MKL-lab1 : SGEMM with Automatic Offload

In this example we use the SGEMM function from Intel MKL, to offloading computation to Intel MIC architecture coprocessor.

1. Objectives and learning goals

- Open in an editor the source code
- Set up the build environment
- Replace the WORK FOR YOU comments with MKL calls
- Build the program
- Run the program with various matrix sizes (e.g. 1000 and 3000)
- Set the following environment variables and repeat the tests:

  export MIC_ENV_PREFIX=MIC
  export MIC_USE_2MB_BUFFERS=16K
  export MIC_KMP_AFFINITY=explicit,granularity=fine,proclist=[1-240:1]
  export KMP_AFFINITY=granularity=fine,compact,1,0

Try now to understand the performance numbers observed for the host and AO execution.