

Hands On #1

- Download and install MPICH on laptop and run cpi example in examples directory
- Copy cpi.c over to Mira, compile, and run
- Compile and run MLife examples from www.cs.illinois.edu/~wgropp/advmpi.tgz on Mira
 - Untar, run configure and make to build executable

Hands On #2

1. Build and run mesh2d on Vesta or your favorite system
 - Run with large enough meshes; e.g., `mfile2d -x 1000 -y 1000`
2. Take the file `mfile2d-pt2pt.c` and modify it to try one of the following (we recommend testing on your laptop first):
 - Persistent sends
 - Ready sends
 - Sendrecv
 - Advanced: Restructure to allow computation during communication
3. Measure the performance (modify `mfile2d.c` to include your new routines)

Hands On #3

- Take the file `mlife2d-pt2ptuv.c` and
 1. Create a new version that that uses `MPI_Type_vector`.
 2. Compare with `mlife2d-pt2pt.c` . How does your version compare?
 3. Advanced: Replace `Type_vector` with `Type_create_resized` with stride to produce a type that can be used for any length vector with the same stride
 4. Advanced: Compare performance with manual packing and with estimate of performance (how fast should it be)?

Hands On #4

- Run `mlife2d` with different sizes of meshes:
 - `Mlife2d -x 2000 -y 2000 -i 100`
 - `Mlife2d -x 4000 -y 4000 -i 100`
- Observe the RMA performance using Fence synchronization
- Create a new version, starting from `mlife2d-fence.c`, that uses `MPI_Win_lock/unlock` synchronization. How does your version perform?
 - Hint: How do you know when you can safely proceed to the next iteration? Or what do you need to do in the code to ensure you can move to the next iteration?

Hands On #5

- Measure performance of halo exchange with different process mappings, using mlife2d
- Advanced (extra credit): compare with expected performance, using simple communication performance model