

CPEG 421/621: Project Assignment: GPGPU vs. Programmable multicore

Issue Date: Dec. 2nd, 2010

Due Date: Dec. 17th, 2010

1. Introduction

In this assignment, you are going to study the two popular high performance architecture classes: GPGPU and general programmable multicore. A representative GPGPU architecture is Nvidia's Fermi; two representative programmable multicore architectures are IBM Cyclops64 and Intel Knights Ferry.

Each study group will focus on one of the three specific architectures.

- Group1 will focus on Nvidia Fermi;
- Group2 will focus on IBM Cyclops64;
- Group 3 will focus on Intel Knights Ferry.

In addition, each group should advocate your architecture – behaving like its technical manager.

You are supposed to read the given materials, do some survey, answer the given questions, and collect high quality articles you believe are good. A final report is required to be submitted by each group. Features of the specific architecture class and the chip listed above should be studied in sufficient detail, including architecture, programming model, system software stack, performance, power consumption, resilience, etc.

The length of the report should not be less than 20 pages, using the Springer LNCS style [1] which is strongly encouraged to follow.

At the end, a list of articles/links are provided which you are supposed to read before answering the questions

2. Questions

Please answer the following questions:

Question1: Please locate/read/survey at least 10-15 good articles. List them and use them in your survey. Also, construct an annotated bibliography.

Question2: Each group should write clearly why you think "your architecture" will win the competition. You should advocate "your architecture" in multiple ways, such as architecture, programming model, software technology, performance, power consumption, resiliency, etc.

3. Ethics for Team Work

It is expected that each member of the group should read all the papers sent out, discuss thoroughly with each other. And each member should contribute a fair amount to the report.

4. References

- [1] <http://www.springer.com/computer/lncs?SGWID=0-164-0-0-0>
- [2] http://sc10.supercomputing.org/schedule/event_detail.php?evid=pan130
- [3] <http://www.hpcwire.com/features/Should-I-Buy-GPGPUs-or-Blue-Gene-106708953.html>
- [4] http://www.nvidia.com/object/fermi_architecture.html
- [5] <http://en.wikipedia.org/wiki/Cyclops64>
- [6] http://www.pcworld.com/businesscenter/article/197620/intel_unveils_new_server_chip_with_32_cores.html
- [7] http://download.intel.com/pressroom/archive/reference/ISC_2010_Skaugen_keynote.pdf
- [8] Debunking the 100X GPU vs. CPU Myth: An Evaluation of Throughput Computing on CPU and GPU, by Victor W Lee, Changkyu Kim, etc, ISCA10.
<http://portal.acm.org/citation.cfm?id=1815961.1816021>
- [9] Understanding throughput-oriented architectures, By Michael Garland and David B. Kirk, Communications of ACM.
<http://portal.acm.org/citation.cfm?id=1839694>