

## Using TAU to Identify Parallel Performance Issues

This exercise will give you experience using performance tools, in this case TAU, to identify parallel performance issues in production-level scientific applications. Although the details differ between different tools, the **process** of performance profiling is similar. The intent of this exercise is to help you become familiar with that process.

If you wish, you are encouraged to use your own scientific application of interest for this exercise. Otherwise, you will work with the LAMMPS molecular dynamics code developed primarily by Steve Plimpton at Sandia National Laboratories (<http://lammps.sandia.gov/>).

**Preparation:** The code and other materials that you need for this exercise are located at:

[http://staff.psc.edu/blood/petascale\\_school\\_2010/LAMMPS\\_QueenBee.tgz](http://staff.psc.edu/blood/petascale_school_2010/LAMMPS_QueenBee.tgz)

Login to the machine of your choice and use `wget` to grab the tarball. Unpack the tarball and read the brief README file in the LAMMPS directory. There are some scripts there to help you set up your environment, compile, and run jobs instrumented with TAU.

### Steps:

1. Compile the LAMMPS code and do a scaling study from 1 to 64 cores to get the baseline performance for this code before instrumentation with TAU. The relevant timing in LAMMPS is the time for the MD loop to complete. You can find this time in seconds by running:

```
grep "Loop time" log.lammps
```

Record your results so you can compare them with the instrumented version.

2. Recompile LAMMPS with TAU instrumentation. Use a Makefile that includes `mpi` and `pdt` (do not include “phase”). Look in the directory given in the example build script for the necessary TAU Makefiles. Run the code again (only one run is necessary, but make sure it is over at least several cores) to measure the overhead of full instrumentation. Compare the timing of this run to the original. You will use this run to create a selective instrumentation file to eliminate most of the overhead.
3. Go to the directory where you ran the instrumented LAMMPS. You should see a number of `profile.*` files corresponding to the number of cores over which you ran. Make sure you have an X server running on your machine (and that you connected via `ssh -X`), and execute the command `paraprof`. Using `paraprof`, create a selective instrumentation file as demonstrated in the lecture. Use this selective instrumentation file to rebuild LAMMPS. Also include `-optKeepFiles` in your `TAU_OPTIONS`. Use the same TAU

Makefile as before and redo the scaling study. This time, however, set the TAU\_CALLPATH variable to 1 and TAU\_CALLPATH\_DEPTH to 50 to collect callpath information.

4. Compare the timings of your selectively instrumented runs to your original runs. The amount of overhead should be reduced. If there is still substantial difference in the timings (greater than 10%) then we will need to assess whether the timings of major functions reported by TAU seem reasonable. You may need to do several runs with and without TAU to determine whether differences are due to TAU instrumentation or random variability between runs.
5. Once you have instrumented runs with low overhead you can use paraprof to explore the profiling data that you have collected. Remember you can use “paraprof --pack” to merge all of your profile\* files into a single .ppk file. You can copy this back to your workstation to use the TAU GUI analysis tools locally. Use paraprof to look at the callpath information for LAMMPS and find the hotspots in the code and the path to the main loop. Some of the instrumented routines are part of the startup or shutdown process. We are only interested in the time corresponding to the MD loop.
6. To isolate the timings that you are interested in, you can use the callpath information and your knowledge of the code to define execution phases. For the purposes of this exercise, this has been done for you. Paste the contents of select\_MDphase.tau (in your LAMMPS directory) at the end of your selective instrumentation file. Rebuild LAMMPS with the new selective instrumentation file, but this time choose a TAU\_MAKEFILE that includes mpi, pdt, and **phase**. Rerun your scaling study with the new executable. This time do *\*not\** set the TAU\_CALLPATH variables in your batch script (phase and callpath profiling interfere with one another in TAU).
7. As demonstrated in the lecture, look at the profiles for the functions in the MD phase. If possible, you will probably want to create \*.ppk files (via “paraprof --pack”) and copy these back to your own machine to look at them with your local version of TAU. For each run, use the filter in the function legend to filter out all functions that are not labeled with the name of the MD phase. For each run, add the mean time to the comparison window. Look at how the functions perform as you increase the number of processors. Also, look at the performance of individual functions across all the processors at high processor counts. Determine which functions are limiting the parallel efficiency of the code and what the problems are.

8. Time permitting, create a selective instrumentation file to include only the top 15 or so major routines from the MD phase. To find these routines, look at a profile that had the TAU\_CALLPATH environment set. You'll want to look at a profile that was done at a high processor count so you see which routines are the time-consuming routines at scale. Don't forget to include routines that **call** the time-consuming routines, especially those that call time-consuming MPI routines. Now rebuild and rerun LAMMPS at the higher processor counts with tracing enabled rather than profiling. You will probably want to edit the "jac.in" input file so that "run" is set to 10-50 steps or so. Use Jumpshot to look at the trace and get a feel for the computation and communication patterns in the code. You may want to check out the "cheat sheet" here:  
<http://www.psc.edu/general/software/packages/tau/TAU-quickref.pdf>

You should also be able to clearly identify from the trace the scaling problems that you saw by looking at the profiles.