



Goddard Space Flight Center



# The Myths and Realities of Cluster Computing

Science Data Processing Session

**Charles D. Norton**

With acknowledgment to members of the HPC Systems Group

Jet Propulsion Laboratory  
California Institute of Technology  
High Performance Computing Systems Group

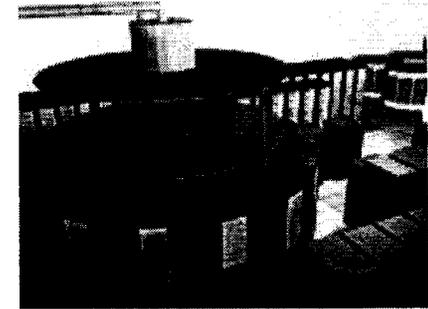
Work sponsored by the Earth Science Technology  
Office/Computational Technologies Program

**4th Quality Mission Software Workshop**  
**Dana Point, CA, May 7-9 2002**

# A Hierarchy of High Performance Systems

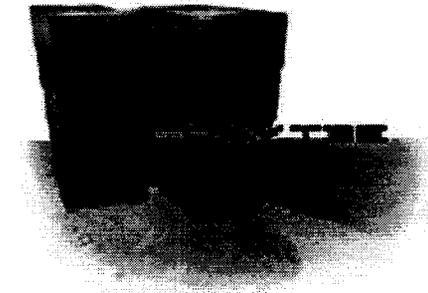
## **Extreme High-End (Petaflops HTMT Architecture)**

- Essential toward advancing the state-of-the-art
- Special purpose, not yet available, highest performance at highest cost



## **Traditional System (Cray, IBM, SGI)**

- Useful, but not highly accessible or configurable
- Benchmark for modern day performance



## **Commodity Cluster (Beowulf-Class)**

- Driven by advances in hardware & software
- Highly available and re-configurable
- Outstanding price/performance



# Impact of Advances in Cluster Components

## Processors ('97)

- 200 Mhz Pentium Pro with 128 MB RAM per node

## Networking and Disk ('97)

- 100 Mb/s Ethernet & 2.6 GB disks per node

## Software ('97)

- Linux, MPI, monitoring tools, compilers, debuggers

## Costs ('97)

- ~2K per node for top-of-line system

## Processors (Now)

- >1 Ghz dual-CPU Processors with >2 GB RAM per node

## Networking and Disk (Now)

- > 2 Gb/s switches & >10 GB disks per node

## Software (Now)

- Incremental enhancements

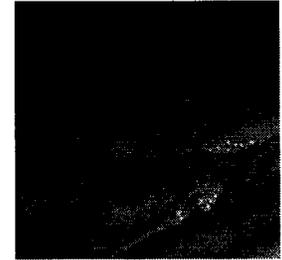
## Costs (Now)

- ~2K per node for top-of-line system and higher based on networking, support, etc...

# Application Development and Migration

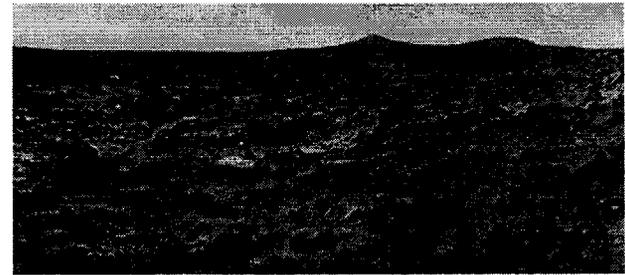
## Task Based (Data Processing)

- Schedule work on processors as needed



## Simple Parallelization (Optimization)

- Run same code on processors, but vary parameters over iterations



## Complete Parallelization (Physics-Based Modeling)

- Distribute both memory and the calculation



Large data sets, real-time results, compute and memory bound applications, risk reduction through design optimization...

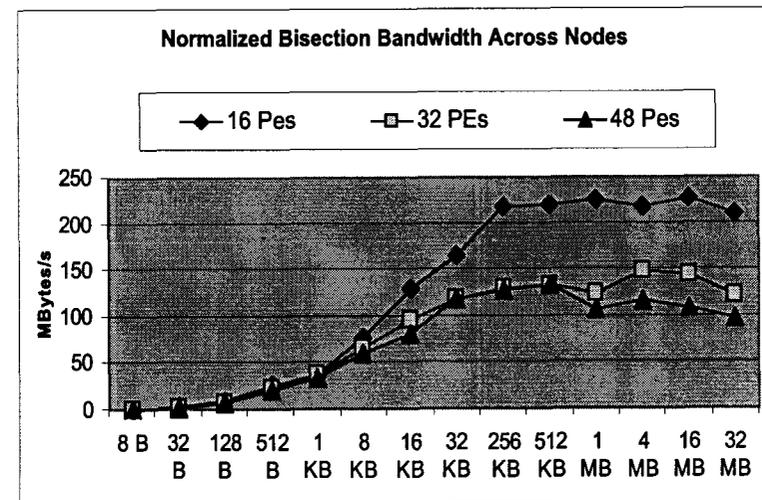
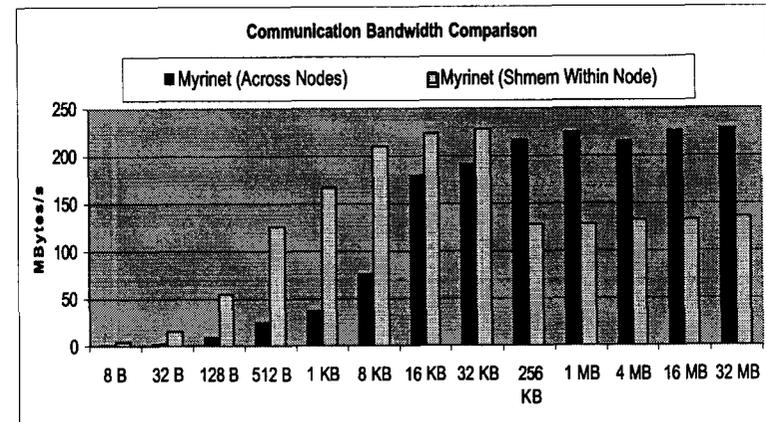
# Cluster Computing Problems and Pitfalls

## Purchasing & Costing Systems

- Support and maintenance costs are significant for large high-end systems and can't be ignored
- State-of-the-art always changes
  - Linux OS, Parallel I/O, Compilers, Cooling

## Knowledge and Experience

- New arena for vendors and users with varying levels of expertise
  - System software, networks, file systems, monitoring and management tools, etc...
- Performance programming can be challenging

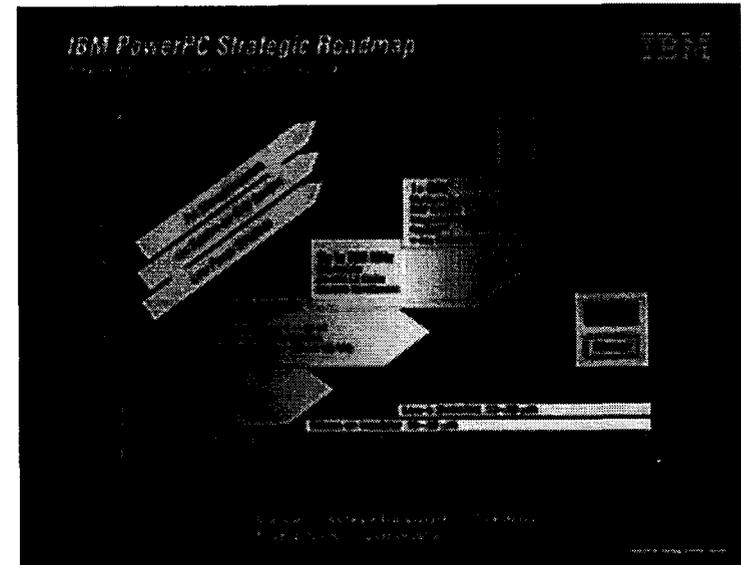


Network/Processor performance implications

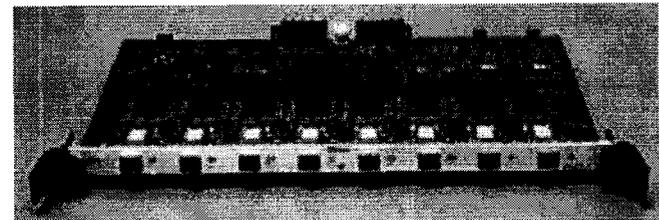
# Cluster System Considerations for SDP

## Hardware, Software, and Maintenance

- Leverage off of industry accepted, and well supported, standards
- Stay abreast of technology advances, but stay away from the leading edge
- System management tools are important and should be integrated with the hardware
- These systems do not run themselves so administration is an important expense



Commodity Processor Roadmap is Strong



Networking Remains the Bottleneck