Accelerating Computational Science and Engineering with Heterogeneous Computing in Louisiana

For Presentation at NVIDIA Booth in SC14

by

Honggao Liu, PhD

Deputy Director of CCT

11/19/2014

LSU | Center for Computation & Technology

Interdisciplinary | Innovative | Inventive

Outline

1. Overview Cyberinfrastructure in Louisiana
2. Trends in accelerator-aided supercomputing
3. Move Louisiana users to a hybrid accelerated environment
4. Early results running on GPU-accelerated HPC clusters

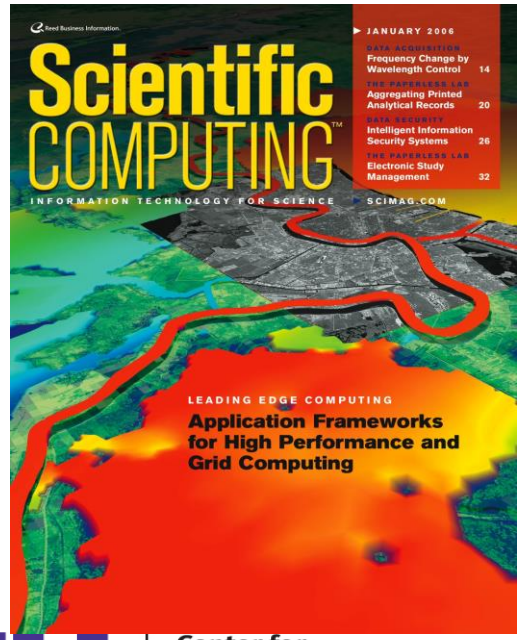

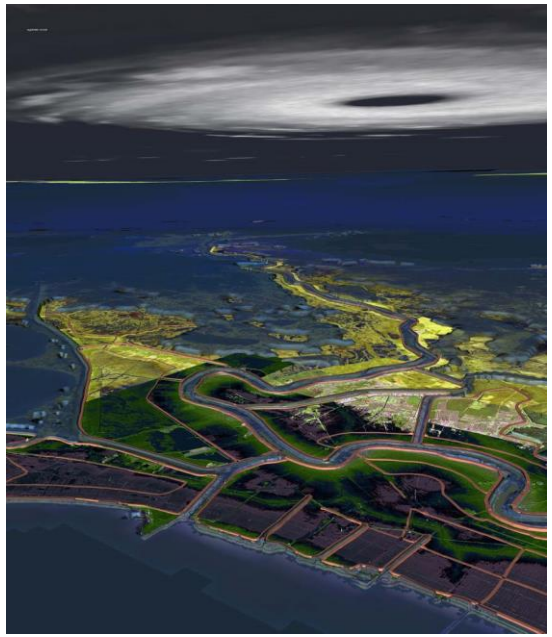

CCT is …

An innovative and interdisciplinary research environment that advances computational sciences and technologies and the disciplines they touch.

- **Faculty lines** – currently, 34 (avg. 50/50 split appointments) across 13 departments and 7 colleges/schools; tenure resides in home department
- **Enablement staff** – currently 15 senior research scientists (non-tenured; mixture of CCT dollars and soft money support) with HPC and scientific visualization expertise who support a broad range of compute-intensive and data-intensive research projects;
- **Education** – Influence design and content of interdisciplinary curricula; for example: (1) computational sciences, (2) visualization, and (3) digital media
- **CyberInfrastructure** – guide LSU's (and state's via LONI) cyber-infrastructure design to support research ➔ high-performance computing (HPC), networking, data storage/management, & visualization; also associated HPC support staff

Louisiana Cyberinfrastructure

3 Layers:

- LONI (Network + HPC)
- LONI Institute
- LA-SiGMA

LSU | Center for Computation & Technology

Interdisciplinary | Innovative | Inventive

Louisiana Cyberinfratructure

- **LONI** base (http://loni.org):
  - A state-of-the-art fiber optics network that runs throughout Louisiana, and connects Louisiana and Mississippi research universities
  - State project since 2005, $40M Optical Network, 4x 10 Gb lambdas
  - $10M Supercomputers installed at 6 sites in 2007, centrally maintained by HPC @ LSU
  - $8M Supercomputer to replace Queen Bee, upgrade network to 100Gbps
- **LONI Institute** (http://institute.loni.org/):
  - Collaborations on top of LONI base
  - $15M Statewide project to recruit computational researchers
- **LA-SiGMA** (http://lasigma.loni.org/):
  - Louisiana Alliance for Simulation-Guided Materials Applications
  - Virtual organization of seven institutions of Louisiana focusing on computational materials science
  - Research and develop tools on top of LONI base and LONI Institute
  - $20M Statewide NSF/EPSCOR Cyberinfrastructure project

Supercomputers in Louisiana Higher Education

LSU CENTER FOR COMPUTATION & TECHNOLOGY

and

INFORMATION TECHNOLOGY SERVICES

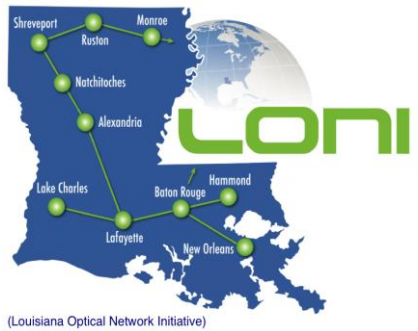

(Louisiana Optical Network Initiative)

2002 : SuperMike : ~ $3M from LSU (CCT & ITS), Atipa Technologies — **17th** in Top500
1024 cores; 3.7 Tflops

2007 : Tezpur : ~ $1.2M from LSU (CCT & ITS), **Dell** — **134th** in Top500
1440 cores; 15.3 Tflops

2007 : Queen Bee : ~ $3M thru BoR/LONI (Gov. Blanco), **Dell** — **23rd** in Top500
5440 cores; 50.7 Tflops; Became NSF-funded node on TeraGrid

2012 : SuperMike-II : $2.65M from LSU (CCT & ITS), **Dell** — **250th** in Top500
7040 cores; 146 + 66 Tflops

2014 : **SuperMIC** : $4.1M from NSF & LSU, **Dell** — **65th** in Top500
7600 cores; 1050 Tflops; Became NSF-funded node on XSEDE

2014 : **QB2** : ~ $6.6M thru BoR/LONI, **Dell** — **46th** in Top500
10080 cores; 1530 Tflops;

DELL

LSU Center for Computation & Technology
Interdisciplinary | Innovative | Inventive

HPC Systems (According to OS)

Linux Clusters

- LSU's HPC
  - SuperMIC (1050 TF)
    NEW in production
  - SuperMike-II (220 TF)
  - Shelob (95 TF)
  - Tezpur (15.3 TF)
    Decommissioned in 2014
  - Philip (3.5 TF)
- LONI
  - QB2 (1530 TF)
    NEW in friendly user mode
  - Queen Bee (50.7 TF)
    Decommissioned in 2014
  - Five (@ 4.8 TF)

AIX Clusters

- LSU's HPC
  - Pandora (IBM P7; 6.8 TF)
  - Pelican (IBM P5+;1.9 TF)
    Decommissioned in 2013
- LONI
  - Five (IBM P5; @ 0.85 TF)
    Decommissioned in 2013

LSU's HPC Clusters

- ✓ SuperMike-II: $2.6M in LSU funding; installed in fall 2012
- ✓ Melete: $0.9M in 2011 NSF/CNS/MRI funding; an interaction-oriented, software-rich cluster w/ tangible interface support
- ✓ **Shelob:** $0.54M in 2012 NSF/CNS funding; a GPU-loaded, heterogeneous, computing platform
- ✓ SuperMIC: $3.92M in 2013 NSF/ACI/MRI funding + $1.7M LSU match; ~ 1PetaFlops HPC system fully loaded w/ Intel Xeon-phi processors

LSU HPC System

- **SuperMike-II** (mike.hpc.lsu.edu)
  - **380 compute nodes**: 16 Intel Sandy Bridge cores @ 2.6GHz, 32GB RAM, 500GB HD, 40Gb/s infiniband, 2x 1Gb/s Ethernet
  - **52 GPU compute nodes**: 16 Intel Sandy Bridge cores @ 2.6GHz, 2 NVIDIA M2090 GPUs, 64GB RAM, 500GB HD, 40Gb/s infiniband, 2x 1Gb/s Ethernet
  - **8 fat compute nodes**: 16 Intel Sandy Bridge cores @ 2.6GHz, 256 GB RAM, 500GB HD, 40Gb/s infiniband, 2x 1Gb/s Ethernet, **Aggregated together by ScaleMP to one big SMP node**
  - **3 head nodes**: 16 Intel Sandy Bridge cores @ 2.6GHz, 64 GB RAM, 2 x 500GB HD, 40Gb/s infiniband, 2x 10Gb/s
  - **1500TB** (scratch + long term) DDN Luster storage

LSU New HPC System

- **SuperMIC** (mic.hpc.lsu.edu)
  - The largest NSF MRI award LSU has ever received ($3.92M with $1.7M LSU match for the project)
  - **Dell is a partner on the proposal, and won the bid!**
  - **360 compute nodes**
    - **2x 10-core 2.8GHz Ivy Bridge CPUs, 2x 7120P PHIs, 64GB Ram**
  - **20 hybrid compute nodes**
    - **2x 10-core 2.8GHz Ivy Bridge CPUs, 1x 7120P PHI, 1x K20X GPU, 64GB Ram**
  - **1 Phi head node, 1 GPU head node**
  - **1 NFS server node,**
  - **1 cluster management node**
  - **960 TB (scratch) Luster storage**
  - **FDR Infiniband**
  - **1.05 PFlops peak performance**

Center for Computation & Technology
Interdisciplinary | Innovative | Inventive

LONI Supercomputing Grid

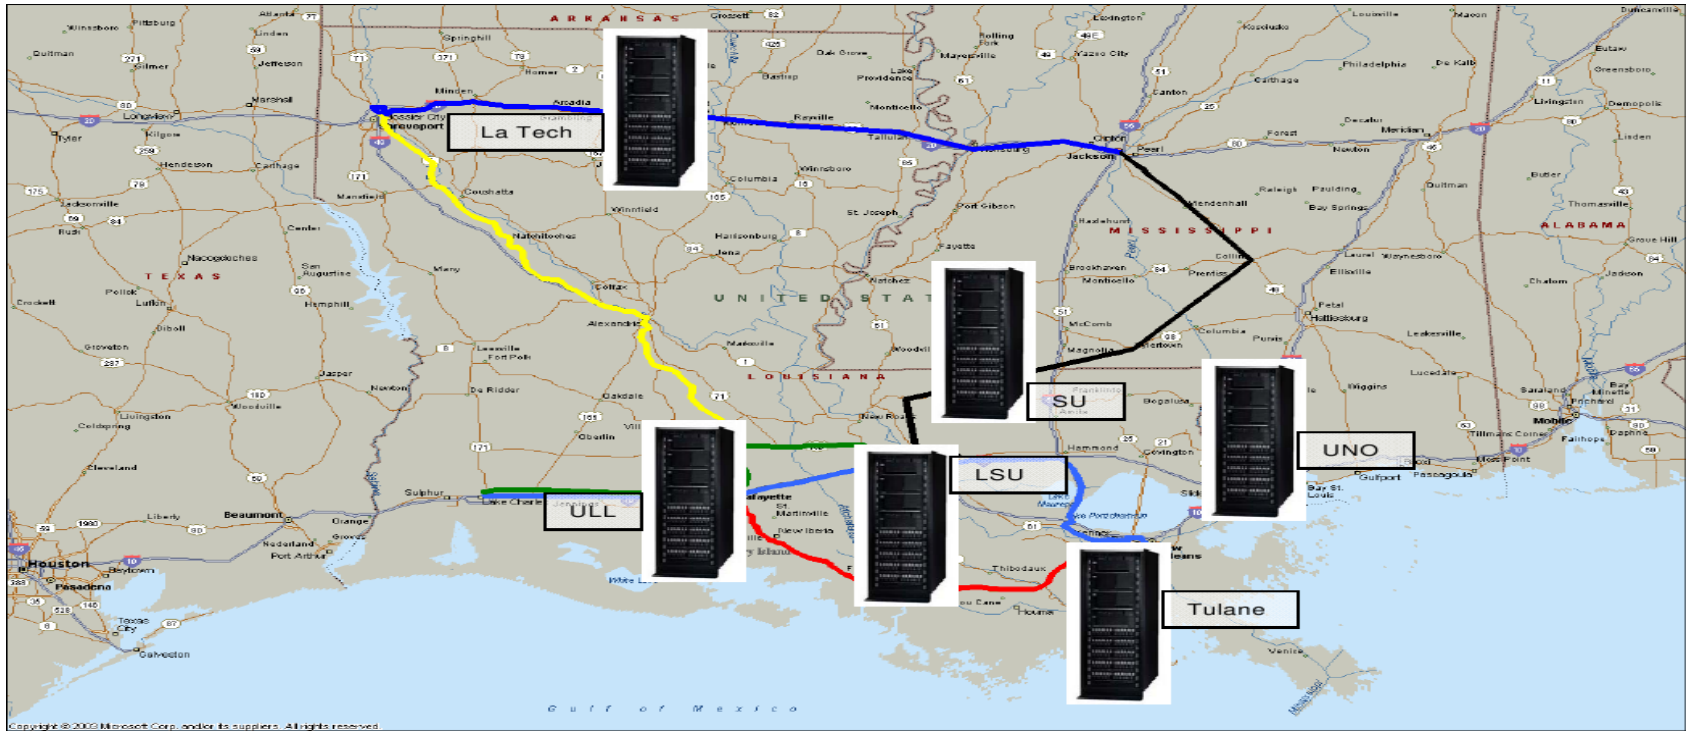

- 6 clusters currently online, hosted at six campuses

LONI's HPC Clusters

- **QB2: 1530 Tflops centerpiece (NEW)**
  - Achieved 1052 TFlops using 476 of 504 compute nodes
  - 480 nodes with NVIDIA K20X
  - 16 nodes 2 Intel Xeon Phi 7120P
  - 4 nodes with NVIDIA K40
  - 4 nodes with 40 Intel Ivy Bridge cores and **1.5 TB** RAM
  - 1600TB DDN storage running Lustre
- **Five 5 TFlops clusters**
  - Online: Eric(LSU), Oliver(ULL), Louie(Tulane), Poseidon(UNO), Painter (LaTech)
  - 128 nodes with 4 Intel Xeons cores@ 2.33 Ghz, 4 GB RAM
  - 9TB DDN storage running Lustre each
- **Queen Bee: 50 Tflops (decommissioned)**
  - 23rd on the June 2007 Top 500 list

LONI New HPC System

- **Queen Bee Replacement (QB2, qb.loni.org)**
  - **Dell won the bid!**
  - **480 GPU compute nodes**
    - **2x 10-core 2.8GHz Ivy Bridge CPUs, 2x K20X GPUs, 64GB Ram**
  - **16 Xeon Phi compute nodes**
    - **2x 10-core 2.8GHz Ivy Bridge CPUs, 2x 7120P PHIs, 64GB Ram**
  - **4 Visualization/compute nodes**
    - **2x 10-core 2.8GHz Ivy Bridge CPUs, 2x K40 GPUs, 128GB Ram**
  - **4 Big Memory compute nodes**
    - **4x 10-core 2.6GHz Ivy Bridge CPUs, 1.5TB Ram**
  - **1 GPU head node and 1 Xeon Phi head node**
  - **1 NFS server node**
  - **2 cluster management nodes**
  - **1600TB (scratch) Luster storage**
  - **FDR Infiniband**
  - **1.53 PFlops peak performance**

Center for Computation & Techn

Interdisciplinary | Innovative | Inventive

Trends in Supercomputing

Multi-core – Many-core
Hybrid processors
    Accelerators for specific kinds of computation
    Co-processors
Application-specific supercomputers

NVIDIA GPU

Intel MIC (Many Integrated Core) –Xeon Phi

Usage of Accelerators in HPC

- Statistics of accelerators in top 500 supercomputers (June 2014 list)

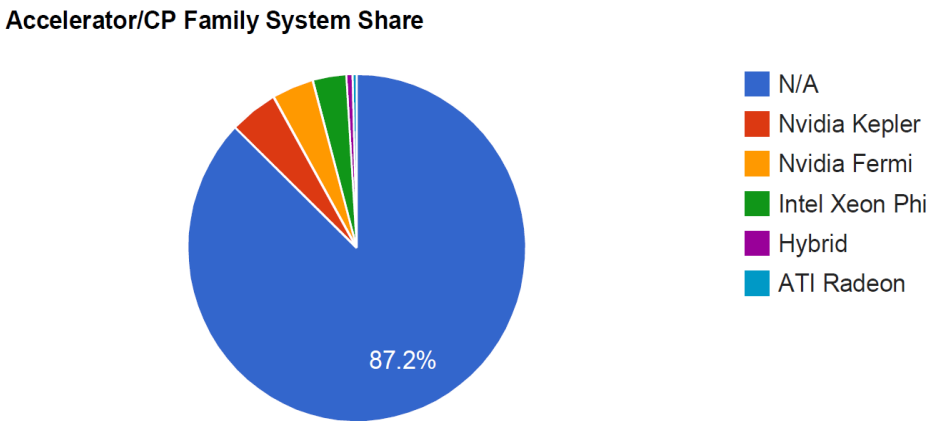

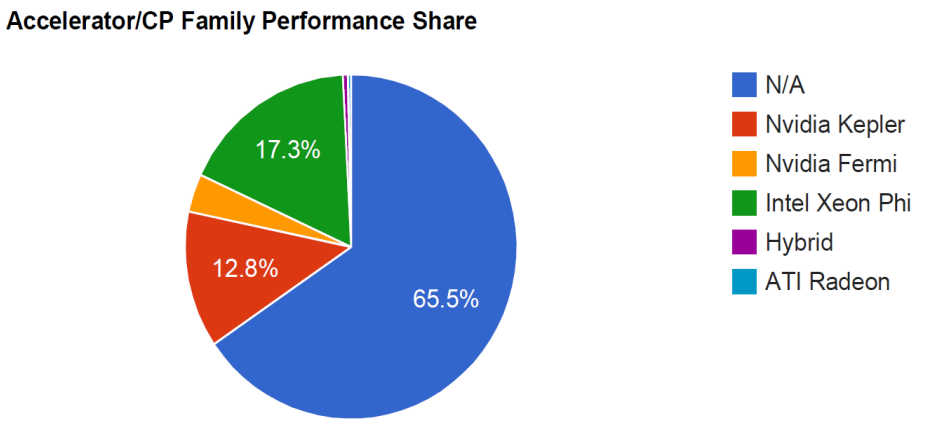

Accelerator/CP Family	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
N/A	436	87.2	179,221,078	252,023,035	15,384,521
Nvidia Kepler	23	4.6	35,133,206	52,529,949	960,416
Nvidia Fermi	20	4	9,813,752	20,123,473	750,056
Intel Xeon Phi	16	3.2	47,390,611	75,176,932	4,234,766
Hybrid	3	0.6	1,373,234	2,018,688	236,284
ATI Radeon	2	0.4	831,900	1,686,749	83,328

Supercomputers in Louisiana Higher Education

LSU CENTER FOR COMPUTATION & TECHNOLOGY

and

INFORMATION TECHNOLOGY SERVICES

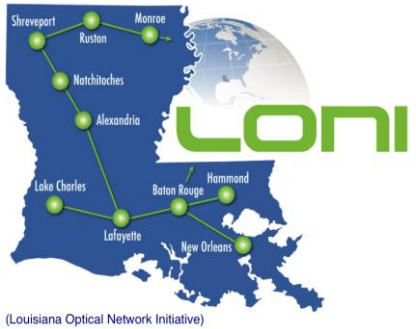

(Louisiana Optical Network Initiative)

2002 : SuperMike : ~ $3M from LSU (CCT & ITS), Atipa Technologies — **17th** in Top500
1024 cores; 3.7 Tflops (1 core/processor)

2007 : Tezpur : ~ $1.2M from LSU (CCT & ITS), **Dell** — **134th** in Top500
1440 cores; 15.3 Tflops (2 cores/processor)

2007 : Queen Bee : ~ $3M thru BoR/LONI (Gov. Blanco), **Dell** — **23rd** in Top500 on TeraGrid
5440 cores; 50.7 Tflops (4 cores/processor)

2012 : SuperMike-II : $2.65M from LSU (CCT & ITS), **Dell** — **250th** in Top500
7040 cores; 146 + 66 Tflops (8 cores/processor, **100 NVIDIA M2090 GPUs**)

2014 : **SuperMIC** : $4.1M from NSF & LSU, **Dell** — **65th** in Top500
7600 cores; 1050 Tflops (10 cores/processor, **740 Intel PHIs + 20 NVIDIA K20X GPUs**)

2014 : **QB2** : ~ $6.6M thru BoR/LONI, **Dell** — **46th** in Top500
10080 cores; 1530 Tflops; (10 cores/processor, **960 NVIDIA K20X + 8 K40 + 32 Intel PHIs**)

DELL

LSU Center for Computation & Technology
Interdisciplinary | Innovative | Inventive

NVIDIA CUDA RESEARCH CENTER

GPU Efforts

- **Why GPU?**
- Spider: 8-node GPU cluster in 2005, visualization group
- A GPU team is formed and funded by LA-SiGMA in 2009
- Renamed as Heterogeneous Computing Team in 2013, and the Technologies for Extreme Scale Computing (TESC) group in 2014
- Is devoted to the development of new computational formalisms, algorithms, and codes optimized to run on heterogeneous computers with GPUs (and Xeon PHIs).
- Develops technologies for next generation supercomputing and big data analytics.
- Fosters interdisciplinary collaborations and trains next generation computational and computer scientists.

TESC Group

- Focus on multiple projects each devoted to the development of different codes, such as codes **for simulations of spin glasses, drug discovery, quantum Monte Carlo Simulations, or classical simulations of molecular system**s.
- A Co-development model, a collaboration of students from different domain sciences or engineering partnered with students from computer science or computing engineering, is ideal for the rapid development of highly optimized codes for GPU or Xeon Phi architectures.
- Includes more than 80 researchers, and its weekly meetings are attended by an average of 40 researchers.
- Also includes the Ste||ar Group developing HPX, the Cactus group, and others at CCT

Education, Outreach and Training

- **Train Louisiana users to a hybrid accelerated environment**
- Training and education at all levels, from primary school through graduate school and beyond, is an essential component of the CCT's year round activities
- **Beowulf Bootcamp**: teaching High Schools about HPC
  - CCT has offered a week-long Beowulf Bootcamp in past 6 years
  - Interactive Lectures, Hands-On with Hardware, Programming
- **Research Experiences for Undergraduates** (REU) & Teachers (RET)

Education, Outreach and Training

- **Computational Sciences Workshops**: over 15 workshops on a broad range of subjects
- **HPC training**: recurring training is regularly provided throughout the year






GeauxDock: Molecular docking package for computer-aided drug discovery

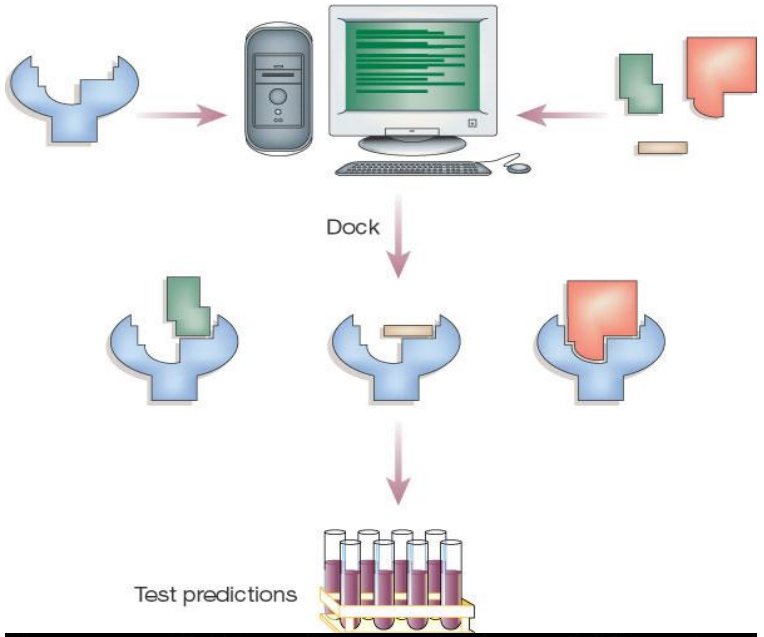

Computational modeling of binding drug to proteins has become an integral component of modern drug discovery pipeline.

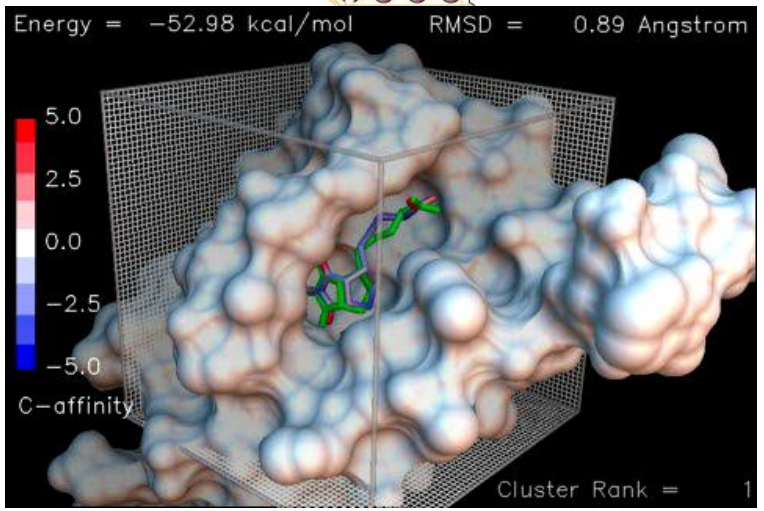

Virtual Screening (VS)

- Ligand based
- Structure based
  - **Ligand-receptor docking**
  - Affinity prediction

Computation Model

Multiple Replica Monte Carlo
Ligand and Protein Conformations

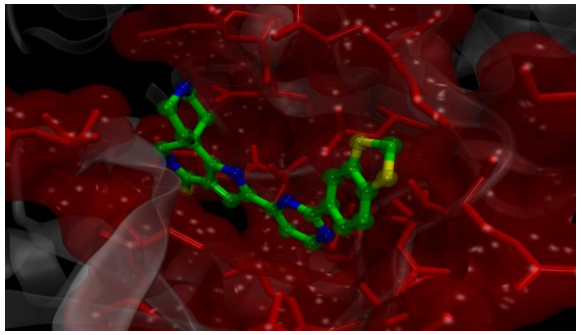

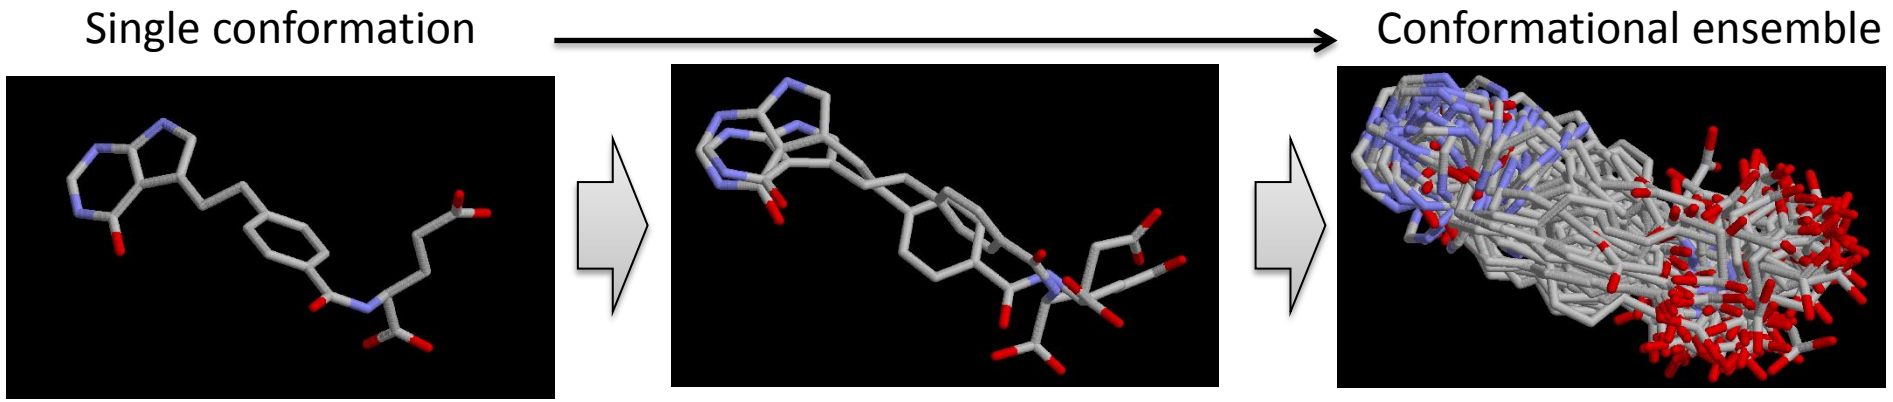

Computer-aided drug development holds a significant promise to speed up the discovery of novel pharmaceuticals at reduced costs.

Docking simulations predict the native pose of the ligand by searching for the global minimum in the energy space.

Computation Model

The Program Outline

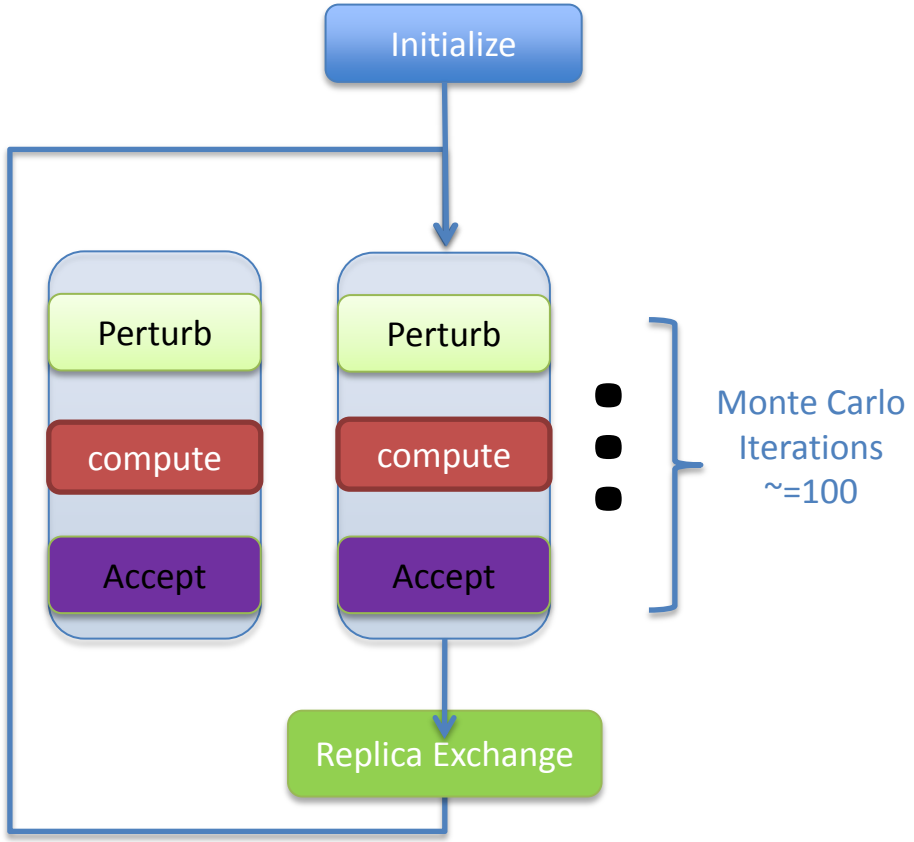

Implementation

Task mapping

	Fine grain	Coarse grain
Domain Model	Pair-wise computation	Replica ensembles
CPU	SIMD	Threads
GPU	threads	Thread

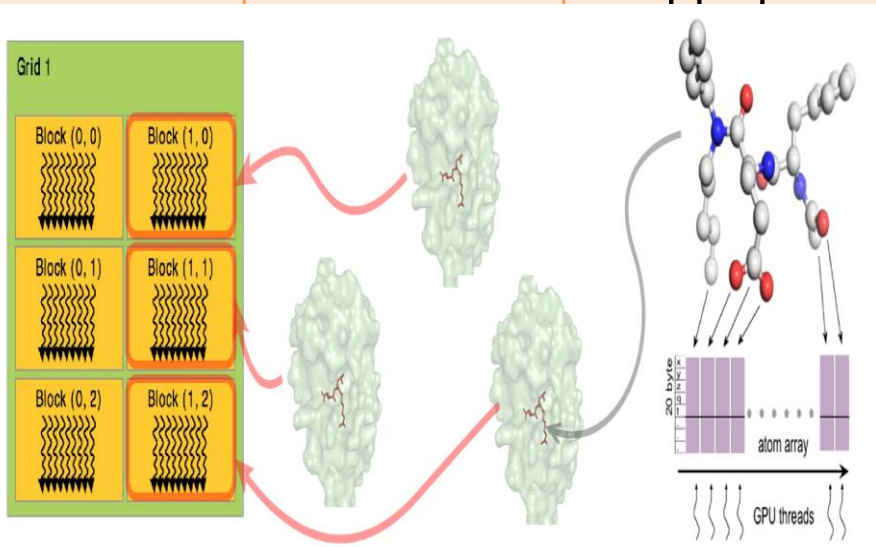

GeauxDock Benchmarks

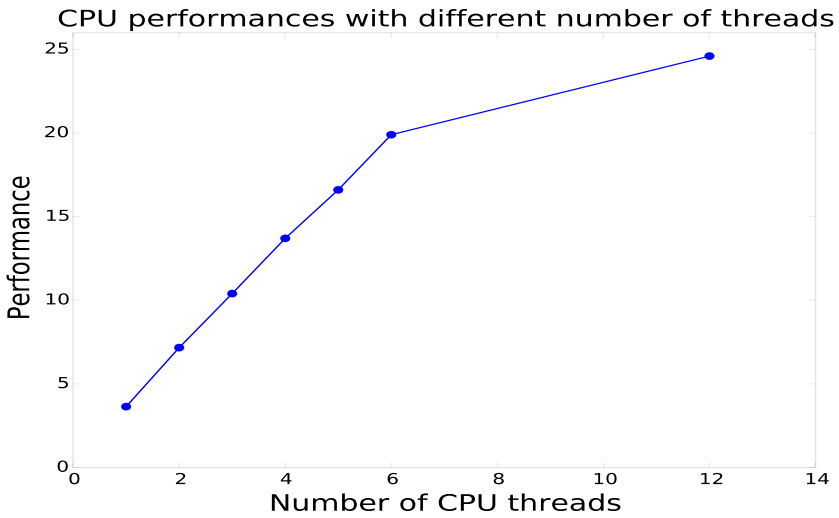

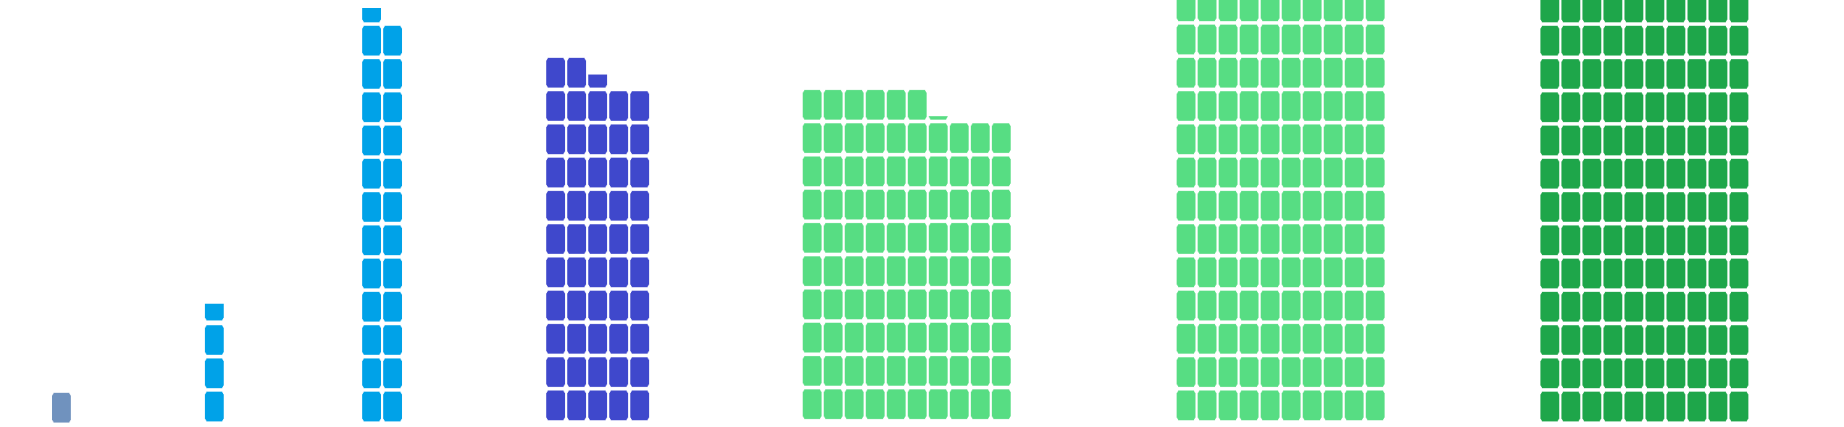

1.0X	3.64X	24.6X	52.4X	96.1X	193.7X	206.4X
C++prototype	CPU 1T	CPU 12T	Xeon Phi	M2090 x1	M2090 x2	GTX980 x1

LSU | Center for Computation & Technology

Interdisciplinary | Innovative | Inventive

Chemora

(**C**omputational **H**ierarchy for **E**ngineering **M**odel-**O**riented **R**e-adjustable **A**pplications)

- A framework for solving systems of PDEs
- Based on Cactus, prominent usage in the computational relativistic astrophysics community
- PDEs are expressed either in a high-level LATEX-like language or in Mathematica
- Discretization stencils are defined separately from equations, and can include Finite Differences, Discontinuous Galerkin Finite Elements (DGFE), Adaptive Mesh Refinement (AMR), and multi-block systems.
- **Use Chemora in the Einstein Toolkit to implement the Einstein Equations on CPUs and GPUs, and study astrophysical systems such as black hole binaries, neutron stars, and core-collapse supernovae**

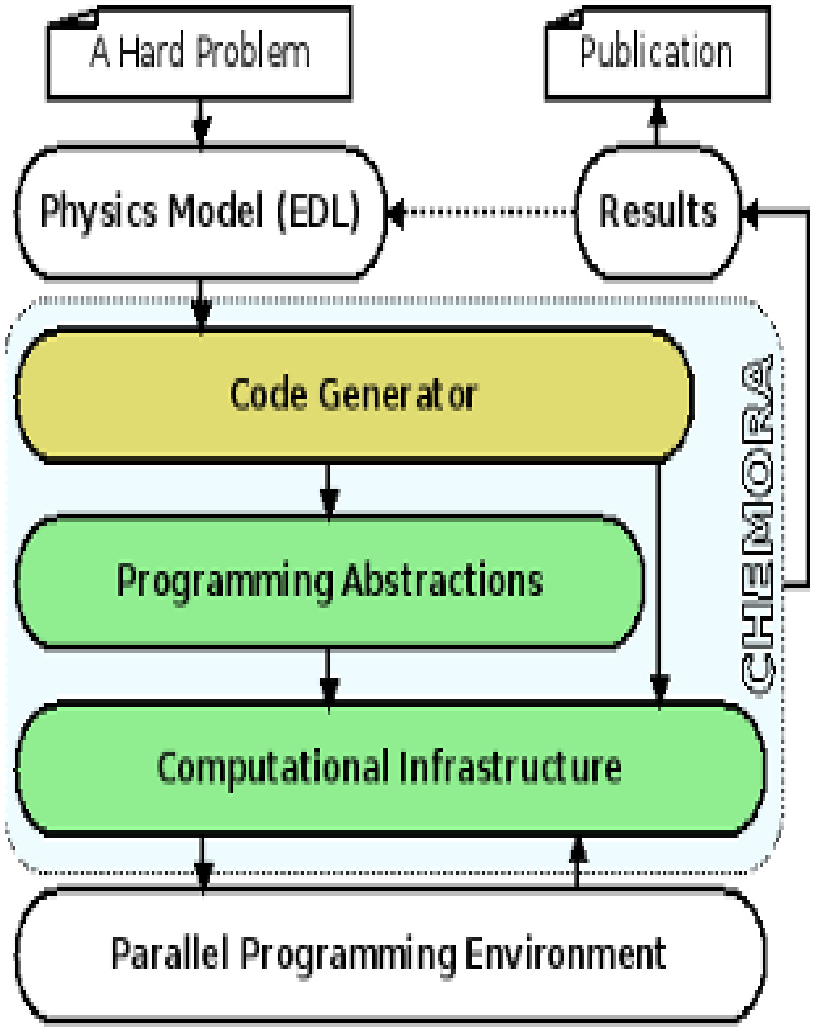

McLachlin Benchmark using Chemora

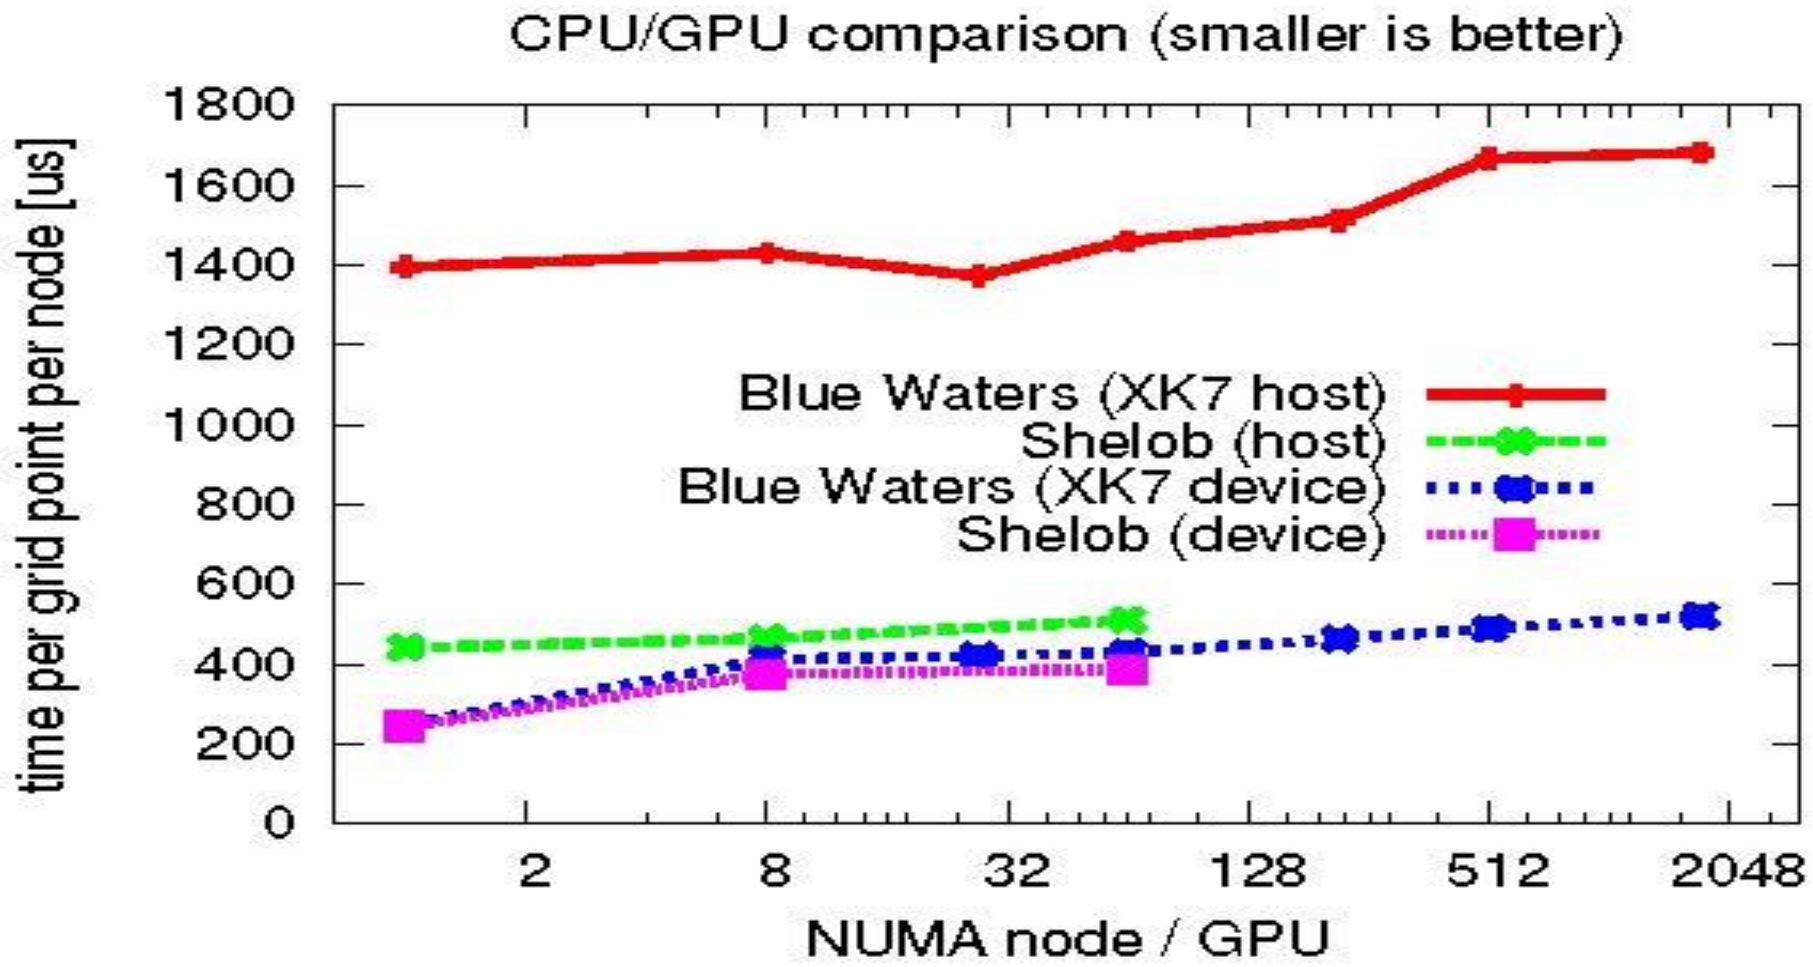

Parallel Tempering Simulation of the 3D Edwards-Anderson Spin Glass System

Design and implement a CUDA code for simulating the random frustrated a 3D Edwards-Anderson Ising model on GPUs.

Our overall design sustains a performance of **33.5** picoseconds per spin flip attempt, with parallel tempering moves.

Fastest GPU implementation for small to intermediate system sizes, comparable to FPGA implementation.

Accelerating Science & Engineering

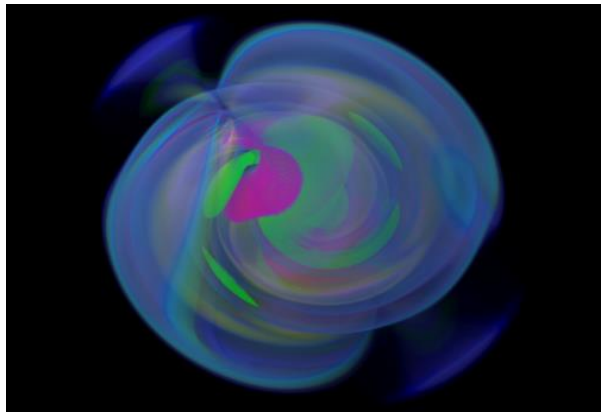

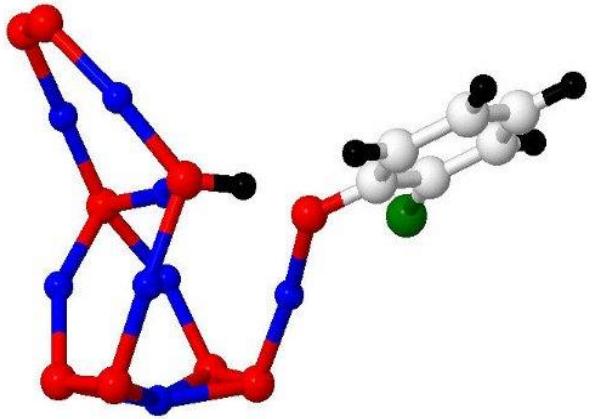

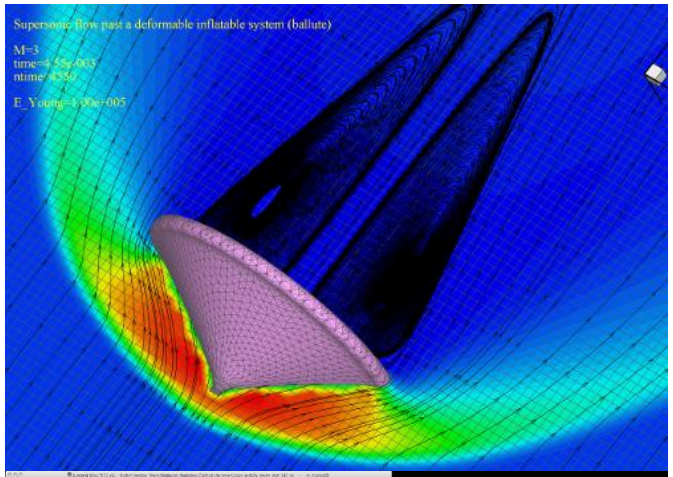

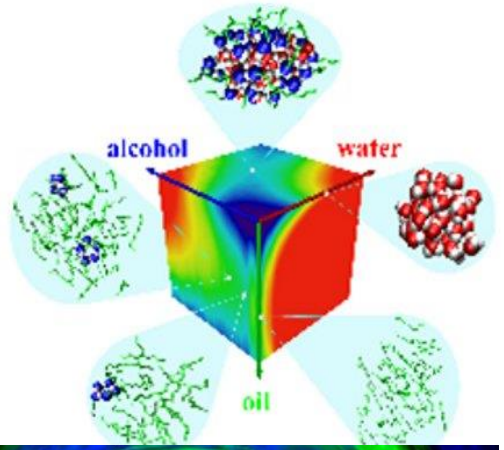

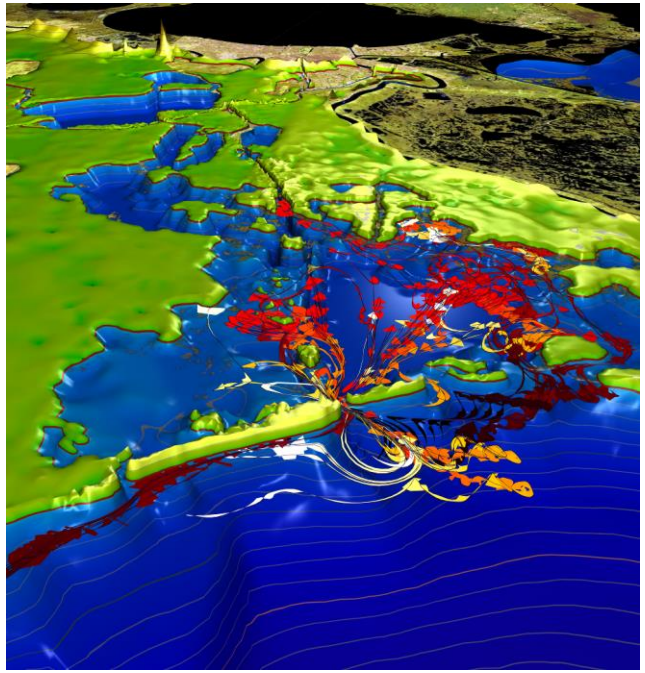

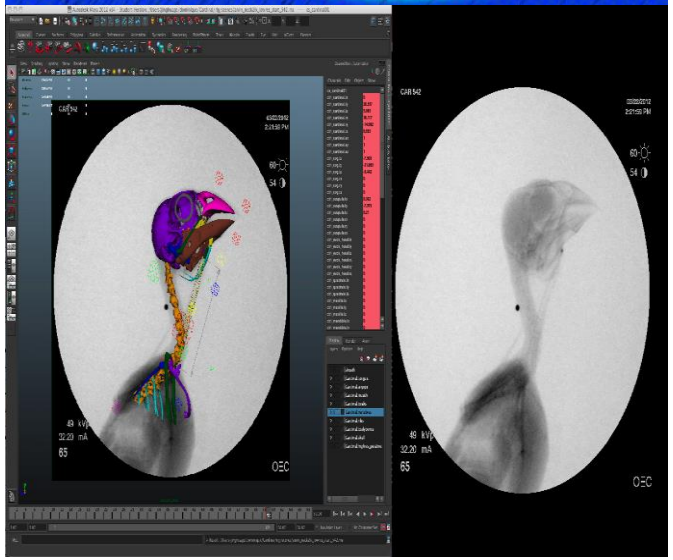

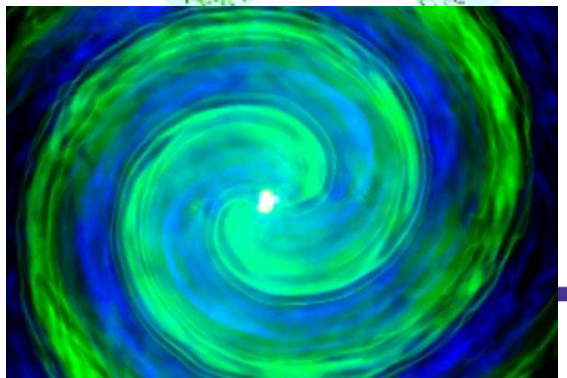

Center for Computation & Technology

Interdisciplinary | Innovative | Inventive

Summary

- Louisiana Cyberinfrastructure is growing tremendously!
- Heterogeneous Computing with GPUs has been enabling computational research and education in Louisiana
- NVIDIA - A long term partner, has helped us to accelerate computational science and engineering discoveries