Multithreaded Optimisation
Exercise Notes

Download and extract the exercise files

Firstly, change directory to make sure you are on the `/work` filesystem on ARCHER.

```
guestXX@archer:~> cd /work/y14/y14/guestXX/
```

`/work` is a high performance parallel file system that can be accessed by both the frontend and compute nodes. **All jobs on ARCHER should be run from the `/work` filesystem.** ARCHER compute nodes cannot access the `/home` filesystem at all: any jobs attempting to use `/home` will fail with an error.
Copy the exercise files archive on ARCHER from `/home/z01/shared` (or use `wget`) and unpack it with the commands

```
cp /home/z01/shared/tpo.tar .
tar xvf tpo.tar
```

1 Exercise

The `TPO/*/MolDyn/` directory contains a not very efficient parallel version of the molecular dynamics code. Build the code using the supplied Makefile.
Now load the performance tools modules with

```
module load perftools
```

Instrument the executable for OpenMP profiling using

```
pat_build -g omp md
```

This will produce a new binary called `md+pat`. Run this using 4 threads. This will produce a file with a `.xf` extension. Generate the profile using

```
pat_report md+pat+???????.xf
```

Now modify the code so that it uses a better loop schedule and/or atomic updates instead of `CRITICAL` — does the performance improve? Redo the profiling.

Extra exercise (Fortran only)

Next modify the code so that it uses an array of locks instead. Write your code so that you can associate more than one particle with a lock.

Extra exercise (Fortran only)

Use a reduction clause for `f` instead. Compare the performance with the other versions.