Development of a GPU-based High-Performance Radiative Transfer Model for the High-spectral Resolution Infrared Sounders

Poster #3

Bormin Huang, Jarno Mielikainen, Hyunjong Oh** and Allen Huang*

CIMSS/SSEC, University of Wisconsin-Madison

*Visiting scholar from Kuopio University, Finland

**Visiting scholar from KMA, Korea

2nd IASI International Conference

25-29January 2010

Sevrier (Haute-Savoie), France

Graphic Processing Unit (GPU)

- Driven by the insatiable market demand for real-time, high-definition 3D graphics, the programmable GPU has evolved into a highly parallel, multithreaded, multicore processor.
- nVidia and ATI (now AMD) are two major manufacturers.

100 teraflop data center:

➢1429 4-CPU servers:

Cost $3.1M & consume 571 KW

Or

➢25 4-GPU servers:

Cost $310K & consume 27KW

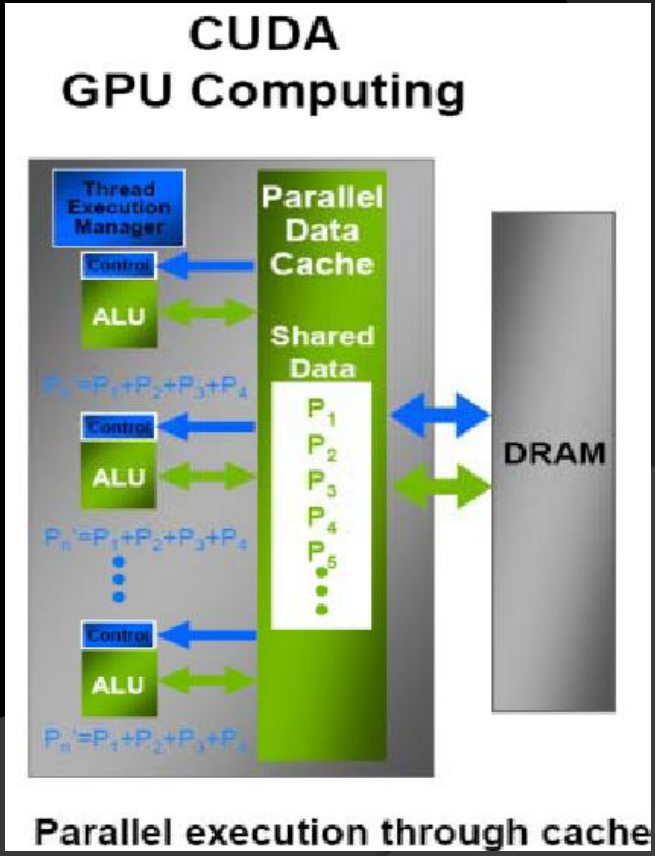

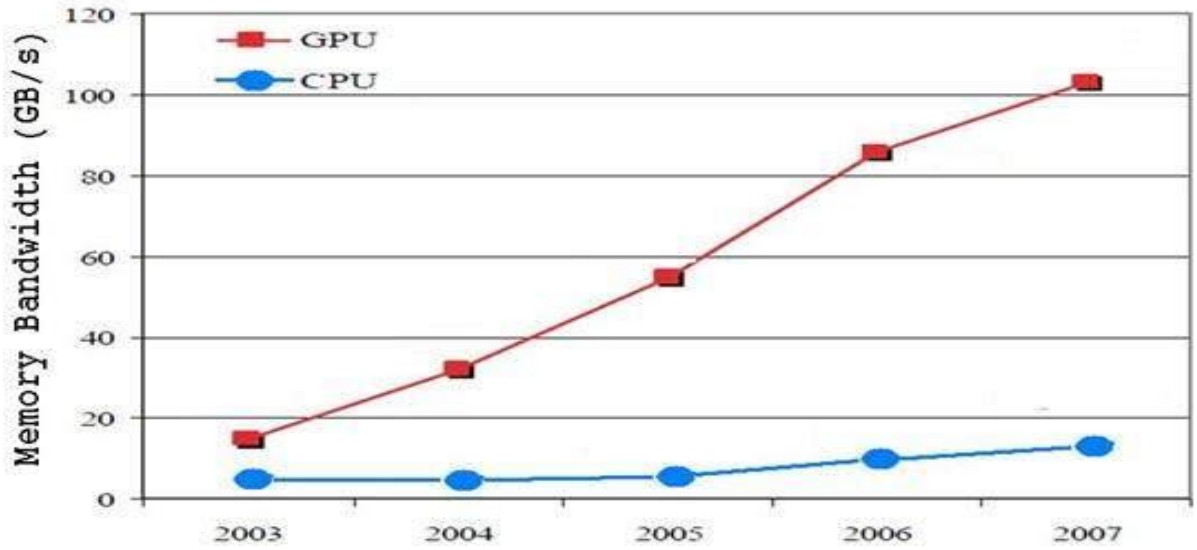

GPU is better than CPU in

- **Performance (GFlops): >250 times**
- **Memory Bandwidth (GB/s): > 5 times**

Our GPU forward model is running on a low-cost personal super computer (~US$7000).
It has a quad-core 2.4 GHz AMD CPU, and 4 Nvidia Tesla 1.3 GHz GPUs with total 960 cores.

Form Factor	10.5" x 4.376", Dual Slot
# of Tesla GPUs	1
# of Streaming Processor Cores	240
Frequency of processor cores	1.3 GHz
Single Precision floating point performance (peak)	933
Double Precision floating point performance (peak)	78
Floating Point Precision	IEEE 754 single & double
Total Dedicated Memory	4 GB GDDR3
Memory Speed	800MHz
Memory Interface	512-bit
Memory Bandwidth	102 GB/sec
Max Power Consumption	187.8 W
System Interface	PCIe x16
Auxiliary Power Connectors	6-pin & 8-pin
Thermal Solution	Active fan sink
Software Development Tools	C-based CUDA Toolkit

Our current benchmark on the $7,000 GPU:

- For the AIRS sounder with 2378 channels: our GPU-based forward model in CUDA obtains 731X speedup over its original Fortran code running on 1-core CPU (0.260 sec).

- For the IASI sounder with 8461 channels: our GPU-based forward model in CUDA obtains 1523X speedup over its original Fortran code running on 1-core CPU (1.003 sec).

- Using the 4-core CPU with OpenMP, the Fortran code speedup is only 2.21x over its 1-core CPU for the AIRS forward model, and only 2.56x for the IASI forward model.

Paper Submitted to Journal of Computational Physics

Development of a GPU-based High-Performance Radiative Transfer Model for the Infrared Atmospheric Sounding Interferometer (IASI)

Future Work

1. **Developing the GPU-based high-performance RTTOV forward model for NWP SAF/Met Office.**
2. **Developing the GPU-based high-performance WRF model for SeaSpace**
3. **Developing the GPU-based high-performance CRTM model for JCSDA.**
4. **Developing the GPU-based full spectrum physical retrieval system.**

Poster #3

Development of a GPU-based High-Performance Radiative Transfer Model for the High-spectral Resolution Infrared Sounders

Poster #3