Utilizing Advanced Computing in Research

Dirk Colbry, Research Specialist
Institute for Cyber-Enabled Research

MICHIGAN STATE UNIVERSITY

ICER

What is Advanced Computing Hardware?

- Anything more advanced than your desktop
- Local resources
  - Lab, Department, Institution (iCER)
- National resources
  - NSF (XSEDE), DOE (Jaguar) , Others
- Commercial Resources (cloud computing)
  - Amazon, Azure, Liquid Web, Others

MICHIGAN STATE UNIVERSITY

ICER

Why use Advanced Computing Hardware?

- Science takes too long
- Computation runs out of memory
- Run out of disk space
- Need licensed software
- Need advanced interface (visualization)

MICHIGAN STATE UNIVERSITY

ICER

1957 MISTIC Mainframe

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

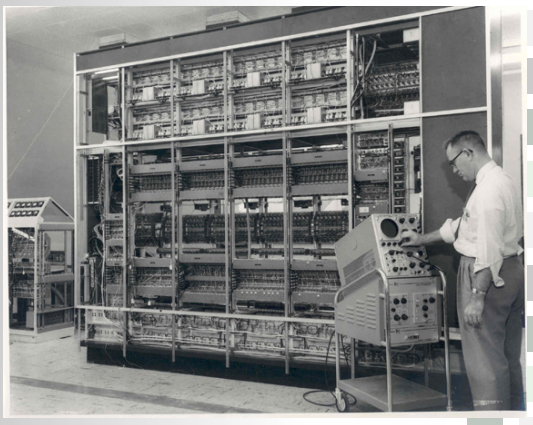

MICHIGAN STATE UNIVERSITY

ICER

iCER and The HPCC

FREE*

- Over 3110 Nodes
- Over 23072 CPU cores
- Over 363 TB of scratch memory space
- 2TB Shared memory machines
- 50GB backed up home directory space
- GPGPU cluster with 64 Tesla Nodes
- High Throughput condor cluster
- Specialized Bioinformatics VMs

MICHIGAN STATE UNIVERSITY

ICER

Available Software

- Center Supported Development Software
  - Intel compilers, openmp, openmpi, mvapich, totalview, mkl, pathscale, gnu...
- Center Supported Research Software
  - MATLAB, R, fluent, abaqus, HEEDS, amber, blast, ls-dyna, starp...
- Customer Software
  - gromacs, cmake, cuda, imagemagick, java, openmm, siesta…
  - For a more up to date list, see the documentation wiki:
    - http://wiki.hpcc.msu.edu/

MICHIGAN STATE UNIVERSITY

ICER

What if I want more?

XSEDE
Extreme Science and Engineering Discovery Environment

NASA

Open Science Grid

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

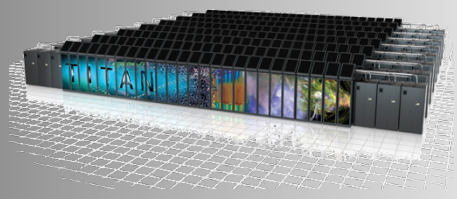

amazon web services™

MICHIGAN STATE UNIVERSITY

ICER

Science and Computation

- Pillars of Science:
  - Theory
  - Experimentation
  - Simulation
  - (Big) Data

Wired Magazine, 2009

MICHIGAN STATE UNIVERSITY

ICER

How much Communication?

- Pleasantly parallel
  - No communication required
- Loosely coupled
  - Typically sync at regular intervals
- Tightly coupled
  - Constant communication

MICHIGAN STATE UNIVERSITY

ICER

Pleasantly Parallel

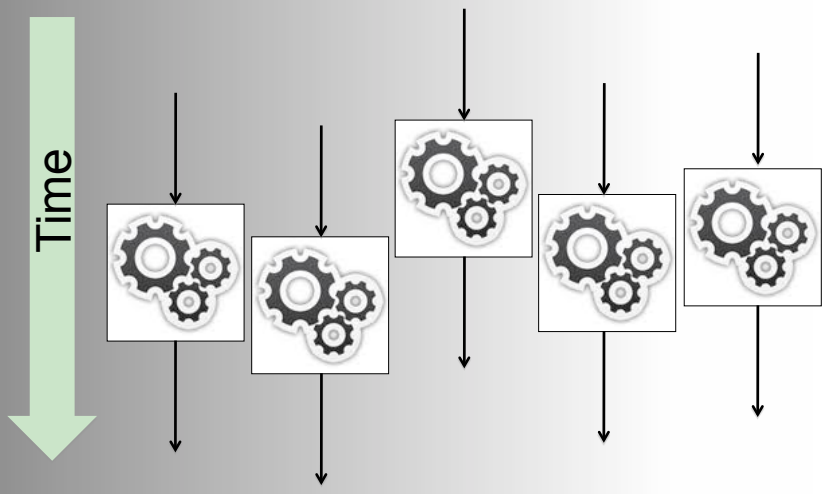

MICHIGAN STATE UNIVERSITY

ICER


Condor
High Throughput Computing
• Job submission system
• Runs like a screen saver
• Steals CPU Cycles
MICHIGAN STATE UNIVERSITY
ICER

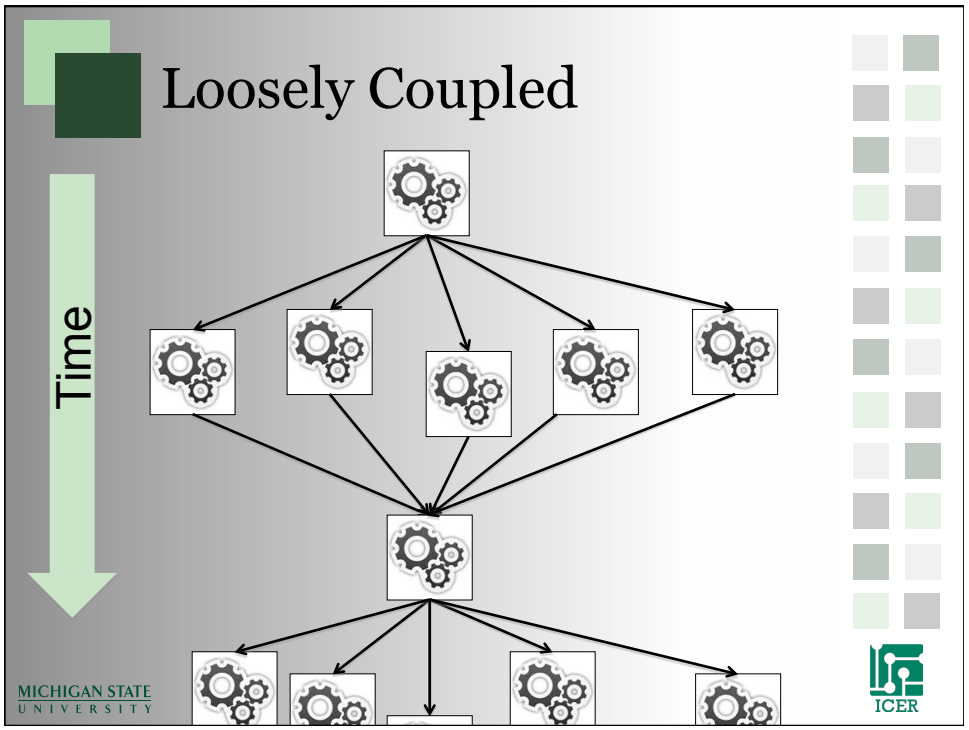
Loosely Coupled
Time
MICHIGAN STATE UNIVERSITY
ICER

Shared Network Communication

- Beowulf Cluster
- Commodity Computers
- High speed network
- MPI
  - Message Passing Interface
  - Programming library
- Parallel File systems
  - Luster

MICHIGAN STATE UNIVERSITY

ICER

Tightly Coupled

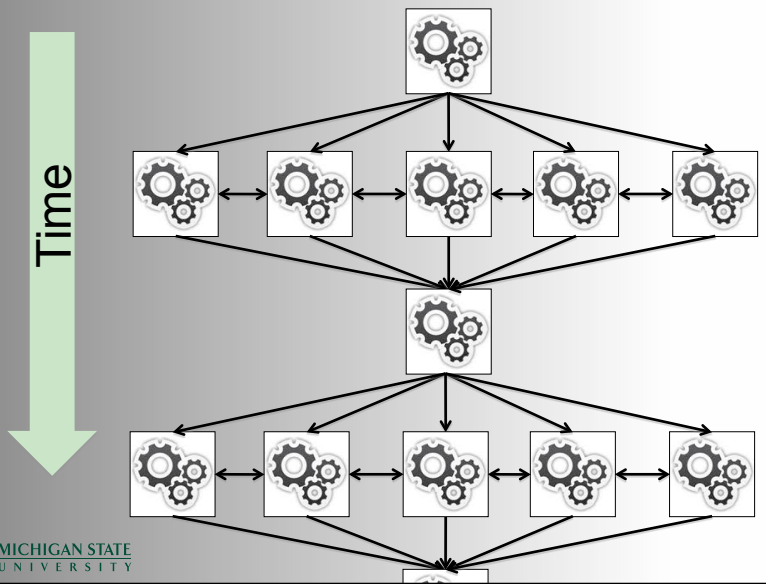

MICHIGAN STATE UNIVERSITY

ICER

Shared Memory Communication

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

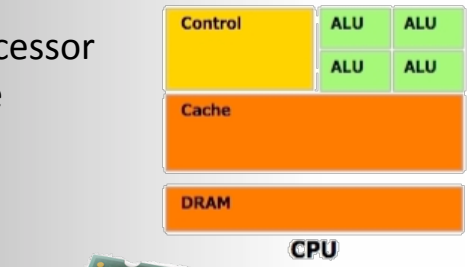

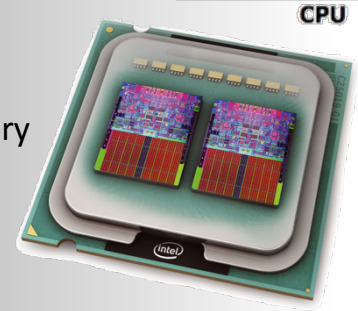

MICHIGAN STATE UNIVERSITY

ICER

GPU

- 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!

NVIDIA CUDA TEACHING CENTER

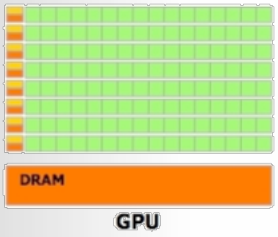

MICHIGAN STATE UNIVERSITY

ICER

Coming SOON
Intel Xeon Phi

- Cross between CPU and GPU
- About 61 Pentium III cores
  - Less cores GPU
  - Easier to use than GPU
- Very new

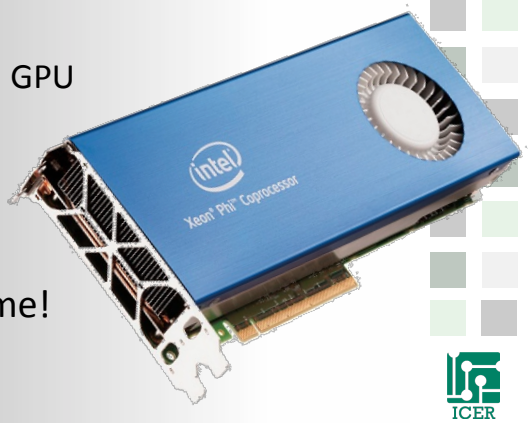

- We are getting some!

MICHIGAN STATE UNIVERSITY

ICER

Summary of Hardware

- Pleasantly parallel
  - Condor
- Loosely Coupled
  - Beowulf cluster
- Tightly Coupled
  - Fat Nodes
  - GPUs
  - Phi

MICHIGAN STATE UNIVERSITY

ICER

Example Problems

- Boundary Simulations
- Data Analysis
- Search

MICHIGAN STATE UNIVERSITY

ICER

Example: Boundary simulations

1. Divide a 2D or 3D simulation space into a grid of cells
2. Define information that is transferred at the boundary of the cells
3. Simulate the dynamics of the cell during a time interval
4. Repeat steps 2 and 3

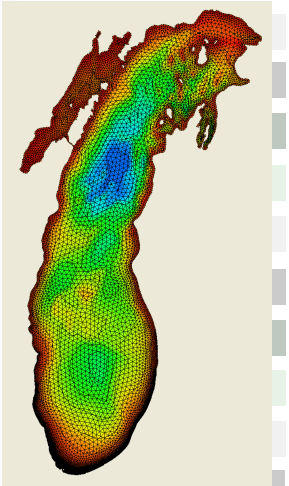

Image Provided by Dr. Mantha Phanikumar, MSU

MICHIGAN STATE UNIVERSITY

ICER

Boundary Simulations

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

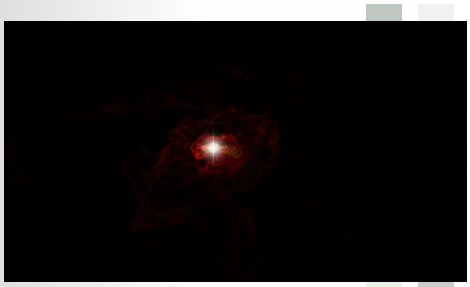

ENZO Simulation, Drs. O'Shea and Smith

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

MICHIGAN STATE UNIVERSITY

ICER

Boundary Simulation

- Tightly to loosely coupled
- Typically solved with MPI
- PDE solutions available for GPU and OpenMP

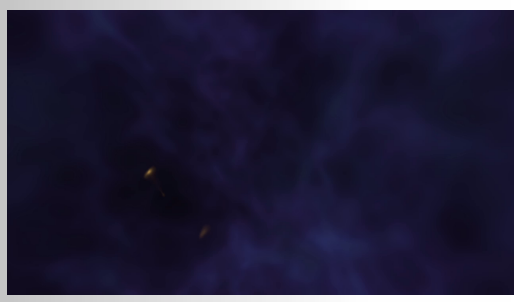

Simulation of the formation of a spiral galaxy
Drs. Turk, O'Shea and Smith

MICHIGAN STATE UNIVERSITY

ICER

Example: Data Analysis

1. Input data file
2. Find features, search or filter data in some way
3. Output Results

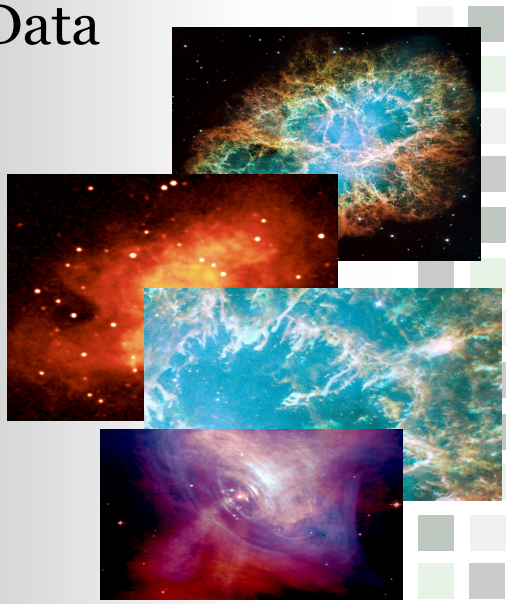

Images from, "Understanding the H_2 Emission from the Crab Nebula", C.T. Richardson, J.A. Baldwin, G.J. Ferland, E.D. Loh, Charles A. Huehn, A.C. Fabian, P.Salomé

MICHIGAN STATE UNIVERSITY

ICER

Data Analysis

- Computer vision tasks
- Bioinformatics
- Astrophysics
- Etc.

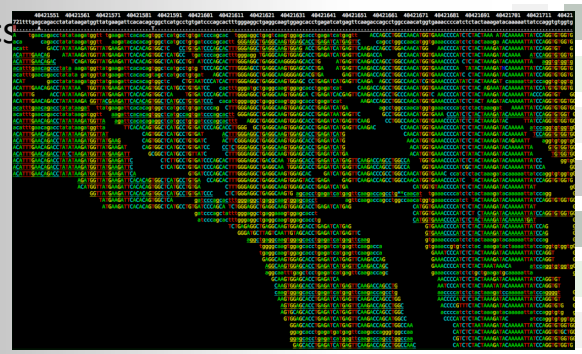

Image generated using SAMtools

MICHIGAN STATE UNIVERSITY

ICER

Data Analysis

- Loosely coupled
- Bulk of computation is typically pleasantly parallel
- Can be I/O bound

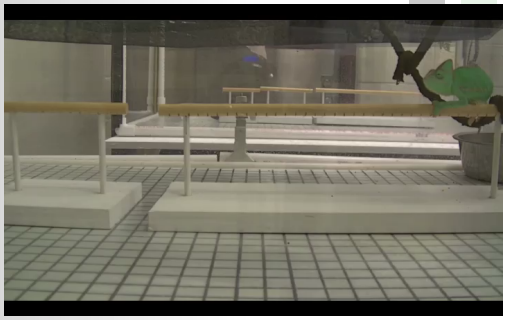

Video Provided by Dr. Fred Dyer

MICHIGAN STATE UNIVERSITY

ICER

Example: Search

- Randomly generate test candidates
- Evaluate the quality of solution
- Repeat until found

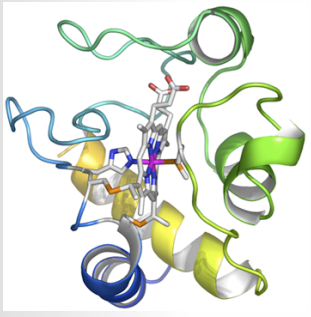

Image Provided by Dr. Warren F. Beck, MSU

MICHIGAN STATE UNIVERSITY

ICER

Search

- Evolution (Avida)
- Genetic Algorithms (Heeds)
- RANSAC
- Monte Carlo
- Etc.

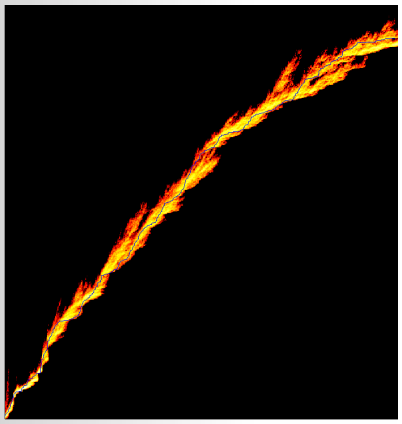

Image Provided by Dr. Charles Ofrea, MSU

MICHIGAN STATE UNIVERSITY

ICER

Search

- Pleasantly parallel
- The more the better
- Typically not I/O bound
- Typically not memory bound

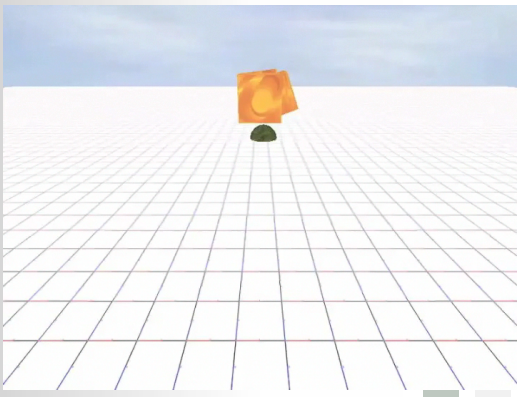

Evolution of an artificial organism that can move and forage for food, Dr. Nicolas Chaumont

MICHIGAN STATE UNIVERSITY

ICER

MSU Seminars in Research and Instructional Technology

May 7 and 8, 2013

- Two days of no-cost seminars to faculty and graduate students on technology topics
  - Morning sessions run from 8:30 to 11:30 am
  - Afternoon sessions run from 1:30 to 4:30 pm
  - Lunch is provided that will feature guest speakers on instructional technology
  - Introduction to HPC (Time TBD)
  - Advanced HPC (Time TBD)

MICHIGAN STATE UNIVERSITY

http://train.msu.edu/faculty/seminars/

ICER

Getting Help

- Documentation and User Manual – wiki.hpcc.msu.edu
- Contact HPCC and iCER Staff for:
  - Reporting System Problems
  - HPC Program writing/debugging Consultation
  - Help with HPC grant writing
  - System Requests
  - Other General Questions
- Primary form of contact - www.hpcc.msu.edu/contact
- HPCC Request tracking system – rt.hpcc.msu.edu
- HPCC Phone – (517) 353-9309
- HPCC Office – 1400 PBS
- Office Hours – Monday – Friday 9am-5pm

MICHIGAN STATE UNIVERSITY

ICER