IEEE Cluster 2013 Tutorial
Programming for the Intel Xeon Phi

PART I - Introduction (1.5 hours)

- Xeon Phi Architecture ~ 30 minutes
- Programming models ~ 30 minutes
  - Native Execution (MPI / Threads / MPI+Threads)
  - MPI on host and Phi
  - MPI on host, offload to Phi
    - Targeted
    - Automatic (MKL)
  - Offload to host from the Phi
- Hands-On Lab ~ 30 minutes
  - Login and explore busybox

PART II - Native Execution (1.5 hours)

- Native Execution ~ 45 minutes
  - Why run native?
  - How to build a native application?
  - How to run a native application?
  - Best practices for running native
    - KMP_AFFINITY
  - Optimization
    - Cache + ALU/SIMD details
    - Vectorization
    - Parallelization
    - Alignment
    - Compiler reports
- Hands-On Lab ~ 45 minutes
  - Interactive exercise using compiler reports
  - Interactive exercise to show affinity settings

PART III - Offload Execution (1.5 hours)

- Offload to Phi ~ 45 minutes
  - What is offloading?
  - Directives
  - Automatic offloading with MKL
  - Compiler assisted offloading
  - Offloading inside a parallel region
- Hands-On Lab ~ 45 minutes
  - Interactive exercise with simple offload and data transfer

PART IV - MPI + Offload Execution (1.5 hours)

- MPI execution ~ 45 minutes
  - Symmetric execution
    - Workload distribution
    - ibrun.sym
    - Correct pinning of MPI tasks on host and coprocessor
    - Interactive exercise showing symmetric at work
  - MPI + offload

- Hands-On Lab ~ 45 minutes
  - Explore symmetric execution