Arctic Region Supercomputing Center

Introduction to ARSC

David Newman (from Tom Logan slides), September 3 2015

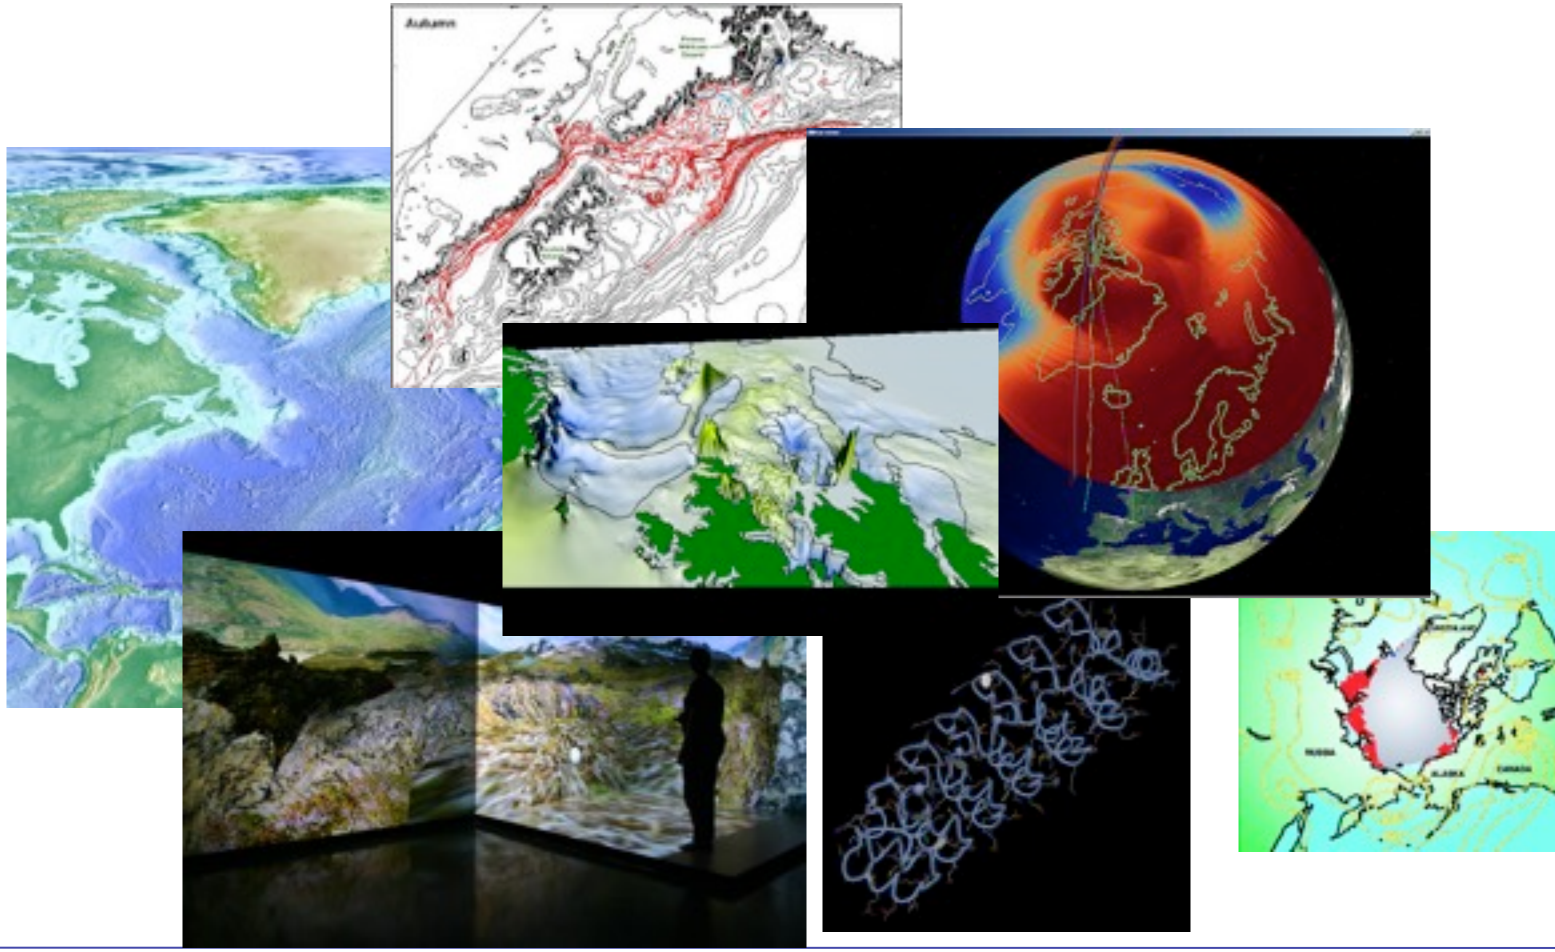

Arctic Region Supercomputing Center

What we do:

- High performance computing, university owned and operated center
- Provide HPC resources and support
- Conduct research locally, globally

Who we are:

Committed to helping scientists seek understanding of our past, present and future by applying computational technology to advance discovery, analysis and prediction

HPC UAF University of Alaska Fairbanks

University Research Center

- ARSC computational, visualization and network resources enable broad opportunities for students and faculty
- HPC training, outreach activities, internships and science application workshops
- HPC facility providing open accessibility for research and education

- Developing, evaluating and using acceleration technologies to drive leading-edge HPC computing
- Employs undergraduates, graduate students , post-doctoral fellows, faculty and staff

New and Current HPC Users in Training

- Introduce new and current users to HPC and visualization resources
- Provides programming skills for successful use of ARSC resources
- In-depth instruction, hands-on assistance developing codes

- Collaborative discussions with other users and parallel computing experts
- Early adoption and assessment of software tools

PHYS 693 - Core Skills in Computational Science

- **Computational Science**
  - High Performance Computing is required
  - Complex computing environment
  - Special skills are necessary
- **Arctic Region Supercomputing Center**
  - Extensive hardware resources
  - Extensive human resources
  - If our users are HPC trained, we rejoice
  - If our users are not, we despair

Hardware Overview

PACMAN –
Penguin Computing Opteron Cluster

- **12 Login Nodes:**
  - 2- Six core 2.2 GHz AMD Opteron Processors; 64 GB memory per node
- **256 Four Core Compute Nodes**
  - 2- Dual core 2.6 GHz AMD Opteron Processors; 16 GB memory per node
- **88 Sixteen Core Compute Nodes**
  - 2 Eight core 2.3 GHz AMD Opteron Processors; 64 GB memory per node
  - QLogic QDR Infiniband; 250 GB local disk
- **20 Twelve Core Compute Nodes**
  - 2 Six core 2.2 GHz AMD Opteron Processors; 32 GB memory per node
- **4 Large Memory Nodes (1 login)**
  - 4 Eight core 2.3 GHz AMD Opteron Processors; 256 GB memory per node
  - QLogic QDR Infiniband; 1 TB local disk; 80 GB solid state drive
- **2 GPU Nodes**
  - 2 Quad core 2.4 GHz Intel CPUs; 64 GB of memory per node
  - QLogic QDR Infiniband; 2 M2050 nVidia Fermi GPU cards

Cray XK6m Cluster

- **2 Login Nodes: "fish1.arsc.edu" and "fish2.arsc.edu"**
  - One six core, 2.6 GHz AMD Istanbul processor per node
  - 16 GB of memory per node.
- **48 GPU Enabled Sixteen Core Compute Nodes**
  - One sixteen core, 2.1 GHz AMD Interlagos Processor per node
  - 64 GB of memory per node (4 GB per core)
  - One nVIDIA Tesla X2090 GPU accelerator with 6GB RDDR5 memory
- **32 Twelve Core Compute Nodes**
  - Two six core, 2.6 GHz AMD Istanbul processors per node
  - 32 GB of memory per node (2.5 GB per core)
- **Cray Proprietary Gemini Interconnect**
- **20 TB Lustre $HOME file system**

** NEWEST SYSTEM **

Fish

Bigdipper

- **Sun SPARC Enterprise T5440 Server that acts as the long term storage server for ARSC academic computational resources.**
  - 4- 8-core T2+ CoolThread processors (32 cores total)
  - 128 GB memory
  - Sun Storage 6780 Array, with 125 TB of general disk cache
  - 2- 10 gigabit ethernet cards
- **Runs SAM-QFS**, which provides high-performance disk storage transparently backed by offline tape storage.
- **IBM TS3500 tape library** with:
- 1,800 3592JC tape cartridges with a raw capacity of 7.2PB utilizing six IBM TS1140 tape drives
- 400 LTO-4 tape cartridges with a raw capacity of 320TB utilizing six IBM TS1050 tape drives
- Total tape slot count of over 2,600 tapes in a mix of LTO and 3592 media

Science at ARSC

Science and Research

- Tsunami
- Oceanography
- Climate
- Ionosphere
- Data

Tsunamis

To ensure reliable tsunami early detection and hazard assessment capabilities, it is essential to create numerical models to forecast future tsunami impact and flooding limits to the coastal regions of Alaska.

Animation demonstrates the run-up in Resurrection Bay, Seward, Alaska as a result of the 1964 earthquake near Kodiak Island.

Animated underwater mud slide generating a tsunami wave.

Tsunamis

To ensure reliable tsunami early detection and hazard assessment capabilities, it is essential to create numerical models to forecast future tsunami impact and flooding limits to the coastal regions of Alaska.

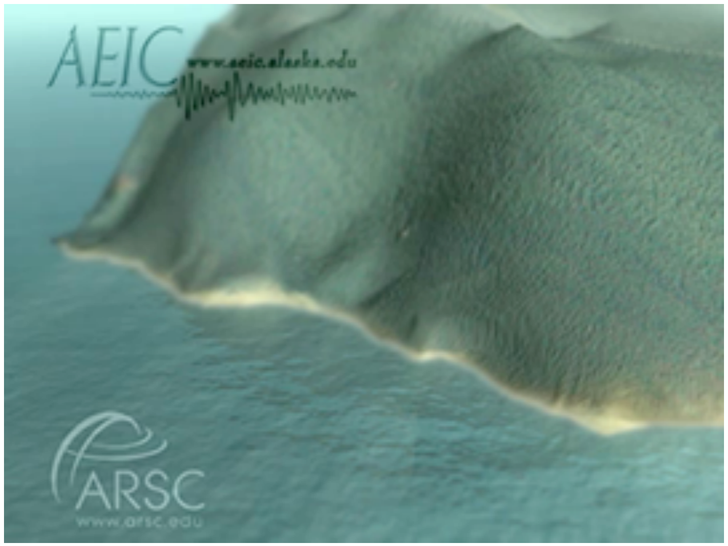

Animated underwater mud slide generating a tsunami wave.

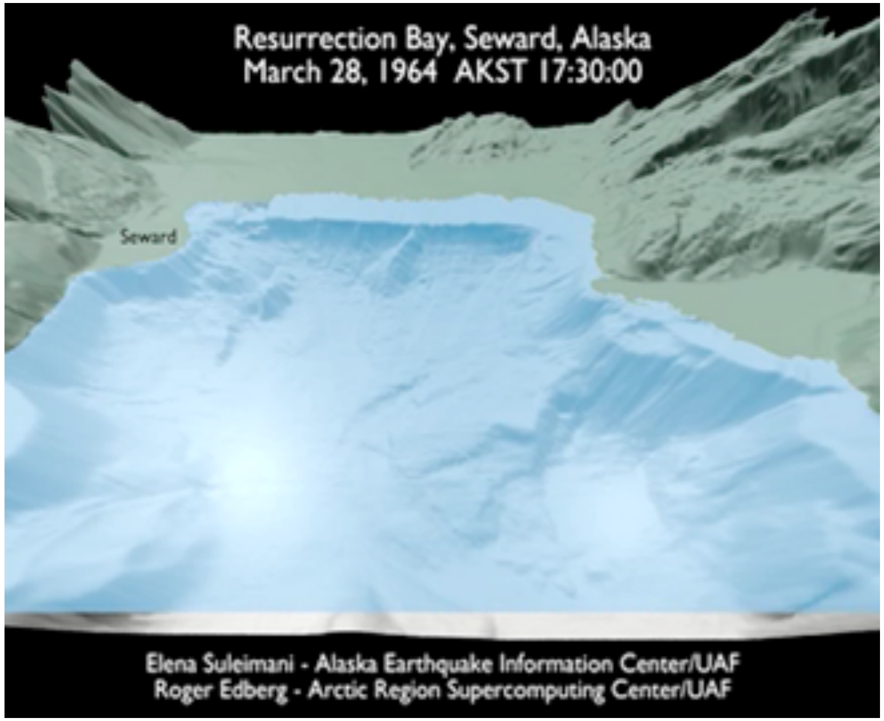

Animation demonstrates the run-up in Resurrection Bay, Seward, Alaska as a result of the 1964 earthquake near Kodiak Island.

Tsunamis

To ensure reliable tsunami early detection and hazard assessment capabilities, it is essential to create numerical models to forecast future tsunami impact and flooding limits to the coastal regions of Alaska.

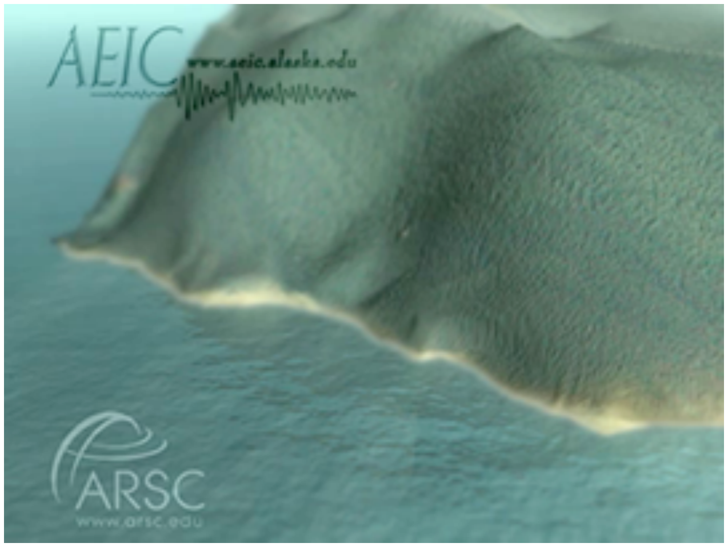

Animated underwater mud slide generating a tsunami wave.

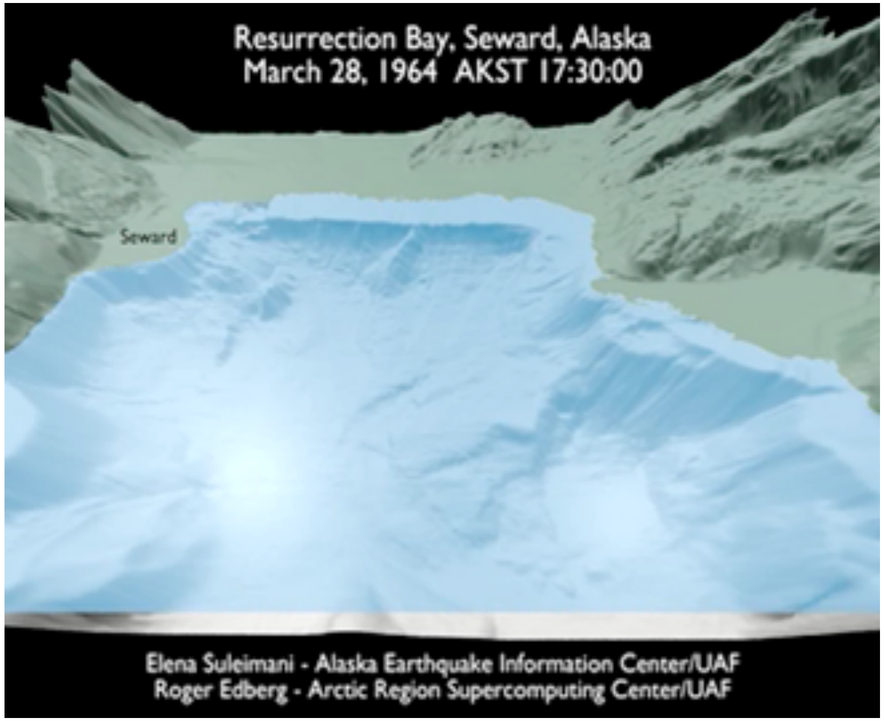

Animation demonstrates the run-up in Resurrection Bay, Seward, Alaska as a result of the 1964 earthquake near Kodiak Island.

Tsunamis

Tsunami Portal

Zygmunt Kowalik and Tom Logan, leads

A website portal where registered users can run models to make predictions on future tsunami impact, with access to to multiple codes for modeling generation, propagation and run-up of tsunamis on multiple computer platforms.

Comparison of two different models from the Tsunami Portal. The scenario is the 1964 earthquake in Alaska.

Modeled sea level pattern generated by the Indian Ocean tsunami, Dec. 26, 2004.

Tsunamis

Tsunami Portal

Zygmunt Kowalik and Tom Logan, leads

A website portal where registered users can run models to make predictions on future tsunami impact, with access to to multiple codes for modeling generation, propagation and run-up of tsunamis on multiple computer platforms.

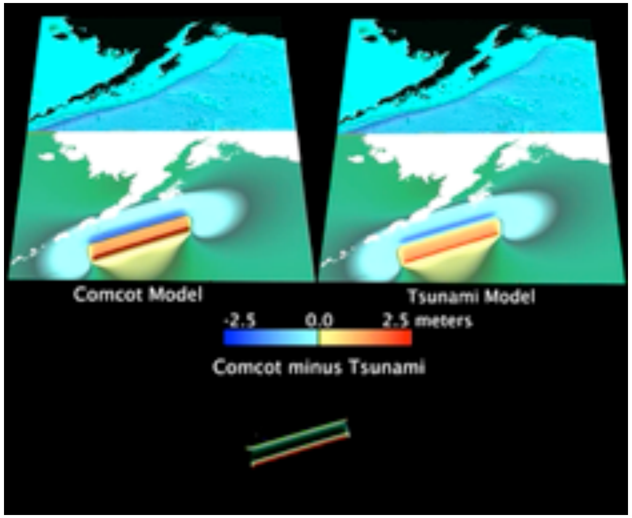

Comparison of two different models from the Tsunami Portal. The scenario is the 1964 earthquake in Alaska.

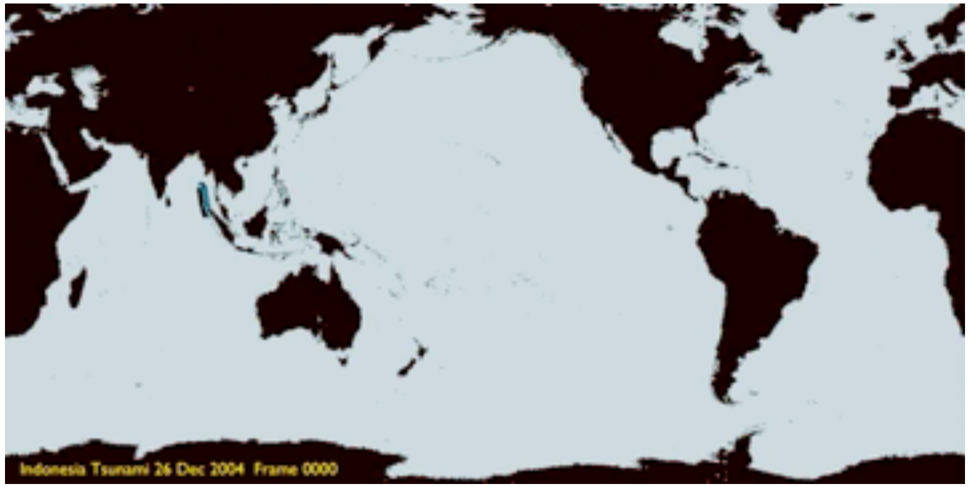

Modeled sea level pattern generated by the Indian Ocean tsunami, Dec. 26, 2004.

Tsunamis

Tsunami Portal

Zygmunt Kowalik and Tom Logan, leads

A website portal where registered users can run models to make predictions on future tsunami impact, with access to to multiple codes for modeling generation, propagation and run-up of tsunamis on multiple computer platforms.

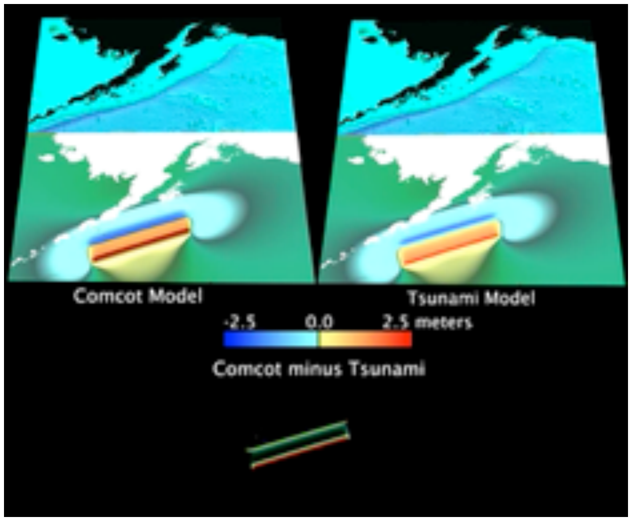

Comparison of two different models from the Tsunami Portal. The scenario is the 1964 earthquake in Alaska.

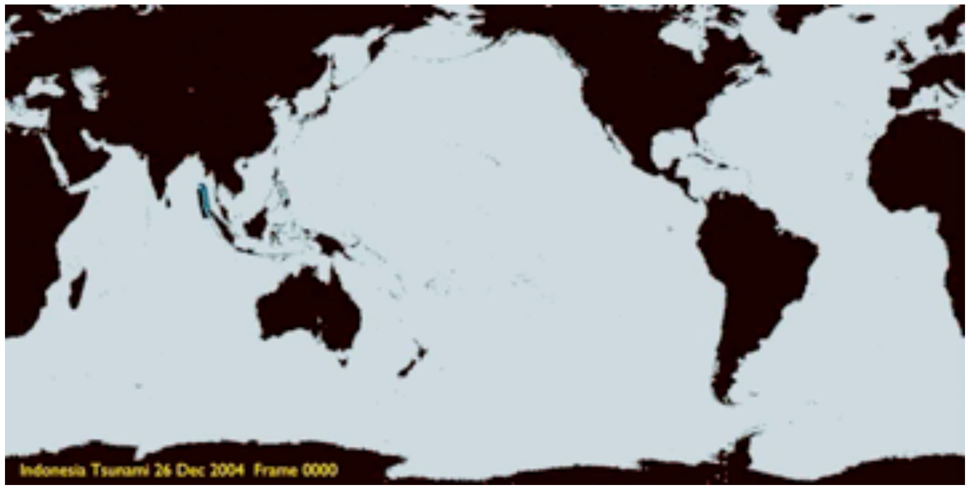

Modeled sea level pattern generated by the Indian Ocean tsunami, Dec. 26, 2004.

Oceanography

Kate Hedstrom and Mark Johnson, leads

Hedstrom (ARSC) is using community models to develop understanding of ocean circulation and ecosystem models with higher spatial resolution and accounting for more natural boundary conditions.

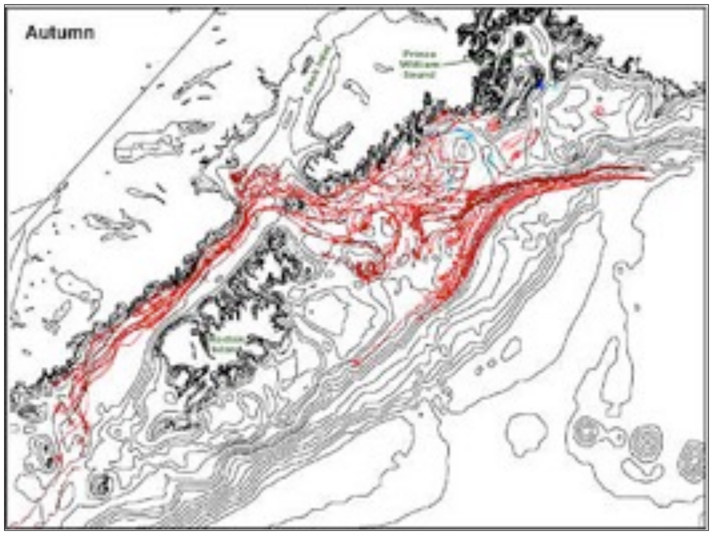

Numerical ocean model shows the isosurface velocity vectors of the Alaska current.

Oceanography

Kate Hedstrom and Mark Johnson, leads

Hedstrom (ARSC) is using community models to develop understanding of ocean circulation and ecosystem models with higher spatial resolution and accounting for more natural boundary conditions.

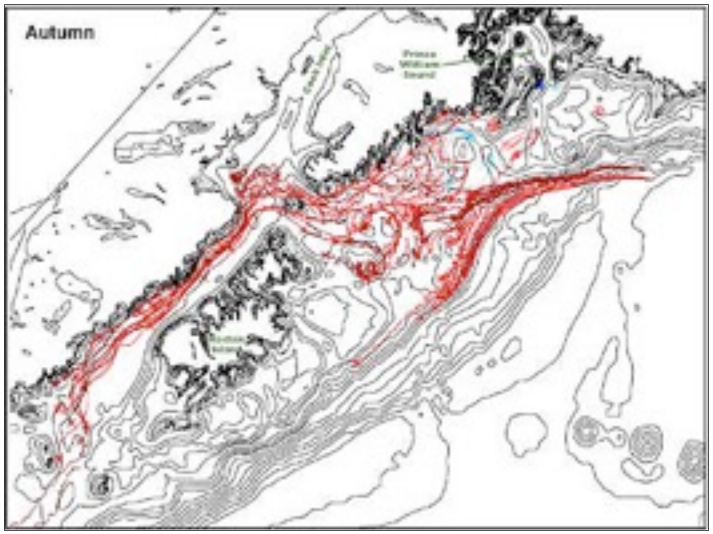

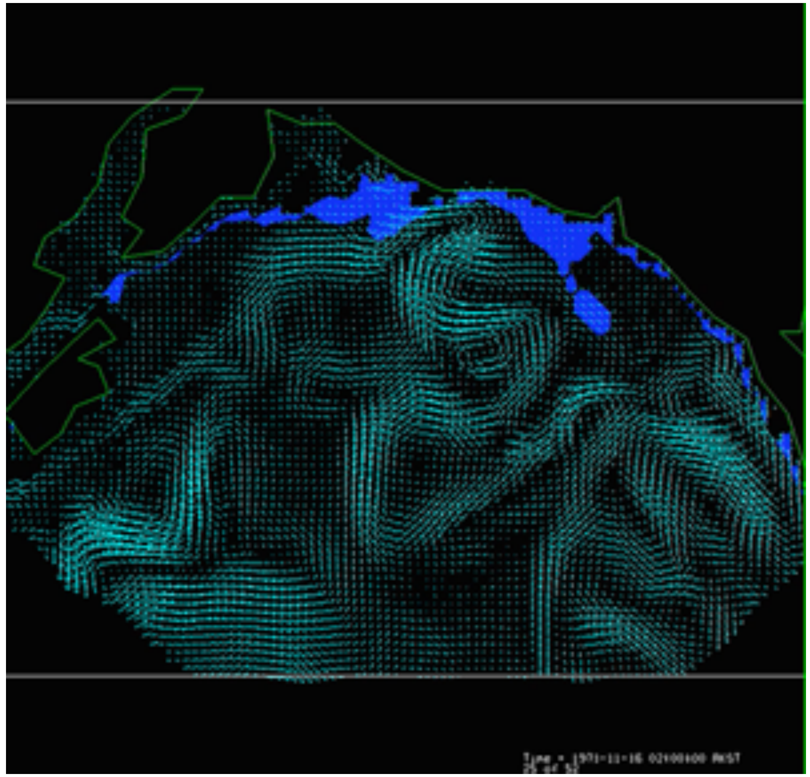

Numerical ocean model shows the isosurface velocity vectors of the Alaska current.

Climate

Don Morton, lead

The Weather Research and Forecast (WRF) model produces high-resolution forecasts of Alaska regions. Twice-daily WRF forecasts are made available to the National Weather Service. To see the WRF web interface, go to weather.arsc.edu and click on the Current Forecast link.

Animation depicts a forecast of wind and cloud fields during a high-wind event of Oct. 2006.

Climate

Don Morton, lead

The Weather Research and Forecast (WRF) model produces high-resolution forecasts of Alaska regions. Twice-daily WRF forecasts are made available to the National Weather Service. To see the WRF web interface, go to weather.arsc.edu and click on the Current Forecast link.

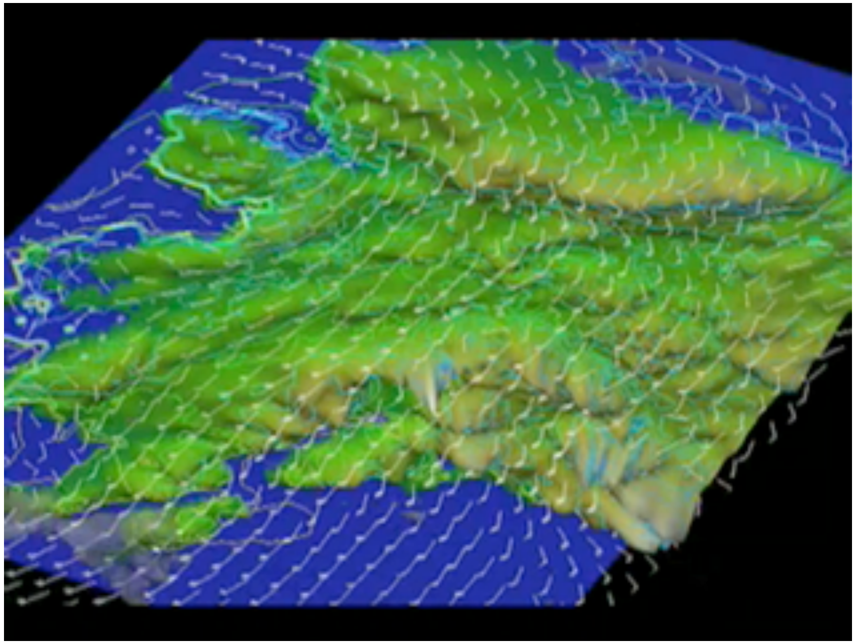

Animation depicts a forecast of wind and cloud fields during a high-wind event of Oct. 2006.

Ionosphere

Brenton Watkins and Sergei Maurits, leads

Maurits (ARSC) and Watkins (GI, Physics) are advancing their world class Eulerian Parallel Polar Ionosphere Model (UAF EPPIM) to significantly improve the model's capability to simulate the ionospheric environment and to predict real-time scintillation, bending, and abnormal propagation of electromagnetic signals in polar regions.

Ionosphere

Brenton Watkins and Sergei Maurits, leads

Maurits (ARSC) and Watkins (GI, Physics) are advancing their world class Eulerian Parallel Polar Ionosphere Model (UAF EPPIM) to significantly improve the model's capability to simulate the ionospheric environment and to predict real-time scintillation, bending, and abnormal propagation of electromagnetic signals in polar regions.

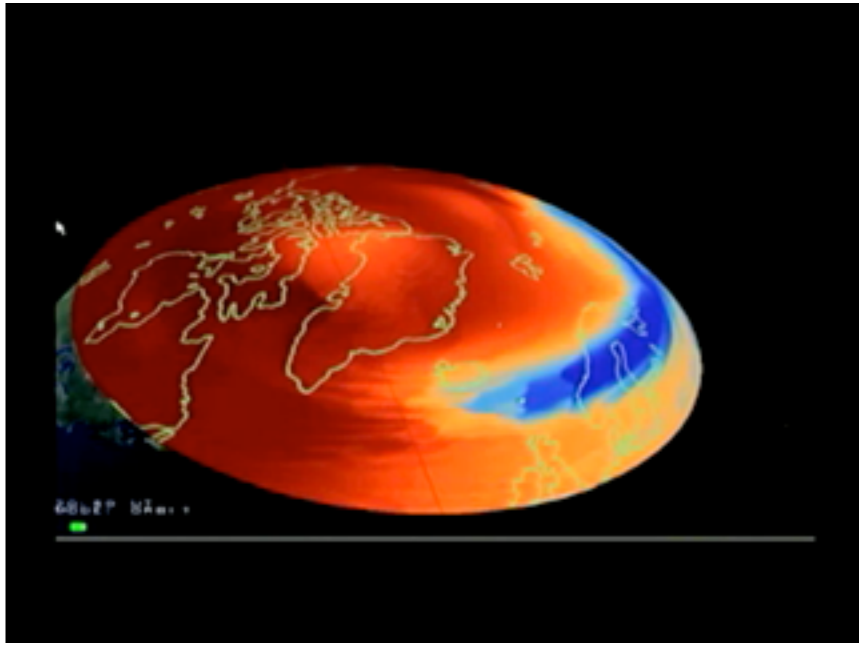

Data

- **Research in data-intensive computing includes:**
  - distributed information retrieval and Web search
  - grid computing with the Globus Toolkit and Web services
  - partnering with faculty members to add metadata and Web-database capability for biological and ecological data
  - remote sensing storage, processing and distribution

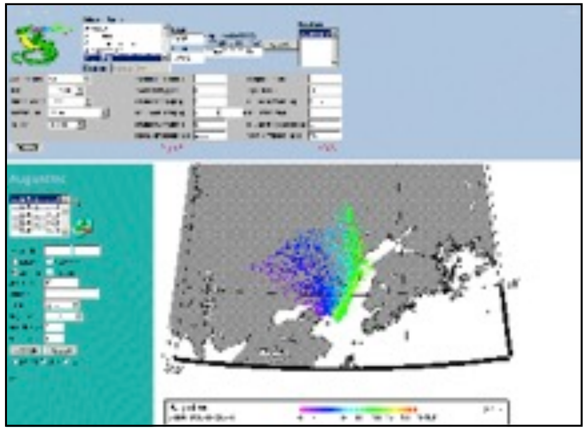

- **New cross-cutting effort with virtual globes (VGs):**
  - general purpose GoogleEarth™ and GINA's SwathViewer VG interaction with geospatially located data
  - VG applications include weather, volcano monitoring, glaciology

Research & Development Culture

- Post Doctoral Fellows
- Graduate Students
- Undergraduate Students
- Faculty
- Enhance/advance research, academic computing and military connections

Staff Resources

- Consultants
- HPC specialists
- Visualization/Data Analysts
- Systems administrators & programmers
- Storage specialists
- Network specialists
- About 20?? staff, plus 10 or so student employees

Supercomputer Architectures

- **They're all Parallel Computers**
- **Three Classes:**
  - Shared Memory
  - Distributed Memory
  - Distributed & Shared Memory
- **They all run Unix/Linux**

Shared Memory Architecture

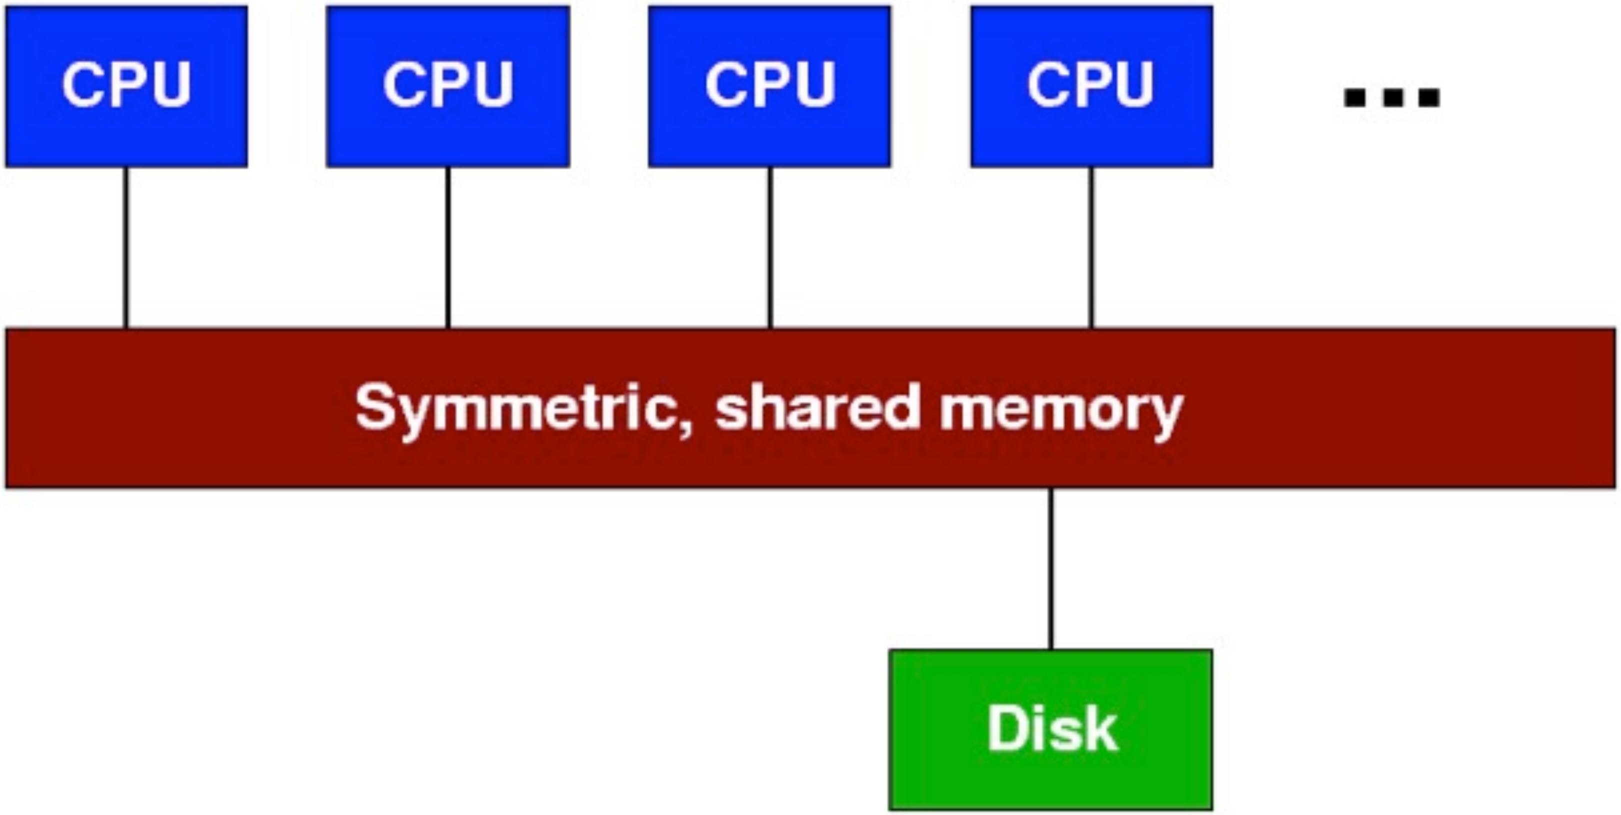

Distributed Memory Architecture

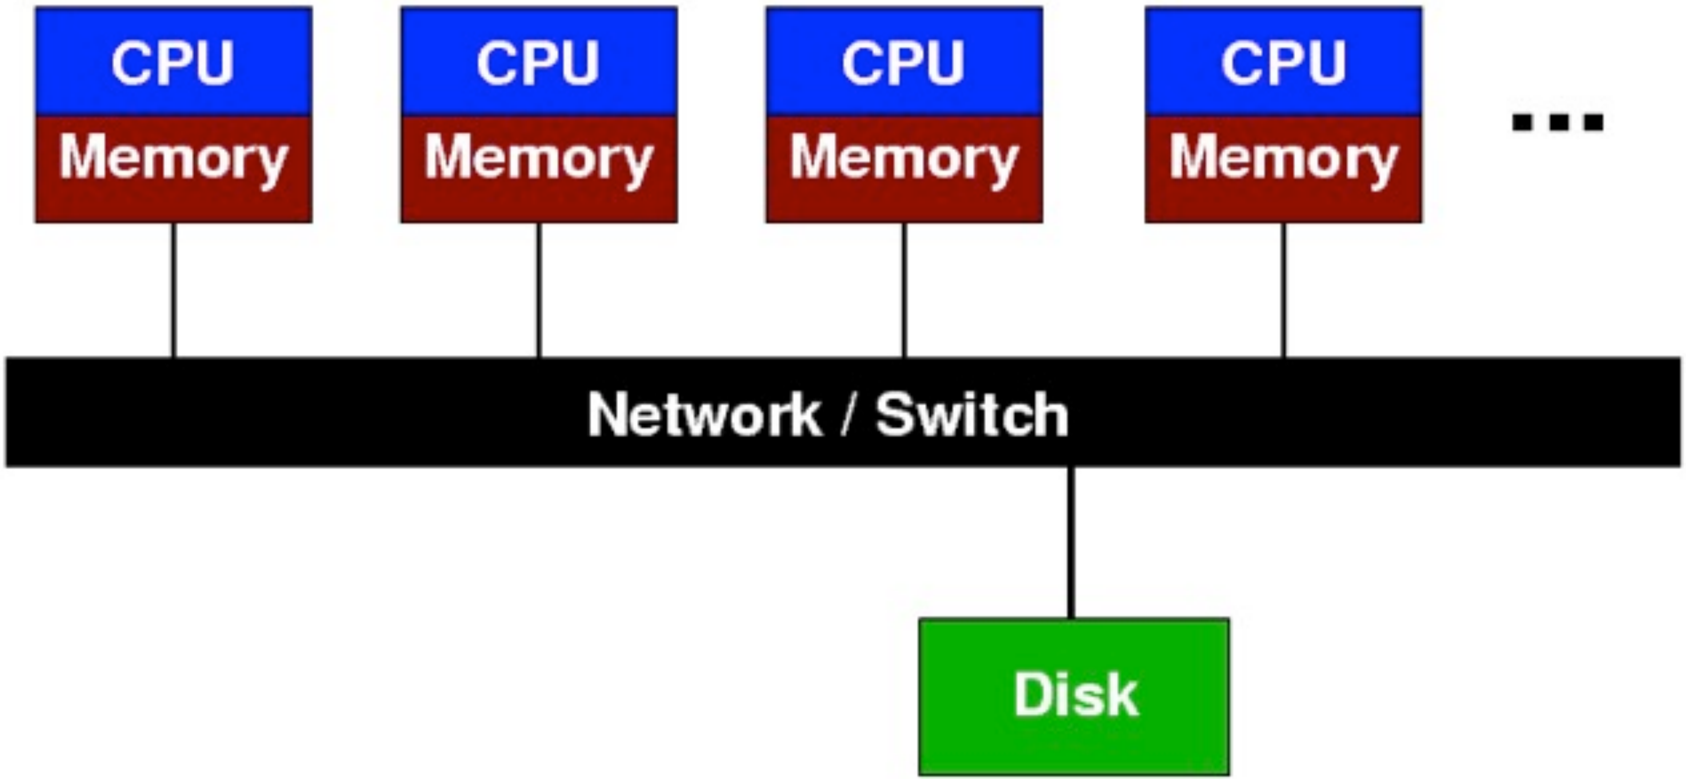

Cluster Architecture

- **Scalable, distributed, shared-memory parallel processor**

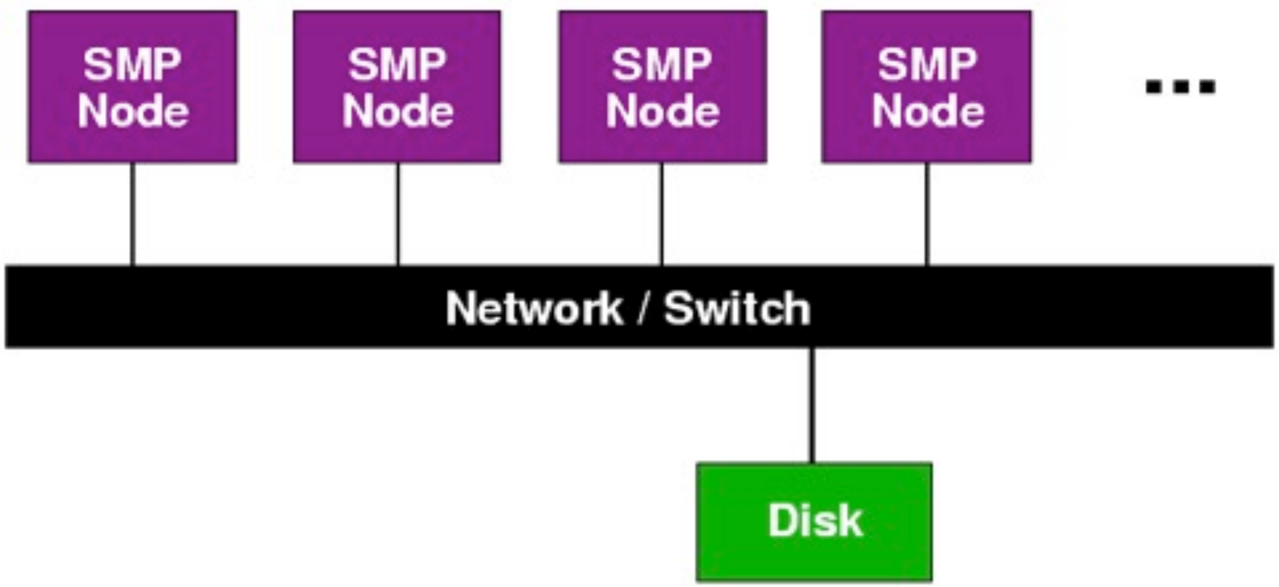

Threaded Programming on Shared-Memory Systems

- **OpenMP**
  - Directives/pragmas added to serial programs
  - A portable standard
- **Other Threaded Paradigms**
  - Java Threads
  - Pthreads
- **Auto-parallel option in the compiler**

Threads Dynamically

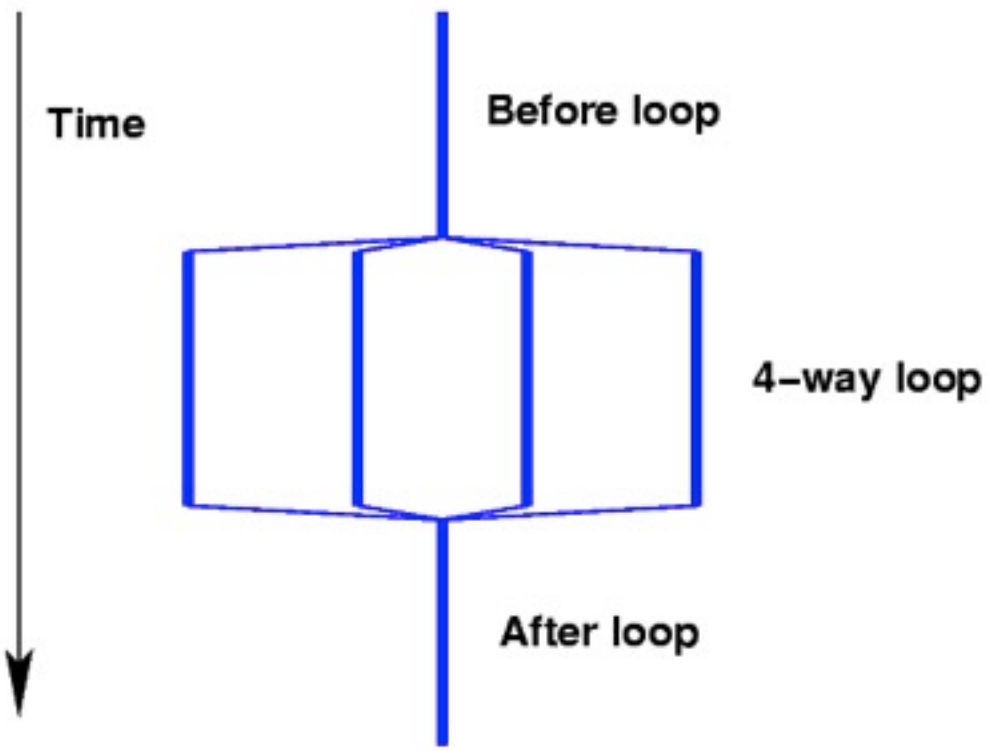

Distributed Processing

Concept:

1) Divide the problem explicitly
2) CPUs Perform tasks concurrently
3) Recombine results
4) All processors may or may not be doing the same thing

Branimir Gjetvaj

Number of Processes

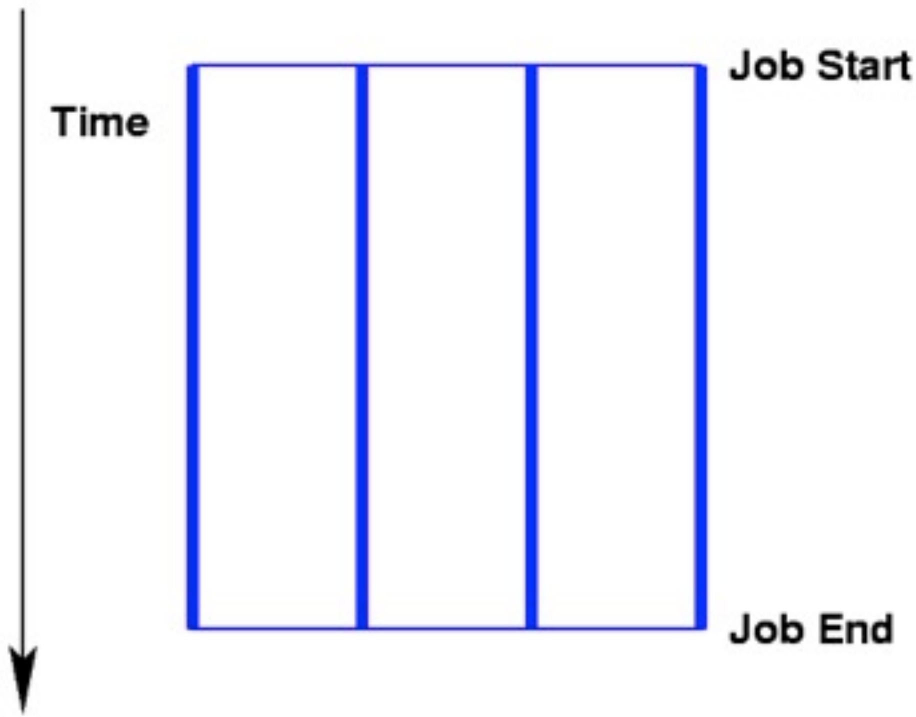

Cluster Programming

- **Shared-memory between processors on one node:**
  - OpenMP, threads, or MPI
- **Distributed-memory methods between processors on multiple nodes**
  - MPI
- **Mixed mode**
  - MPI distributes to nodes, OpenMP within node

Agenda for your Class (syllabus)

West Ridge Research Building

- Moving to WRRB 009 Classroom

- Restrooms across the hall
- Emergency Exit
- Main ARSC area upstairs
- Access Cards

Questions?