Massively-Parallel Heterogeneous Computing for Solving Large Problems



Wen-mei Hwu

Professor and Sanders-AMD Chair, ECE, NCSA University of Illinois at Urbana-Champaign

With contributions from

Mert Hidayetoglu, Abdul Dakkak, Carl Pearson, and Simon Garcia







A major paradigm shift

- In the 20th Century, we were able to understand, design, and manufacture what we can measure
 - Physical instruments and computing systems allowed us to see farther, capture more, communicate better, ...





A major paradigm shift

- In the 20th Century, we were able to understand, design, and manufacture what we can measure
 - Physical instruments and computing systems allowed us to see farther, capture more, communicate better, understand natural processes, control artificial processes...
- In the 21st Century, we are able to understand, design, and create what we can compute
 - Computational models are allowing us to see even farther, going back and forth in time, learn better, test hypothesis that cannot be verified any other way, ...





Examples of Paradigm Shift

20th Century

- Small mask patterns
- Electronic microscope and Crystallography with computational image processing
- Anatomic imaging with computational image processing
- Optical telescopes
- Teleconference
- GPS

ECE ILLINOIS

.

- Optical proximity correction
- Computational microscope with initial conditions from Crystallography
- Metabolic imaging sees disease before visible anatomic change

21st Century

- Gravitational wave telescopes
- Tele-emersion augmented reality
- Self-driving cars



What is powering the paradigm shift?

- Large clusters (scale out) allow solving realistic problems
 - 1.5 Peta bytes of DRAM in Illinois Blue Waters
 - E.g., 0.5 Å (0.05 nm) grid spacing is needed for accurate molecular dynamics
 - interesting biological systems have dimensions of mm or larger
 - Thousands of nodes are required to hold and update the grid points.
- Fast nodes (scale up) allow solution at realistic time scales
 - Simulation time steps at femtosecond (10⁻¹⁵ second) level needed for accuracy
 - Biological processes take milliseconds or longer

- Current molecular dynamics simulations progress at about one day for each 100 microseconds of the simulated process.
- Interesting computational experiments take weeks (used to be months)



What types of applications are demanding computing power today?

- First-principle-based models
 - Problems that we know how to solve accurately but choose not to because it would be "too expensive"
 - High-valued applications with approximations that cause inaccuracies and lost opportunities
 - Medicate imaging, earthquake modeling, weather modeling, astrophysics modeling, precision digital manufacturing, combustion modeling,
- Applications that we have failed to program
 - Problems that we just don't know how to solve
 - High-valued applications with no effective computational methods
 - Computer vision, natural language dialogs, stock trading, fraud detection, ...

ILLINOIS

2013: Blue Waters and Titan Computing Systems

System Attribute	NCSA Blue Waters	ORNL Titan
Vendors Processors	Cray/AMD/NVIDIA Interlagos/Kepler	Cray/AMD/NVIDIA Interlagos/Kepler
Total Peak Performance (PF) Total Peak Performance (CPU/GPU)	12.5 7.1/5.4	27.1 2.6/24.5
Number of CPU Chips Number of GPU Chips	49,504 4,224	18,688 18,688
Amount of CPU Memory (TB)	1600	584
Interconnect	3D Torus	3D Torus
Amount of On-line Disk Storage (PB)	26	13.6
Sustained Disk Transfer (TB/sec)	>1	0.4-0.7
Amount of Archival Storage	300	15-30
Sustained Tape Transfer (GB/sec)	100	7

ILLINOIS



Next Wave of Machines

• To be deployed in 2018-2019

ECE ILLINOIS

- The NSF leadership machine will be 2-3X of Blue Waters in computing throughput
- This does not increase the problem size much if your algorithm has complexity of O(N²) or even O(N log(N))
- The nodes will be limited by the memory bandwidth for most lowcomplexity algorithms

Most Blue Waters Job use less than or equal to 4,000 nodes (1/6 of the machine) today.



We know what we want but don't know how to build it.



Slide courtesy of Steve Oberlin, NVIDIA

ECE ILLINOIS

Some different modalities of Real-world Data



This seems to be a combinational logic design problem.





Combinations Logic Specification – Truth Table

Input			
а	b	С	output
0	0	0	0
0	0	1	1
0	1	0	1
0	1	1	0
1	0	0	1
1	0	1	0
1	1	0	0
1	1	1	1





What if we did not know the truth table?

- Look at enough observation data to construct the rules
 - 000 \rightarrow 0
 - 011 \rightarrow 0
 - 100 \rightarrow 1
 - 110 \rightarrow 0

• If we have enough observational data to cover all input patterns, we can construct the truth table and derive the logic!

I L L I N O¹² S

LeNet-5, a convolutional neural network for handwritten digit recognition.







Two Important ML Modalities in Science

- Comparing Theoretical Predictions and Observational Data
 - Using theory to generate training data
 - Use trained model to look for relevant patterns in the observational data
 - Tradeoff training algorithm complexity for reduced data analysis algorithm complexity (LIGO example)
- Using ML to identify the key dimensionalities in simulation or observational data
 - Based on more traditional statistical approaches such as Eigen analysis
 - High algorithm complexity of data analysis



Conclusion and Outlook

- Applications have very large and increasing appetite for more computing power
 - Both larger scale clusters and faster nodes
 - Key to solving large, real world problems

- We can expect only about 2-3X increase in machine capability every 5 years moving forward
 - Low-complexity algorithms will continue to gain importance
- Heterogeneity has become the norm for all hardware systems
 - HPC community are currently seeing about 2-3x application speedup from GPUs
 - Recent positive spiral between deep learning and GPU computing
 - We expect much more heterogeneity in both computing and memory devices in the next generation

