OSU

CS475

Project 5

Author
By JAKE GOODWIN

Email:
goodwjak@oregonstate.edu

May 18, 2022

1 Introduction

This lab write up describes the performance of a numeric integration problem when using a varying number of nodes in addition to changing the number of threads.

2 Hardware

This was ran on the the OSU Rabbit server. I ran it on the rabbit server because I know that it's currently working. I have also included the up time from the rabbit server below.

- CPU: Intex Xenon E5-2630
- Speed: 2.40Ghz
- Max Speed: 3.20Ghz

```
17:40:16 up 154 days,  4:15, 62 users,  load average: 5.86, 9.92, 12.18
```

3 Tables and graphs

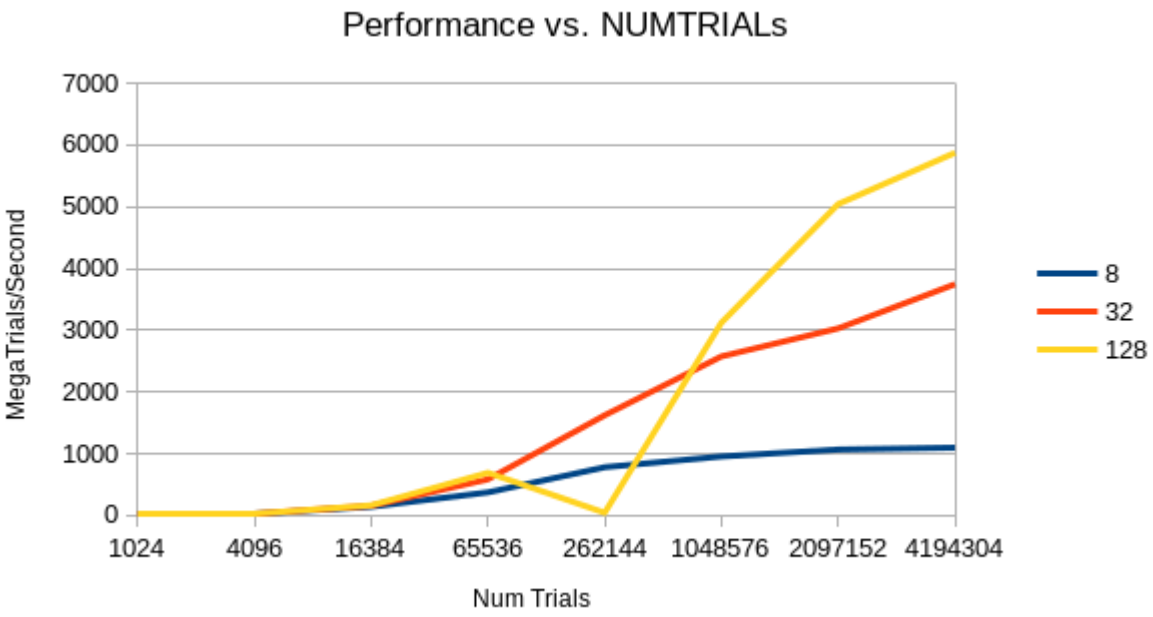

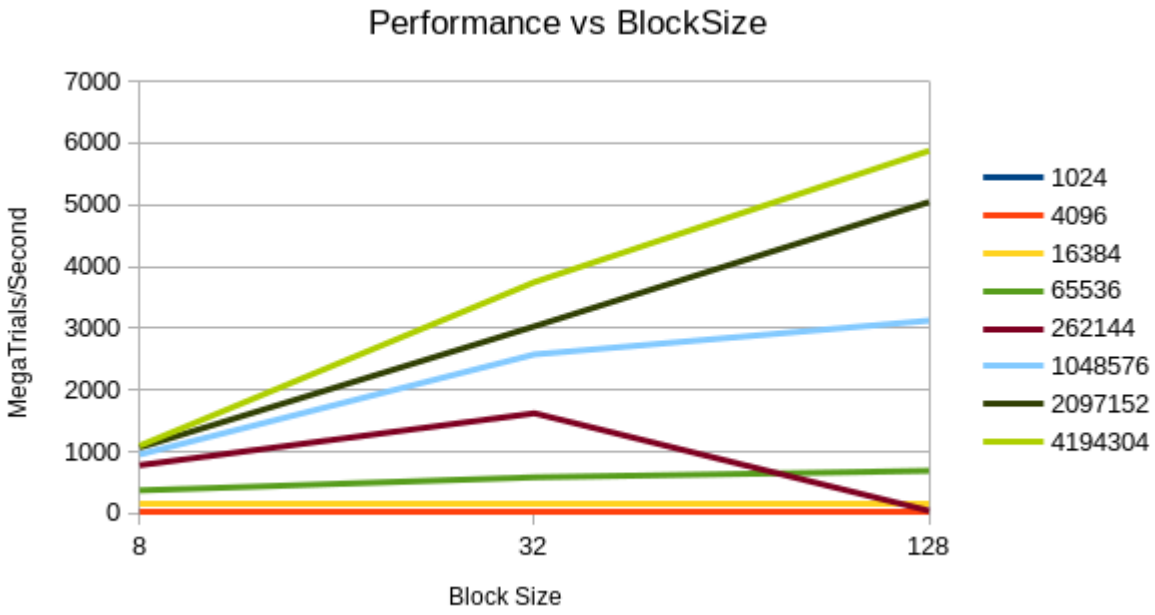

4 Performance patterns

For the second graph where each of the lines on the graph represents the number of trials, it appears that all of them benefit from higher block sizes until 32 at which point it starts to change for a lot of them.

5 Why of patterns

I think that around 32 blocks for a lot of the trial sizes the overhead starts to reduce the performance of the MegaTrials per second.

6 Blocksize hypothesis

7 A block size of 8 severely limits how much can be done at once. I think it prevents full usage of the hardware.

8 project 1 and 5 comparison

The performance using cuda made it to almost 6000 MegaTrials per second compared to my project 1 results of approximately 140 using a 4 core cpu.

9 Summary of GPU in parallel computing

For programs that have a very good parallel fraction and a large data set GPUs help improve performance by large margins.