Benchmarking NWP Kernels on Multi- and Many-core Processors

John Michalakes, NCAR
Manish Vachharajani, University of Colorado

AGU December 16, 2008

NCAR

Colorado
University of Colorado at Boulder

Motivation

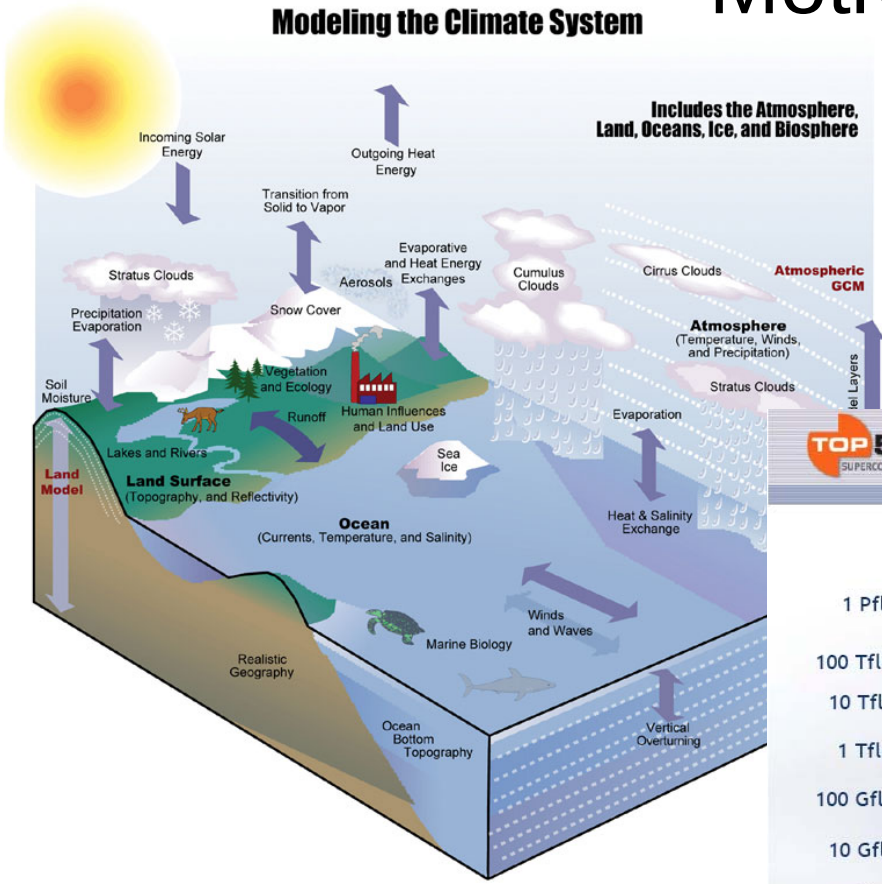

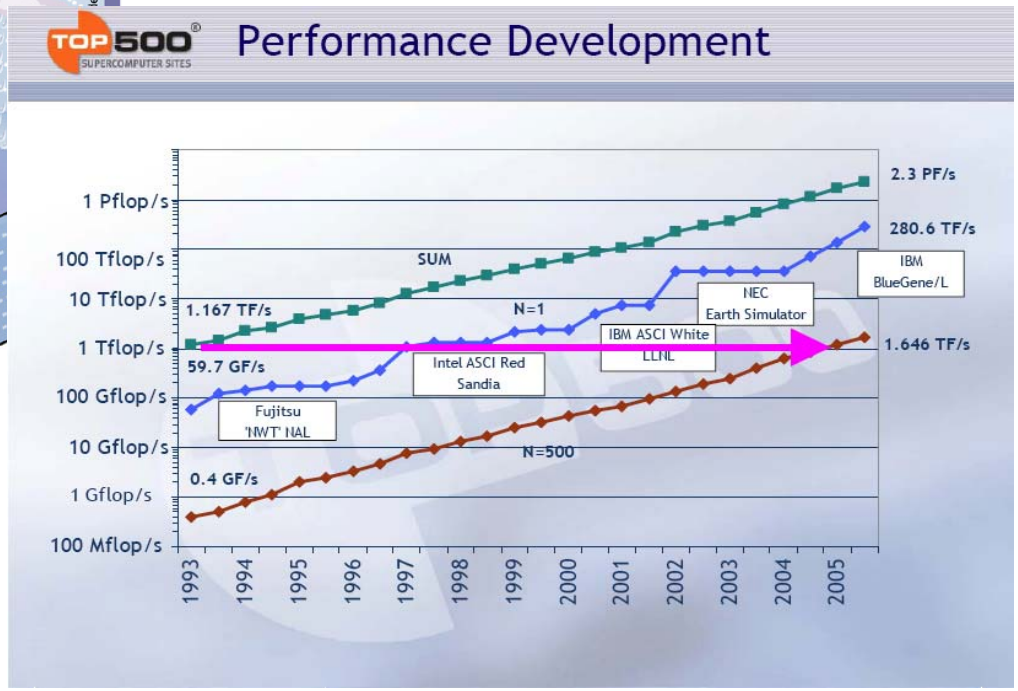

"At this rate, $10^{10} - 10^{12}$ improvement will take 40 years" – Rich Loft, NCAR

Hardware Acceleration

- Recover parallelism being wasted
- Multi-core and many-core
  - Higher core-count CPUs
  - Micro-SIMD and Vector instructions
  - Graphics Processing Units (GPUs)
  - Cell BE
  - FPGA
- Metrics: Flops, Flops/dollar, Flops/watt
- How do we judge effectiveness for earth system applications?

Application Benchmarks

- Objectives
  - Provide basis for comparing and evaluating hardware
  - Develop understanding of factors affecting performance
  - Feed back kernels into source models
- Kernel benchmarks
- Whole code benchmarks
- Make codes and results publicly available

WRF Application Benchmarks

- WSM5 cloud microphysics
  - Original, GPU, & multi-core
- 5th Order Positive-definite advection/diffusion
  - Original, GPU
- Radiation (LW and SW)
  - Original only
- KPP chemistry kinetics
  - Under development for GPU, Cell BE (VA. Tech)

NCAR

Colorado
University of Colorado at Boulder

WSM5 Microphysics

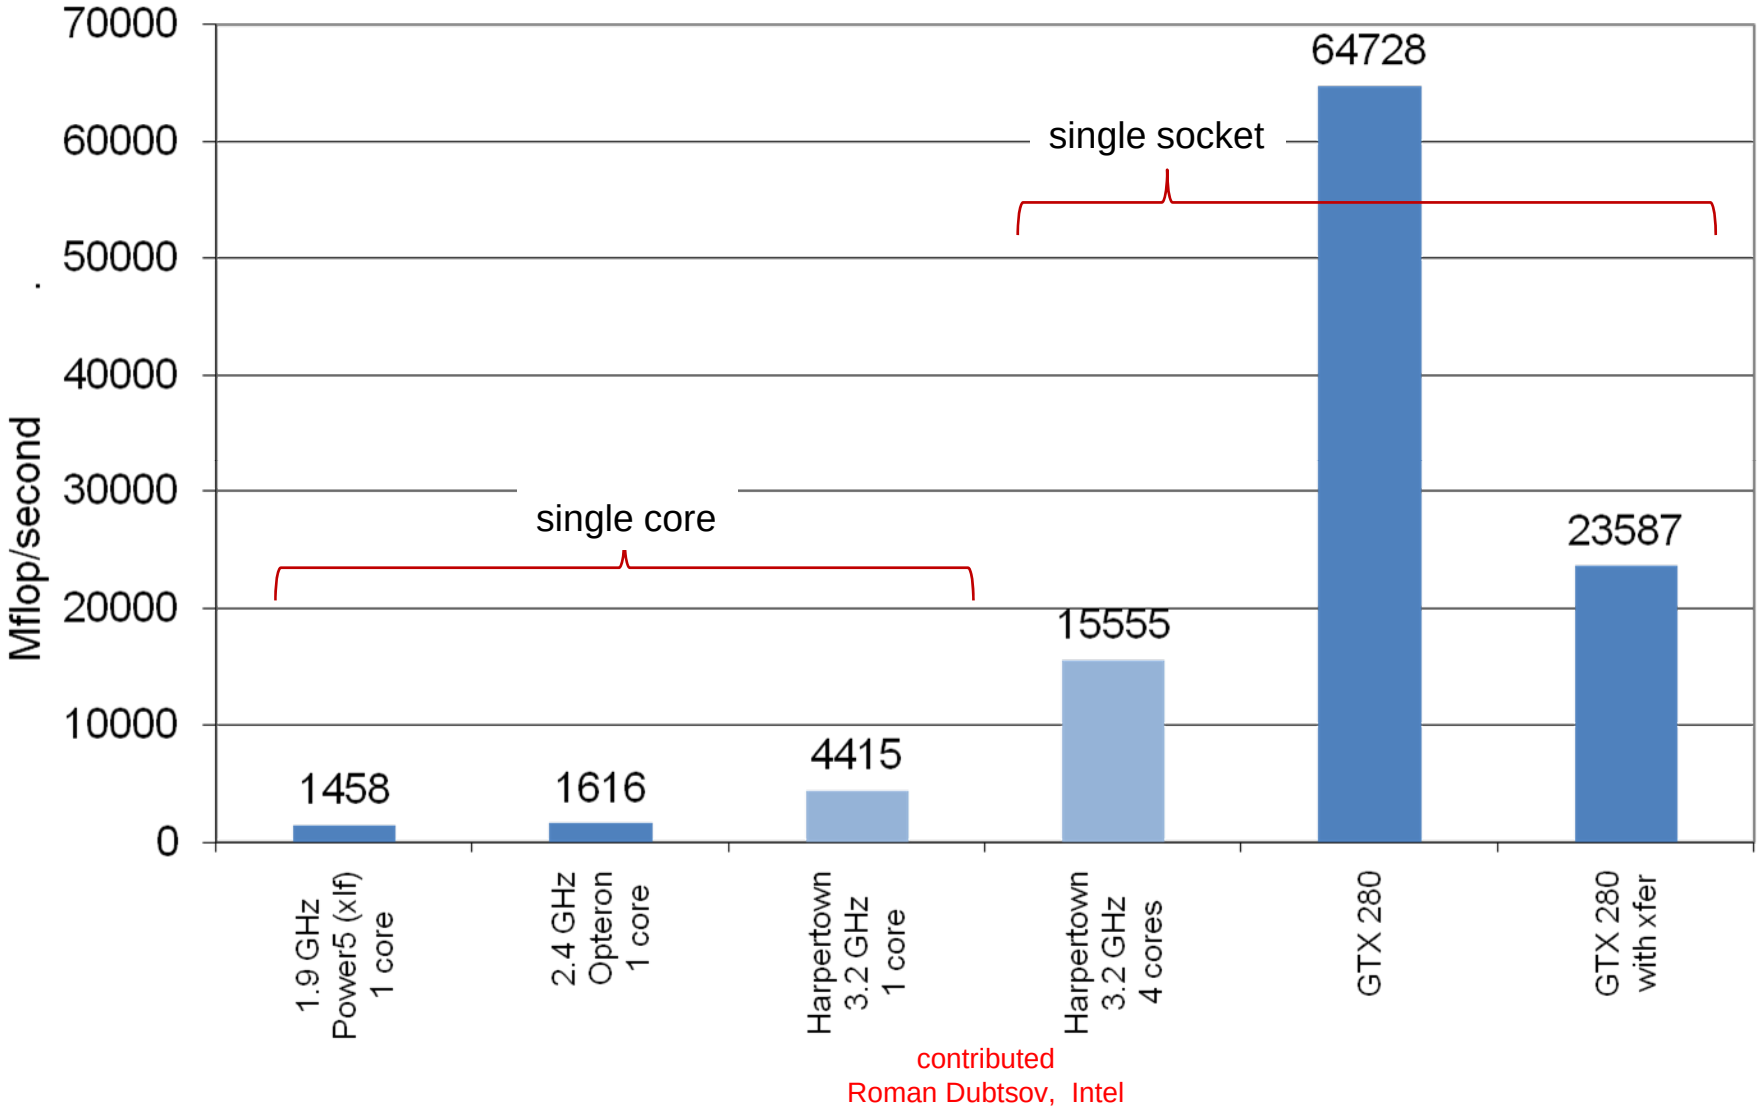

contributed
Roman Dubtsov, Intel

Advection/Diffusion

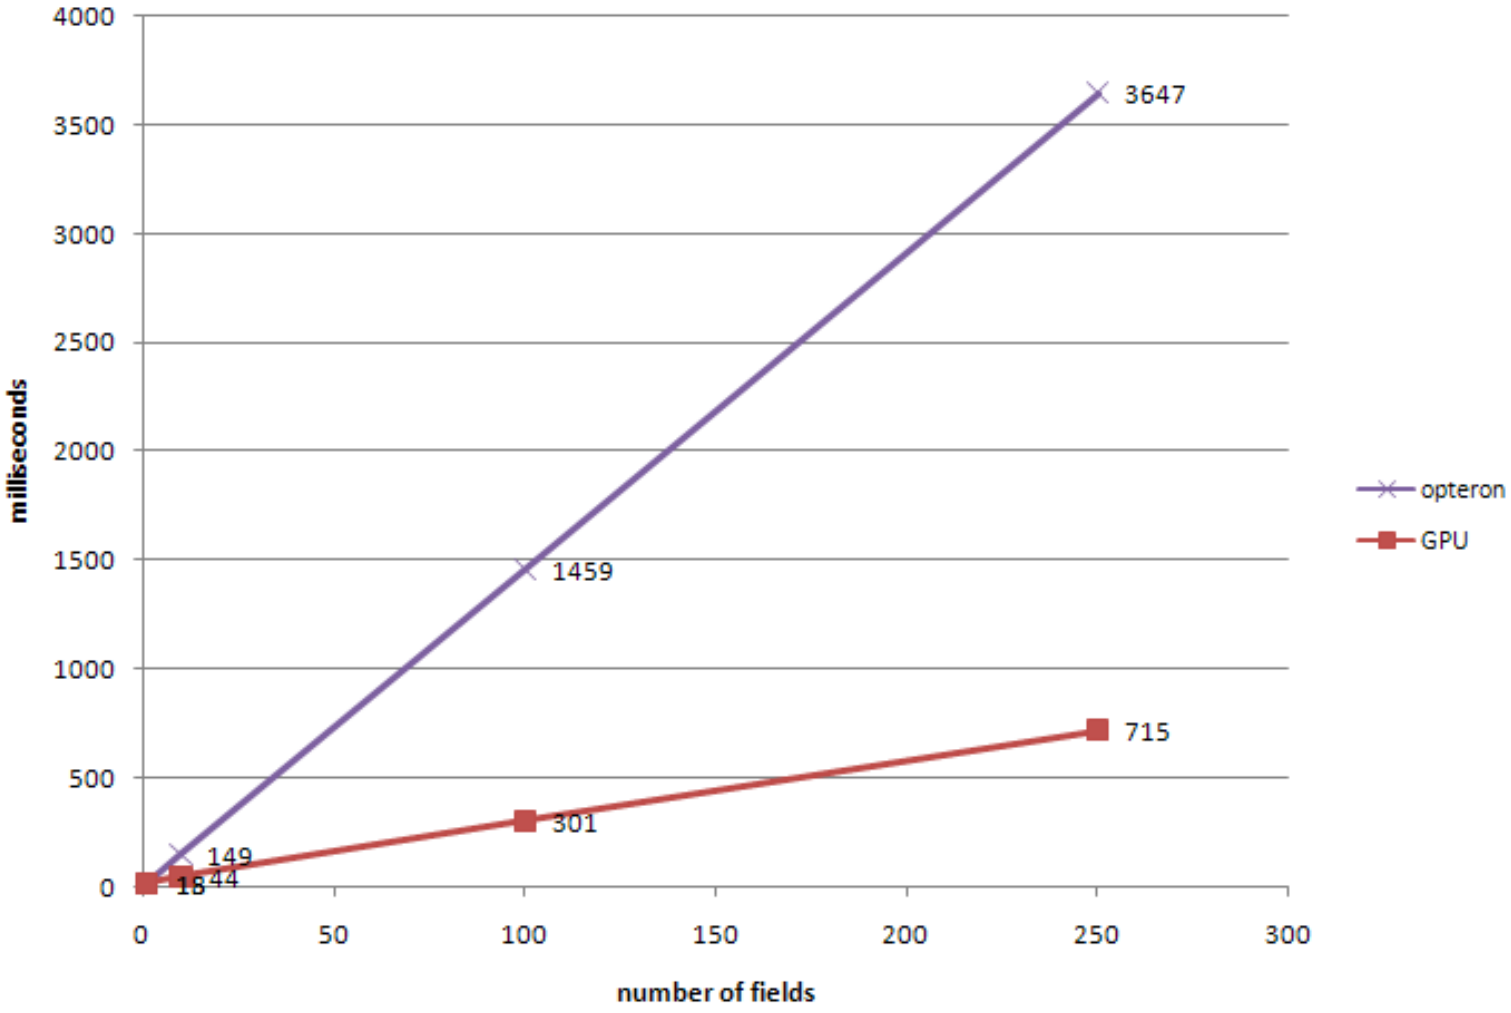

Chemistry (still working on this one)

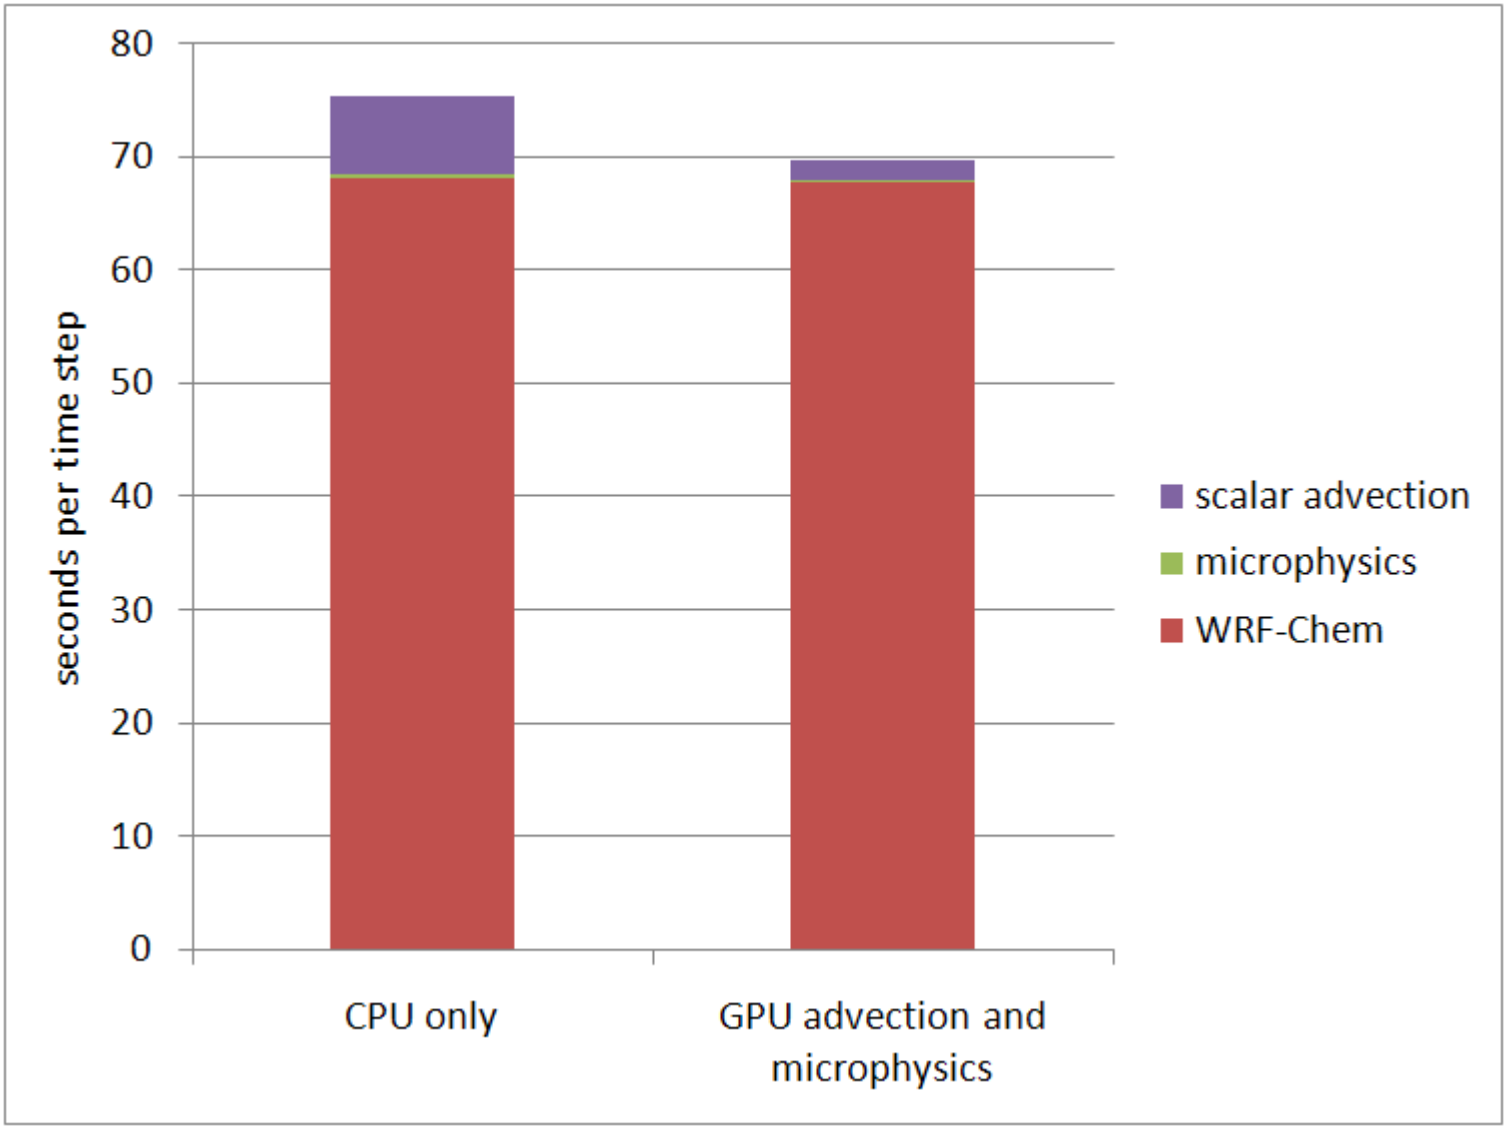

Computational Footprints

General[1]	WSM5	Adv./Diffus.	Chemistry[2]
FP operations	2702	301	609,226
Load/Stores	1224	396	1,021,227
Ops. per word	2.2	0.76	0.600
Ops. per byte	0.55	0.19	0.075

CUDA on GPU	WSM5	Adv./Diffus.	Chemistry
Registers per thread	60	40	84
Occupancy	17%	28%	19%
Speedup	40[3]	4[4]	2[4]

[1] per cell per step
[2] double precision
[3] GTX280 vs. single core 2.4 GHz AMD Opteron
[4] Tesla vs. single core 2.83 GHz Intel Xeon

Summary

- Ongoing work
  - Whole code benchmarks
  - Additional kernels, other models
  - Message passing from device memories
- Other information
  - Michalakes, J., M. Vachharajani. GPU Acceleration of Numerical Weather Prediction. *Parallel Processing Letters*. Volume 18, Issue: 4 (December 2008). World Scientific. pp 531-548.
  - http://www.mmm.ucar.edu/wrf/WG2/GPU

NCAR

Colorado
University of Colorado at Boulder