High Performance Computing at ESRL

Supporting the Computational Needs of ESRL's Scientists and Modelers

Craig Tierney[2], Leslie Hart[1], Pam Weber[1], Lee Cohen[2], Chris Harrop[2], Ed Moxley[3], Forrest Hobbs[4], Nathan Dauchy[4], Eric Schnepp[4], Glen Pankow[5]

[1]NOAA – Earth System Research Laboratory, [2]Cooperative Institute for Research in Environmental Sciences, [3]Riverside Technology , Inc., [4]Raytheon, [5]Cooperative Institute for Research in the Atmosphere

Overview

We provide high performance computing systems and support necessary for the development of tomorrow's weather and hurricane models across ESRL and NOAA. Our research into emerging technologies and experienced computer scientists and architects allow us to provide cost effective HPC resources to our community.

History of HPC at ESRL

- 1992 Intel Paragon – The first massively parallel computer in NOAA
- 2000 Jet – First in a series of clustered high-performance computing systems
  - First HPC cluster procured by government through competitive acquisition
  - First cluster system with integrated storage area network and archive system
- 2002 – **Eighth fastest system in the world**
  - **First in power efficiency of the Top Ten fastest systems**
- 2008 – Acquired first GPU cluster
- 2009 – Only major HPC center to adopt multiple major HPC file systems at scale

Current HPC Resources

- 77 Teraflops total performance
- 19.1 Terabytes of RAM
- 900 Terabytes of high-speed disk
- 2 PB tape archive

User Base

Our diverse user base includes:

NOAA: GSD, PSD, CSD, GMD, AOML, NCEP, PMEL, GFDL, SWPC, NSSL, NOS;, NGDC, GLERL, NWS
University: FSU, NC State, U. Illinois, and, NCAR

Workflow Management

ESRL has developed a fault-tolerant workflow management engine to facilitate the reliable and efficient execution of numerical models for real-time and retrospective experiments.

- End-to-end automation of complex workflows
- Automated enforcement of interdependencies between workflow components, data, and temporal events

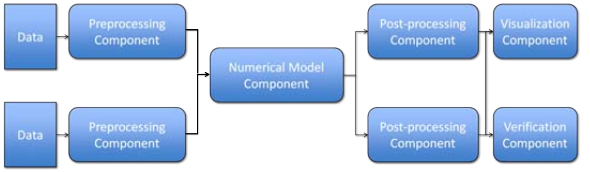

GPU Computing

Continuing our tradition of experimenting with and adoption of new technologies, we have started to research the viability of Graphical Processing Units for weather modeling.

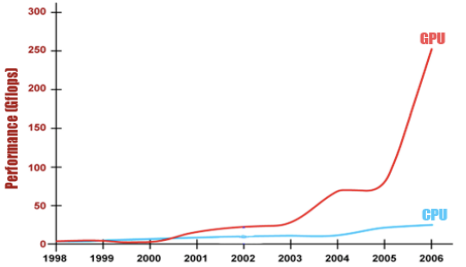

Current GPU Cluster

- 16 nodes – 32 Teraflops
- Early performance results show **34X speedup** for NIM over existing CPU technology

Partnerships

We work with different groups to augment our existing resources to meet their computational needs. By leveraging existing resources and staff additional resources can be provided for minimal additional cost.

HFIP – The goal of the Hurricane Forecast Improvement Project (HFIP) is to improve one-to-five day hurricane forecasts in terms of intensity, track, and storm surge, as well as reduce uncertainty.

HRRR – The High-Resolution Rapid Refresh (HRRR) weather model is being developed by ESRL scientists to support prediction of severe storms for the FAA.

Future of ESRL High Performance Computing

2010

- Identify power saving measures through equipment updates
  - Reduce power consumption up to 75 percent
  - No additional up-front cost to NOAA
- Double the size of the HFIP System
  - Total Jet System performance over 140 Teraflops

The Future

- In 2012, migrate traditional high performance computing workload to the ARRA-funded NOAA consolidated site
- Continue to be the site for testing and adopting emerging computational technologies and paradigms
  - GPUs, low-power processors, new-emerging technologies
- Continue supporting dedicated computational modeling projects such as HFIP and FAA

ESRL researches and leverages emerging technologies to maximize return on investment to best support the development of tomorrow's weather and hurricane models.