Introduction to iCER and the HPCC

Dr. Dirk Colbry
Research Specialist
Institute for Cyber-Enabled Research

MICHIGAN STATE UNIVERSITY

ICER

Agenda

- Introduction to HPCC/iCER
- Bioinformatics Support at iCER
- Doing more faster - Pleasantly Parallel

MICHIGAN STATE UNIVERSITY

ICER

Science and Computation

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

Wired Magazine, 2009

MICHIGAN STATE UNIVERSITY

ICER

What is Advanced Computing Hardware?

- Anything more advanced than your desktop
- Local resources
  - Lab, Department, Institution (HPCC)
- National resources
  - NSF (XSEDE, Blue Waters), 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
- Need licensed software
- Need advanced interface (visualization)

MICHIGAN STATE UNIVERSITY

ICER

Types of problems

- CPU bound
  - Lots of computing (think simulation)
- Memory bound
  - Requires lots of memory (think genomics)
- I/O bound
  - Requires lots of data (think astronomy)

** Many problems fall in more than one category

MICHIGAN STATE UNIVERSITY

ICER

Types of Systems

Local Resources

- Special Use Systems
  - SMP – Symmetric Multiprocessor (fat node)
  - GPGPU (General Purpose Graphics Processing Unit)
- Typical HPC cluster
  - Commodity computers
  - High speed backbone
  - High speed network storage

National and Commercial

- Advanced HPC
  - Specialty hardware
  - High speed backbone
  - High speed storage
- Grid
  - Many HPC systems linked together by high speed network
- Cloud
  - Lots of definitions
  - Typically refers to computing as a service using highly flexible virtual machines

MICHIGAN STATE UNIVERSITY

ICER

1957 MISTIC Mainframe

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

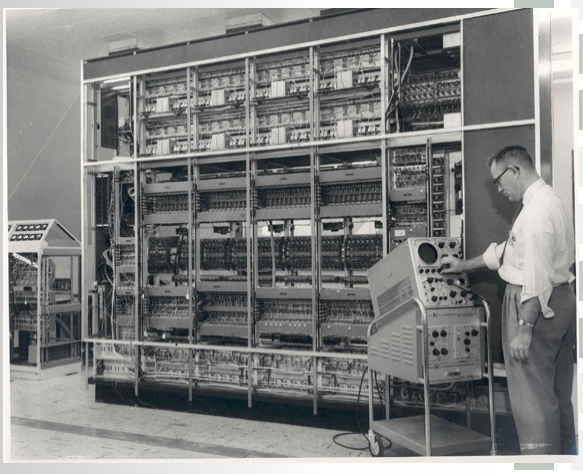

MICHIGAN STATE UNIVERSITY

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

2005 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

10

ICER

HPC Systems

FREE*

- Large Memory Nodes (up to 2TB!)
- GPU Accelerated cluster (K20, M1060)
- PHI Accelerated cluster (5110p)
- Over 540 nodes, 7000 computing cores
- Access to high throughput condor cluster
- 363TB high speed parallel scratch file space
- 50GB replicated file spaces
- Access to large open-source software stack and specialized bioinformatics VMs

MICHIGAN STATE UNIVERSITY

ICER

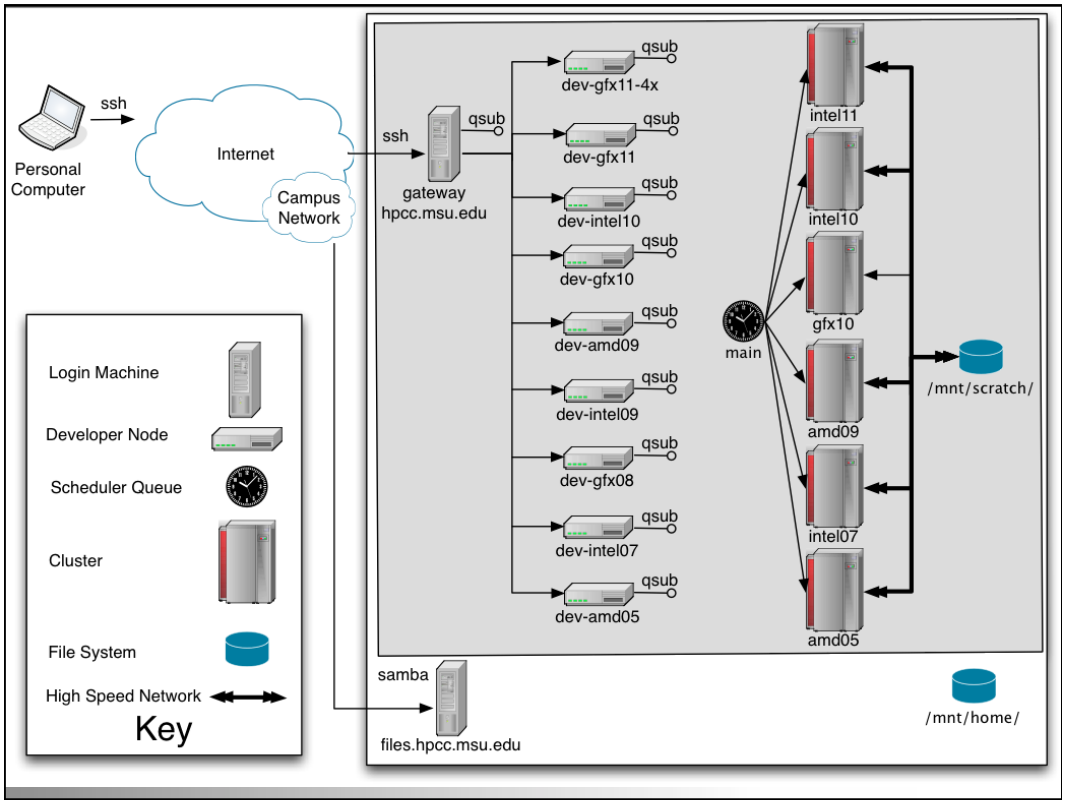

2009 iCER

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

Bigger Science

- The goal of iCER is NOT:
  - Kflops
- Instead, the goal of iCER IS:
  - KSciences / second
- Doing More Science, Faster
  - Reducing the "Mean time to Science"
- iCER is designed to help researchers do their science and when appropriate scale them up to one of the national labs

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

Accounts

- Each user has 50Gigs of backed-up personal hard drive space.
  - /mnt/home/username/
- Users have access to more than 363TB of high speed parallel scratch space.
  - /mnt/scratch/username/
- Shared group space is also available upon request.

MICHIGAN STATE UNIVERSITY

ICER

Buy-In Opportunities

- Every 6 months to a year
- We will maintain your computers for you
- Researchers get exclusive use of their nodes within 4 hours of submitting a job
- Buy-in jobs will automatically overflow into the general resources.

MICHIGAN STATE UNIVERSITY

ICER

2013 Buy-in

- In progress now!
  - $4,200* - General purpose base nodes
    - 20 cores, 64 gb
  - $5,500* - 128gb nodes
  - $8,200* - 128gb and 2xK20 GPU nodes
  - $9,200* - 128gb and 2xPhi Card nodes

* Estimate only, other costs may apply. Please talk to me if you would like more details

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

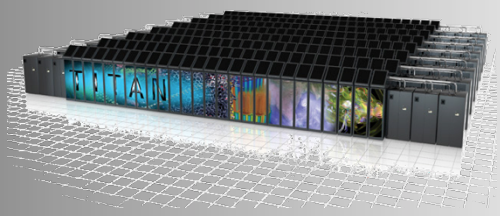

amazon web services™

MICHIGAN STATE UNIVERSITY

ICER

More than Just Hardware

- Education
- Grant Writing / Support
- Advancing Computational Science

MICHIGAN STATE UNIVERSITY

ICER

User Training, Education and Support

- Local Workshops
  - Software carpentry
  - Introduction to Linux and HPCC
  - Advanced HPCC
- Remote Training
  - VSCSE – Virtual School for Computer Science Education
  - XSEDE training Workshops

software carpentry

XSEDE
Extreme Science and Engineering Discovery Environment

vscse

MICHIGAN STATE UNIVERSITY

ICER

2013
CYBER INFRASTRUCTURE DAYS
OCTOBER 24-25
Save The Date!

http://tech.msu.edu/CI-Days/

MICHIGAN STATE UNIVERSITY

ICER

People

- (Future) iCER Director
  - Dr. Kenneth M Merz
- HPCC Director
  - Dr. Bill Punch
- Administrative Assistant
  - Kelly Osborn



MICHIGAN STATE UNIVERSITY

ICER

HPC Administrators

- Andrew Keen
- Jim Leikert
- Gregory Mason
- Jason Muffett




MICHIGAN STATE UNIVERSITY

ICER

iCER Research Specialist

Ben Ong and Me

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

MICHIGAN STATE UNIVERSITY

ICER

iCER Bioinformatics Support

MICHIGAN STATE UNIVERSITY

ICER

Bioinformatics Outreach

- Biological analysis software
- Common data set repositories
- Database services
- Virtual Machines
- Research-based storage space
- Seminars
- One-on-one Consulting
- Programming and scripting assistance

wiki.hpcc.msu.edu/display/Bioinfo/Bioinformatics+Support+at+MSU

MICHIGAN STATE UNIVERSITY

ICER

Some Examples

- Multithreaded BLAST – shared memory
- mpiBLAST – distributed memory
- Velvet Assembler – multithreaded shared mem
- MAKER2 – MPI, distributed memory
- Parallel Solutions: OpenMP, OpenMPI, MVAPICH, MVAPICH2, CUDA

MICHIGAN STATE UNIVERSITY

ICER

Virtual Machines

- Virtual “Servers” expressed in software
- Available for research labs/working groups
- Flavors currently available:
  - Galaxy (single node only)
  - BLAST (web browser based)
  - UCSC/GBrowse Genome Browsers
  - MySQL database services (postgres can be provided also)

MICHIGAN STATE UNIVERSITY

ICER

Tutorials

- Titus Brown's ANGUS-NGS tutorials, converted for using examples on HPC instead of Amazon
- Using UCSC for certain tasks
- mpiBLAST
- Velvet and Oases
- Others being developed...

MICHIGAN STATE UNIVERSITY

ICER

Bioinformatics Contact

- Ticket requests:
  - https://rt.hpcc.msu.edu/index.html
  - Please include "Bioinformatics Help" in the subject to more quickly route your request.

MICHIGAN STATE UNIVERSITY

ICER

Doing more faster

Pleasantly Parallel Techniques

MICHIGAN STATE UNIVERSITY

ICER

Single Thread Jobs

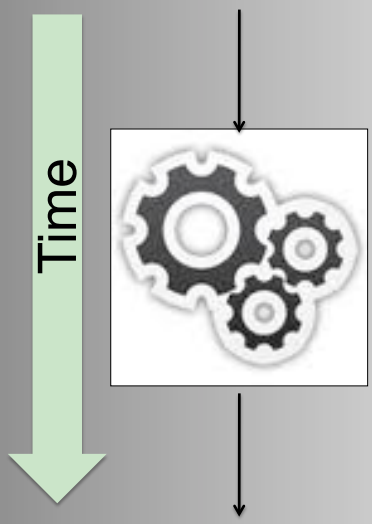

One CPU can only run one thing at a time. (sort of)

MICHIGAN STATE UNIVERSITY

ICER

Communication

- Shared Memory
  - Ex. OpenMP
- Shared Network
  - Ex. MPI
- Distributed Network
  - Ex. Map-Reduce
- Dedicated Accelerators
  - Ex. GPGPU and Phi
- Hybrid Systems

MICHIGAN STATE UNIVERSITY

ICER

Pleasantly Parallel

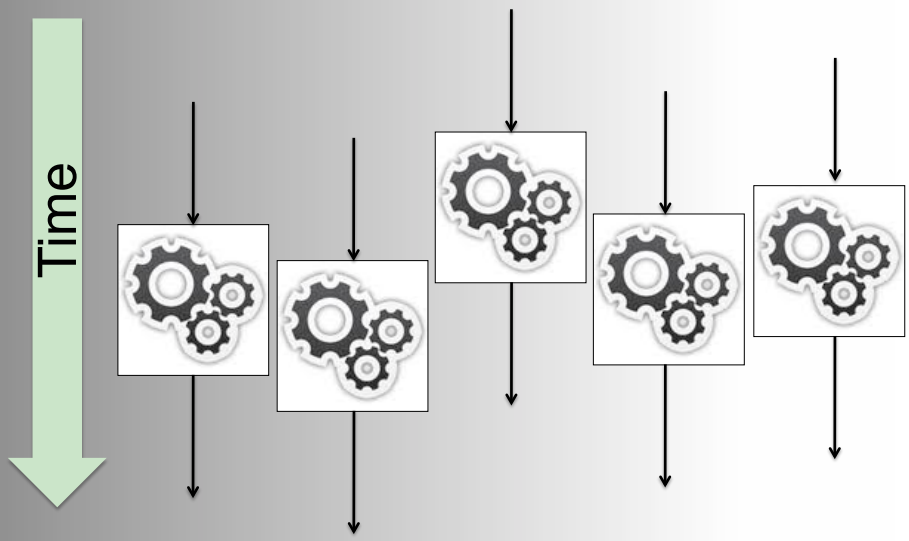

MICHIGAN STATE UNIVERSITY

ICER

How fast can we go?

- T - How long does each operation take?
- N - How many operations do you need to run?
- CPUs – Number of Cores job will run on.

- Single CPU time estimate:
  - TxN
- Best possible Pleasantly parallel time:
  - (TxN)*overhead/CPUs

MICHIGAN STATE UNIVERSITY

ICER

Who are you? -- Biometrics

MICHIGAN STATE UNIVERSITY

ICER

Pairwise-All Problem

- Database of faces
- Compare everything to everything else
- Calculate a Matching score to use for identification http://tech.msu.edu/CI-Days/

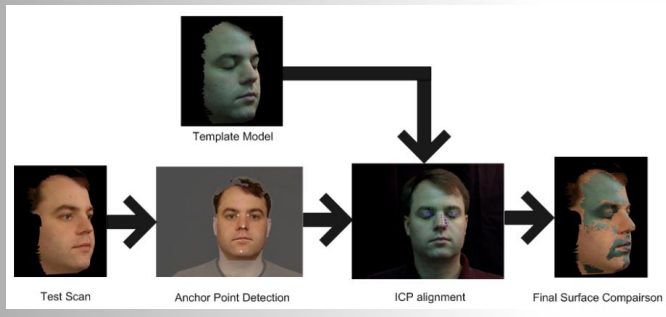

MICHIGAN STATE UNIVERSITY

ICER

943 x 943 Similarity Matrix

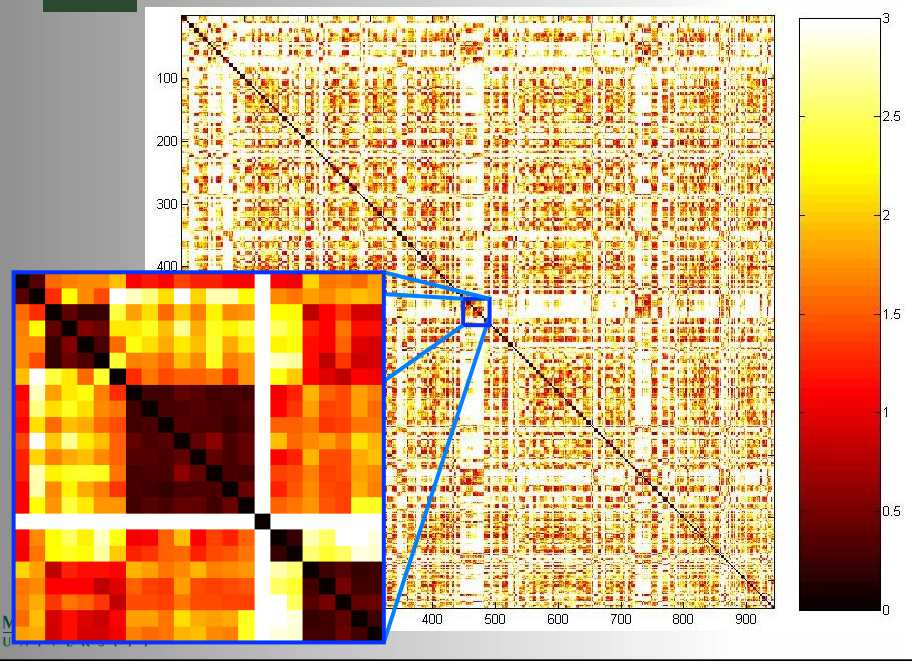

ICER

Estimated Calculation Times

- Preprocessing
  - 943 * 12 (seconds) ≈ 189 Minutes
- Matching
  - 943 * 943 * 5 (seconds) ≈ 103 Days
- Scans matched to themselves always result in 0 mm
  - (943 * 943 – 943) * 5 (seconds) ≈ 103 Days
- The Proposed Alignment Algorithm is symmetric.
  - (943 * 943 – 943)/2 * 5 (seconds) ≈ 51.5 Days
- We also load models once per row instead of every time
  - (943*943-943)/2 * 3 (seconds) + 943 * 2 (seconds) ≈ 31 Days

MICHIGAN STATE UNIVERSITY

40

ICER

Calculation Time for Full Similarity Matrix

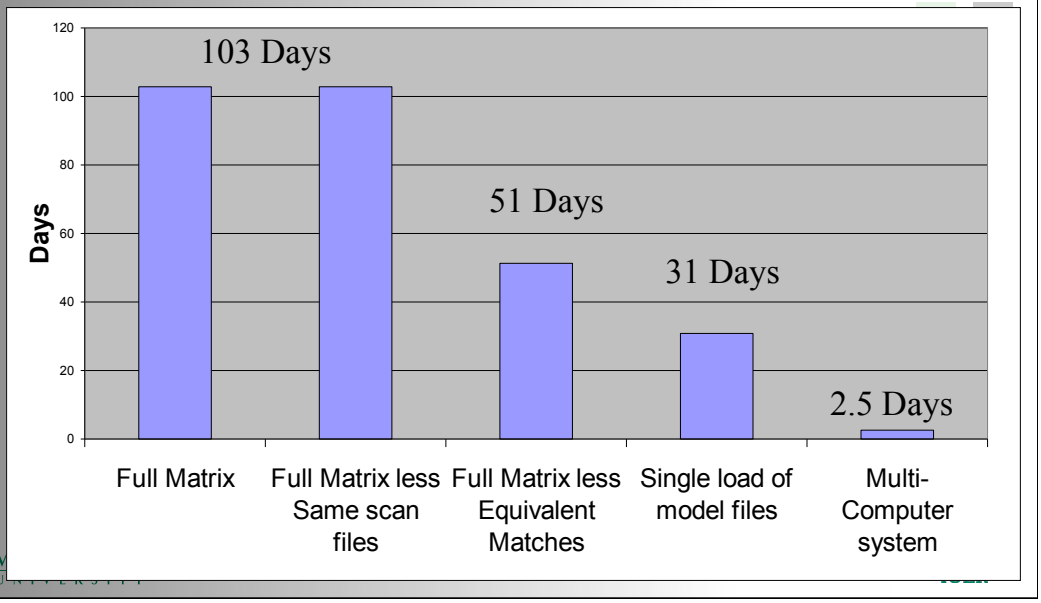

How do we go even bigger?

- 5000 scans.
  - 1.5 years on a single processor computer
  - 13 days on our ad-hoc cluster.
  - 1.5 days on the HPCC

MICHIGAN STATE UNIVERSITY

ICER

Step to Pleasantly Parallel

- Figure out command line
- Estimate single job time:
  - Should be > 5 minutes
  - Should be < 1 week
  - Best if < 4 hours
- Make a submissions script
- Submit Job

MICHIGAN STATE UNIVERSITY

ICER

Example

- Folder full of input files:

1.in	5.in	9.in	13.in	17.in
2.in	6.in	10.in	14.in	18.in
3.in	7.in	11.in	15.in	19.in
4.in	8.in	12.in	16.in	

- Want folder full of output files:

1.out	5.out	9.out	13.out	17.out
2.out	6.out	10.out	14.out	18.out
3.out	7.out	11.out	15.out	19.out
4.out	8.out	12.out	16.out	

- Command Syntax:
  - ./myprogram inputfile > outputfile

MICHIGAN STATE UNIVERSITY

ICER

PBS Job Arrays

- One submission script copied many times
- Uses the PBS –t option
  - Ranges: 1-10
  - Lists: 2,4,100,3
  - Combination: 1-10,20,50,100
- Distinguish between jobs by using the PBS_ARRAYID environment variable

MICHIGAN STATE UNIVERSITY

ICER

Simple Job Array

```
#!/bin/bash -login
#PBS -l walltime=00:05:00,mem=2gb
#PBS -l nodes=1:ppn=1,feature=gbe
#PBS -t 1-100

cd ${PBS_O_WORKDIR}

./myprogram ${PBS_ARRAYID}.in > ${PBS_ARRAYID}.out

qstat -f ${PBS_JOBID}
```

MICHIGAN STATE UNIVERSITY

ICER

Example: Job Arrays

- Get the bleder_farm example:

```
>module load powertools
>getexample blender_farm
>cd ./blender_farm
```

- Look at the qusb file, using “less” command

```
>less blender_farm.qsub
```

- Submit the job

```
>qsub blender_farm.qsub
```

MICHIGAN STATE UNIVERSITY

ICER

HPCC Job array limitations

- Can not have more than 144 cores running
- Can not submit more than 256 jobs at once

- Lots of ways to work around this problem.

MICHIGAN STATE UNIVERSITY

ICER

Job array numbers

- All numbers in a job array have the same base number
  - 7478210
- Each PBS_ARRAYID is show in square brackets
  - 7478210[1]
  - 7478210[2]
- Delete all jobs using one command
  - qdel 7478210[]

MICHIGAN STATE UNIVERSITY

ICER

Large Shared Memory Systems (Fat Nodes)

- 256gb,512gb,1tb and 2tb nodes
- 32 – 64 cores

MICHIGAN STATE UNIVERSITY

ICER

MSU HTCondor Cluster

- Runs like a screen saver and Scavenges CPU cycles:
  - Approximately 400 nodes
  - Approximately 800 cores
  - WindowsXP

MICHIGAN STATE UNIVERSITY

ICER

Getting Help

- 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
- Documentation and User Manual – wiki.hpcc.msu.edu
- HPCC Phone – (517) 353-9309
- HPCC Office – 1400 PBS
- Office Hours – Monday – Friday 9am-5pm

MICHIGAN STATE UNIVERSITY

ICER