

⋮

High-End Computing Systems Group



Department of Computer and Info. Science
The Ohio State University

<http://nowlab.cis.ohio-state.edu/systems/>

⋮

⋮

High-End Computing and Its Benefits

- Computation spreaded over hundreds and thousands of processors
 - Provides far-reaching benefits to society: e.g. new drugs, safer and fuel-efficient vehicles, environmental modeling, nuclear stockpile safety, scientific discoveries in a broad range of disciplines
 - Essential to science and engineering research
 - Enabling element of the United States national security program
 - New applications of high-end computing are ripe for exploration: e.g. intelligent systems, design, transportation, crisis management
- ⋮

•
•
•

Our Vision

- Cover the entire range, from the Applications level, through the Systems Software level, to the Networking and Communications level, in developing high-end systems and enabling their use in application areas:
 - Data-Intensive Computing
 - Datamining and Database Systems
 - High-Performance Scientific Computing
 - Middleware and Compilers
 - Network-Based Computing
- Carry out research in an integrated manner (systems, networking, and applications)

• • • • • • • •

•
•
•

Collaboration within Department - An Integrated Team

- Systems
 - D. K. Panda (Architecture, Commn., & Networking)
 - P. Sadayappan (Programming Env. & Resource Management)
 - Mario Lauria (Communication and I/O)
 - Srini Parthasarathy (Data-intensive computing & Data Mining)
 - Hakan Ferhatosmanoglu (Databases)
 - Gagan Agrawal (Grid Computing, Middleware, Compilers)
 - Joel Saltz (High Performance Computing Software, Bioinformatics)
- Applications
 - R. Crawfis (volume rendering)
 - R. Machiraju (streaming visualization)
 - H.-W. Shen (large-scale time-varying visualization)

• • • • • • • •

•
•
•

Collaboration

- Industry
 - IBM TJ Watson, Intel, HP
- National Laboratories and Consortia
 - Sandia, Los Alamos, Pacific Northwest Labs, Argonne, Oak Ridge, NPACI, Wright Patterson, and Army Research Lab
- Several departments at OSU
 - College of Engineering, Maths and Physical Science, Life Sciences, Biomedical Informatics, and Ohio Supercomputer Center (OSC)
 - Internet 2 Technology Evaluation Center (ITEC-Ohio)

•
•
•

Research Funding

- National Science Foundation
- Department of Energy
- Department of Defense
- National Library of Medicine
- Ohio Board of Regents
- IBM
- Sandia National Lab
- Pacific Northwest National Lab
- Los Alamos National Lab
- Ameritech

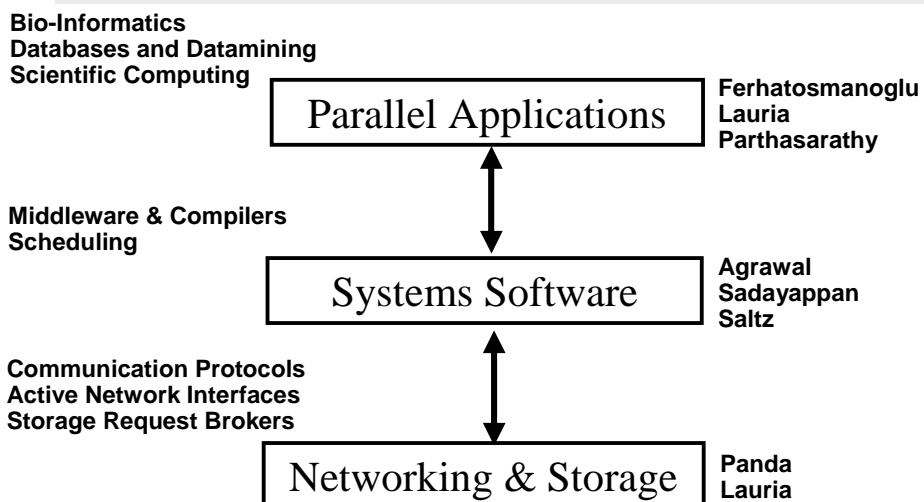
⋮

The Research Challenge

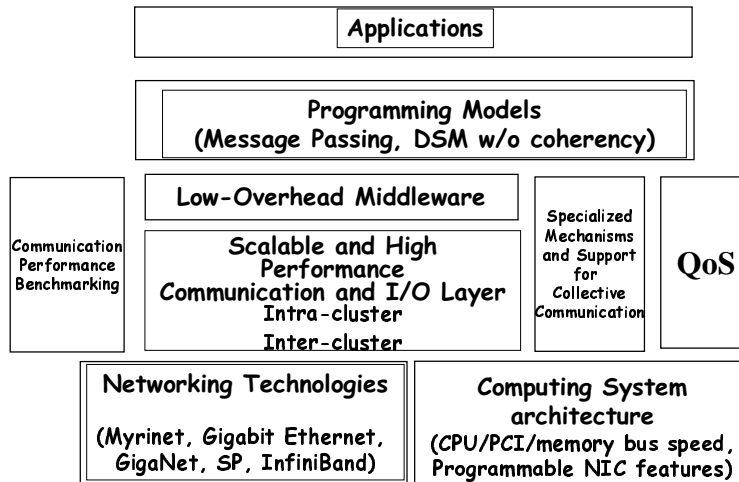
- The cost/performance ratio of computing, networking, and storage components has improved dramatically, making the aggregate available "raw" computing power very high
 - However, the fraction of the potential power that is actually utilized is getting smaller every year
 - Two broad challenges:
 - Make the large aggregate power of distributed computing resources available in a transparent manner to high-end applications
 - Create high-level approaches to ease the development of high-end applications
- ⋮

⋮

Research Areas Covered



Overview of Research Projects (Panda)




Projects

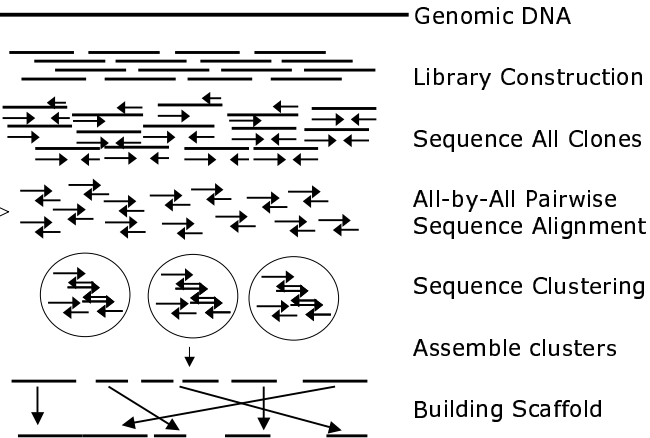
- Developing low-latency point-to-point communication in clusters with modern interconnects
 - Gigabit Ethernet, Myrinet, VIA, Quadrics, and InfiniBand
- Taking advantage of Active NIC for
 - Collective communication (broadcast, multicast, reduction)
 - QoS (bandwidth allocation, policing, monitoring)
- Efficient Implementation of Programming Model Layers with better communication support
 - Message Passing Interface (MPI), Distributed Shared Memory (DSM), Get/Put, Sockets
- Design of scalable data centers

Projects (Mario)


Whole Shotgun Assembly: a case study in data intensive computing



~12,000,000 Reads
(6.2 GB)



Cluster architecture and data intensive computing




Research topic #1 – high perf I/O inside the cluster

- 1 TB of disk @ 1K\$ possible using IDE disks
- 1 GB/s aggregate bandwidth on small cluster
- Research topics:
 - File System vs. Parallel Library
 - how to deliver this perf to the apps?

Research topic #2 – high perf data transfer between clusters over the GRID

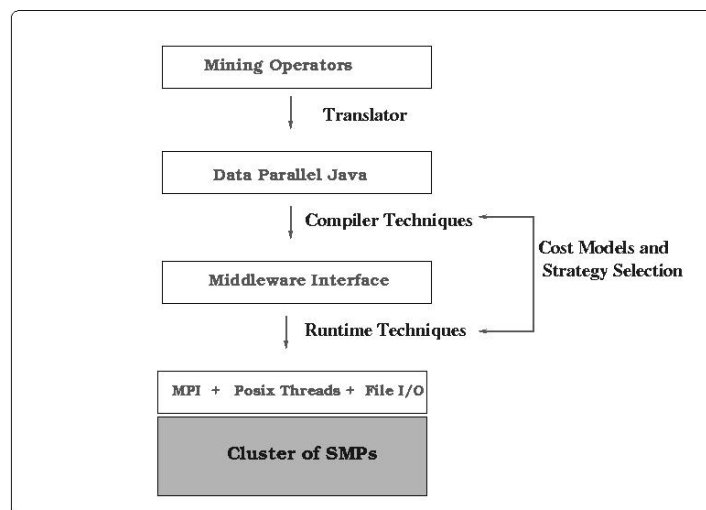
- 1 Gb/s WAN it's here – we are doing it on campus
- 1 Gb/s cheap& easy – just use Gb Ethernet
- Research topics:
 - how to move data fast in/out of the cluster?
 - 1 Gb/s on the wire => ? to the application



Projects (Agrawal)

- Overall focus: System Support for Commercial and Scientific Data Intensive Applications
- Includes:
 - Support at Middleware / Runtime and Compiler level
 - Support for Scientific Data Intensive Applications (analysis of data from simulations and sensors) and commercial data intensive applications (data mining, OLAP)
 - Target Parallel Environments (Clusters, Parallel Machines) and Distributed / Grid Environment

An Example Project



⋮

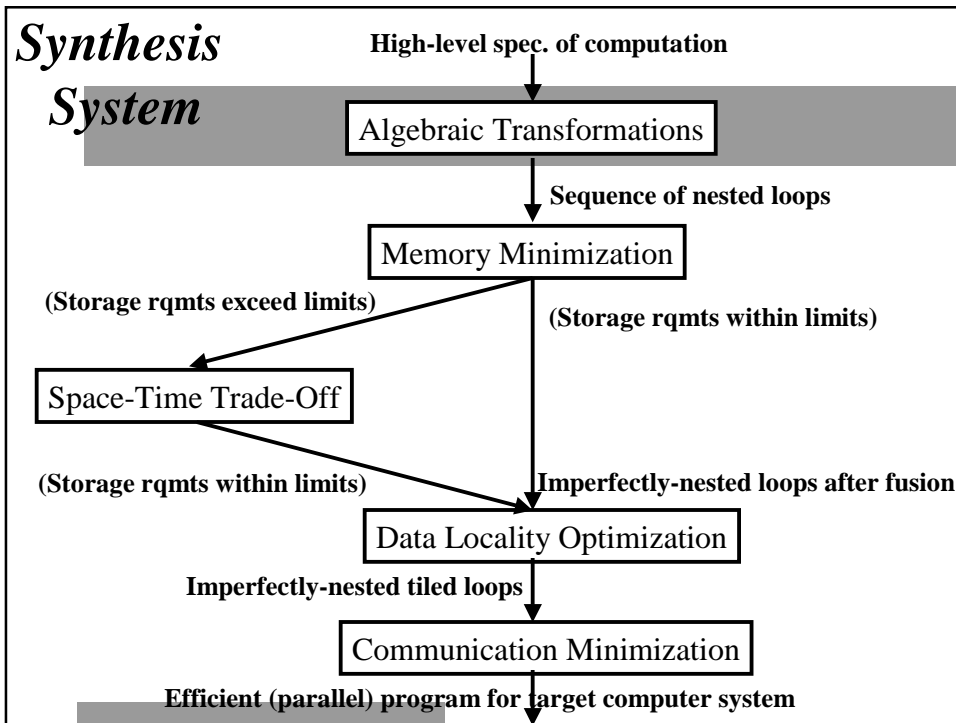
Other Projects

- **Compiling for a Grid Environment**
- **Compiling scientific data intensive applications**
- **Performance modeling and prediction**
- **Algorithms for sequential and parallel data cube construction**

⋮

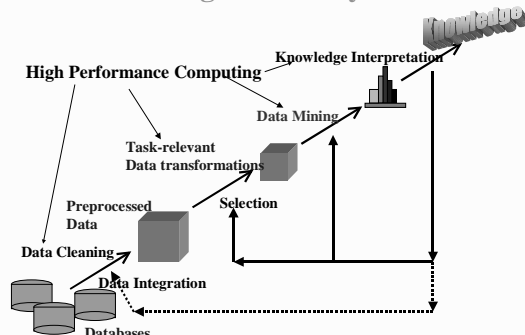
Research Interests (Sadayappan)

- **Network-Based Computing (with Panda)**
 - Resource Optimization and Scheduling
- **High-Perf. Scientific Computing (with Baumgartner)**
 - Automatic synthesis and optimization of parallel programs for a class of quantum chemistry computations



Projects (Srini) High Performance Data Mining

Knowledge Discovery Process



- KDD process
 - Interactive
 - Iterative
 - Compute and I/O Intensive
- Research
 - Novel DM algorithms
 - Parallel algorithms
 - Incremental algorithms
 - Systems support
 - Applications
 - Intrusion Detection
 - Bioinformatics
 - Web/Text Mining

- Sample Project: Active Mining on Streaming Data
 - Crucial Issue: Data Influx Rate Exceeds Processing Rate
 - Problem for time-critical applications (e.g. Network Intrusion Detection)

- Sample Project: Looking for Structure in Scientific Data:
 - Protein Structure Detection
 - Important for protein function analysis
 - Distance based structure representation
 - Frequency based pattern identification
 - Alpha Helix-related substructures (figure)

Alpha-helix-related substructures

Original protein

Projects (Hakan)

Database Support for Modern Applications

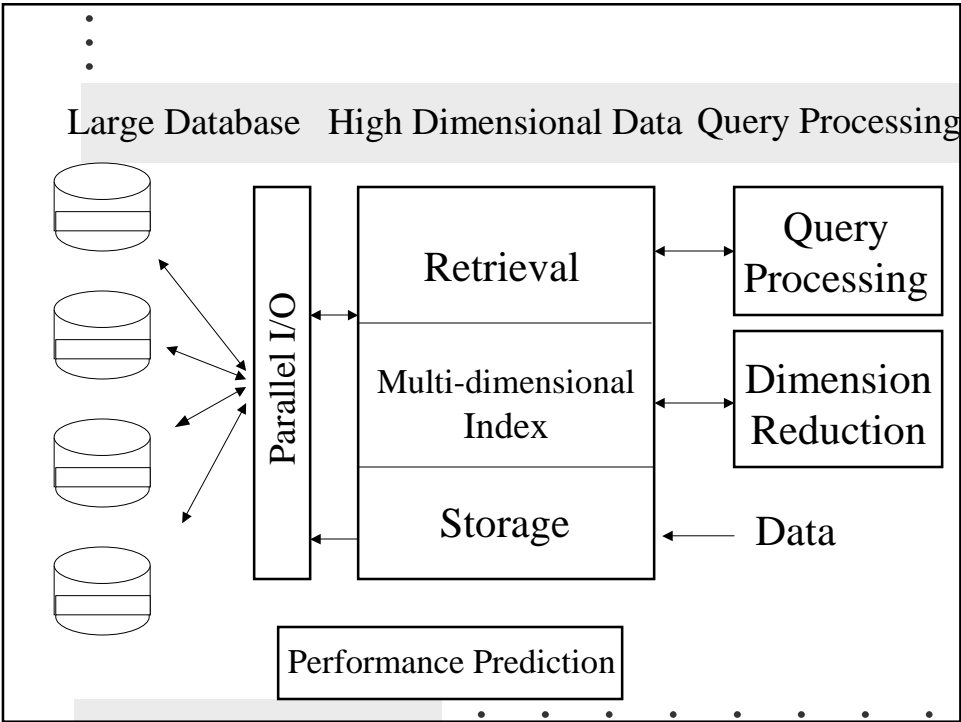
- Multimedia databases (image, audio, video, document)
 - *IBM QBIC, Muscle Fish, CNN* video server
- Spatial databases (GIS):
 - *ESRI* Spatial Data Engine, *Oracle* Universal Server with Spatial Data, *Informix* Spatial Datablades
- Time-series DBs (stock market): *IBM Quest*
- Biomedical and String DBs (genetics, networking data)
- **Problem:** *Massive* amount of *Multi-dimensional* data
- **Goal:** *Scalable* systems for *Efficient* queries

⋮

Our Approach

- "Dimensionality"
 - High dimensional storage and retrieval
 - Vector space and string data
- "Scalability" (w.r.t. size & performance)
 - Multi-disk/processor architectures
 - Parallel I/O
 - Internet, global computing
- "Query Processing"
 - Support a wide range of queries
 - Define new kinds of useful queries and techniques for them

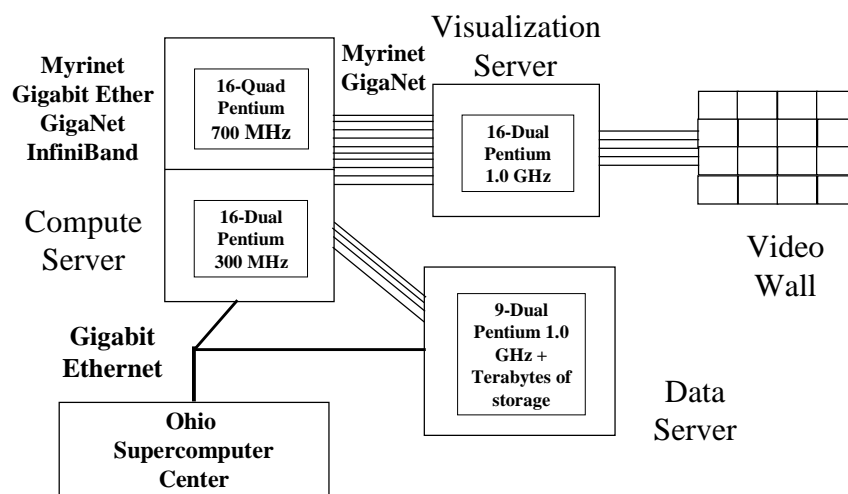
⋮



Examples of Large Scale Projects in the Group

- Active Data Repository and Datacutter (Saltz, Agrawal; DARPA, DOE, NSF): Systems software for efficient representation. Manipulation and querying of multi-dimensional multi-resolution data
- Scalable Communication Support with Modern Networking Technologies for Large Clusters (Panda; DOE SciDAC):
- Synthesis of High-Performance Algorithms for Quantum Chemistry (Sadayappan, Baumgartner; NSF ITR)

Distributed Cluster Environment



•
•
•

Students and Accomplishments

- 25 graduate students
- First employment of graduated students
 - IBM TJ Watson, IBM Research, SGI, Compaq/Tandem, Pacific Northwest National Lab, Fore Systems, Microsoft, Lucent, Citrix,
- Several of the past and current students
 - OSU Graduate Fellowship
 - OSU Presidential Fellowship
 - IBM Co-operative Fellowship
 - CIS annual research award
 - Best Papers at conferences
- Big demand for summer internships every year
 - IBM TJ Watson, IBM Almaden, Argonne, Sandia, Los Alamos, Pacific Northwest National Lab, Intel, Dell, ..

•
•
•

Systems Seminar Series

- Held every Friday afternoon followed by a pizza party
- External and Internal (faculty and students) speakers
- Attended by all systems faculty and students
- Provides an environment for exchange of information (technical and social)

⋮

Summary

- Addressing cutting-edge research challenges with focus on multi-disciplinary applications
- Significant research funding from federal, industrial and state sources
- Synergistic group with significant growth in the last few years
- Plans for continue to grow during the next few years

⋮

Web/E-mail Pointers

<http://nowlab.cis.ohio-state.edu/systems>

E-mails:

{panda,saday,agrawal,lauria,srini,hakan,saltz}
@cis.ohio-state.edu

⋮