

# LAMMPS Guide

Last updated 9 February, 2023 • 4 min read

From the LAMMPS README file:

*LAMMPS is a classical molecular dynamics simulation code designed to run efficiently on parallel computers. It was developed at Sandia National Laboratories, a US Department of Energy facility, with funding from the DOE. It is an open-source code, distributed freely under the terms of the GNU Public License (GPL).*

## LAMMPS MPI and package install

LAMMPS has various options for building and installing different settings and packages. Due to LAMMPS having a lot of custom builds and various package options, we recommend users to install LAMMPS locally under their /home directories or team's /shared folder for ease of use and the current needs for the user or team. Reading the documentation pages for LAMMPS, we can compile LAMMPS with CMake or with regular make commands. We will use the GCC compiler in this example:

 Step 1: Create directory project and download LAMMPS:



```
> # Create a directory for our LAMMPS project.
1
2 mkdir -pv ~/software/lammps
3 cd ~/software/lammps
4
5 # Download the latest Lammps Linux tarball file from their do
6
7 Go to: https://www.lammps.org/download.html
8
9 # Copy the download link for the LAMMPS stable release and pa
10 wget pasteDownloadLinkforLammpsTarballHere
11
12 #Unpack the tarball and cd into the new lammps directory that
13 #This example will show the LAMMPS 23Jun2022 release
14
15 tar -xvf lammps-stable.tar.gz
16 cd lammps-23Jun2022
17
18 # Directory structure should look like the following:
19
20 ~/software/lammps/lammps-23Jun2022
21
22
23
```

## Step 2: Edit LAMMPS makefile

```

> #LAMMPS allows for CMake or regular make builds
1 #In this example we will show the regular make build and inst
2 #Going to build the MPI version of LAMMPS
3
4 cd lammeps-23Jun2022/src/MAKE
5
6 #Make a copy of the Makefile.mpi and rename it to Makefile.or
7
8 cp Makefile.mpi Makefile.origmpi
9
10 #Copy the Makefile.omp file from the src/MAKE/OPTIONS folder
11
12 cp OPTIONS/Makefile.omp Makefile.mpi
13 vi Makefile.mpi
14 Press i to go into Insert mode in the VI editor
15
16 #Go down to the MPI Library section and update the paths to p
17
18 MPI_INC = -I/gpfs/sharedfs1/admin/hpc2.0/apps/openmpi/4.1.
19 MPI_PATH = -L/gpfs/sharedfs1/admin/hpc2.0/apps/openmpi/4.1.
20 MPI_LIB = -lmpi -lpthread
21
22 #Most software installs do not use custom edited Makefiles ar
23 #If the build would need fftw3, the FFT library section woul
24 #Now we need to add some Link flags at the top to have LAMMPS
25 #If the build would need blas or lapack libraries, they woul
26
27 FC = mpifort
28 CC = mpicc
29 CXX = mpicxx
30 CCFLAGS = -g -O3 -fopenmp -std=c++11
31 SHFLAGS = -fPIC
32 DEPFLAGS = -M
33
34 LINK = mpicxx
35 LINKFLAGS = -g -O3 -fopenmp -std=c++11 -L/gpfs/sharedfs1/admi
36 LIB = -lstdc++ -lgfortran
37 SIZE = size
38
39 ARCHIVE = ar
40 ARFLAGS = -rc
41 SHLIBFLAGS = -shared -rdynamic
42
43 #Once entered, save the Makefile with the following key strok
44
45 ESC key (to get out of Insert mode)
46 :wq!
47
48
49

```

### Step 3: Build LAMMPS packages

```
> #Go back one directory to the /src folder:
1
2 cd ..
3
4 #Should show the following path:
5
6 ~/software/lammps/lammps-23Jun2022/src/
7
8 #Load the gcc/11.3.0 and openmpi/4.1.4 modules before building
9
10 module load gcc/11.3.0 openmpi/4.1.4
11
12 # Install needed LAMMPS packages for run with the following c
13
14 make yes-packagenamehere
15
16 # Here are some example Commands to install the openmp, qeq,
17
18 make yes-openmp
19
20 make yes-qeq
21
22 make yes-meam
23
24 make yes-reaxff
25
26 # After installing the needed packages, the following command
27
28 make ps
29
30 # Or if a list of installed packages is needed without showing
31
32 make pi
33
34
35
36
```

## Step 4: Build LAMMPS

```

> # Build the LAMMPS MPI executable, the -j flag tells mpi to b
1
2 make -j4 mpi
3
4 # LAMMPS should start building the executable with the packag
5 # Once LAMMPS finishes building under the /src folder, LAMMPS
6 lmp_mpi
7
8 #Go back to the main lammps folder
9
10 cd ~/software/lammps/lammps-23Jun2022
11
12 #Create a /bin folder to move the lmp_mpi executable into
13
14 mkdir bin
15
16 #Copy the lmp_mpi exectuable over to this new /bin folder
17
18 rsync -a --progress ~/software/lammps/lammps-23Jun2022/src/ln
19
20 #Continue to the next step to create a module file for the ne
21 #If a module is not needed, the path can be specified in a su
22
23 ~/software/lammps/lammps-23Jun2022/bin/lmp_mpi -i restOfLammp
24
25

```

## Step 5: Create a Module file for LAMMPS

We can create a module file for LAMMPS that so that we can conveniently load LAMMPS and it's dependencies. The name that you choose for your module file is important as that is what module uses to reference it. We will make our name different by adding the "-mine" suffix to help separate it from the system installed vasp.

```

> mkdir -p ~/mod/lammps
1 cd ~/mod/lammps
2 vi 23Jun2022-mine
3
4
5

```

```

> #%Module1.0
1
2 # Throw an error if any of these modules are loaded.
3 conflict lammps
4
5 # Load the particular GCC and openmpi versions we used for th
6 module load gcc/11.3.0
7 module load openmpi/4.1.4
8
9 # Modify the PATH to use our compiled LAMMPS. Do not use a t
10 prepend-path PATH ~/software/lammps/lammps-23Jun2022/bin/
11
12
13

```

If you are interested, in learning about module files you can read `man modulefile`

Finally, make sure that `module` knows to look in your `~/mod` directory for your module files by setting the `MODULEPATH` environmental variable:

```

> nano ~/.bashrc # Add the lines below.
1
2
3

```

```

> # My modules
1 source /etc/profile.d/modules.sh
2 MODULEPATH=${HOME}/mod:${MODULEPATH}
3
4
5

```

Reload your `~/.bashrc` file in your current shell:

```

> 1 source ~/.bashrc
1 2 # Finally Now we can load and run our VASP module
2 3 module load lammps/23Jun2022-mine
3 4 which lmp_mpi
4 5 lmp_mpi -h
5
6

```

To install packages AFTER lammps is installed, go to the LAMMPS installer `/src` folder:

```

> cd ~/software/lammps/lammps-23Jun2022/src
1

```

Then enter the following depending on which packages that would like to be installed:

Example Commands to install the qeq, meam, and reaxff LAMMPS packages:

```
> make yes-qeq
1
2 make yes-meam
3
4 make yes-reaxff
5
6
```

After installing the needed packages, the following command can confirm if the packages installed:

```
> make ps
1
```

Or if a list of installed packages is needed without showing the ones that are not installed, the following command can help:

```
> make pi
1
```

Once installing the needed packages, lammps would need to be rebuilt from the previous build steps to invoke the changes.

```
> make mpi
1
```

Once LAMMPS rebuilds with the new packages, the rebuilt `lmp_mpi` executable would need to be moved back to the `/bin` folder

```
> #Copy the lmp_mpi executable over to this new /bin folder
1 rsync -a --progress ~/software/lammps/lammps-23Jun2022/src/ln
2
3
```

If the `lammps/23Jun2022-mine` module is currently loaded, it would need to be unloaded and reloaded to have the changes take into effect.

```
> module unload lammps
1
2 module load lammps/23Jun2022-mine
3
4
```

The `Imp_mpi` MPI build above should automatically look for the `openmpi/4.1.4` HPC module.

The `openmpi/4.1.4` module would need to be loaded when the `Imp_mpi` executable is called.

Running LAMMPS with MPI would need specific MPI command line options to bypass issues when ignoring the old openib fabric and running on the new UCX framework.

These settings are automatically loaded when the `openmpi/4.1.4` module is loaded in the environment.

This feature of the module allows the options to not be specified in the `mpirun` or `mpiexec` commands.

The following example will show how to run the LAMMPS MPI executable by spawning 4 MPI threads and run the specified *lammpsInputFileHere* input file and save the data output to a file called *OutPutFileName.txt*

```
mpirun -np 4 Imp_mpi -in lammpsInputFileHere > OutPutFileName.txt
```

## Step 7: The updated openmpi/4.1.4 modules [🔗](#)

For extra information, the options that the **openmpi/4.1.4** and **openmpi/4.1.4-ics** modules set will be explained here (and are not needed when running the `Imp_mpi` command in a submission script):

```
mpirun --mca opal_warn_on_missing_libcuda 0 -mca pm1 ucx --mca btl  
^vader,tcp,openib,uct -x UCX_NET_DEVICES=mlx5_0:1 -np 4  
~/software/lammps/lammps-23Jun2022/bin/Imp_mpi -in  
restOfLammpsCommandHere
```

Lets break down the mpi command line options:

The `--mca opal_warn_on_missing_libcuda 0` option disables the CUDA warning message

The `-mca pm1 ucx` tells mpi to run using the UCX framework for the point-to-point message layer

The `--mca btl ^vader,tcp,openib,uct` tells MPI to not run on the vader,tcp,openib, or uct frameworks

The `-x UCX_NET_DEVICES=mlx5_0:1` option tells MPI to target the Infiniband card `mlx5_0` and connect to the card on port 1

The `-np 4` flag tells MPI to run on 4 MPI threads

The `~/software/lammps/lammps-23Jun2022/bin/Imp_mpi` path tells MPI to



run the `Imp_mpi` executable under the given path, but if there is a local module loaded, the path to the `Imp_mpi` executable would not be needed. Instead the `Imp_mpi` command can replace the full path above.

```
-np 4 Imp_mpi
```

The `-i restOfLammpsCommandHere` indicates the command line options needed for LAMMPS, in this case the `-i` refers to the LAMMPS input file for LAMMPS to process.

Or a different way to pass the input file and generate an output file from the given input, can be the following command:

```
Imp_mpi -in lammpsInputFileHere > OutPutFileName.txt
```

## Running LAMMPS

If possible, you should always first run your code on your local machine just to ensure that your code is correct. You can do it on a small dataset and a small configuration (single processor, etc.). This way you would be able to catch any errors not related to the cluster even before submitting your job.

Below we show a step-by-step example of how to run a simple LAMMPS simulation in the cluster. We have used one of the examples bundled with LAMMPS distribution, namely, `flow`.

## Copy your code and data to the cluster

We are assuming that you are using the terminal to copy your data. If you are using a GUI client, you should be able to do it in a visual way.

Open a terminal to [connect to the cluster](#) and create a directory for the experiment.

```
> mkdir lammptest && cd lammptest
1
```

The code from the example can be downloaded using the chunk of code below.

```
> ## cloning lammps project
1 git clone https://github.com/lammps/lammps.git
2 ## moving the files from the "flow" example to our working di
3 mv lammps/examples/flow* .
4 ## deleting the lammps project
5 rm -rf lammps
6
```

Now, your `lammptest` directory should have the following files

```

> ls -lt
1 total 40
2 -rw-r--r-- 1 netid domain users 5443 Nov 10 09:58 log.27Nov18
3 -rw-r--r-- 1 netid domain users 5444 Nov 10 09:58 log.27Nov18
4 -rw-r--r-- 1 netid domain users 5445 Nov 10 09:58 log.27Nov18
5 -rw-r--r-- 1 netid domain users 5441 Nov 10 09:58 log.27Nov18
6 -rw-r--r-- 1 netid domain users 1503 Nov 10 09:58 in.flow.poi
7 -rw-r--r-- 1 netid domain users 1505 Nov 10 09:58 in.flow.cou
8

```

## SLURM script

**SLURM** is the scheduler program for our cluster. In the cluster, we need to create a simple script that would tell **SLURM** how to run your job. For details see the [SLURM Guide](#).

You can either create this script in the terminal using any editor such as **nano**, or you can create it in your local machine and use the **scp** command to copy it into the cluster. We can put this script in the **lampstest** directory, and it would contain the following lines:


```

> $ cd ~/lampstest
1 $ cat lammps_job.sh
2
3 #!/bin/bash
4 #SBATCH -N 1
5 #SBATCH -n 36
6 #SBATCH --constraint='skylake'
7 #SBATCH -o lammps_sim_out-%J.txt
8 #SBATCH -e lammps_sim_out-%J.txt
9 #SBATCH --mail-type=ALL
10 #SBATCH --mail-user=user@uconn.edu
11 module load openmpi/4.1.4
12 mpiexec lmp_mpi < in.flow.couette
13

```

This script is telling how many processors we need as well as which files the output (and errors) should be written to. Basically, the lines starting with **#SBATCH** provide the switches for the **sbatch** command, which submits a job to **SLURM**. Note that we have told **SLURM** to email us at every event for this job such as **begin / queued / end / error** etc.

The last line is the command that would be run as the job. It invokes the **lammps** module with the input **~/lampstest/flow/in.flow.couette**.

 If you are using **mpirun** in your submission scripts, it is recommended to use the following command syntax

```

> mpirun lmp_mpi -in inputfile
1

```

The `lammps` documentation site mentioned this note if `mpirun` is being used with the `<` operator:

*The redirection operator `<` will not always work when running in parallel with `mpirun`; for those systems the `-in` form is required.*

## Submitting your job


Before you submit your job make sure that the `LAMMPS` module is loaded, as described in the first part of this guide. When you are ready, simply do the following:

```
> sbatch lammps_job.sh
1
```

## Checking output

When the job is done we would get email notifications. You can also check your job status using the `sjobs` command. We can check on the `lammps` output itself using `tail`:

```
> cat lammps_sim_out-JOBID.txt
1
```

 Note that you should replace `JOBID` with the id associated to the submitted job.