



KNL Performance Comparison: *GaitSym*

March 2017

1. Compilation, Setup and Input

Compilation

This code has currently not been released but is available on ARCHER at:

```
/home/b16/b16/wisb16/Unix/cvs/UnifiedBootstrap2017
```

On Xeon I use the g++ environment and the -static -O3 flags; on KNL I use -static -O3 again with g++

Setup

A typical run on KNL has been 512 cores, 4 threads per core (aprun -n 2048 -j 4). I'm using select=8:aoe=quad_100 options. A typical run on Xeon is bigger. Typically 3072 cores with hyperthreading enabled (aprun -n 6144 -j 2). I'm using select=128.

Input

This isn't terribly straightforward because there are multiple input files. You can find typical examples in:

```
/home/b16/b16/wisb16/work/05-03a_giraffatitan_archer  
/home/b16/b16/wisb16/knl_work/05-03a_giraffatitan_knl
```

2. Performance Data

I'm afraid I haven't been terribly systematic about this but for the settings given about a run of 1 million iterations takes 41,952 s on KNL and 1,842 s on Xeon. This means that the Xeon is 22.8 times faster when running on 6 times as many cores, so the performance per core is 3.80 times higher on the Xeon. That seems moderately believable to me given the relative processor speeds. My code should run very well on KNL because the memory footprint is small and it is mostly CPU bound.

3. Summary and Conclusions

My feeling is that I get very similar performance on both systems. It is slower per KNL core but I think that is what I would expect given the clock frequency of the different technologies. In the end it all depends on what the relative pricing of the different systems is.