ARSC Host Site Presentation

Virginia L. Bedford, Arctic Region Supercomputing Center, Fairbanks, Alaska

ABSTRACT: *The Arctic Region Supercomputing Center (less than three years old) includes a large memory YMP (8/1024), a T3D, SGI visualization laboratories, mass storage system connected by various networking media. This talk will present an overview of these components and their support of arctic and other scientific research. Some experiences and recommendations for setting up a new site will be discussed.*

Introduction

The Arctic Region Supercomputing Center (ARSC) was funded by the Department of Defense on August 21, 1992, as part of the Strategic Environmental Research and Development Program (SERDP). The Y-MP M98, named Denali, became operational in January 1993. One of its features is 1 GigaWord of memory. In February 1994 a 128PE T3D was installed and this past January the processors on the Y-MP were increased from four to eight, and the memory of the T3D was increased from 2MW to 8MW per Processing Element.

Mission: The Arctic Region Supercomputing Center (ARSC) supports high-performance computing and networking activities associated with the Department of Defense and University of Alaska. The ARSC facilitates national and international high latitude and other research by providing high-performance visualization services to the University of Alaska faculty, researchers, and students and to other national and international scientists while fostering the growth and diversification of the Alaskan economy.

Facilities and Resources

ARSC supercomputing resources include a Y-MP M98 with 8 processors and 1 GW of memory, a T3D with 128 PEs and 8GB total memory, a StorageTek Powerhorn robotic tape silo holding about 1.1 TB, 89GB of disk storage and an out of silo tape library. These are tied with visualization resources including multiple Silicon Graphics servers, workstations and disk space.

The hardware and technical staff of ARSC are located on the University of Alaska Fairbanks campus in the Butrovich Building. Although our CRAY Y-MP and T3D have been located here since original installation in the spring of 1993, the technical staff and developing Video Production Lab moved to the Butrovich facility just two months ago.

ARSC provides a training lab in the Butrovich Building for regular classes taught by staff on various topics, including programming and optimization on the YMP and T3D, visualization with AVS, and ABAQUS. In addition this lab provides facilities for offering courses from outside vendors and other special workshops on programs such as AVS and Wavefront.

Staff use a mix of Silicon Graphics Indy workstations and Macintosh computers.

ARSC provides most information about our resources and policies on our ARSC WWW server (http://www.arsc.edu). As this is intended to be our central repository for all information and user policies, it will soon also be accessible from our supercomputers via lynx for our users who do not have direct access to a GUI web browser.

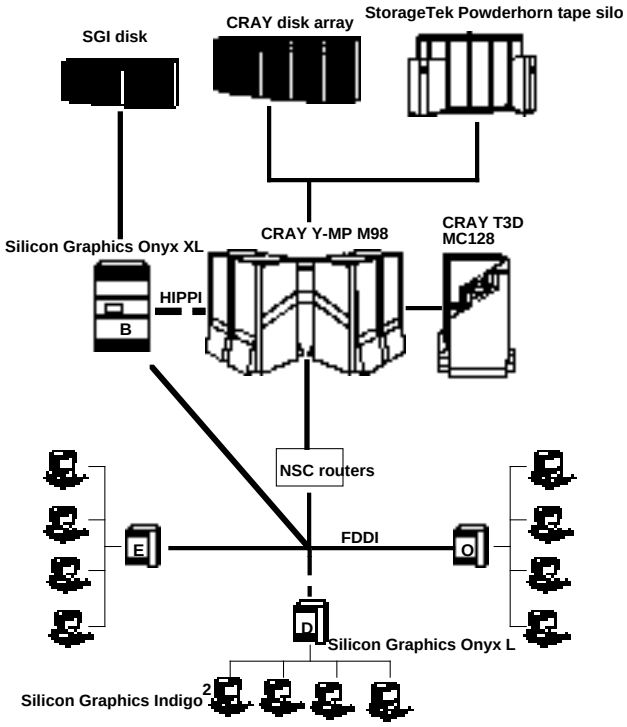

Figure 1: ARSC resources

Visualization Labs

In addition to our CRAY supercomputers, there are three ARSC visualization laboratories on the University of Alaska Fairbanks campus, located in buildings of the School of Engineering, Institute of Marine Sciences and Geophysical Institute. Each lab is equipped with an Silicon Graphics Onyx L server, and four Indigo^2 visualization workstations. A master server is located in the Butrovich Building. The ARSC Video Production Laboratory is also located in the Butrovich Building.

The visualization labs and production laboratory will soon be connected to each other by FDDI, and by a HiPPI connection, to the Y-MP. These projects should be completed later this fall.

The lab managers in each lab work partly for User Services, and partly for the host departments, answer questions about the fundamental use of the equipment. Each serves as a liaison between the particular users of the lab and as ARSC advocates for the use of visualization tools for their particular department users.

Users and Research

ARSC users (approximately 400) include researchers from the University of Alaska and other Universities, as well as branches of the military. Commercial businesses are also represented through the Technology Affiliates Program.

Research includes climate forecasting models, such as ocean circulation models and sea-ice-land-atmosphere interaction, reservoir modeling, scientific visualization, SAR signal processing, volcanic hazard monitoring and prediction atmospheric attenuation of communication satellite signals, strategic mineral management and high performance computing and networking

Staff and Infrastructure

ARSC's staff includes the Director and Administrative Services group, Development, Technical Services and User Services as well as several research scientist associations. Also included are four onsite CRI engineers and analysts. Technical Services has Service Agreements with University of Alaska Operations and Network Services groups for support in those areas.

ARSC User Services

Mission

The mission of the ARSC User Services Group is to provide the highest quality tools and training to enable researchers and computational scientists to efficiently compute, solve problems, and analyze their results through visualization and other techniques. In addition to providing support on current systems, User Services introduces new hardware and software tools for the users and provides persuasive evidence that new techniques result in better science. User Services also engages clients in dialog and provides advocacy within ARSC for user needs for computational resources, training and documentation. The customers of the User Services Group are external ARSC users, the Director, the Allocations Committee and ARSC Administrative Services.

Responsibilities

Support of external Web server, training, applications software support, user account administration (client account database support), problem tracking and resolution, representation and promotion of, and advocacy for, ARSC client interests through intimate knowledge of user needs and problems, assistance in proposal development, user allocations, billing (with technical services) and addressing contractor statement of work task requirements.

ARSC Technical Services

Mission

The mission of the ARSC Technical Services Group is to support the mission and users of ARSC by providing well-managed and efficiently run computer systems, with clearly defined policies and procedures known and followed by all staff and communicated to users. The ARSC Technical Services Group develops and follows processes that help ARSC respond quickly to the growth of and change in ARSC; it balances the need to be leading edge with the need to provide a robust system. The customers of the Technical Services Group are external ARSC users and all ARSC staff.

Responsibilities

Planning, evaluation, development, operations, monitoring and management of supercomputer operations, system integrity, disaster planning and recovery, operating systems and compiler maintenance and upgrades, security, storage management, accounting, supercomputer networking, performance evaluation and tuning, capacity planning, technology assessment, technical consulting, problem solving, status reporting, quality assurance, third party software, internal staff server (temporarily assigned to Technical Services), staff work station support, and billing (with User Services).

Operations is responsible for providing the facilities and support necessary for the efficient and effective running of the computer systems. This includes providing computer room space, meeting electrical and cooling requirements, fire protection, environmental monitoring, coordination with physical plant, tape media storage, coordination of the service agreements for the tape silo. Assistance with hardware acquisition and installation planning is part of this area as well. Day to day operations also include console monitoring and tape mounting, if needed.

External Network

ARSC currently uses a satellite T1 link to the lower 48, using Northwestnet as our internet provider. This link is also used by the University. In addition we are currently testing a terrestrial frame-relay T1 link, which is used by selected users who are participating in the tests. As expected, the terrestrial link has significantly improved customer satisfaction.

Both response time (due to the satellite delays) and bandwidth are a significant issue for us. Providing improved service for our users who are located Outside of Alaska is a high current priority. Several changes should be completed this fall to improve response: ARSC will switch to a terrestrial link, and the University will purchase its own link. These will be run in parallel and shared by both. Planned connections for Fall 1995 include the Defense Research and Engineering Network (DREN), and install several other dedicated links to specific users.

Software Configuration

Operating Systems

We are currently running Unicos 8.0.2.2, plan to upgrade to Unicos 8.0.4 this fall, and to Unicos 9.0 next year. One primary advantage for us in upgrading Unicos this fall will be to allow for further MAX upgrades. We are interested in being more current on MAX, and in taking advantage of job rollout and checkpointing for the T3D. Another big plus will be the ability to take advantage of the new MSP for DMF which will provide data compression.

NQS

Most of the work done on the Y-MP is done through NQS batch jobs, with some interactive work. The T3D is similarly characterized. A project planned for this fall is total rework of our queue structure. The T3D queue structure has been modified several times in an attempt to manually balance workload, and interactive and batch use. For example, within the past few months we added a 128PE queue which runs over the weekend, to allow for big longer jobs.

Storage Management

ARSC has 89 GigaBytes of disk storage on DD-60s and DD-62s. This is supplemented by a StorageTek Powderhorn robotic tape silo which can currently hold about 1.1 TB of data. Cray ReeL Librarian (CRL 2.0.6) and the Data Migration Facility (DMF 2.2.6) are used for data storage.

All tapes are managed by CrayReeLLibrarian. Of the possible 5,546 slots within the silo we have allocated 4,000 tapes to extend our user and /tmp file systems by the Data Migration Facility (DMF). The remainder are used by users directly and for file system backups.

Four user file systems and one temporary file system are divided is based upon user function. These five file systems are exported read-only to the visualization and staff workstations.

Users are given a disk quota and encouraged to use CRL and /tmp for large files. The temporary file system is extended somewhat atypically by DMF and data on it is kept for 21 days, if possible. Although this is an unusual treatment of temporary file storage which is more usually flushed, it has been useful for some of our users.

All DMF tapes reside within the silo. A close eye is kept on the available DMF tape threshold through various tools. The response to reaching this limit is to work with users to remove big old, unaccessed data to avoid lowering our retention period on /tmp data, a last resort. We reserve a small threshold of available slots for importing tapes into the silo. This fall, after conversion to Unicos 8.0.4, we will add use of the new DMF MSP for compression.

To maintain a pool of available CRL tapes, those CRL tapes which have not been accessed within some period of time are moved out of the silo and users are notified. The tapes are still accessible, but require a manual tape mount by operators. Users are requested to notify us whenever they would like to reactivate their data.

Our goal is to maintain active data available for immediate use within the extended file system and to provide a location for temporarily inactive and archival data.

Users running large models often also have large amounts of data, and developing a long term appropriate storage plan is a high priority.

Future plans include implementing some sort of dual copy DMF process, possibly integrated with helical scan technology.

Account Management

Userids and groups, quotas and resource allocation are all managed within the User Services group. Because we use NFS to export our file systems, exact uid and gid matching is required.

Our resource management system, known as “resman”, was originally developed at the San Diego Supercomputing Center, and has been modified for our use. With support of a Cray grant, an alternate system has been developed, which is based on a customized web server which maintains state. Once implemented this new resource management system is expected to be able to do all account management. In addition to allowing project managers to maintain and query their own account balances it will allow for some system management .

Service Units are allocated to users by ARSC and a User Allocation committee on a quarterly basis. The accounting algorithm is intended to reflect some relative cost as well as influence user behaviour towards use of desirable and available resources.

```
SUs = 1.0   * Y-MP CPU [hr]
  +   0.01 * T3D CPU  [hr]
  +   0.003 * (Disk*real time) [GB-hr]
  +   0.003 * (DMFSilo*real time)[GB-hr]
  +   0.005 * (Y-MP Mem*Y-MP CPU) [MW- hr]
  +   0.0   *  CRL Tape
```

Figure 2: ARSC accounting algorithm

T3D

ARSC has a 128 PE T3D, one of the first. This past January memory was upgraded to 8MW per PE, with a total 8 GB available. We have a parallel systems specialist who teaches courses on use of the T3D, puts out a regular newsletter which is also available on the ARSC web pages, and identifies and works with

possibilities for code which may be appropriately run on the T3D.

The tight coupling of the T3D and Unicos operating systems have made it difficult for us to stay at the latest version of MAX. For example, we would like to take advantage of job rollout. However, this required upgrading to Unicos 8.0.3/4, and we have not had the resources to do this. We plan to upgrade to Unicos 8.0.4 this fall, and to upgrade MAX at the same time.

This past year, with the help of CRI, we evaluated the possibility of implementing Phase 2 I/O to improve I/O speed, and determined that it would not be appropriate for our site.

ARSC and the WWW

ARSC has provided information via the web for just over a year. This has been quite successful in documenting site information as well as policies. We are in the process of installing a line mode web browser (lynx) on our Y-MP so that the web server can provide full information to our users.

About six months ago we extended this server by creating an internal, staff only server, which contains detailed operational procedures for Operations, Network, User Services and Technical Services staff. It has proven extremely useful for documenting standard methods for current staff, providing information for newly hired staff and as background for a system audit.

In addition to detailed checklists for some standard procedures a software and hardware inventory is being developed. These integrate links to other useful sites such as the CRINFORM pages and third party product home sites.

ARSC Internal Support Information

The ARSC Staff Web server is intended to provide a private, central repository for policies, processes, procedures and internal ARSC information.

- ARSC Missions
- ARSC Staff Projects
- Administrative Information
- Meeting Notes, Status Reports and Agendas
  - Agenda for next Ops meeting
- System Software (What, When, Where, Why, Who and How)
- Storage Management
- Security Policies and Guidelines
- Resource Accounting
- Technical Services
- User Services
- Public web server
  - Document submission/revision procedure
  - Personal home page policy
- Problem Management
- Service Agreements
  - Operations Support
  - Network Support
- Macintosh Support
- CRInform
- CRI Service Bulletins

Figure 3: Internal Support home page

ARSC System Software: What, When, Where, Why, Who, How

- System and Local Software
  - Local Systems Software Layout (Crays)
  - Software Installation Request Process (Crays and SGIs)
  - Cray Software Versions, History, Details
- Upgrade Policies and Processes
  - Compiler Upgrades
  - T3D MAX Operating System Upgrades
- System Testing Policies and Processes
  - "It was the best of times, it was the worst of times."
  - Testing Checklists

Figure 4: System software internal web page

User Services

Procedures

- Consult Desk procedure
- Account maintenance
  - Denali
  - SGI workstations
- Revising a web server document (external server)
- Remote tape drive operating procedures
  - chmotd, tool for changing the motd on denali
  - ednews, tool for adding/revising news items on denali
- Library update procedures

Tasks

- Training tasks for new consultants
- Current tasks and projects for consultants and interns

References

- Man page information

Figure 5: User Services internal page