I. High performance computing with OpenMPI

Often the resources (number of cores and available memory) of a single workstation are not sufficient to perform calculations of realistic quantum fields on lattices. Therefore we need to learn how to combine resources of multiple machines via the OpenMPI framework.

When starting the computation several instances of the program are started both locally and on remote machines, which are then able to communicate via OpenMPI routines. (The details depend on the setup of the computing cluster used). This is necessary to synchronize memory boundaries and e.g. in the computation of global observables.

a.) Go through the provided source code and understand its content. Compile the program using mpicc and submit locally with four processes via mpirun –n 4 ./PROGNAME

b.) Based on the explanation of the MPI functions MPI_Scatter and MPI_Gather
http://www.mpich.org/static/docs/v3.1/www3/MPI_Scatter.html
http://www.mpich.org/static/docs/v3.1/www3/MPI_Gather.html
write a program in which the root process generates a set of random numbers (here simply use drand48()) and distributes it evenly to the nodes, with the nodes all computing their sub average of the numbers and transmitting the outcome back to root for the final average over the sub averages. (e.g. 3 Processes 15 random numbers in total)

c.) Instead of MPI_Gather use the MPI function MPI_Reduce to directly calculate the average of the random numbers on the root process

https://computing.llnl.gov/tutorials/mpi/#Collective_Communication_Routines

http://www.mpich.org/static/docs/v3.1/www3/MPI_Reduce.html