

Computational Tools for Multiscale Simulations - CTCI

J. Ramanujam and S. Dua

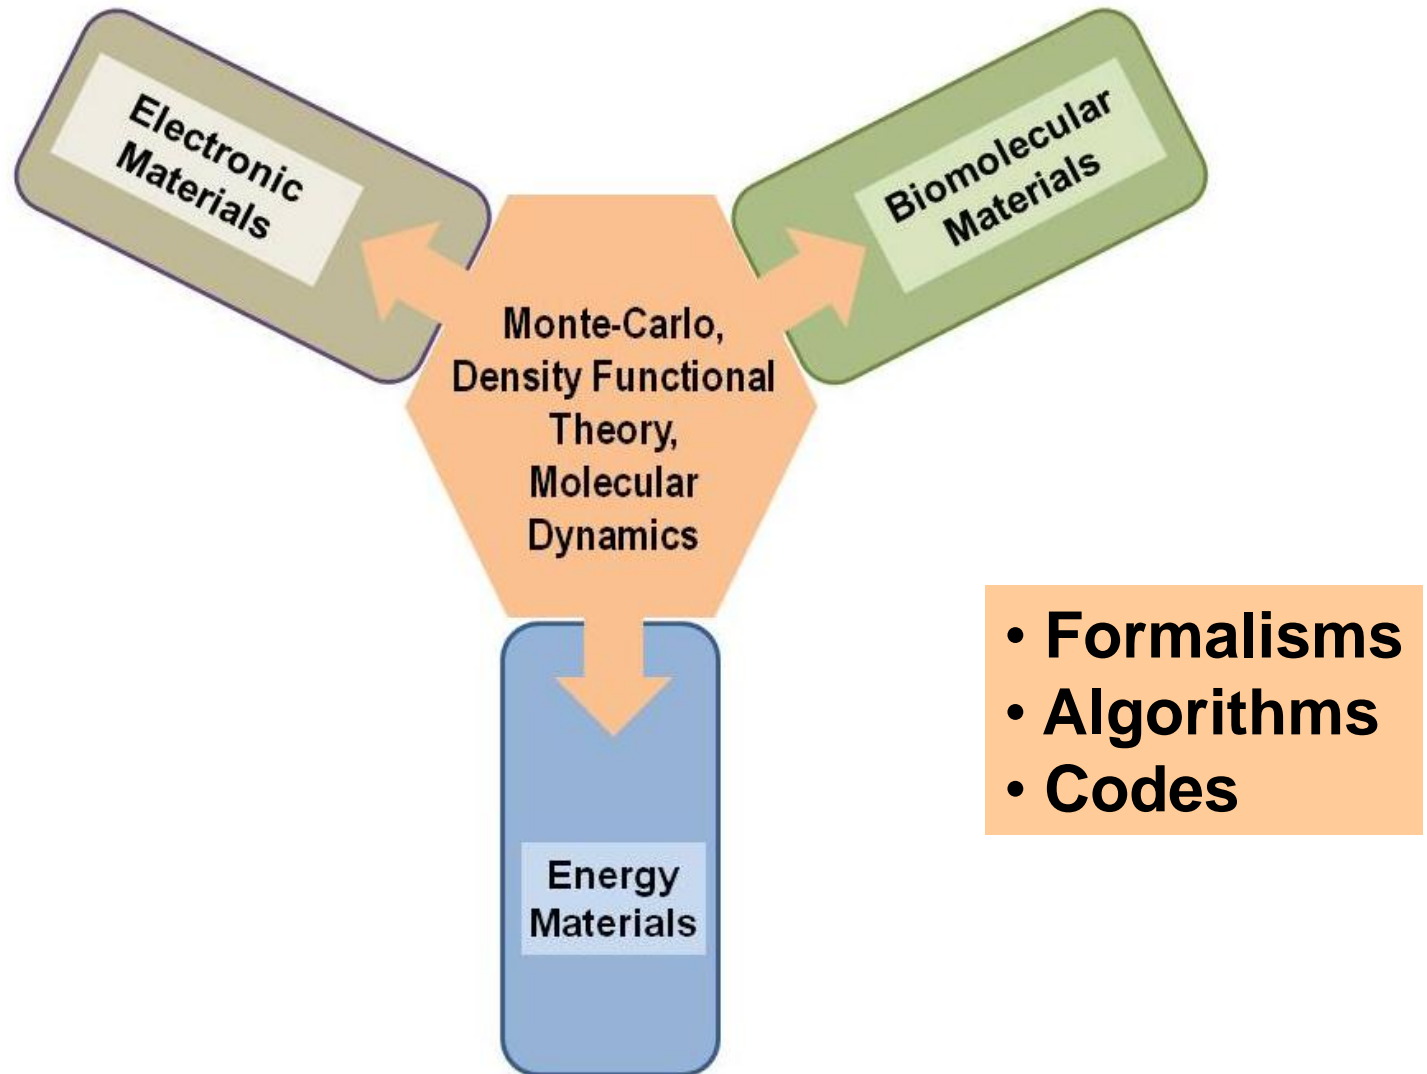


All-Hands Meeting, Baton Rouge: May 28, 2011

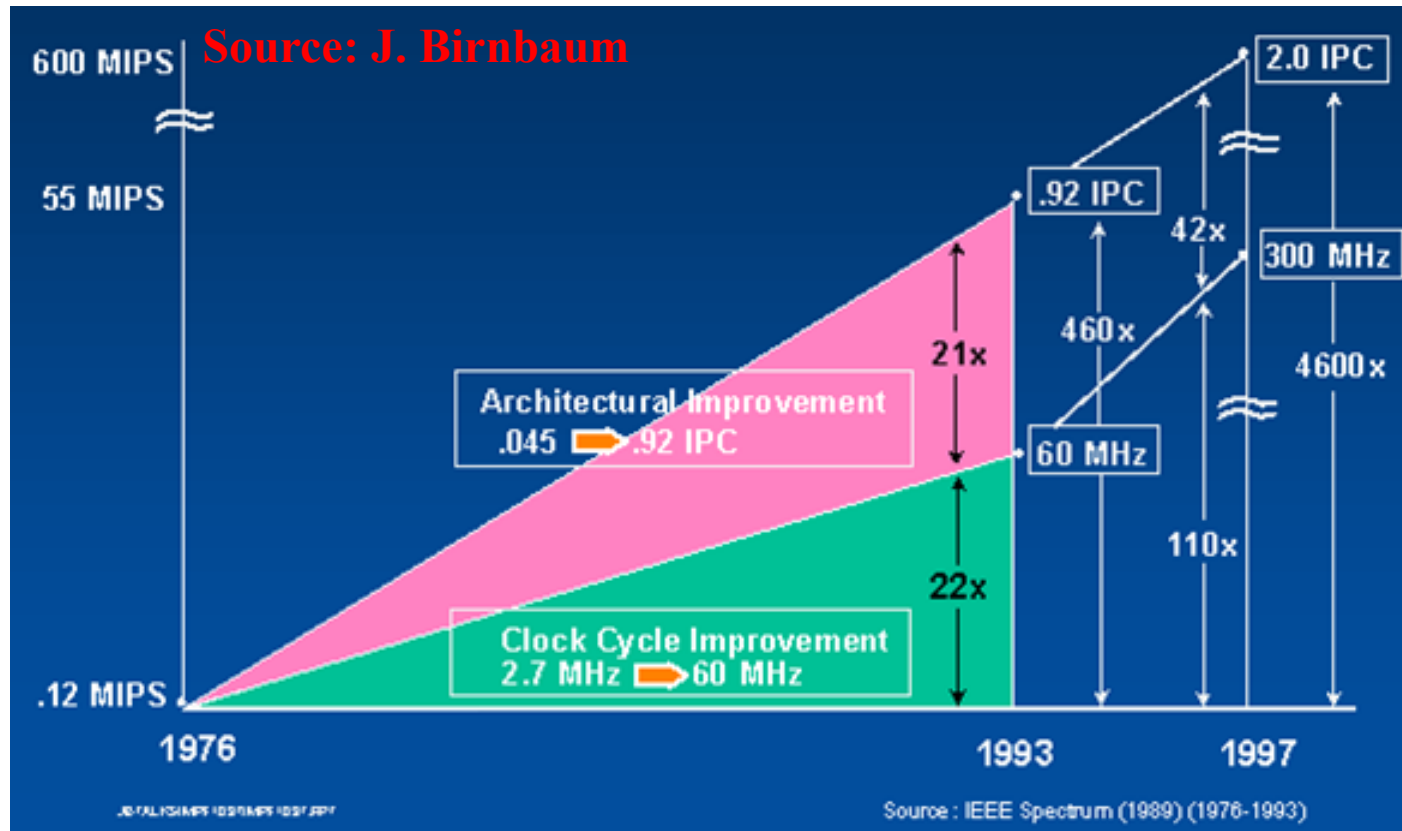
CTCI: CyberTools and CyberInfrastructure

- Goal: Build transformational common toolkits around three core formalisms/algorithms indispensable for computational materials science:
 - Next Generation Monte Carlo Codes
 - Massively Parallel Density Functional Theory and Force Field Methods
 - Large-scale Molecular Dynamics

CTCI: CyberTools and CyberInfrastructure




The Good Old Days for Software



- Single-processor performance experienced dramatic improvements from clock, and architectural improvement (Pipelining, Instruction-Level-Parallelism)
- Applications experienced automatic performance improvement

Revolution is Happening Now

- Chip density is continuing increase $\sim 2x$ every 2 years
 - Clock speed is not
 - Number of processor cores may double
- There is little or no more hidden parallelism (ILP) to be found
- **Parallelism must be exposed to and managed by software**



Source: Intel, Microsoft (Sutter) and Stanford (Olukotun, Hammond)

Challenges to the Future of Supercomputing

- Exaflops-scale computing by 2020
- 100X power consumption improvement
- 1000X increase in parallelism
- 10X improvement in efficiency
- Multicore and GPU processor architectures
- User programming productivity
- Fault tolerance for reliability
- Graph processing for STEM and knowledge management applications

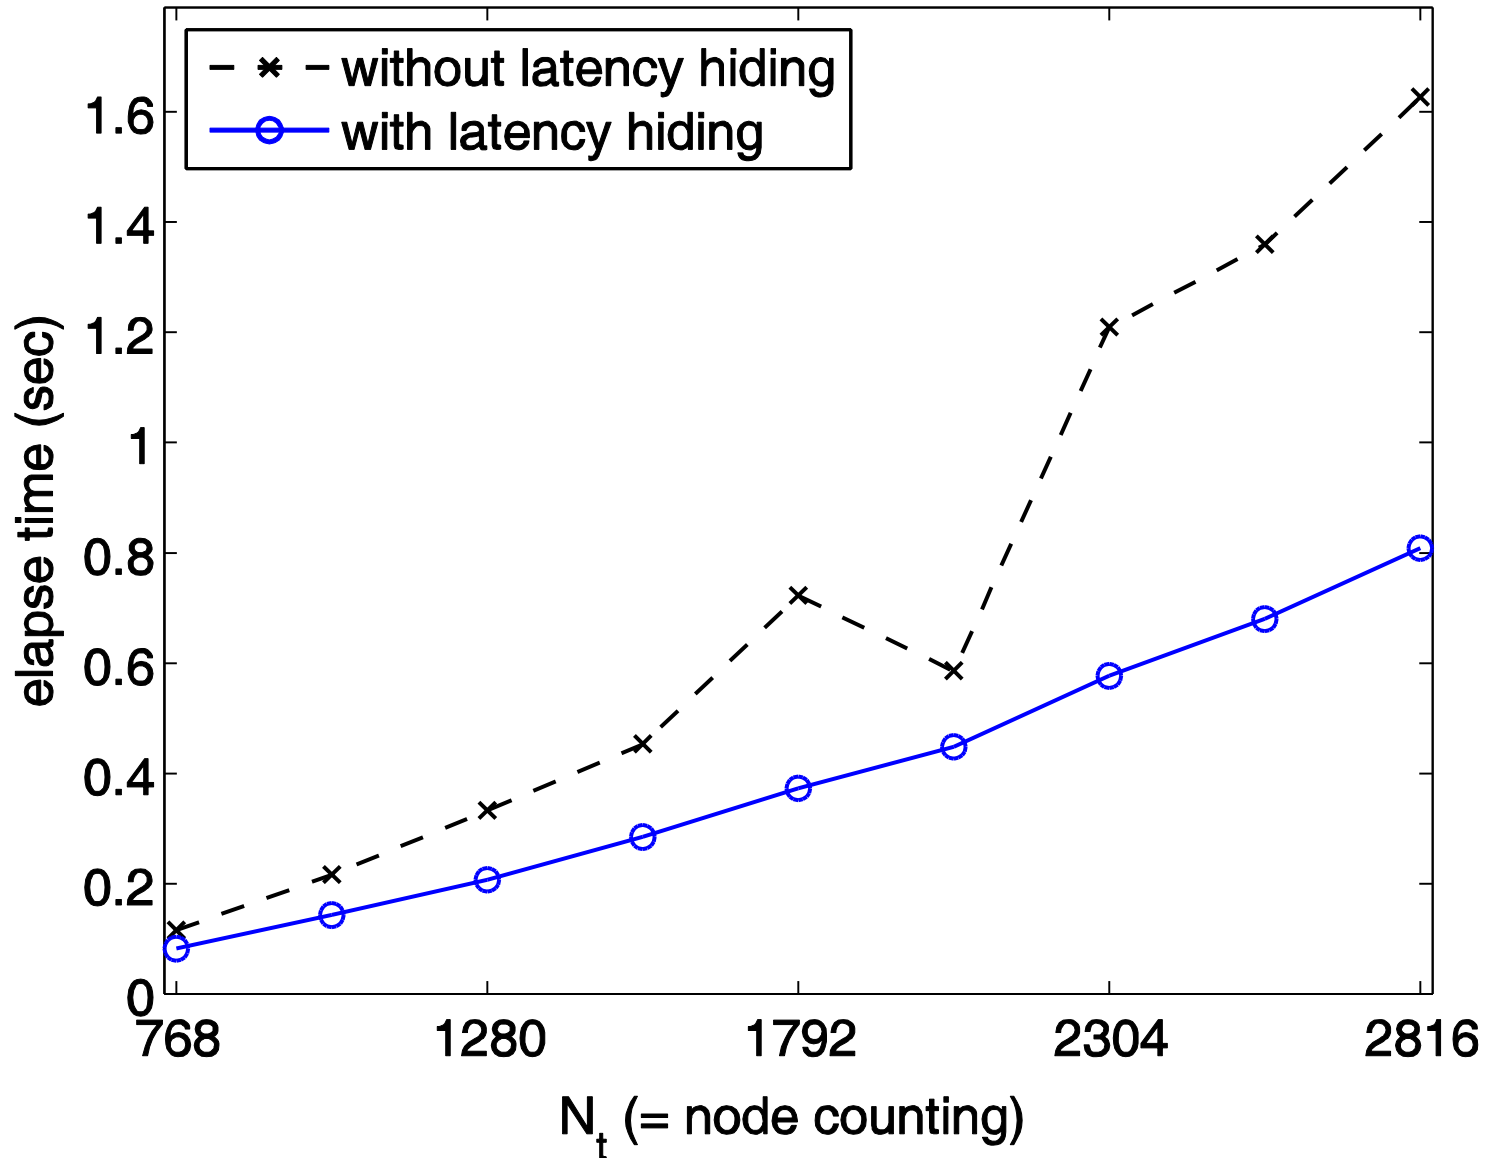
Scaling many-body codes to 30,000 processors [Lee, Ramanujam, Jarrell, Moreno (LSU)]

- Multi-Scale Many Body (MSMB) methods are required to more accurately model competing phases, including magnetism, superconductivity, insulating behavior etc. of strongly-correlated materials
 - In these multi-scale approaches, explicit methods are used only at the shortest length scales.
 - Two particle many-body methods, which scale algebraically, are used to approximate the weaker correlations at intermediate length scales.

Scaling many-body codes to 30,000 procs

- Advanced latency hiding methods, together with new national leadership class machines such as NSF's Kraken, have brought two-particle many-body theories to the point where they may be simulated.
- Hide the interprocessor communication in local calculations
- Latency hiding is used with tensor operations (which are the inner loops in these codes)

Scaling many-body codes to 30,000 procs



Constraint-based Synchronization for Extreme Scale Execution (Sterling, LSU)

Strategy

- Catalyze revolution through paradigm shift
- Develop new execution model for codesign and interoperability of system layers
- Attack problems of starvation, latency, overhead, and contention
- Implement experimental proof of concept software (HPX, ParalleX)

Accomplishments

- Derive ParalleX execution model
- Implement HPX runtime system software
- Application experiments demonstrate 2-3X performance gain
- Sponsored research by NSF, DARPA, NSA, DOE, Army & Microsoft

The Execution Model Imperative

- HPC in 6th Phase Change
 - Driven by technology opportunities and challenges
 - Historically, catalyzed by paradigm shift
- Guiding principles for governing system design and operation
 - Semantics, Mechanisms, Policies, Parameters, Metrics
- Enables holistic reasoning about concepts and tradeoffs
 - Serves for Exascale the role of *von Neumann architecture* for sequential
- Essential for co-design of all system layers
 - Architecture, runtime and operating system, programming models
 - Reduces design complexity from $O(N^2)$ to $O(N)$
- Empowers discrimination, commonality, portability
 - Establishes a phylum of UHPC class systems
- Decision chain
 - For reasoning towards optimization of design and operation



Game Changer – Runtime System


- Runtime system
 - is: ephemeral, dedicated to and exists only with an application
 - is not: the OS, persistent and dedicated to the hardware system
- Moves us from *static* to *dynamic* operational regime
 - Exploits situational awareness for causality-driven adaptation
 - Guided-missile with continuous course correction rather than a fired projectile with fixed-trajectory
- Based on foundational assumption
 - Untapped system resources to be harvested
 - More computational work will yield reduced time and lower power
 - Opportunities for enhanced efficiencies discovered only in flight
 - New methods of control to deliver superior scalability
- “Undiscovered Country” – adding a dimension of systematics
 - Adding a new component to the system stack
 - Path-finding through the new trade-off space



HPX: Adaptive Mesh Refinement

Wallclock time ratio MPI/HPX


(Depending on levels of AMR refinement - LoR, pollux.cct.lsu.edu, 32 cores)



GPUs

- GPUs are specialized for compute-intensive, data-parallel computations (e.g., rendering in graphics)
- More transistors on a GPU are devoted to processing data rather than flow control and data caching
- Provide significant application-level speedup over general purpose single processor execution
- General Purpose computation on GPU: GPGPU
 - Accelerate critical path of applications
 - Typically use data-parallel algorithms
 - Applications: image processing, game effects or physics processing, modeling, scientific computing, linear algebra, convolution, sorting, ...

Graphics Processing Units (GPUs)



Source: NVIDIA

Jack Dongarra on GPUs

“GPUs have evolved to the point where many real world applications are easily implemented on them and run significantly faster than on multi-core systems. Future computing architectures will be hybrid systems with parallel-core GPUs working in tandem with multi-core CPUs.”

Jack Dongarra

Professor, University of Tennessee

(LINPACK, BLAS, Top 500, ...)

Parallel Programming Tools (Leangsuksun, LaTech)

- Collaboration with Drs. Alex Safarenko and Sven-Bodo Scholz, University of Hertfordshire, UK.
- Parallel Programming Tool Development based on SAC (Single Assignment C) toolset that enables parallel application developers expressing their problems in a high-level language
- Tool set is capable to generate codes in various HPC architectures including multicore and GPU.
- The collaborative enhancement will introduce fault tolerance to SAC and its parallel applications.

Novel Architectures and Compilation

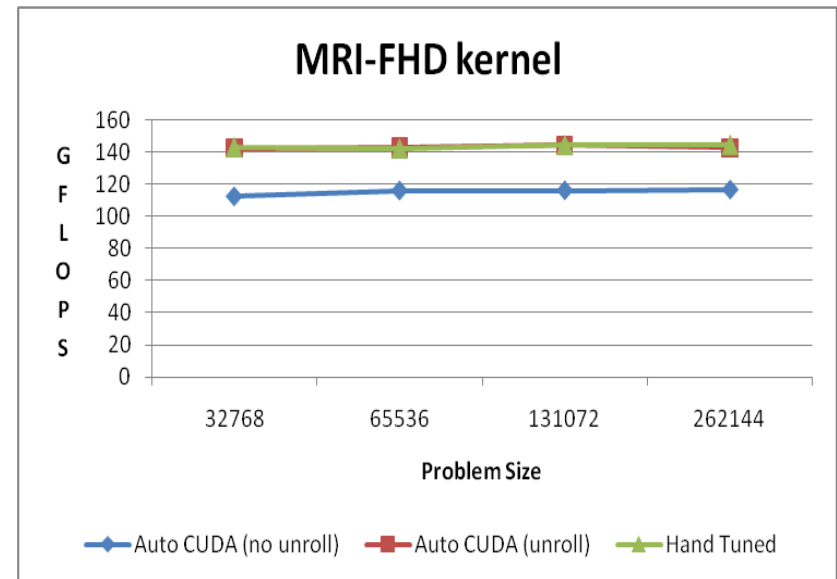
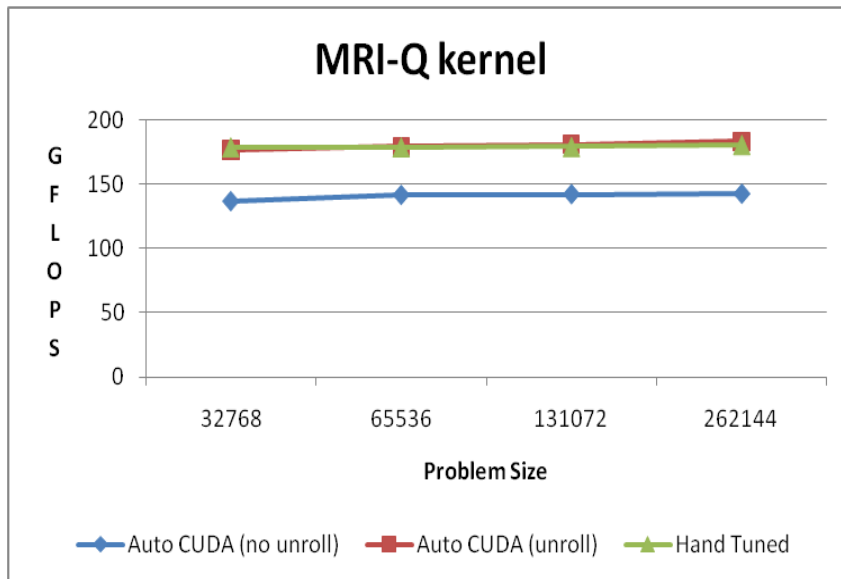
Techniques (Ramanujam, LSU)

- GPUs are cost-effective and powerful processors for general-purpose computing
- Programming GPUs (for example, the NVIDIA GPUs using CUDA or OpenCL) is still tedious:
 - The performance of the GPU is highly sensitive to the formulation of the kernel; needs significant experimentation
 - Programmers may like this low level of control and this is appropriate for library development; unlike the case of general-purpose processors, compilers and tools are not quite helpful here
- **Strategy:** Develop a system (PLUTO) for automatic transformation of sequential input programs into efficient parallel CUDA/OpenCL programs

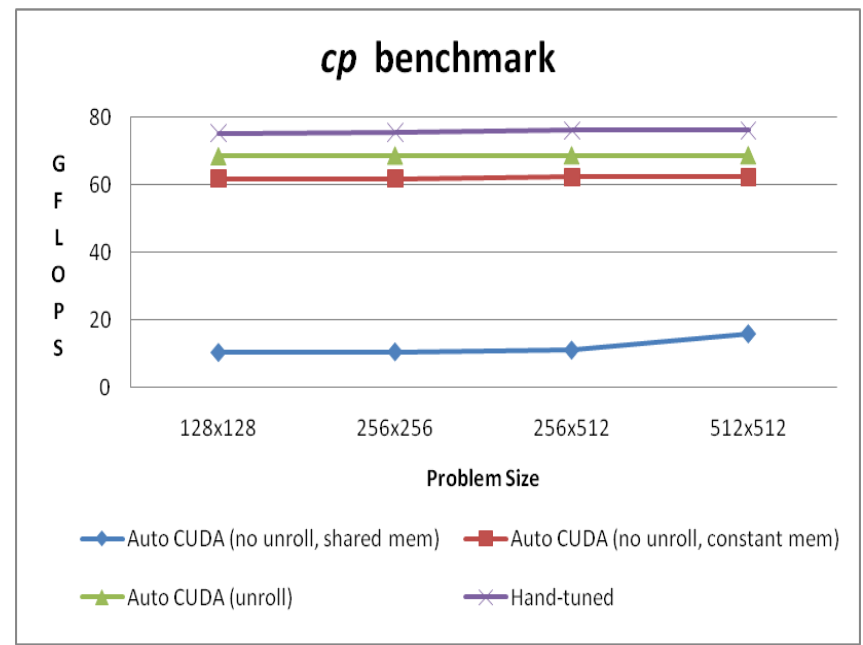
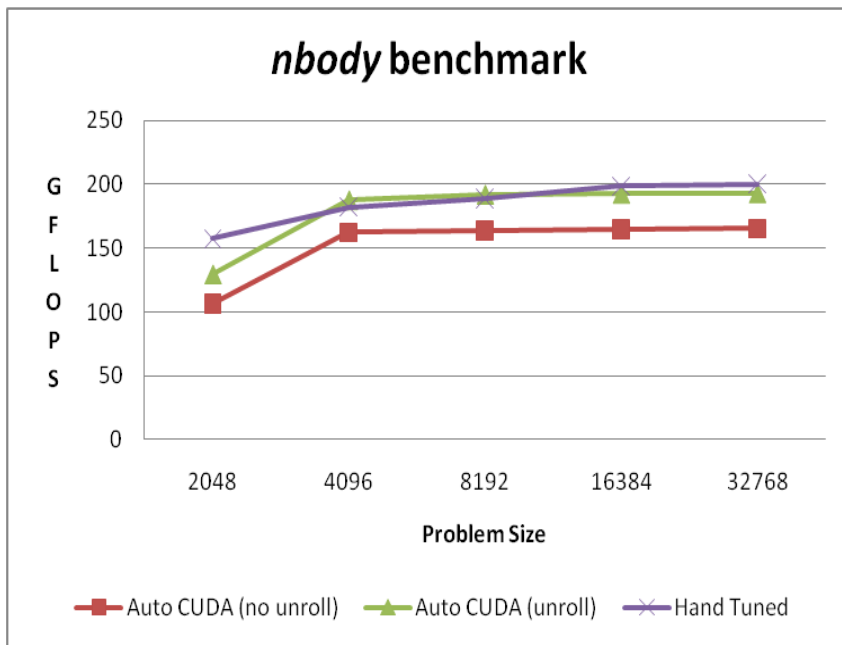
CUDA Model

- CUDA views the GPU as a compute *device* that is a coprocessor to the CPU (*host*):
 - GPU (the *device*) executes data-parallel parts (*kernels*) of an application
 - The GPU executes a *kernel* on many threads in parallel
 - GPU threads have very small thread creation overhead (extremely lightweight)
 - The GPU requires thousands of threads for full efficiency (in contrast to multicore CPUs, which need only a few)
 - GPU has its own memory DRAM (*device memory*)

Automatic Transformer (PLUTO): Experimental Evaluation - 1



Automatic Transformer (PLUTO): Experimental Evaluation - 2



Execution Management Tools (Jha, LSU)

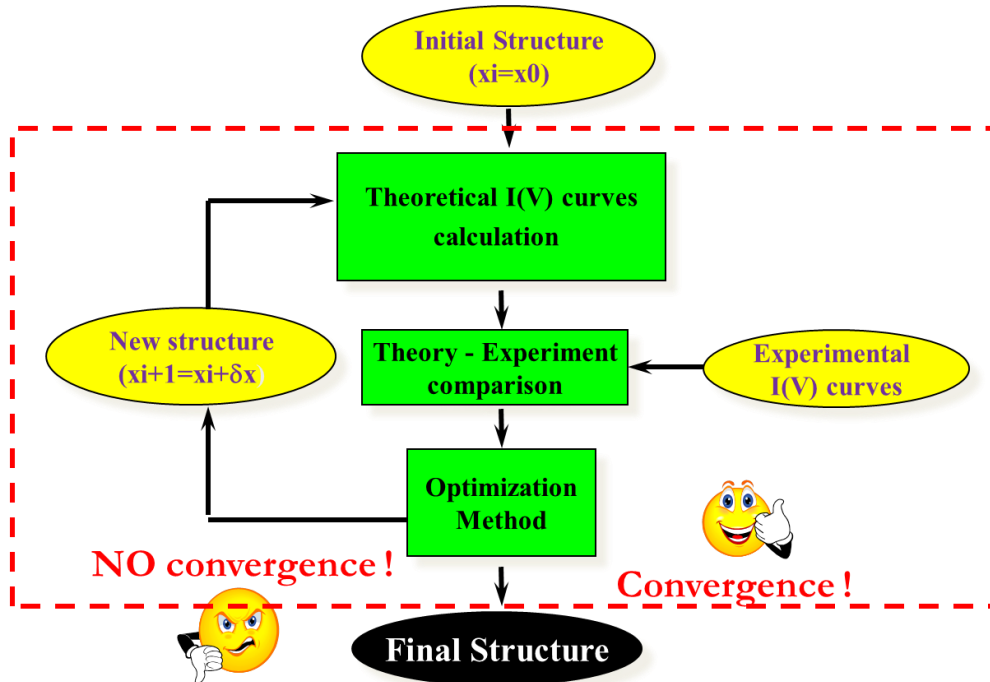
- Often, there is a need to support the execution of multiple codes with data input-output relationship among them in a clean way
- Primary focus on integrating codes with the SAGA Pilot-Job mechanism on LONI machines
- In progress: Extending this capability to Gaussian and other group-codes, that have different runtime profiles and adaptivity
- Related techniques have been applied to support the cross-site runs of an ensemble of Molecular Dynamics (MD) runs for approximately 6000 ensembles in conjunction with Tom Bishop (Tulane).

Inverse Problem in Structure Determination by LEED (Von Braun Nascimento, Ward Plummer, Hannah Manuel: LSU)

- Multiple scattering forces LEED analysis to be indirect
- Quantitative comparison theory-experiment
- Search for the best fitted structural model: Search Problem
- Hard task: locate the global minimum in an N-dimensional parameters space;

Using the results of actual observations to infer the values of the parameters characterizing the system under investigation

Search Process



Complex Transition Metal Oxides :
complex structure → many structural parameters to be optimized;

Necessary Improvements

- 1) Global Search Methods;
- 2) Faster Multiple Scattering calculations
- 3) **Direct Methods (inverse problem)**, surface structure directly from experimental data → **Holy Grail of LEED !**

New Paradigm in Developing Force Fields (Wick, LaTech; and Rick, UNO)

- Developed new force fields relevant for catalysis and biological simulations – these methods rely on two relevant characteristics:
 - reactivity in aqueous environments, and
 - Intermolecular charge transfer
- Trade-off between accuracy and computational efficiency
- These models have accuracies as good or better than many of the commonly used models for water, without any additional significant computational demand.
- Work in progress: Developing charge transfer models for ions. Charge transfer between the solvent, ions and nanotubes is important based on the results of Pratt and Hoffman at Tulane (SD2)

Temporal Pattern Discovery in High-Throughput Data Domains: Dua, LaTech

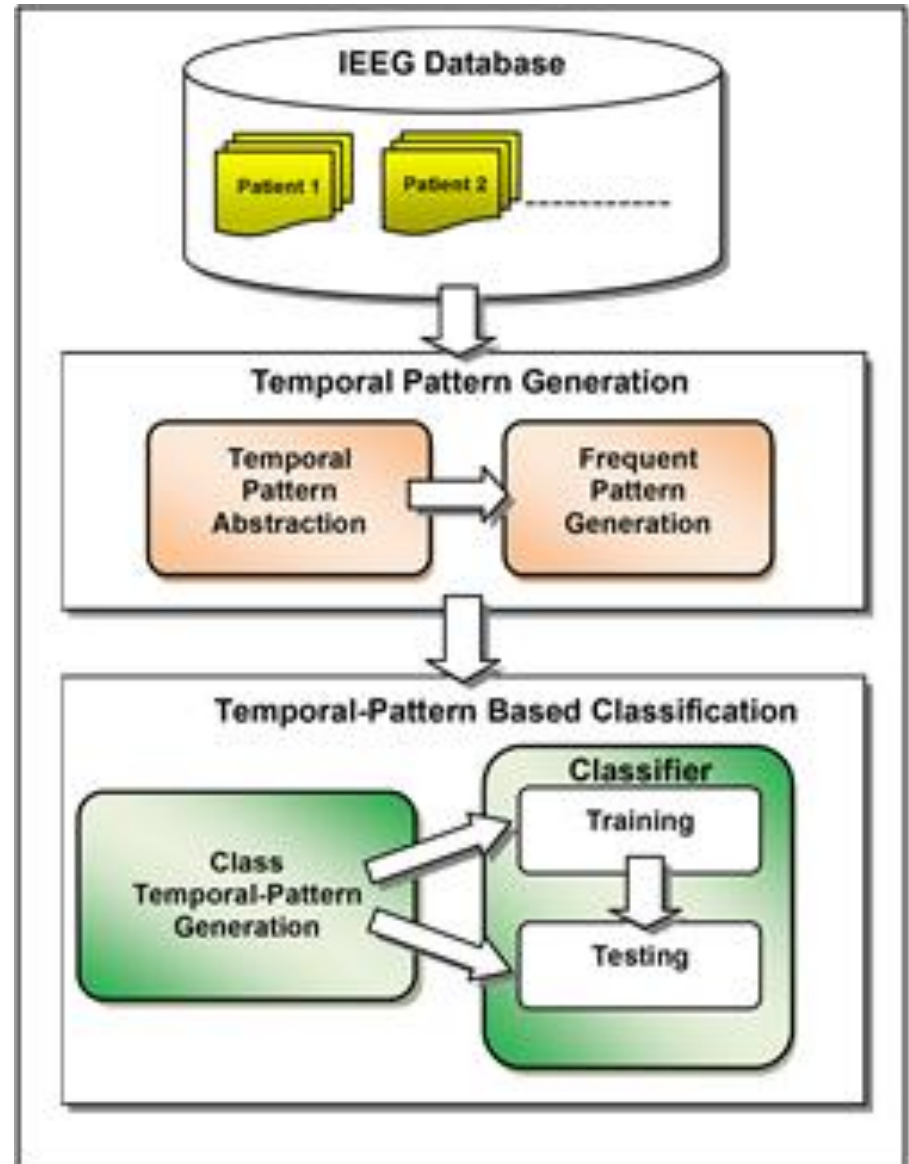
- Multivariate temporal data are collections of contiguous data values that reflect complex temporal changes over a given duration.
- Technological advances have resulted in huge amounts of such data in high-throughput disciplines, including data from MD simulations
- Traditional data mining techniques do not utilize the complex behavioral changes in multivariate time series data over a period of time and are not well suited to classify multivariate temporal data

Temporal Pattern Discovery (Dua)

- Temporal pattern mining on multivariate time series data can uncover interesting patterns, which can be further utilized to extract complex temporal patterns from a variety of time series data to build an effective automated decision system.
- We have developed an automated system that builds a temporal pattern-based framework to classify multivariate time series data.
- Comparative study shows that this method produces overall superior results.

Temporal Pattern Discovery (Dua)

- Two algorithms developed:
 - one for temporal pattern generation; and
 - one for multivariate time-series classification.
- This approach has higher accuracy (86%) compared to other methods (around 72%)



Education, Outreach, Collaborations

- **Courses**

- High Performance Computing, Computational Science, Compiler Optimizations

- **Outreach**

- Beowulf Boot Camp for High School Students and Teachers, 2011
- GPU Programming Models, Optimizations and Tuning, tutorial: International Symp. on Code Gen. & Opt., April 2011
- GPUs & General-Purpose Multicores: Programming Models, Optimizations & Tuning, tutorial: Intl. Conf. Supercomp., June 2011

- **External Collaborations**

- Sandia National Labs, Pacific Northwest National Labs, Oak Ridge National Labs, Lawrence Berkeley Natl. Labs, Ohio State University, U. Delaware, U. Hertfordshire, ...

- **CTCI Seminars**

- Sven-Bodo Scholz (Univeristy of Hertfordshire, UK), *High-Productivity and High-Performance with Multicores*
- Zhang Le, (Michigan Tech University), *Using multi-scale and multi-resolution model with Graphics processing Unit (GPU) to simulate real-time actual cancer progression*

Summary

- Scaling of MSMB methods to 30,000 cores
- Runtime system for exascale computing
- GPUs
 - Parallel programming tools
 - Compilation technology
 - Implementation of high-performance codes
- Execution management
- Inverse problem (LEED)
- New force-field development
- New analysis techniques for data from MD simulations
- Strong collaborations with Science Drivers