-for-the-intel-xeon-phi-coprocessor.pdf>].

[17] *Using the Intel MPI Library on Intel Xeon Phi Coprocessor Systems*, <http://software.intel.com/en-us/articles/using-the-intel-mpi-library-on-intel-xeon-phi-coprocessor-systems> [<http://software.intel.com/en-us/articles/using-the-intel-mpi-library-on-intel-xeon-phi-coprocessor-systems>].

[18] *Debugging on Intel Xeon Phi Coprocessor Use Case Overview*, <http://software.intel.com/en-us/articles/debugging-on-intel-xeon-phi-coprocessor-use-case-overview> [<http://software.intel.com/en-us/articles/debugging-on-intel-xeon-phi-coprocessor-use-case-overview>].

[19] *Intel Xeon Phi Coprocessor Instruction Set Reference Manual*, <http://software.intel.com/sites/default/files/forum/278102/327364001en.pdf> [<http://software.intel.com/sites/default/files/forum/278102/327364001en.pdf>].

[20] *An Overview of Programming for Intel Xeon processors and Intel Xeon Phi coprocessors*, http://software.intel.com/sites/default/files/article/330164/an-overview-of-programming-for-intel-xeon-processors-and-intel-xeon-phi-coprocessors_1.pdf [http://software.intel.com/sites/default/files/article/330164/an-overview-of-programming-for-intel-xeon-processors-and-intel-xeon-phi-coprocessors_1.pdf].

[21] *Programming and Compiling for Intel Many Integrated Core Architecture*, <http://software.intel.com/en-us/articles/programming-and-compiling-for-intel-many-integrated-core-architecture> [<http://software.intel.com/en-us/articles/programming-and-compiling-for-intel-many-integrated-core-architecture>].

[22] *Building a Native Application for Intel Xeon Phi Coprocessors*, <http://software.intel.com/en-us/articles/building-a-native-application-for-intel-xeon-phi-coprocessors> [<http://software.intel.com/en-us/articles/building-a-native-application-for-intel-xeon-phi-coprocessors>].

[23] “Using Intel Math Kernel Library on Intel Xeon Phi Coprocessors” section in the *MKL User's Guide*, http://software.intel.com/sites/products/documentation/doclib/mkl_sa/11/mkl_userguide_lnx/index.htm [http://software.intel.com/sites/products/documentation/doclib/mkl_sa/11/mkl_userguide_lnx/index.htm].

- [24] “Support Functions for Intel Many Integrated Core Architecture” section in the MKL Reference Manual, http://software.intel.com/sites/products/documentation/doclib/mkl_sa/11/mklman/index.htm [http://software.intel.com/sites/products/documentation/doclib/mkl_sa/11/mklman/index.htm].
- [25] Intel Math Kernel Library on the Intel Xeon Phi Coprocessor, <http://software.intel.com/en-us/articles/intel-mkl-on-the-intel-xeon-phi-coprocessors> [<http://software.intel.com/en-us/articles/intel-mkl-on-the-intel-xeon-phi-coprocessors>].
- [26] Intel Compiler 13.0 User Guide and Reference Manual, <http://software.intel.com/sites/products/documentation/doclib/stdxe/2013/composerxe/compiler/cpp-lin/index.htm> [<http://software.intel.com/sites/products/documentation/doclib/stdxe/2013/composerxe/compiler/cpp-lin/index.htm>].
- [27] "Using Intel Trace Analyzer and Collector for Intel Many Integrated Core Architecture" in Intel Cluster Studio 2013 Tutorial, http://software.intel.com/sites/products/documentation/hpc/ics/ics2013/ics_tutorial/index.htm#Linux_itac.htm [http://software.intel.com/sites/products/documentation/hpc/ics/ics2013/ics_tutorial/index.htm#Linux_itac.htm].
- [28] Advanced Optimizations for Intel MIC Architecture, <http://software.intel.com/en-us/articles/advanced-optimizations-for-intel-mic-architecture> [<http://software.intel.com/en-us/articles/advanced-optimizations-for-intel-mic-architecture>].
- [29] Optimization and Performance Tuning for Intel Xeon Phi Coprocessors - Part 1: Optimization Essentials, <http://software.intel.com/en-us/articles/optimization-and-performance-tuning-for-intel-xeon-phi-coprocessors-part-1-optimization> [<http://software.intel.com/en-us/articles/optimization-and-performance-tuning-for-intel-xeon-phi-coprocessors-part-1-optimization>].
- [30] Optimization and Performance Tuning for Intel Xeon Phi Coprocessors, Part 2: Understanding and Using Hardware Events, <http://software.intel.com/en-us/articles/optimization-and-performance-tuning-for-intel-xeon-phi-coprocessors-part-2-understanding> [<http://software.intel.com/en-us/articles/optimization-and-performance-tuning-for-intel-xeon-phi-coprocessors-part-2-understanding>].
- [31] Requirements for Vectorizable Loops, <http://software.intel.com/en-us/articles/requirements-for-vectorizable-loops/> [<http://software.intel.com/en-us/articles/requirements-for-vectorizable-loops/>].