

Intel Xeon Phi

A guide to Intel Xeon Phi usage

Intel Xeon Phi can be programmed in several modes. The default mode on Anselm is offload mode, but all modes described in this document are supported.

Intel Utilities for Xeon Phi

To get access to a compute node with Intel Xeon Phi accelerator, use the PBS interactive session

```
$ qsub -I -q qmic -A NONE-0-0
```

To set up the environment module “Intel” has to be loaded

```
$ module load intel/13.5.192
```

Information about the hardware can be obtained by running the micinfo program on the host.

```
$ /usr/bin/micinfo
```

The output of the “micinfo” utility executed on one of the Anselm node is as follows. (note: to get PCIe related details the command has to be run with root privileges)

MicInfo Utility Log

Created Mon Jul 22 00:23:50 2013

System Info

HOST OS	: Linux
OS Version	: 2.6.32-279.5.2.el6.Bull.33.x86_64
Driver Version	: 6720-15
MPSS Version	: 2.1.6720-15
Host Physical Memory	: 98843 MB

Device No: 0, Device Name: mic0

Version

Flash Version	: 2.1.03.0386
SMC Firmware Version	: 1.15.4830
SMC Boot Loader Version	: 1.8.4326
uOS Version	: 2.6.38.8-g2593b11
Device Serial Number	: ADKC30102482

Board

Vendor ID	: 0x8086
Device ID	: 0x2250
Subsystem ID	: 0x2500
Coprocessor Stepping ID	: 3
PCIe Width	: x16
PCIe Speed	: 5 GT/s
PCIe Max payload size	: 256 bytes
PCIe Max read req size	: 512 bytes
Coprocessor Model	: 0x01
Coprocessor Model Ext	: 0x00
Coprocessor Type	: 0x00
Coprocessor Family	: 0x0b
Coprocessor Family Ext	: 0x00
Coprocessor Stepping	: B1
Board SKU	: B1PRQ-5110P/5120D
ECC Mode	: Enabled
SMC HW Revision	: Product 225W Passive CS

Cores

Total No of Active Cores	: 60
Voltage	: 1032000 uV
Frequency	: 1052631 kHz

Thermal

Fan Speed Control	: N/A
Fan RPM	: N/A
Fan PWM	: N/A
Die Temp	: 49 C

GDDR

GDDR Vendor	: Elpida
GDDR Version	: 0x1
GDDR Density	: 2048 Mb
GDDR Size	: 7936 MB
GDDR Technology	: GDDR5
GDDR Speed	: 5.000000 GT/s
GDDR Frequency	: 2500000 kHz
GDDR Voltage	: 1501000 uV

Offload Mode

To compile a code for Intel Xeon Phi a MPSS stack has to be installed on the machine where compilation is executed. Currently the MPSS stack is only installed on compute nodes equipped with accelerators.

```
$ qsub -I -q qmic -A NONE-0-0
$ module load intel/13.5.192
```

For debugging purposes it is also recommended to set environment variable “OFFLOAD_REPORT”. Value can be set from 0 to 3, where higher number means more debugging information.

```
export OFFLOAD_REPORT=3
```

A very basic example of code that employs offload programming technique is shown in the next listing. Please note that this code is sequential and utilizes only single core of the accelerator.

```
$ vim source-offload.cpp
```

```
#include <iostream>

int main(int argc, char* argv[])
{
    const int niter = 100000;
    double result = 0;

    #pragma offload target(mic)
    for (int i = 0; i < niter; ++i) {
        const double t = (i + 0.5) / niter;
        result += 4.0 / (t * t + 1.0);
    }
    result /= niter;
    std::cout << "Pi ~ " << result << '\n';
}
```

To compile a code using Intel compiler run

```
$ icc source-offload.cpp -o bin-offload
```

To execute the code, run the following command on the host

```
./bin-offload
```

Parallelization in Offload Mode Using OpenMP

One way of parallelization a code for Xeon Phi is using OpenMP directives. The following example shows code for parallel vector addition.

```
$ vim ./vect-add
```

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

```
typedef int T;
```

```

#define SIZE 1000

#pragma offload_attribute(push, target(mic))
T in1[SIZE];
T in2[SIZE];
T res[SIZE];
#pragma offload_attribute(pop)

// MIC function to add two vectors
__attribute__((target(mic))) add_mic(T *a, T *b, T *c, int size) {
    int i = 0;
    #pragma omp parallel for
        for (i = 0; i < size; i++)
            c[i] = a[i] + b[i];
}

// CPU function to add two vectors
void add_cpu (T *a, T *b, T *c, int size) {
    int i;
    for (i = 0; i < size; i++)
        c[i] = a[i] + b[i];
}

// CPU function to generate a vector of random numbers
void random_T (T *a, int size) {
    int i;
    for (i = 0; i < size; i++)
        a[i] = rand() % 10000; // random number between 0 and 9999
}

// CPU function to compare two vectors
int compare(T *a, T *b, T size ){
    int pass = 0;
    int i;
    for (i = 0; i < size; i++){
        if (a[i] != b[i]) {
            printf("Value mismatch at location %d, values %d and %dn",i, a[i], b[i]);
            pass = 1;
        }
    }
    if (pass == 0) printf ("Test passedn"); else printf ("Test Failedn");
    return pass;
}

int main()

```

```

{
    int i;
    random_T(in1, SIZE);
    random_T(in2, SIZE);

    #pragma offload target(mic) in(in1,in2) inout(res)
    {

        // Parallel loop from main function
        #pragma omp parallel for
        for (i=0; i<SIZE; i++)
            res[i] = in1[i] + in2[i];

        // or parallel loop is called inside the function
        add_mic(in1, in2, res, SIZE);

    }

    //Check the results with CPU implementation
    T res_cpu[SIZE];
    add_cpu(in1, in2, res_cpu, SIZE);
    compare(res, res_cpu, SIZE);

}

```

During the compilation Intel compiler shows which loops have been vectorized in both host and accelerator. This can be enabled with compiler option “-vec-report2”. To compile and execute the code run

```
$ icc vect-add.c -openmp_report2 -vec-report2 -o vect-add
```

```
$ ./vect-add
```

Some interesting compiler flags useful not only for code debugging are:

Debugging `openmp_report[0|1|2]` - controls the compiler based vectorization diagnostic level `vec-report[0|1|2]` - controls the OpenMP parallelizer diagnostic level

Performance optimization `xhost - FOR HOST ONLY` - to generate AVX (Advanced Vector Extensions) instructions.

Automatic Offload using Intel MKL Library

Intel MKL includes an Automatic Offload (AO) feature that enables computationally intensive MKL functions called in user code to benefit from attached Intel Xeon Phi coprocessors automatically and transparently.

Behavioral of automatic offload mode is controlled by functions called within the program or by environmental variables. Complete list of controls is listed here.

The Automatic Offload may be enabled by either an MKL function call within the code:

```
mkl_mic_enable();
```

or by setting environment variable

```
$ export MKL_MIC_ENABLE=1
```

To get more information about automatic offload please refer to “Using Intel® MKL Automatic Offload on Intel® Xeon Phi™ Coprocessors” white paper or Intel MKL documentation.

Automatic offload example

At first get an interactive PBS session on a node with MIC accelerator and load “intel” module that automatically loads “mkl” module as well.

```
$ qsub -I -q qmic -A OPEN-0-0 -l select=1:ncpus=16
$ module load intel
```

Following example show how to automatically offload an SGEMM (single precision - g dir=“auto”>eneral matrix multiply) function to MIC coprocessor. The code can be copied to a file and compiled without any necessary modification.

```
$ vim sgemm-ao-short.c
```

```
#include <stdio.h>
#include <stdlib.h>
#include <malloc.h>
#include <stdint.h>

#include "mkl.h"

int main(int argc, char **argv)
{
    float *A, *B, *C; /* Matrices */

    MKL_INT N = 2560; /* Matrix dimensions */
    MKL_INT LD = N; /* Leading dimension */
    int matrix_bytes; /* Matrix size in bytes */
    int matrix_elements; /* Matrix size in elements */

    float alpha = 1.0, beta = 1.0; /* Scaling factors */
    char transa = 'N', transb = 'N'; /* Transposition options */
```

```

    int i, j; /* Counters */

    matrix_elements = N * N;
    matrix_bytes = sizeof(float) * matrix_elements;

    /* Allocate the matrices */
    A = malloc(matrix_bytes); B = malloc(matrix_bytes); C = malloc(matrix_bytes);

    /* Initialize the matrices */
    for (i = 0; i < matrix_elements; i++) {
        A[i] = 1.0; B[i] = 2.0; C[i] = 0.0;
    }

    printf("Computing SGEMM on the host\n");
    sgemm(&transa, &transb, &N, &N, &N, &alpha, A, &N, B, &N, &beta, C, &N);

    printf("Enabling Automatic Offload\n");
    /* Alternatively, set environment variable MKL_MIC_ENABLE=1 */
    mkl_mic_enable();

    int ndevices = mkl_mic_get_device_count(); /* Number of MIC devices */
    printf("Automatic Offload enabled: %d MIC devices present\n", ndevices);

    printf("Computing SGEMM with automatic workdivision\n");
    sgemm(&transa, &transb, &N, &N, &N, &alpha, A, &N, B, &N, &beta, C, &N);

    /* Free the matrix memory */
    free(A); free(B); free(C);

    printf("Done\n");

    return 0;
}

```

Please note: This example is simplified version of an example from MKL. The expanded version can be found here: `$MKL_EXAMPLES/mic_ao/blas/source/sgemm.c**`

To compile a code using Intel compiler use:

```
$ gcc -mkl sgemm-ao-short.c -o sgemm
```

For debugging purposes enable the offload report to see more information about automatic offloading.

```
$ export OFFLOAD_REPORT=2
```

The output of a code should look similar to following listing, where lines starting with [MKL] are generated by offload reporting:

```

Computing SGEMM on the host
Enabling Automatic Offload
Automatic Offload enabled: 1 MIC devices present
Computing SGEMM with automatic workdivision
[MKL] [MIC --] [AO Function]      SGEMM
[MKL] [MIC --] [AO SGEMM Workdivision] 0.00 1.00
[MKL] [MIC 00] [AO SGEMM CPU Time]      0.463351 seconds
[MKL] [MIC 00] [AO SGEMM MIC Time]      0.179608 seconds
[MKL] [MIC 00] [AO SGEMM CPU->MIC Data] 52428800 bytes
[MKL] [MIC 00] [AO SGEMM MIC->CPU Data] 26214400 bytes
Done

```

Native Mode

In the native mode a program is executed directly on Intel Xeon Phi without involvement of the host machine. Similarly to offload mode, the code is compiled on the host computer with Intel compilers.

To compile a code user has to be connected to a compute with MIC and load Intel compilers module. To get an interactive session on a compute node with an Intel Xeon Phi and load the module use following commands:

```
$ qsub -I -q qmic -A NONE-0-0
```

```
$ module load intel/13.5.192
```

Please note that particular version of the Intel module is specified. This information is used later to specify the correct library paths.

To produce a binary compatible with Intel Xeon Phi architecture user has to specify “-mmic” compiler flag. Two compilation examples are shown below. The first example shows how to compile OpenMP parallel code “vect-add.c” for host only:

```
$ icc -xhost -no-offload -fopenmp vect-add.c -o vect-add-host
```

To run this code on host, use:

```
$ ./vect-add-host
```

The second example shows how to compile the same code for Intel Xeon Phi:

```
$ icc -mmic -fopenmp vect-add.c -o vect-add-mic
```


Execution of the Program in Native Mode on Intel Xeon Phi

The user access to the Intel Xeon Phi is through the SSH. Since user home directories are mounted using NFS on the accelerator, users do not have to copy binary files or libraries between the host and accelerator.

To connect to the accelerator run:

```
$ ssh mic0
```

If the code is sequential, it can be executed directly:

```
mic0 $ ~/path_to_binary/vect-add-seq-mic
```

If the code is parallelized using OpenMP a set of additional libraries is required for execution. To locate these libraries new path has to be added to the LD_LIBRARY_PATH environment variable prior to the execution:

```
mic0 $ export LD_LIBRARY_PATH=/apps/intel/composer_xe_2013.5.192/compiler/lib/mic:$LD_LIBRARY_PATH
```

Please note that the path exported in the previous example contains path to a specific compiler (here the version is 5.192). This version number has to match with the version number of the Intel compiler module that was used to compile the code on the host computer.

For your information the list of libraries and their location required for execution of an OpenMP parallel code on Intel Xeon Phi is:

```
/apps/intel/composer_xe_2013.5.192/compiler/lib/mic
```

```
libiomp5.so libimf.so libsvml.so libirng.so libintlc.so.5
```

Finally, to run the compiled code use:

```
$ ~/path_to_binary/vect-add-mic
```

OpenCL

OpenCL (Open Computing Language) is an open standard for general-purpose parallel programming for diverse mix of multi-core CPUs, GPU coprocessors, and other parallel processors. OpenCL provides a flexible execution model and uniform programming environment for software developers to write portable code for systems running on both the CPU and graphics processors or accelerators like the Intel® Xeon Phi.

On Anselm OpenCL is installed only on compute nodes with MIC accelerator, therefore OpenCL code can be compiled only on these nodes.

```
module load openc1-sdk openc1-rt
```

Always load “opencl-sdk” (providing devel files like headers) and “opencl-rt” (providing dynamic library libOpenCL.so) modules to compile and link OpenCL code. Load “opencl-rt” for running your compiled code.

There are two basic examples of OpenCL code in the following directory:

```
/apps/intel/opencl-examples/
```

First example “CapsBasic” detects OpenCL compatible hardware, here CPU and MIC, and prints basic information about the capabilities of it.

```
/apps/intel/opencl-examples/CapsBasic/capsbasic
```

To compile and run the example copy it to your home directory, get a PBS interactive session on of the nodes with MIC and run make for compilation. Make files are very basic and shows how the OpenCL code can be compiled on Anselm.

```
$ cp /apps/intel/opencl-examples/CapsBasic/* .
$ qsub -I -q qmic -A NONE-0-0
$ make
```

The compilation command for this example is:

```
$ g++ capsbasic.cpp -lOpenCL -o capsbasic -I/apps/intel/opencl/include/
```

After executing the complied binary file, following output should be displayed.

```
./capsbasic
```

```
Number of available platforms: 1
```

```
Platform names:
```

```
[0] Intel(R) OpenCL [Selected]
```

```
Number of devices available for each type:
```

```
CL_DEVICE_TYPE_CPU: 1
```

```
CL_DEVICE_TYPE_GPU: 0
```

```
CL_DEVICE_TYPE_ACCELERATOR: 1
```

```
** Detailed information for each device **
```

```
CL_DEVICE_TYPE_CPU[0]
```

```
CL_DEVICE_NAME: Intel(R) Xeon(R) CPU E5-2470 0 @ 2.30GHz
```

```
CL_DEVICE_AVAILABLE: 1
```

```
...
```

```
CL_DEVICE_TYPE_ACCELERATOR[0]
```

```
CL_DEVICE_NAME: Intel(R) Many Integrated Core Acceleration Card
```

```
CL_DEVICE_AVAILABLE: 1
```

```
...
```

More information about this example can be found on Intel website: <http://software.intel.com/en-us/vcsample/samples/caps-basic/>

The second example that can be found in “/apps/intel/opencl-examples” >directory is General Matrix Multiply. You can follow the the same procedure to download the example to your directory and compile it.

```
$ cp -r /apps/intel/opencl-examples/* .
$ qsub -I -q qmic -A NONE-0-0
$ cd GEMM
$ make
```

The compilation command for this example is:

```
$ g++ cmdoptions.cpp gemm.cpp ../common/basic.cpp ../common/cmdparser.cpp ../common/oclobject.
```

To see the performance of Intel Xeon Phi performing the DGEMM run the example as follows:

```
./gemm -d 1
Platforms (1):
  [0] Intel(R) OpenCL [Selected]
Devices (2):
  [0] Intel(R) Xeon(R) CPU E5-2470 0 @ 2.30GHz
  [1] Intel(R) Many Integrated Core Acceleration Card [Selected]
Build program options: "-DT=float -DTILE_SIZE_M=1 -DTILE_GROUP_M=16 -DTILE_SIZE_N=128 -DTILE_G
Running gemm_nn kernel with matrix size: 3968x3968
Memory row stride to ensure necessary alignment: 15872 bytes
Size of memory region for one matrix: 62980096 bytes
Using alpha = 0.57599 and beta = 0.872412
...
Host time: 0.292953 sec.
Host perf: 426.635 GFLOPS
Host time: 0.293334 sec.
Host perf: 426.081 GFLOPS
...
```

Please note: GNU compiler is used to compile the OpenCL codes for Intel MIC. You do not need to load Intel compiler module.

MPI

Environment setup and compilation

Again an MPI code for Intel Xeon Phi has to be compiled on a compute node with accelerator and MPSS software stack installed. To get to a compute node with accelerator use:

```
$ qsub -I -q qmic -A NONE-0-0
```

The only supported implementation of MPI standard for Intel Xeon Phi is Intel MPI. To setup a fully functional development environment a combination of Intel compiler and Intel MPI has to be used. On a host load following modules before compilation:

```
$ module load intel/13.5.192 impi/4.1.1.036
```

To compile an MPI code for host use:

```
$ mpiicc -xhost -o mpi-test mpi-test.c
```

To compile the same code for Intel Xeon Phi architecture use:

```
$ mpiicc -mmic -o mpi-test-mic mpi-test.c
```

An example of basic MPI version of “hello-world” example in C language, that can be executed on both host and Xeon Phi is (can be directly copy and pasted to a .c file)

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

int main (argc, argv)
    int argc;
    char *argv[];
{
    int rank, size;

    int len;
    char node[MPI_MAX_PROCESSOR_NAME];

    MPI_Init (&argc, &argv);          /* starts MPI */
    MPI_Comm_rank (MPI_COMM_WORLD, &rank);      /* get current process id */
    MPI_Comm_size (MPI_COMM_WORLD, &size);      /* get number of processes */

    MPI_Get_processor_name(node,&len);

    printf( "Hello world from process %d of %d on host %s\n", rank, size, node );
    MPI_Finalize();
    return 0;
}
```

MPI programming models

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

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.)**

Coprocessor-only model** - all MPI ranks reside only on the coprocessors.

Symmetric model** - the MPI ranks reside on both the host and the coprocessor.
Most general MPI case.

Host-only model

In this case all environment variables are set by modules, so to execute the compiled MPI program on a single node, use:

```
$ mpirun -np 4 ./mpi-test
```

The output should be similar to:

```
Hello world from process 1 of 4 on host cn207
Hello world from process 3 of 4 on host cn207
Hello world from process 2 of 4 on host cn207
Hello world from process 0 of 4 on host cn207
```

Coprocessor-only model

There are two ways how to execute an MPI code on a single coprocessor: 1.) lunch the program using “**mpirun**” from the coprocessor; or 2.) lunch the task using “**mpiexec.hydra**” from a host.

Execution on coprocessor**

Similarly to execution of OpenMP programs in native mode, since the environmental module are not supported on MIC, user has to setup paths to Intel MPI libraries and binaries manually. One time setup can be done by creating a “**.profile**” file in user’s home directory. This file sets up the environment on the MIC automatically once user access to the accelerator through the SSH.

```
$ vim ~/.profile
```

```
PS1='[u@h W]$ '
```

```
export PATH=/usr/bin:/usr/sbin:/bin:/sbin
```

```
#OpenMP
```

```
export LD_LIBRARY_PATH=/apps/intel/composer_xe_2013.5.192/compiler/lib/mic:$LD_LIBRARY_PATH
```

```
#Intel MPI
```

```
export LD_LIBRARY_PATH=/apps/intel/impi/4.1.1.036/mic/lib:$LD_LIBRARY_PATH
```

```
export PATH=/apps/intel/impi/4.1.1.036/mic/bin/:$PATH
```

Please note: - this file sets up both environmental variable for both MPI and OpenMP libraries. - this file sets up the paths to a particular version of Intel

MPI library and particular version of an Intel compiler. These versions have to match with loaded modules.

To access a MIC accelerator located on a node that user is currently connected to, use:

```
$ ssh mic0
```

or in case you need specify a MIC accelerator on a particular node, use:

```
$ ssh cn207-mic0
```

To run the MPI code in parallel on multiple core of the accelerator, use:

```
$ mpirun -np 4 ./mpi-test-mic
```

The output should be similar to:

```
Hello world from process 1 of 4 on host cn207-mic0
Hello world from process 2 of 4 on host cn207-mic0
Hello world from process 3 of 4 on host cn207-mic0
Hello world from process 0 of 4 on host cn207-mic0
```

Execution on host

If the MPI program is launched from host instead of the coprocessor, the environmental variables are not set using the “profile” file. Therefore user has to specify library paths from the command line when calling “mpiexec”.

First step is to tell mpiexec that the MPI should be executed on a local accelerator by setting up the environmental variable “I_MPI_MIC”

```
$ export I_MPI_MIC=1
```

Now the MPI program can be executed as:

```
$ mpiexec.hydra -genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/ -host mic0 -n 4 ~/mpi-test-mic
```

or using mpirun

```
$ mpirun -genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/ -host mic0 -n 4 ~/mpi-test-mic
```

Please note: - the full path to the binary has to be specified (here: “>~/mpi-test-mic”) - the LD_LIBRARY_PATH has to match with Intel MPI module used to compile the MPI code

The output should be again similar to:

```
Hello world from process 1 of 4 on host cn207-mic0
Hello world from process 2 of 4 on host cn207-mic0
Hello world from process 3 of 4 on host cn207-mic0
Hello world from process 0 of 4 on host cn207-mic0
```

Please note that the “mpiexec.hydra” requires a file “>pmi_proxy” from Intel MPI library to be copied to the MIC filesystem. If the file is missing please

contact the system administrators. A simple test to see if the file is present is to execute:

```
$ ssh mic0 ls /bin/pmi_proxy  
/bin/pmi_proxy
```

Execution on host - MPI processes distributed over multiple accelerators on multiple nodes

To get access to multiple nodes with MIC accelerator, user has to use PBS to allocate the resources. To start interactive session, that allocates 2 compute nodes = 2 MIC accelerators run qsub command with following parameters:

```
$ qsub -I -q qmic -A NONE-0-0 -l select=2:ncpus=16
```

```
$ module load intel/13.5.192 impi/4.1.1.036
```

This command connects user through ssh to one of the nodes immediately. To see the other nodes that have been allocated use:

```
$ cat $PBS_NODEFILE
```

For example:

```
cn204.bullx  
cn205.bullx
```

This output means that the PBS allocated nodes cn204 and cn205, which means that user has direct access to “**cn204-mic0**” and “**cn-205-mic0**” accelerators.

Please note: At this point user can connect to any of the allocated nodes or any of the allocated MIC accelerators using ssh: - to connect to the second node :
** \$ ssh cn205 - **to connect to the accelerator on the first node from the first node:** \$ ssh cn204-mic0** or \$ ssh mic0 - to connect to the accelerator on the second node from the first node: **\$ ssh cn205-mic0**

At this point we expect that correct modules are loaded and binary is compiled. For parallel execution the mpiexec.hydra is used. Again the first step is to tell mpiexec that the MPI can be executed on MIC accelerators by setting up the environmental variable “I_MPI_MIC”

```
$ export I_MPI_MIC=1
```

The launch the MPI program use:

```
$ mpiexec.hydra -genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/  
-genv I_MPI_FABRICS_LIST tcp  
-genv I_MPI_FABRICS shm:tcp  
-genv I_MPI_TCP_NETMASK=10.1.0.0/16  
-host cn204-mic0 -n 4 ~/mpi-test-mic  
: -host cn205-mic0 -n 6 ~/mpi-test-mic
```

or using mpirun:

```
$ mpirun -genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/
-genv I_MPI_FABRICS_LIST tcp
-genv I_MPI_FABRICS shm:tcp
-genv I_MPI_TCP_NETMASK=10.1.0.0/16
-host cn204-mic0 -n 4 ~/mpi-test-mic
: -host cn205-mic0 -n 6 ~/mpi-test-mic
```

In this case four MPI processes are executed on accelerator cn204-mic and six processes are executed on accelerator cn205-mic0. The sample output (sorted after execution) is:

```
Hello world from process 0 of 10 on host cn204-mic0
Hello world from process 1 of 10 on host cn204-mic0
Hello world from process 2 of 10 on host cn204-mic0
Hello world from process 3 of 10 on host cn204-mic0
Hello world from process 4 of 10 on host cn205-mic0
Hello world from process 5 of 10 on host cn205-mic0
Hello world from process 6 of 10 on host cn205-mic0
Hello world from process 7 of 10 on host cn205-mic0
Hello world from process 8 of 10 on host cn205-mic0
Hello world from process 9 of 10 on host cn205-mic0
```

The same way MPI program can be executed on multiple hosts:

```
$ mpiexec.hydra -genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/
-genv I_MPI_FABRICS_LIST tcp
-genv I_MPI_FABRICS shm:tcp
-genv I_MPI_TCP_NETMASK=10.1.0.0/16
-host cn204 -n 4 ~/mpi-test
: -host cn205 -n 6 ~/mpi-test
```

Symmetric model

In a symmetric mode MPI programs are executed on both host computer(s) and MIC accelerator(s). Since MIC has a different architecture and requires different binary file produced by the Intel compiler two different files has to be compiled before MPI program is executed.

In the previous section we have compiled two binary files, one for hosts “**mpi-test**” and one for MIC accelerators “**mpi-test-mic**”. These two binaries can be executed at once using mpiexec.hydra:

```
$ mpiexec.hydra
-genv I_MPI_FABRICS_LIST tcp
-genv I_MPI_FABRICS shm:tcp
-genv I_MPI_TCP_NETMASK=10.1.0.0/16
-genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/
-host cn205 -n 2 ~/mpi-test
```



```
: -host cn205-mic0 -n 2 ~/mpi-test-mic
```

In this example the first two parameters (line 2 and 3) sets up required environment variables for execution. The third line specifies binary that is executed on host (here cn205) and the last line specifies the binary that is execute on the accelerator (here cn205-mic0).

The output of the program is:

```
Hello world from process 0 of 4 on host cn205
Hello world from process 1 of 4 on host cn205
Hello world from process 2 of 4 on host cn205-mic0
Hello world from process 3 of 4 on host cn205-mic0
```

The execution procedure can be simplified by using the mpirun command with the machine file as a parameter. Machine file contains list of all nodes and accelerators that should be used to execute MPI processes.

An example of a machine file that uses 2 >hosts (**cn205** and **cn206**) and 2 accelerators (**cn205-mic0** and **cn206-mic0**) to run 2 MPI processes on each of them:

```
$ cat hosts_file_mix
cn205:2
cn205-mic0:2
cn206:2
cn206-mic0:2
```

In addition if a naming convention is set in a way that the name of the binary for host is “**bin_name**” and the name of the binary for the accelerator is “**bin_name-mic**” then by setting up the environment variable **I_MPI_MIC_POSTFIX** to “**-mic**” user do not have to specify the names of both binaries. In this case mpirun needs just the name of the host binary file (i.e. “mpi-test”) and uses the suffix to get a name of the binary for accelerator (i.e. “mpi-test-mic”).

```
$ export I_MPI_MIC_POSTFIX=-mic
```

>To run the MPI code using mpirun and the machine file “hosts_file_mix” use:

```
$ mpirun
-genv I_MPI_FABRICS shm:tcp
-genv LD_LIBRARY_PATH /apps/intel/impi/4.1.1.036/mic/lib/
-genv I_MPI_FABRICS_LIST tcp
-genv I_MPI_FABRICS shm:tcp
-genv I_MPI_TCP_NETMASK=10.1.0.0/16
-machinefile hosts_file_mix
~/mpi-test
```

A possible output of the MPI “hello-world” example executed on two hosts and two accelerators is:

```
Hello world from process 0 of 8 on host cn204
Hello world from process 1 of 8 on host cn204
Hello world from process 2 of 8 on host cn204-mic0
Hello world from process 3 of 8 on host cn204-mic0
Hello world from process 4 of 8 on host cn205
Hello world from process 5 of 8 on host cn205
Hello world from process 6 of 8 on host cn205-mic0
Hello world from process 7 of 8 on host cn205-mic0
```

Please note: At this point the MPI communication between MIC accelerators on different nodes uses 1Gb Ethernet only.

Using the PBS automatically generated node-files

PBS also generates a set of node-files that can be used instead of manually creating a new one every time. Three node-files are generated:

Host only node-file: - `/lscratch/PBS_JOBID/nodefile-cnMIConlynode-file`
: `ä - /lscratch/{PBS_JOBID}/nodefile-mic` Host and MIC node-file: -
`/lscratch/${PBS_JOBID}/nodefile-mix`

Please note each host or accelerator is listed only per files. User has to specify how many jobs should be executed per node using “-n” parameter of the mpirun command.

Optimization

For more details about optimization techniques please read Intel document Optimization and Performance Tuning for Intel® Xeon Phi™ Coprocessors