Institute for Cyber-Enabled Research: Regional Organization to Promote Computation in Science

Dr. Dirk Colbry
colbrydi@msu.edu
Institute for Cyber Enabled Research

MICHIGAN STATE UNIVERSITY

ICER

Agenda

- Introduction
- Problems we are solving
- Hardware available at MSU

MICHIGAN STATE UNIVERSITY

ICER

1957 MISTIC Mainframe

- MSU's first mainframe
- Hand built by grad students
  - Dick Reid
  - Glen Keeney

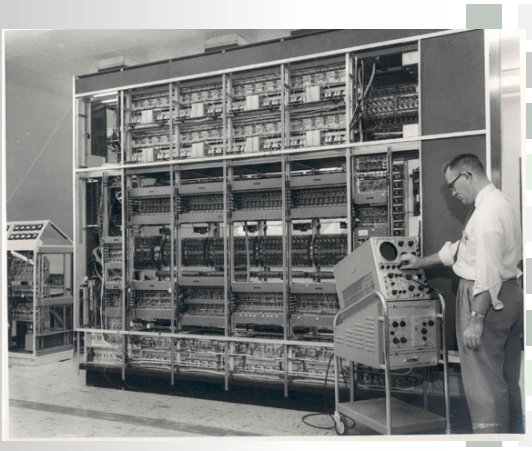

MICHIGAN STATE UNIVERSITY

ICER

After MISTIC

- 1957 MISTIC
- 1963-1973 CDC 3600
- 1967 Computer Science Department
- 1968 CDC 6500
- 1971 MERIT
- 1978 Cyber 750
- **2004 HPCC**
- **2009 ICER**

MICHIGAN STATE UNIVERSITY

ICER

2004 MSU HPCC

- Provide a level of performance beyond what you could get and reasonably maintain as a small group
- Provide a variety of technology, hardware and software, that would allow for innovation not easily found

MICHIGAN STATE UNIVERSITY

ICER

2009 iCER

The Institute for Cyber Enabled Research(iCER) at Michigan State University (MSU) was established to coordinate and support multidisciplinary resource for computation and computational sciences. The Center's goal is to enhance MSU's national and international presence and competitive edge in disciplines and research thrusts that rely on advanced computing.

MICHIGAN STATE UNIVERSITY

ICER

iCER Research Specialist

- Me
  - Research Consulting
  - HPCC Programming
  - Proposal Writing
  - Training and Education
  - Outreach

MICHIGAN STATE UNIVERSITY

ICER

Agenda

- Introduction
- **Problems we are solving**
- Hardware available at MSU

MICHIGAN STATE UNIVERSITY

ICER

Single Thread Jobs

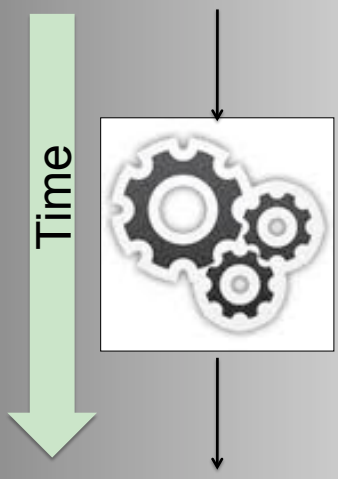

One CPU can only run one thing at a time. (sort of)

MICHIGAN STATE UNIVERSITY

ICER

Pleasantly Parallel

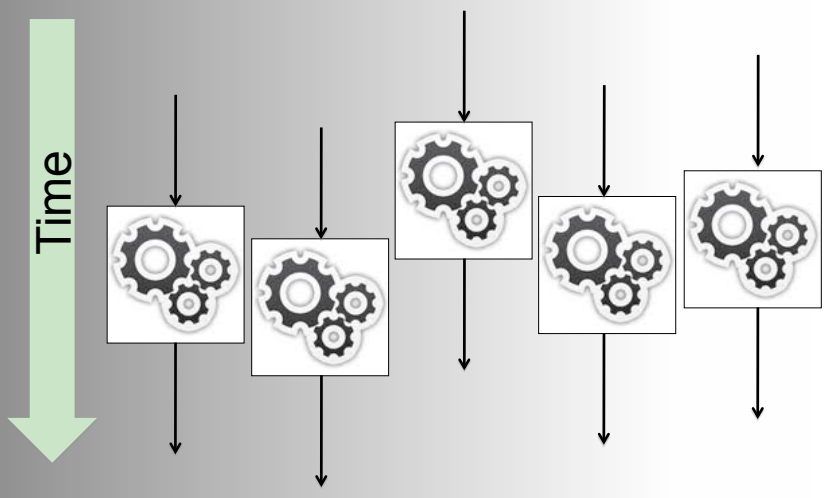

MICHIGAN STATE UNIVERSITY

ICER

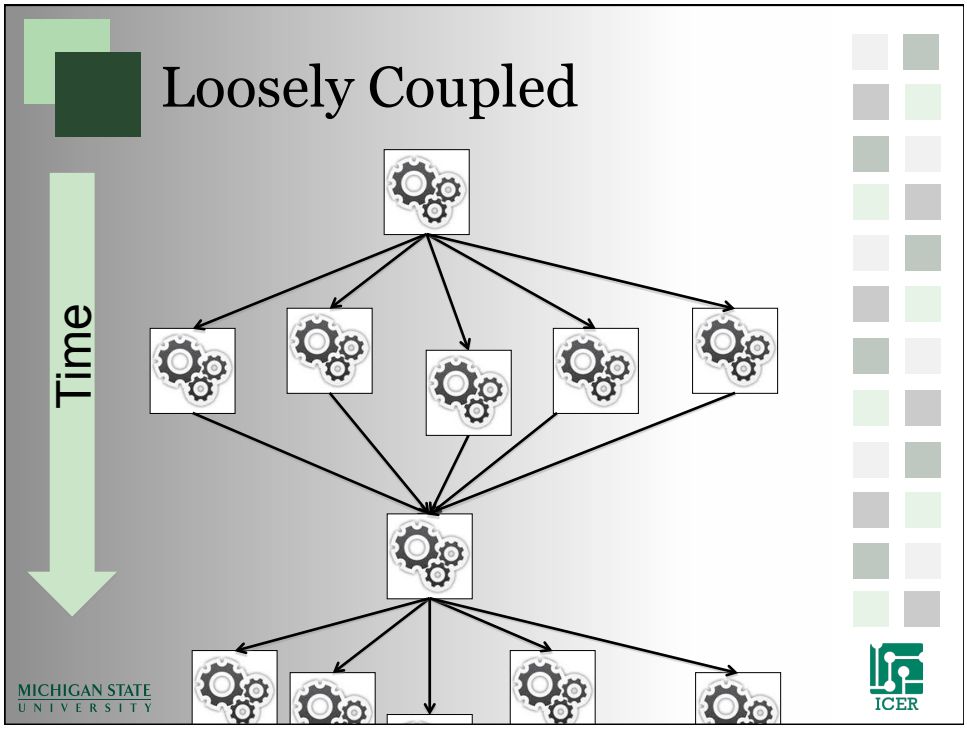
Loosely Coupled
Time
MICHIGAN STATE UNIVERSITY
ICER

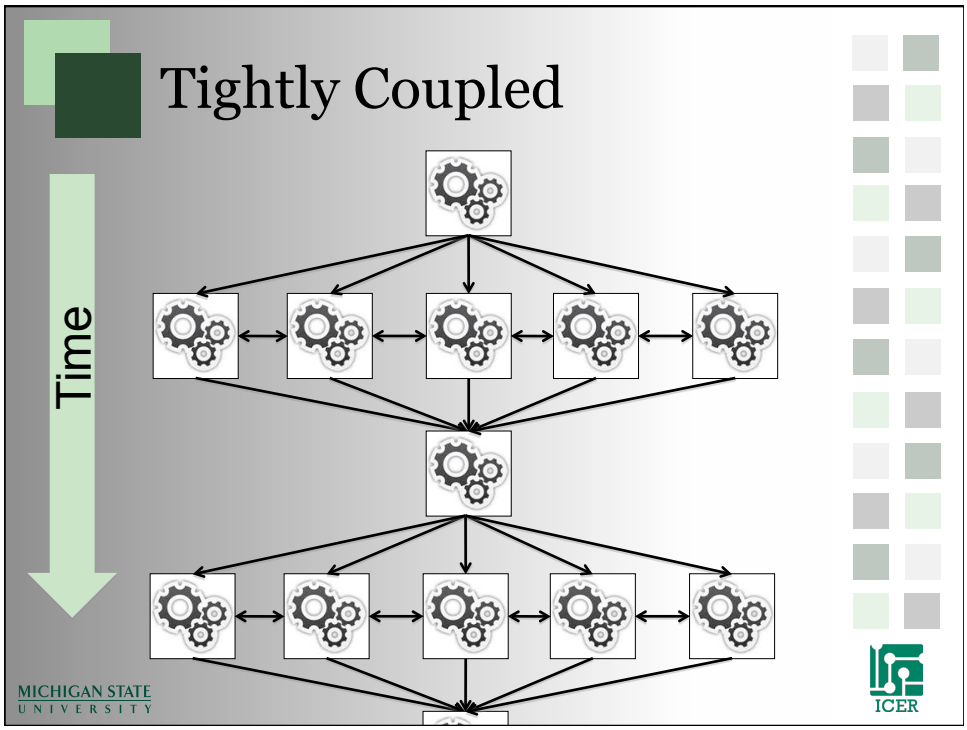
Tightly Coupled
Time
MICHIGAN STATE UNIVERSITY
ICER

What problems are we solving?

- Boundary Simulations
- Data Analysis
- Search

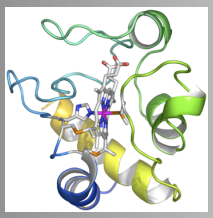

Image Provided by Dr. Warren F. Beck, MSU

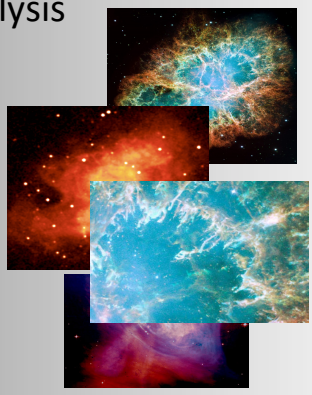

Images from, "Understanding the H_2 Emission from the Crab Nebula", C.T. Richardson, J.A. Baldwin, G.J. Ferland, E.D. Loh, Charles A. Huehn, A.C. Fabian, P.Salomé

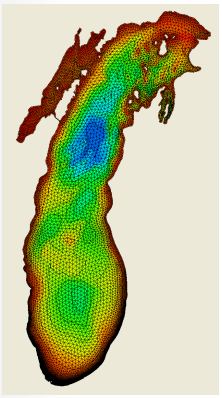

Image Provided by Dr. Mantha Phanikumar, MSU

MICHIGAN STATE UNIVERSITY

ICER

Boundary Simulations

- Fluid dynamics
- Finite element analysis
- Molecular dynamics
- Weather
- Etc.

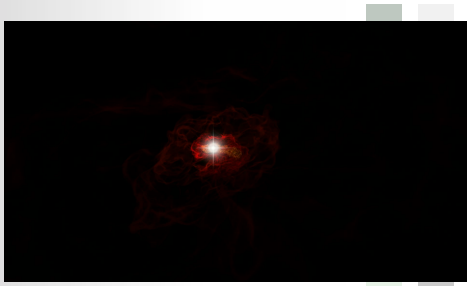

ENZO Simulation, Drs. O'Shea and Smith

- System of PDE (Partial Differential equations)
- Mathematically equivalent to inverse of a matrix

MICHIGAN STATE UNIVERSITY

ICER

Data Analysis

- Computer vision tasks
- Some Bioinformatics
- Astrophysics
- Etc.

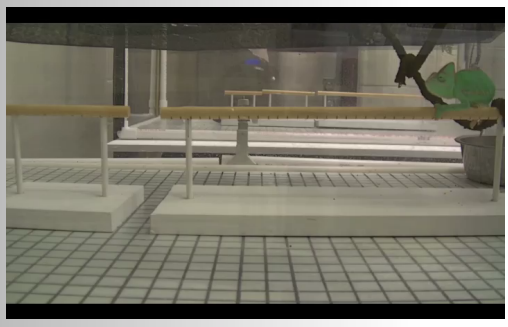

Video Provided by Dr. Fred Dyer

MICHIGAN STATE UNIVERSITY

ICER

Search

- Genome sequencing
- Analytics
- Optimization
- Etc.

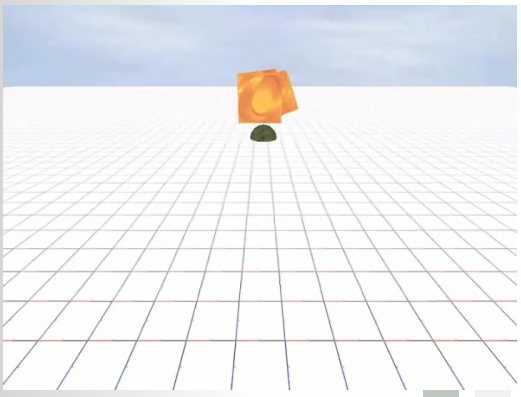

Evolution of an artificial organism that can move and forage for food, Dr. Nicolas Chaumont

MICHIGAN STATE UNIVERSITY

ICER

Agenda

- Introduction
- Problems we are solving
- **Hardware available at MSU**

MICHIGAN STATE UNIVERSITY

ICER

Communication

- Shared Memory
- Shared Network
- Distributed Network
- Dedicated Accelerators
- Hybrid Systems

MICHIGAN STATE UNIVERSITY

ICER

General Purpose Clusters

MICHIGAN STATE UNIVERSITY

ICER

Commodity Cluster

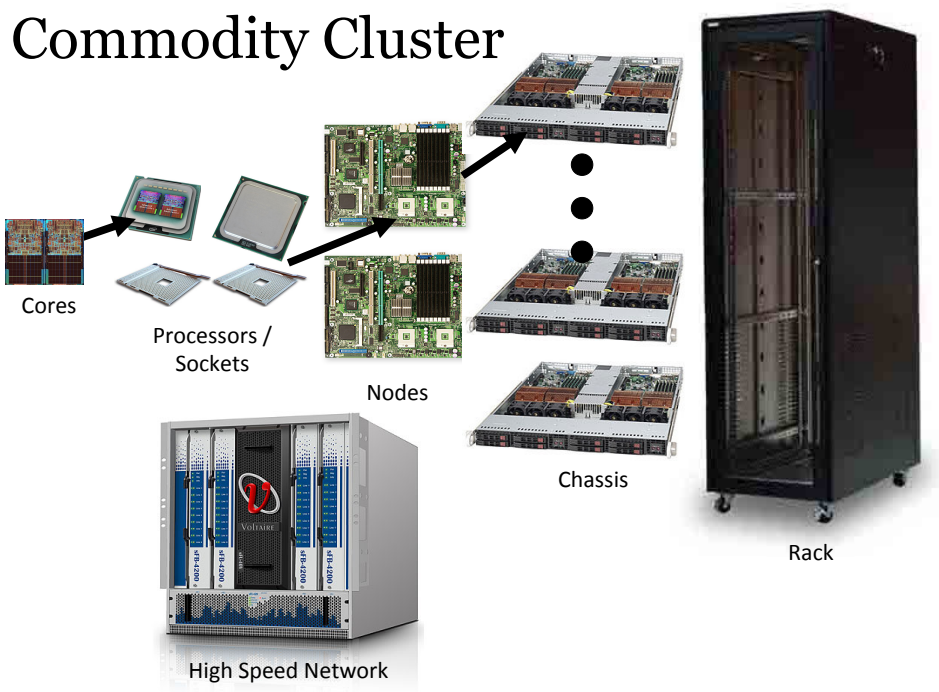

Large Shared Memory Systems (Fat Nodes)

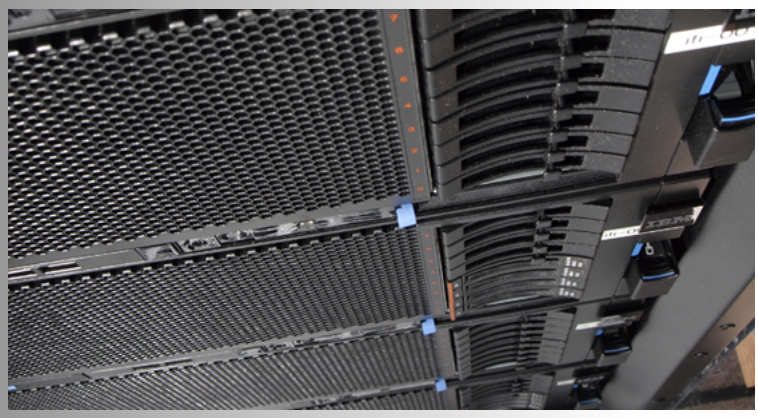

MICHIGAN STATE UNIVERSITY

ICER

Shared Memory Communication

- Cores on a processor share the same memory
- OpenMP
- Fat nodes
  - 64 cores
  - 2TB of memory

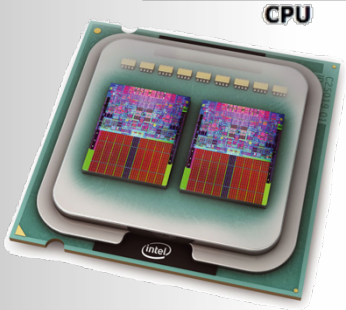

MICHIGAN STATE UNIVERSITY

ICER

General Purpose GPU Accelerated Systems

MICHIGAN STATE UNIVERSITY

ICER

GPUs

- Cards used to render graphics on a computer
- Hundreds of cores
- Not very smart cores
- But, if you can make your research look like graphics rendering you may be able to run really fast!

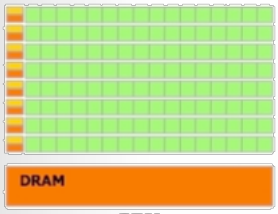

MICHIGAN STATE UNIVERSITY

ICER

Pros and Cons

- Benefits
  - Lots of processing cores.
  - Works with the CPU as a co-processor
  - Very fast local memory bandwidth
  - Large online community of developers
- Drawbacks
  - Can be difficult to program.
  - Memory Transfers between GPU and CPU are costly (time).
  - Cores typically run the same code.
  - Errors are not detected (on older cards)
  - Double precision calculations are slow (On older cards)

MICHIGAN STATE UNIVERSITY

ICER

Intel Xeon Phi

- Cross between CPU and GPU
- About 60 Pentium I cores
  - Less cores than GPU
  - Easier to use than GPU
    - OpenMP
    - MPI
- Very new
  - January 2013

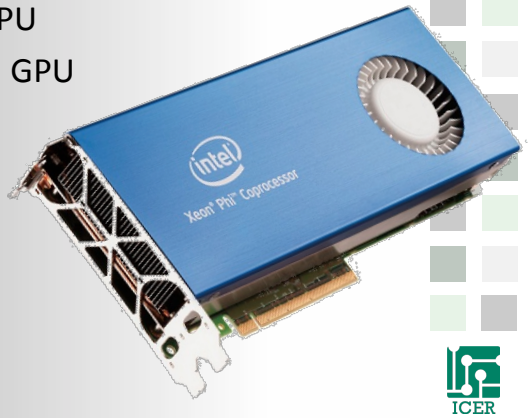

MICHIGAN STATE UNIVERSITY

ICER

High Throughput HTCondor Cluster

MICHIGAN STATE UNIVERSITY

ICER

Condor
High Throughput Computing

- Job submission system
- Runs like a screen saver
- Steals CPU Cycles

MICHIGAN STATE UNIVERSITY

ICER

Which approach is the best?

- Depends on what you are doing?
- Depends on how much communication you need.
- Depends on what hardware you have.
- Depends on how much time you have.

MICHIGAN STATE UNIVERSITY

ICER

Questions?

MICHIGAN STATE UNIVERSITY

ICER