A Practical Demonstration of Running a First LAMMPS Simulation on Newton

David Keffer
Department of Materials Science & Engineering
University of Tennessee, Knoxville
date begun: January 27, 2016
date last updated: February 4, 2016

Table of Contents

I. Purpose of Document

The purpose of this document is to provide a complete and unambiguous demonstration of running a simple LAMMPS simulation on the Newton Cluster at the University of Tennessee. This document includes specific instructions and sample files.

II. Accessing Newton

II.A. Download and Install Necessary Software

You will need software to allow you to access the cluster remotely from your own computer. There are two software that you must download, a terminal emulator and a file transfer software. Some codes perform both these actions.

The Newton system website contains links free versions of this software from the following URL: https://newton.utk.edu/bin/view/Main/Workshop1Connecting.

In this example, I am going to use a windows machine. I download and install the PuTTY SSH (secure socket handler) client and the FileZilla SFTP (secure file transfer protocol) client.

The necessary login information for both software is given in Table 1.

hostname	login.newton.utk.edu
username	your UT NetID
password	your UT NetID password
remote port	22

Table 1. Login information for the Newton cluster.

II.B. Login to Newton

Use Putty to login to Newton. You are presented with a command line prompt. We are going to proceed using a series of commands, all of which are summarized in Table A in Appendix A.

- Locate where you are using the `pwd` command.
- Examine the contents of the current directory, using the `ls -al` command.
- Make a directory for this class, using the `mkdir mse614` command.
- Change to the newly created directory, using the `cd mse614` command.
- Make a directory within mse614 for this class, using the `mkdir example_01_lj` command.
- Change to the newly created directory, using the `cd example_01_lj` command.

II.B. Create LAMMPS Input File

The Linux environment has a default editor, `vi`. It is a line editor, which predates the existence of a computer mouse. While `vi` is cherished by programmers, many non-programmers find it a very onerous file editor. There are numerous `vi` tutorials online.

In the age of rapid file transmission, there is an acceptable and more user-friendly alternative, which we will use in this demonstration.

We will use the free software, Notepad++. This software can be downloaded from the following URL: https://notepad-plus-plus.org/ . We are going to open Notepad++ and create a new empty file titled `in_lj_v01.txt`. This will be our LAMMPS input file.

In Appendix B of this document is a sample input file to perform a simulation of a Lennard-Jones fluid. For the purpose of this demonstration, we are going to simply copy the entire contents of Appendix B, beginning with line "`# define units`" and ending with the line "`run 2000`", into our input file `in_lj_v01.txt`, opened in Notepad++. Once this file is saved, we have our input file.

Note that in the input file, all lines that begin with the "#" sign (formerly called a 'pound' sign and more recently called a 'hash' sign) are comments. They are not required for the code to run. They are included to help make the purpose of each line more clear. All blank lines are also unnecessary. Blank lines are included to make the organization of the input file more apparent to the eye. Any other line begins with a LAMMPS command and is followed by inputs for that particular command.

The LAMMPS manual is online at the following URL: http://lammps.sandia.gov/ . The easiest way to find out what a command means and what options are available to it is google. For example, googling "lammps units" directs us to the appropriate page within the lammps manual, http://lammps.sandia.gov/doc/units.html , with all of the relevant information on the units command.

II.C. Transfer Input File to Newton

In order to run the LAMMPS code, we must transfer the input file to the cluster. To this end, we use the FileZilla software. FileZilla requires the same login information as PuTTy, which was presented in Table 1.

Open Filezilla and have the local window correspond to the location where you have saved your input file. Navigate the remote window to the desired destination directory, `~/mse614/example_01_lj`. Transfer the input file.

The combination of Notepad++ and FileZilla allows us to edit any file locally in a user-friendly software and then transfer it immediately to the remote destination, where it can be used.

III. Running LAMMPS

III.A. Load LAMMPS Module

At the command line prompt, identify the available modules, using the `module avail` command. At the time of writing we observe that the most recent version of lammps that appears is in a module called `lammps/13May14`. We load the module using the following command: `module load lammps/13May14`.

III.B. Create Job File for the Queue

We are going to run the job through the queue. In a previous lecture, we discussed the contents of the job file. In this example we simply present the job file that we are going to use in Table 2. We will call this file `job.sge`. We can get this file onto Newton by copying the contents of

the appendix into a file in Notepad++ and then transferring it to the cluster using FileZilla, exactly as we did for the input file.

```
#$ -N Test
#$ -q short*
#$ -cwd
#$ -pe openmpi* 2
module load lammps/13May14
lmp < in_lj_v01.txt
```

Table 2. Job file for cluster queue.

In this job file, we are running LAMMPS (`lmp` command) with `in_lj_v01.txt` as the input file. We are running in the short queue with two processors.

III.C. Submit Job File for the Queue

We next submit the job file to the queue.

- Submit the job using the `qsub job.sge` command.
- Check the status of the job, using the `qstat` command.
- List the contents of the directory, using the `ls -al` command.

The content of the directory are now:

```
[dkeffer@sigma00 example_01_lj]$ ls -al
total 320
drwxr-xr-x 2 dkeffer cmrg   4096 Jan 27 12:13 .
drwxr-xr-x 3 dkeffer cmrg   4096 Jan 26 12:31 ..
-rw-r--r-- 1 dkeffer cmrg   2844 Jan 27 11:44 in_lj_v01.txt
-rw-r--r-- 1 dkeffer cmrg     99 Jan 27 12:10 job.sge
-rw-r--r-- 1 dkeffer cmrg   9058 Jan 27 12:13 log.lammps
-rw-r--r-- 1 dkeffer cmrg 290716 Jan 27 12:13 output.xyz
-rw-r--r-- 1 dkeffer cmrg      0 Jan 27 12:13 Test.e6397478
-rw-r--r-- 1 dkeffer cmrg   6278 Jan 27 12:13 Test.o6397478
-rw-r--r-- 1 dkeffer cmrg      0 Jan 27 12:13 Test.pe6397478
-rw-r--r-- 1 dkeffer cmrg      0 Jan 27 12:13 Test.po6397478
```

The file `log.lammps` contains a copy of the LAMMPS input file and the LAMMPS output. The file `Test.o6397478` contains the redundant output that would have been printed to the screen had it been an interactive job. The file `output.xyz` contains the atomic coordinates in an xyz format.

In this case, we transfer these file to our local directory for further examination using Notepad++.

III.D. Examining the Output of the LAMMPS Simulation

The file `Test.o6397478` contains the redundant output that would have been printed to the screen had it been an interactive job. This file is reproduced in Appendix C.

The most relevant information is the table of thermodynamic properties. There are three such tables corresponding to the NVE equilibration, the NVT equilibration and the NVT data production runs. The most relevant information is the table of thermodynamic properties from the NVT data production. This information is summarized in Table 3. This is a toy simulation.

For a longer simulation, the equilibration table would be studied to ensure that the system had actually reached equilibration. Once that was validated the data production run could be used to generate publishable data.

It is important to note that the thermodynamic properties reported in this table are instantaneous values. All analysis of averages and fluctuations must be performed as part of the post-processing. This analysis can be performed in any software of your choice (Excel or Matlab) simply by copying the relevant table into the software.

Step	PotEng	KinEng	TotEng	Temp	Press	Density
2000	-1821.594	756.59549	-1064.9985	0.98707827	-0.3111245	0.5
2100	-1805.8618	777.9501	-1027.9117	1.0149382	-0.1243004	0.5
2200	-1801.1839	797.4162	-1003.7677	1.0403342	-0.10290561	0.5
2300	-1780.5184	757.9926	-1022.5258	0.98890098	-0.19909778	0.5
2400	-1803.8874	788.00112	-1015.8863	1.028051	-0.30086952	0.5
2500	-1806.3675	770.407	-1035.9605	1.0050972	-0.23416908	0.5
2600	-1818.0796	770.41401	-1047.6656	1.0051063	-0.19240302	0.5
2700	-1821.7818	770.85111	-1050.9307	1.0056766	-0.34659051	0.5
2800	-1811.3153	782.00793	-1029.3074	1.0202321	-0.36878276	0.5
2900	-1818.8302	785.93459	-1032.8956	1.025355	-0.28715067	0.5
3000	-1841.4016	807.48686	-1033.9147	1.0534727	-0.24091218	0.5
3100	-1838.2069	843.90237	-994.30449	1.1009816	-0.04616304	0.5
3200	-1838.4076	774.9594	-1063.4482	1.0110364	-0.31697307	0.5
3300	-1842.646	772.30297	-1070.343	1.0075707	-0.15439713	0.5
3400	-1834.814	767.92509	-1066.889	1.0018592	-0.12814917	0.5
3500	-1850.819	793.94239	-1056.8766	1.0358022	-0.33874609	0.5
3600	-1849.1511	754.52376	-1094.6273	0.98437543	-0.26298136	0.5
3700	-1879.2309	769.83769	-1109.3933	1.0043545	-0.35762976	0.5
3800	-1865.7572	689.74132	-1176.0159	0.89985821	-0.34823303	0.5
3900	-1849.0607	732.9395	-1116.1212	0.95621591	-0.24712437	0.5
4000	-1861.6776	759.24558	-1102.4321	0.99053566	-0.36298469	0.5

Table 3. Table of Thermodynamic properties from the file `Test.o6397478` or `log.lammps`.

The file `output.xyz` contains the atomic coordinates in an xyz format. This too can be viewed in Notepad++, at least for a small file such as was generated in this example. This file contains 10,974 lines composed of 21 blocks of 514 lines each. The 514 lines compose a configuration including two header lines (the first reporting the number of atoms in the configuration, 512 in this case, and the second reporting the time associated with the configuration. The last 512 lines are the atomic symbol and coordinates of each atom. There are 21 frames in this file because we ran for 2000 steps and saved every 100 steps, resulting in 20 saved configurations. The additional configuration is the initial configuration at the beginning of the data production run. A sample of the contents of this file are shown in Table 4.

These configurations can be viewed independently or viewed as a movie using the visualization software of your choice. Virtually all visualization software can read this standard xyz format.

line #	file content
1	`512`
2	`Atoms. Timestep: 2000`
3	`Ar 6.68092 2.05039 2.30096`
4	`Ar 0.117642 7.8707 0.718617`
...	skipped 509 lines
514	`Ar 7.05901 6.59419 1.06683`
515	`512`
516	`Atoms. Timestep: 2100`
517	`Ar 6.1127 1.94847 1.98184`
518	`Ar 0.109745 8.72332 0.639257`
...	skipped 509 lines
1028	`Ar 6.8389 6.24782 0.519626`
1029	`512`
1030	`Atoms. Timestep: 2200`
1031	`Ar 5.98165 1.90162 1.44947`
...	and so on

Table 4. Table of coordinates from `output.xyz`.

IV. I Did Exactly What You Said and It Doesn't Work!

It has been demonstrated by students that it is possible to exactly follow the instructions in this handout and not meet with the expected results. To date, two culprits are responsible, which are described below. As others are found, they will be added to this list.

IV.A. Add an empty line to the end of all files

Linux and LAMMPS may not execute the final line of a file if it is not followed by a carriage return (new line). Therefore you need an empty line at the end of the file. This applies to both the `job.sge` file and the LAMMPS input file, which in this example is `in_lj_v01.txt`.

IV.B. Run the dos2unix command on all files

Different operating systems insert different invisible characters into the file to represent such things as tabs and carriage returns (end of line). The "end of line" character is particularly troublesome. The command `dos2unix` converts these invisible characters from the dos (Windows) default to the unix (Linux) default. It is executed as shown in Appendix A. If this problem impacts you, you will likely need to run this conversion command on both the `job.sge` file and the LAMMPS input file, which in this example is `in_lj_v01.txt`. The converted file replaces the old file.

Appendix A. Common Linux Commands

`cat` *file*	view contents of file
`cd` *dir*	change to a directory located in current directory
`cd ..`	change to the parent directory
`cp` *oldfile newfile*	copy a file from location to a new location
`dos2unix` *file.txt*	convert text file from dos to unix format
`lmp <` *inputfile*	run LAMMPS
`ls`	list contents of directory
`ls -al`	list contents of directory including hidden files
`man` *command*	access manual (help file) for linux command
`mkdir` *dir*	make a new directory
`module avail`	list available modules
`module load` *modulename*	load module
`pwd`	print working directory
`qsub` *jobfile*	submit a job to the queue
`qstat`	check on the status of your jobs in queue
`rm -i` *file*	remove a file permanently (There is no restore.)
`rmdir` *dir*	remove a directory
`vi` *file*	edit file

Table A. Useful Linux commands

Appendix B. Sample LAMMPS Input File

The following file is a sample LAMMPS input file to simulate a Lennard-Jones Fluid. For the purpose of this example, this file is named `in_lj_v01.txt`.

```
#  define units
units       lj

#  specify periodic boundary conditions
boundary p p p

#  define atom_style
#  full covers everything
atom_style  full

#  define simulation volume
#  If I want N = 512 atoms
#  and I want a density of rho = 0.5 atoms/lj-sigma^3
#  Then I can determine the size of a cube by
#  size = (N/rho)^(1/3)
variable side      equal 10.0793684
region      boxid block 0.0 ${side} 0.0 ${side} 0.0 ${side}
create_box  1 boxid

#  specify initial positions of atoms
#  sc = simple cubic
#  0.5 = density in lj units
lattice     sc 0.50


#  place atoms of type 1 in boxid
create_atoms    1 box

#   define mass of atom type 1
mass        1 1.0

#  specify initial velocity of atoms
#  group = all
#  reduced temperature is T = 1.0 = lj-eps/kb
#  seed for random number generator
#  distribution is gaussian (e.g. Maxwell-Boltzmann)
velocity    all create 1.0 87287 dist gaussian

#  specify interaction potential
#  pairwise interaction via the Lennard-Jones potential with a cut-off at 2.5 lj-sigma
pair_style  lj/cut 2.5


#  specify parameters between atoms of type 1 with an atom of type 1
#  epsilon = 1.0, sigma = 1.0, cutoff = 2.5
pair_coeff  1 1 1.0 1.0 2.5

# add long-range tail correction
pair_modify tail yes

#  specify parameters for neighbor list
#  rnbr = rcut + 0.3
neighbor    0.3 bin
```

```
#  specify thermodynamic properties to be output
#  pe = potential energy
#  ke = kinetic energy
#  etotal = pe + ke
#  temp = temperature
#  press = pressure
#  density = number density
#  output every thousand steps
#  norm = normalize by # of atoms (yes or no)
thermo_style custom step pe ke etotal temp press density

# report instantaneous thermo values every 100 steps
thermo 100

# normalize thermo properties by number of atoms (yes or no)
thermo_modify norm no

#  specify ensemble
#  fixid = 1
#  atoms = all
#  ensemble = nve or nvt
fix     1 all nve

#  define time step
timestep 0.005

# run 1000 steps in the NVE ensemble
# (this equilibrates positions)
run 1000

#  stop fix with given fixid
#  fixid = 1
unfix 1

#  specify ensemble
#  fixid = 2
#  atoms = all
#  ensemble = nvt
#  temp = temperature
#  initial temperature = 1.0
#  final temperature = 1.0
#  thermostat controller gain = 0.1 (units of time, bigger is less tight control)
fix     2 all nvt temp 1.0 1.0 0.1

# run 1000 steps in the NVT ensemble
# (this equilibrates thermostat)
run     1000

#   save configurations
#   dumpid = 1
#   all atoms
#   atomic symbol is Ar
#   save positions every 100 steps
#   filename = output.xyz
#
dump    1       all xyz 100 output.xyz
dump_modify 1 element Ar

# run 1000 more steps in the NVT ensemble
# (this is data production, from which configurations are saved)
run     2000
```

Appendix C. Sample LAMMPS Output File

The following file is the output file generated by LAMMPS. For the purpose of this example, this file is named `Test.o6397478`.

```
Warning: no access to tty (Bad file descriptor).
Thus no job control in this shell.
LAMMPS (13 May 2014)
Created orthogonal box = (0 0 0) to (10.0794 10.0794 10.0794)
  1 by 1 by 1 MPI processor grid
Lattice spacing in x,y,z = 1.25992 1.25992 1.25992
Created 512 atoms
Setting up run ...
Memory usage per processor = 5.52089 Mbytes
Step PotEng KinEng TotEng Temp Press Density
       0   -1736.2204        766.5   -969.72039            1   -2.1168568          0.5
     100    -1782.228    787.19826   -995.02974    1.0270036  -0.26892657          0.5
     200   -1820.0792    823.21778   -996.86139    1.0739958  -0.37460293          0.5
     300   -1803.9399    807.18409   -996.75579    1.0530777  -0.33098281          0.5
     400   -1797.5689    800.31689   -997.25205    1.0441186  -0.31234687          0.5
     500   -1786.2936    789.84338   -996.45025    1.0304545  -0.18389744          0.5
     600   -1804.5412    807.17997   -997.36127    1.0530724  -0.36063593          0.5
     700   -1818.1448    821.35838   -996.78644      1.07157  -0.30655591          0.5
     800   -1816.5595    819.69929   -996.86024    1.0694055  -0.34280504          0.5
     900    -1814.498    817.04784   -997.45017    1.0659463  -0.24186121          0.5
    1000   -1819.7291    822.24597   -997.48311    1.0727279  -0.25796404          0.5
Loop time of 0.301349 on 1 procs for 1000 steps with 512 atoms

Pair  time (%) = 0.211025 (70.0268)
Bond  time (%) = 9.05991e-05 (0.0300645)
Neigh time (%) = 0.0689821 (22.8911)
Comm  time (%) = 0.0128846 (4.27564)
Outpt time (%) = 9.03606e-05 (0.0299854)
Other time (%) = 0.00827622 (2.74639)

Nlocal:    512 ave 512 max 512 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Nghost:    1406 ave 1406 max 1406 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Neighs:    11676 ave 11676 max 11676 min
Histogram: 1 0 0 0 0 0 0 0 0 0

Total # of neighbors = 11676
Ave neighs/atom = 22.8047
Ave special neighs/atom = 0
Neighbor list builds = 100
Dangerous builds = 100
Setting up run ...
Memory usage per processor = 5.52089 Mbytes
Step PotEng KinEng TotEng Temp Press Density
    1000   -1819.7291    822.24597   -997.48311    1.0727279  -0.25796404          0.5
    1100   -1834.0435    798.54785   -1035.4956    1.0418106  -0.42770213          0.5
    1200   -1815.9261    787.50572   -1028.4204    1.0274047  -0.33194761          0.5
    1300   -1821.6024    773.44741    -1048.155    1.0090638  -0.13189678          0.5
    1400   -1836.4965    816.57257    -1019.924    1.0653262  -0.18939315          0.5
    1500   -1835.5686    788.19329   -1047.3753    1.0283017   -0.2739253          0.5
    1600   -1826.7094    759.49849   -1067.2109   0.99086561  -0.08375573          0.5
    1700   -1810.4568    744.76995   -1065.6869   0.97165029  -0.29001831          0.5
    1800   -1823.9328    740.44112   -1083.4917   0.96600276  -0.42210131          0.5
```

```
    1900   -1836.8422    745.76233   -1091.0799   0.97294499    -0.455944          0.5
    2000    -1821.594    756.59549   -1064.9985   0.98707827   -0.3111245          0.5
Loop time of 0.257463 on 1 procs for 1000 steps with 512 atoms

Pair  time (%) = 0.177809 (69.0618)
Bond  time (%) = 0.000102282 (0.0397267)
Neigh time (%) = 0.0588109 (22.8425)
Comm  time (%) = 0.0105119 (4.08287)
Outpt time (%) = 6.91414e-05 (0.0268549)
Other time (%) = 0.0101602 (3.94628)

Nlocal:    512 ave 512 max 512 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Nghost:    1379 ave 1379 max 1379 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Neighs:    11684 ave 11684 max 11684 min
Histogram: 1 0 0 0 0 0 0 0 0 0

Total # of neighbors = 11684
Ave neighs/atom = 22.8203
Ave special neighs/atom = 0
Neighbor list builds = 100
Dangerous builds = 100
Setting up run ...
Memory usage per processor = 6.56581 Mbytes
Step PotEng KinEng TotEng Temp Press Density
    2000    -1821.594    756.59549   -1064.9985   0.98707827   -0.3111245          0.5
    2100   -1805.8618     777.9501   -1027.9117    1.0149382   -0.1243004          0.5
    2200   -1801.1839     797.4162   -1003.7677    1.0403342  -0.10290561          0.5
    2300   -1780.5184     757.9926   -1022.5258   0.98890098  -0.19909778          0.5
    2400   -1803.8874    788.00112   -1015.8863     1.028051  -0.30086952          0.5
    2500   -1806.3675      770.407   -1035.9605    1.0050972  -0.23416908          0.5
    2600   -1818.0796    770.41401   -1047.6656    1.0051063  -0.19240302          0.5
    2700   -1821.7818    770.85111   -1050.9307    1.0056766  -0.34659051          0.5
    2800   -1811.3153    782.00793   -1029.3074    1.0202321  -0.36878276          0.5
    2900   -1818.8302    785.93459   -1032.8956     1.025355  -0.28715067          0.5
    3000   -1841.4016    807.48686   -1033.9147    1.0534727  -0.24091218          0.5
    3100   -1838.2069    843.90237   -994.30449    1.1009816  -0.04616304          0.5
    3200   -1838.4076     774.9594   -1063.4482    1.0110364  -0.31697307          0.5
    3300    -1842.646    772.30297    -1070.343    1.0075707  -0.15439713          0.5
    3400    -1834.814    767.92509    -1066.889    1.0018592  -0.12814917          0.5
    3500    -1850.819    793.94239   -1056.8766    1.0358022  -0.33874609          0.5
    3600   -1849.1511    754.52376   -1094.6273   0.98437543  -0.26298136          0.5
    3700   -1879.2309    769.83769   -1109.3933    1.0043545  -0.35762976          0.5
    3800   -1865.7572    689.74132   -1176.0159   0.89985821  -0.34823303          0.5
    3900   -1849.0607     732.9395   -1116.1212   0.95621591  -0.24712437          0.5
    4000   -1861.6776    759.24558   -1102.4321   0.99053566  -0.36298469          0.5
Loop time of 0.557075 on 1 procs for 2000 steps with 512 atoms

Pair  time (%) = 0.38108 (68.4072)
Bond  time (%) = 0.000164747 (0.0295736)
Neigh time (%) = 0.125276 (22.4881)
Comm  time (%) = 0.0218763 (3.927)
Outpt time (%) = 0.00712967 (1.27984)
Other time (%) = 0.021549 (3.86824)

Nlocal:    512 ave 512 max 512 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Nghost:    1413 ave 1413 max 1413 min
Histogram: 1 0 0 0 0 0 0 0 0 0
Neighs:    11781 ave 11781 max 11781 min
Histogram: 1 0 0 0 0 0 0 0 0 0
```

```
Total # of neighbors = 11781
Ave neighs/atom = 23.0098
Ave special neighs/atom = 0
Neighbor list builds = 200
Dangerous builds = 200
```