Intel Xeon Phi
Exercises

29.06.2016
Fabio Affinito

Preliminaries:
Exercises will be executed on the EURORA cluster. For any other technical reference look at
http://www.hpc.cineca.it/content/eurora-user-guide

In order to connect to this cluster you should use ssh from a terminal:
`ssh <USERNAME>@login.eurora.cineca.it`

where the usernames and the passwords will be given by the lab assistant.

In order to compile your files, you need to load modules and to source some files::

```
module load intel
module load mkl (if required)
source $INTEL_HOME/bin/compilervars.sh intel64
```

You can execute your files on the compute nodes using a job script or in an interactive session.

If you want to use a job script, you should put at the begin of your job script the following lines:

```
#PBS -l walltime=0:10:00
#PBS -l select=1:ncpus=1:nmics=1
#PBS -A train_chyb2015
#PBS -q parallel
#PBS -W group_list=train_chyb2015
```

If you want to use an interactive session (this is better for testing and debugging purposes), you can request an interactive session by specifying

```
qsub -I  -l walltime=0:10:00  -l select=1:ncpus=1:nmics=1 -A train_chyb2015 -W
group_list=train_chyb2015 -q parallel
```

Fortran version:

Exercise 1

- Copy `omp_offload_start.F90` to `omp_offload.F90`
- Edit `omp_offload.F90` and add code to offload the OpenMP section and to offload the test to check whether or not the code is running on the coprocessor
- Compare `omp_offload.F90` to `omp_offload_ours.F90` to make sure you got everything
- Make sure that the number of threads is unconstrained (unset OMP_NUM_THREADS)
- Build the result for host-only and check the vectorization messages
  `ifort -vec-report=3 -openmp -no-offload omp_offload.F90 main.F90`
- Build the result for offload and compare the vectorization messages with the case of the host compilation
  `ifort -vec-report=3 -openmp omp_offload.F90 main.F90`
- Run the result with different numbers of threads on the coprocessor so that you can see the scaling
- What sort of scaling do you see?

Exercise 2

- Make a copy of `mCarlo_offload_start.F90`:

```
cp mCarlo_offload_start.F90 mCarlo_myoffload.F90
```

- Add code to offload the “do_calculation” subroutine and write code in order to test whether or not the code is running on the coprocessor. The code to be offloaded was placed in a subroutine to simplify the creation of the streams on the coprocessor rather than on the processor.
- Build the result:
  ```
  ifort -mkl -openmp mCarlo_myoffload.F90
  ```
- It could happen that for the last recent of the compiler, the VSL_METHOD_DGAUSSIAN_BOXMULLER2 is no longer present. You can replace it with the new method name VSL_RNG_METHOD_GAUSSIAN_BOXMULLER2
- Compare your result to `mCarlo_offload_ours.F90`

Exercise 3

“Native” Intel® Xeon PhiTM coprocessor applications treat the coprocessor as a standalone multicore computer. Once the binary is built on your host system, it is copied to the “filesystem” on the coprocessor along with any other binaries and data it requires. The program is then run from the ssh console. (On Cineca machines you can reach a MIC card typing ssh $HOSTNAME-mic0 or $HOSTNAME-mic1).

- Build our sample application with the –mmic flag. The sample code is a single-file version of the matrix multiply code we previously worked with:
  ```
  ifort –mmic –vec-report=3 –openmp omp_offload_native.F90
  ```
- Once you are in a compute node (accessed in interactive mode) you can log on the coprocessor
  ```
  ssh $HOSTNAME-mic0
  ```
  or
  ```
  ssh $HOSTNAME-mic1
  ```
  There you should find on your filesystem the file that you compiled on the host side. Actually the filesystem is a the host filesystem mounted with NFS on the MIC cards
- Try to run your code
  ```
  ~ # ./a.out 2048
  ```
- As you noted from the error message, we are missing the OpenMP runtime library needed to run this application. So, when you're logged on the Xeon Phi, export the proper library directory:
  ```
  export LD_LIBRARY_PATH=
  /cineca/prod/compilers/intel/cs-xe-2013/binary/composerxe/lib/mic/
  ```
- Try to run again.

Exercise 4

Code of any complexity tends to do things in stages. This can complicate things when multiple stages need to execute on a coprocessor, and you need the results from one stage to persist until the next call. In this section, we will explore how this is done.

- Take a look at `omp_offload_ours.F90` and note how the data transfer and work happen in a single offload call. Let us artificially change this into three stages and observe what happens.
- Start with `omp_3stageoffload_nopersist.F90`. Build it and observe what happens when it runs:
  ```
  ifort -O3 omp_3stageoffload_nopersist.F90 -o mmul_nopersist
  ./mmul_nopersist 2048
  ```
- You will see an error message.
- Now compare `omp_3stageoffload_nopersist.F90` to `omp_3stageoffload_persist.F90`
- Build and run `omp_3stageoffload_persist.F90`:
  ```
  ifort -O3 omp_3stageoffload_persist.F90 -o mmul_persist
  ./mmul_persist 2048
  ```
- Did you get the expected result?
- Make sure you understand how the alloc_if, free_if, and nocopy qualifiers are used in the offload statement. Refer to the compiler reference manual.

Exercise 5

Codes often operate on blocks of data which require the data block to be moved to the coprocessor at the start of the computation and back to the host at the end. Such codes benefit by the use of asynchronous data transfers where the coprocessor computes one block of data while another block is being transferred from the host. Asynchronous transfers can also improve performance for codes requiring multiple data transfers between the host and the coprocessor.

- Take a look at `do_offload subroutine` in `async_start.F90` and notice how the two arrays are processed one after the other using offload statements.
- Change this code so that you transfer one array while the other one is computing. Modify the `do_async` function to use asynchronous data transfers.
- Build and run the program.
  ```
  ifort -o async.out async_start.F90
  ./async.out
  ```
- Notice that the `do_async` function is faster compared to the `do_offloads` function.
- Make sure you understand how the signal and wait qualifiers are used in the offload statements. Refer to the compiler reference manual for more details.

C/C++ version

Exercise 1

- Copy `omp_offload_start.cpp` to `omp_offload.cpp`
- Edit `omp_offload.cpp` and add code to offload the OpenMP section and to offload the test for whether or not the code is running on the coprocessor
- Compare `omp_offload.cpp` to `omp_offload_ours.cpp` to make sure you got everything
- Make sure that the number of threads is unconstrained (unset OMP_NUM_THREADS)
- Build the result for host-only and check the vectorization messages
  ```
  icc -vec-report=3 -openmp -no-offload omp_offload.cpp main.cpp
  ```
- Build the result for offload and compare the vectorization messages wrt to the offload version:
  ```
  icc -vec-report=3 -openmp omp_offload.cpp main.cpp
  ```
- Run the result with different numbers of threads on the coprocessor so that you can see the scaling
- What sort of scaling do you see?

Exercise 2

- Make a copy of `mCarlo_offload_start.cpp`:
  ```
  cp mCarlo_offload_start.cpp mCarlo_myoffload.cpp
  ```
- Add code to offload the OpenMP section and write code in order to test whether or not the code is running on the coprocessor. Note how we had to move the VSLStreamStatePtr definitions within the offload statement block (compare to `mCarlo_offload_ours.cpp`).
- Build the result:
  ```
  icc -mkl -openmp mCarlo_myoffload.cpp
  ```
- It could happen that for the last recent of the compiler, the `VSL_METHOD_DGAUSSIAN_BOXMULLER2` is no longer present. You can replace it with the new method name: `VSL_RNG_METHOD_GAUSSIAN_BOXMULLER2`
- Compare your result to mCarlo_offload_ours.cpp

Exercise 3

“Native” Intel® Xeon PhiTM coprocessor applications treat the coprocessor as a standalone multicore computer. Once the binary is built on your host system, it is copied to the “filesystem” on the coprocessor along with any other binaries and data it requires. The program is then run from the ssh console. (On Cineca machines you can reach a MIC card typing ssh $HOSTNAME-mic0 or $HOSTNAME-mic1).

- Build our sample application with the –mmic flag. The sample code is a single-file version of the matrix multiply code we previously worked with:
  `icc -mmic -vec-report=3 -openmp omp_offload_native.cpp`
- Once you are in a compute node (accessed in interactive mode) you can log on the coprocessor
  `ssh $HOSTNAME-mic0`
  or
  `ssh $HOSTNAME-mic1`
  There you should find on your filesystem the file that you compiled on the host side. Actually the filesystem is a the host filesystem mounted with NFS on the MIC cards
- Try to run your code
  `~ # ./a.out 2048`
- As you noted from the error message, we are missing the OpenMP runtime library needed to run this application. So, when you're logged on the Xeon Phi, export the proper library directory:
  `export LD_LIBRARY_PATH=`
  `/cineca/prod/compilers/intel/cs-xe-2013/binary/composerxe/lib/mic/`
- Try to run again.

Exercise 4

Code of any complexity tends to do things in stages. This can complicate things when multiple stages need to execute on a coprocessor, and you need the results from one stage to persist until the next call. In this section, we will explore how this is done. Take a look at `omp_offload_ours.cpp` and note how the data transfer and work happen in a single offload call. Let us artificially change this into three stages and observe what happens.

- Start with `omp_3stageoffload_nopersist.cpp`. Build it and observe what happens when it runs:
  `icc -O3 -openmp omp_3stageoffload_nopersist.cpp -o mmul_nopersist`
  `./mmul_nopersist 2048`
- You will see an error message.
- Now compare `omp_3stageoffload_nopersist.cpp` to `omp_3stageoffload_persist.cpp`
- Build and run `omp_3stageoffload_persist.cpp`:
  `icc -O3 -openmp omp_3stageoffload_persist.cpp -o mmul_persist`
  `./mmul_persist 2048`
- Did you get the expected result?
- Make sure you understand how the alloc_if, free_if, and nocopy qualifiers are used in the offload statement. Refer to the compiler reference manual.

Exercise 5

Codes often operate on blocks of data which require the data block to be moved to the coprocessor at the start of the computation and back to the host at the end. Such codes benefit by the use of asynchronous data transfers where the coprocessor computes one block of data while another block is being transferred from the host. Asynchronous transfers can also improve performance for codes requiring multiple data transfers between the host and the coprocessor.

- Take a look at `do_offload` function in `async_start.cpp` and notice how the two arrays are processed one after the other using offload statements.
- Change this code so that you transfer one array while the other one is computing. Modify the `do_async` function to use asynchronous data transfers.
- Build and run the program.
  `icc -o async.out async_start.cpp`
  `./async.out`
- Notice that the `do_async` function is faster compared to the `do_offloads` function.
- Make sure you understand how the signal and wait qualifiers are used in the offload statements. Refer to the compiler reference manual for more details.