

LAMMPS Best Practices for Intel® Cluster Ready



BEST PRACTICES

1. Introduction:

The following best practices document is provided as courtesy of the HPC Advisory Council.

2. Version Information:

Download LAMMPS at:

<http://lammps.sandia.gov/tars/>

Download FFTW from:

<http://www.fftw.org/>

3. Prerequisites:

The instructions from this best practice have been tested with the following configuration:

3.1 Hardware:

- Intel Cluster Ready, Dell PowerEdge M610 14-node cluster
- Intel Xeon X5670 CPUs
- Memory: 24GB per node
- Mellanox ConnectX-2 QDR InfiniBand Adapters
- Mellanox QDR InfiniBand Switch

3.2 Software:

- OS: CentOS 5 Update 4
- Application: LAMMPS
- Compilers: Intel compilers, GNU compilers
- MPI: Intel MPI 4, Open MPI 1.4.2, MVAPI-CH2-1.5, Platform MPI 7.1
- Benchmark workload: rhodo.scaled

4. Building LAMMPS

Extract LAMMPS

```
% tar xvfz lammps-30Aug10.tar.gz
```

Modify the Makefile according to MPI, compiler and library used

```
% cd lammps-30Aug10/src/MAKE  
% vim Makefile.{name_of_target}  
% make {name_of_target}
```

Makefile changes for using GNU compilers and FFTW library

```
CC = mpicxx  
CCFLAGS = -g -O  
DEPFLAGS = -M  
LINK = mpicxx  
LINKFLAGS = -g -O  
LIB =  
FFT_INC = -I/application/fftw-2.1.5-gnu/include -DFFT_  
FFTW  
FFT_PATH = -L/application/fftw-2.1.5-gnu/lib  
FFT_LIB = -lfftw
```

Makefile changes for using Intel compilers and MKL library

```
CC = icc  
CCFLAGS = -O3 -fno-alias -ip -unroll0  
DEPFLAGS = -M  
LINK = icc  
LINKFLAGS = -O -L/application/intel/11.1/064/mkl/lib/  
em64t  
LIB = -lstdc++ -lpthread -mkl:cluster -lguide  
FFT_INC = -DFFT_FFTW -I/application/intel/11.1/064/  
mkl/include/fftw  
FFT_PATH =  
FFT_LIB = /application/intel/11.1/064/mkl/lib/em64t/  
libfftw2xf_intel.a
```

Makefile changes for using Intel MPI 4

```
MPI_INC = -I/application/intel/impi/include64 -DMPICH_  
IGNORE_CXX_SEEK  
MPI_PATH = -L/application/intel/impi/lib64  
MPI_LIB = -lpthread -lmpi
```

Makefile changes for using Open MPI 1.4.2

```
MPI_INC = -I/application/openmpi-1.4.2-gnu/include  
-DMPICH_IGNORE_CXX_SEEK
```

```
MPI_PATH = -L/application/openmpi-1.4.2-gnu/lib
```

```
MPI_LIB = -lpthread -lmpi
```

Makefile changes for using MVAPICH2-1.5

```
MPI_INC = -I/application/mvapich2-1.5-gnu/include -DM-  
PICH_IGNORE_CXX_SEEK
```

```
MPI_PATH = -L/application/mvapich2-1.5-gnu/lib
```

```
MPI_LIB = -lpthread -lmpich
```

Makefile changes for using Platform MPI 7.1

```
MPI_INC = -I/opt/platform_mpi/include -DMPICH_IG-  
NORE_CXX_SEEK
```

```
MPI_PATH = -L/opt/platform_mpi/lib/linux_amd64
```

```
MPI_LIB = -lpthread -lmpi
```

5. Running LAMMPS

Running with Intel MPI

```
% mpdboot -r ssh -f /home/pak/mpd.hosts.ib.14 -n 14
```

```
% mpiexec -np 168 -IB -print-rank-map ~/lammmps-  
30Aug10-gnu-openmpi-1.4.2/src/lmp_maia_intel_mpi -in  
in.rhodo.scaled -var x 4 -var y 6 -var z 7
```

```
%mpdallexit
```

Running with Open MPI

```
% mpirun -np 168 -hostfile ~/hostfile-ompi.14 -mca  
mpi_paffinity_alone 1 -mca btl self,sm,openib ~/lammmps-  
30Aug10-gnu-openmpi-1.4.2/src/lmp_maia_gnu_openmpi  
-in in.rhodo.scaled -var x 4 -var y 6 -var z 7
```

Running with MVAPICH2 MPI

```
%mpdboot -r ssh -f /home/pak/mpd.hosts.ib.14 -n 14
```

```
% mpiexec -np 168 ~/lammmps-30Aug10-gnu-openm-  
pi-1.4.2/src/lmp_maia_intel_mvapich2 -in in.rhodo.scaled  
-var x 4 -var y 6 -var z 7
```

```
% mpdallexit
```

Running with Platform MPI

```
% mpirun -np 168 -IBV -cpu_bind -hostfile ~ /hostfile-  
hpmi-ib14 ~/lammmps-30Aug10-gnu-openmpi-1.4.2/src/  
lmp_maia_intel_platform -in in.rhodo.scaled -var x 4 -var  
y 6 -var z 7
```

6. Profiling LAMMPS with IPM

```
export IPM_KEYFILE=/application/ipm-impi-intel/include/  
ipm_key
```

```
export IPM_PARALLEL_IO=no
```

```
export IPM_REPORT=full
```

```
export IPM_LOG_DIR=/home/ipm
```

```
export IPM_LOGDIR=/home/ipm
```

```
% mpdboot -r ssh -f /home/pak/mpd.hosts.ib.14 -n 14
```

```
% LD_PRELOAD=/application/ipm-impi-intel/lib/libipm.so  
mpiexec -np 168 -IB -print-rank-map ~/lammmps-30Aug10-  
gnu-openmpi-1.4.2/src/lmp_maia_intel_mpi -in in.rhodo.  
scaled -var x 4 -var y 6 -var z 7
```

```
% mpdallexit
```

