Interactive Ocean Model Demonstration Over a National ATM Network

Rick O'Connor and *William Gardner*, Grumman Data Systems and Services, Inc., *Dan Fox*, Naval Research Laboratory, Stennis Space Center, *Peter Flynn*, Planning Systems Inc., *Robert Moorehead*, Mississippi State University, and *Bob Knesel*, Naval Oceanographic Office, Stennis Space Center

ABSTRACT: *During Super Computing '95, the Naval Oceanographic Office (NAVOCEANO) and Naval Research Laboratory (NRL) conducted a Live Test Demonstration of a 3-Dimensional, interactive Ocean Model running on a Silicon Graphics, Inc.,(SGI) ONYX and a CRAY C916. The ONYX and C916 were interconnected from a Hybrid Legacy LAN and Local OC-3 based ATM Network at San Diego, California through the ACTS ATM Internetwork (AAI) DS-3 based ATM Wide Area Network (WAN) to a Local OC-3 based ATM Network at Stennis Space Center, Mississippi. This paper presents the networking methodology, results obtained and lessons learned, including subsequent projects planned for near-term implementation.*

1 Introduction

The Supercomputer Center at the Naval Oceanographic Office (NAVOCEANO) administers and operates one of the four Department of Defense (DoD) High Performance Computing (HPC) Major Shared Resource Centers (MSRC), established October 1, 1994, under the auspices of the DoD HPC Modernization Program. The NAVOCEANO DoD MSRC, located at Stennis Space Center (SSC), Mississippi, is responsible for providing oceanographic products and services to all elements of the Department of Defense. NAVOCEANO collects and processes oceanographic, acoustic, and mapping, charting, & geodesy (MC&G) data, and its Naval Ice Center collects and processes ice features using ships, aircraft, satellites, and data buoys. Data bases and tactical products are produced and disseminated to ensure safe and accurate navigation and effective employment of tactical and strategic weapons systems and sensor systems.

The Supercomputer Center consists of a CRAY C916, a CRAY YMP8, and a CRAY YMP2E. These Supercomputers and associated Subsystems provide world ocean and atmospheric data at the accuracy and resolution required and improved predictions of environmental parameters which effect the performance of naval weapons and system sensors at sea. Communications links provide connectivity from T-1 (1.544 Mbits/sec) to T-3 (45 Mbits/sec) speeds for DoD users. The Center features advanced, intelligent Switched Ethernet and Fiber Distributed Data Interface (FDDI) Local Area Networks (LAN) and the foremost Network Management System (NMS), SPECTRUM, available today.

NAVOCEANO provides HPC capabilities to support DoD basic research and exploratory advanced development and operational efforts of the DoD Science and Technology Program. NAVOCEANO's T-3 connection to the Advanced Communications Technology Satellite (ACTS) Asynchronous Transfer Mode (ATM) Internetwork (AAI), and involvement with the Information Wide Area Year (IWAY) SuperComputing'95 (SC'95) event afforded the opportunity to demonstrate a High-Speed Application still under development over an experimental ATM Network in real time. NAVOCEANO's sponsorship and computer and manpower resource support made the project possible.

2 Demonstration Model

The Naval Research Laboratory (NRL) at Stennis Space Center (SSC) has developed Ocean Models and Data Assimilation Systems that provide realistic depictions and forecasts of the ocean environment. One of the models under development was chosen as a demonstration for the Supercomputing '95 event in San Diego. Ocean Models of this type typically run as batch jobs at SuperComputing sites. Output from a model is usually plotted and animated several days or weeks later. Consequently, the development cycle that includes fine-tuning the performance of the model can be quite extensive.

In order to test the concept of Interactive Model Development, a model of the Sea of Japan was modified to replace its output data files with a UNIX socket connected to a remote, separately running Visualization process. The model's input routine was also expanded to allow interactive changes to some critical parameters.

The Visualization Program entitled Interactive Structure Time Varying Visualization (ISTV), under development at Mississippi State University (MSU), was modified to read the information from the remotely run model. ISTV allowed the researcher to view three dimensional (3-D) ocean structure interactively and change the view point as desired. Information that was displayed included layer interface depths and actual ocean current vectors in each layer.

The circulation model ran on the CRAY C916 at NAVOCEANO, with the output viewed in real time at SC'95 using ISTV running on the ONYX and the C916, while critical parameters were adjusted remotely. ISTV was modified to include routines from the SC'95 Virtual Reality CAVE Library. This process provided a "walk-through" of the running model in addition to the screen view. The importance of the ATM connection was to exhibit a meaningful response to the application executing at SSC but displayed at SC'95.

The model, running at NAVOCEANO, provided ocean current data over an experimental ATM Network to the 3-D Virtual Reality environment at San Diego. Researchers, students, and observers were able to watch the ocean model displayed in real time and the user was able to steer around the Sea Of Japan, from the ocean floor up to various altitudes in the air over the surrounding terrain. The ocean current model data was displayed as a series of connected spheres for different depths in the sea.

3 Network Methodology

The ONYX and C916 were interconnected from a Hybrid Legacy LAN and Local Optical Carrier-3 (OC-3, 155 Mbits/sec) based ATM Network at SC'95 San Diego, CA, through the AAI Digital Signal level 3, (DS-3, 45 Mbits/sec) based ATM Wide Area Network (WAN), to a Local OC-3 based ATM Network at NAVOCEANO SSC.

AAI, consisting of FORE Systems ATM switches interconnected throughout the United States, provided the WAN link between the two sites. The WAN represented a pure ATM Network due to the fact that Classical Internet Protocol (IP) was converted to ATM cells at the local switches and was transported by ATM technology and protocols. The "ATM" portion of the Network Architecture went one step further at each site.

At NAVOCEANO, Classical IP over ATM Adaptation Layer 5 (AAL5) was generated from a Bus-Based Gateway (BBG) to the site ATM Switch that was connected to AAI. The BBG provided the CRAY C916 with an ATM OC-3 Network Interface operating at 155 Mbits/sec. The CRAY side of the BBG was a High Performance Parallel Interface (HiPPI) connection operating at 800 Mbps.

Organization	Location	Responsibility
NAVOCEANO Naval Oceanographic Office	Stennis Space Center, Mississippi	CRAY C916 Running Intercative Sea Of Japan Ocean Model Server Program CRAY C916 BBG ATM Interface Config SSC ATM Network Config
NRL Naval Research Laboratory	Super Computing '95 San Diego, California	SGI ONYX Running Interactive Sea Of Japan Ocean Model Client w/ Real Time Display
MSU Mississsippi State University	Super Computing '95 San Diego, California	SGI ONYX Running ISTV Visualization Program SGI ONYX Running CAVE Library Routines
AAI ACTS ATM Internetwork	Super Computing '95 San Diego, California AAI TIOC Kansas City, Missouri Stennis Space Center, Mississippi	AAI ATM WAN Config Network Connectivity To SC'95 And Stennis Space Center
SCInet IWAY Contractor	Super Computing '95 San Diego, California	SC '95 Hybrid Legacy LAN / Local ATM LAN Config

Figure 1: SC'95 DEMONSTRATION ORGANIZATIONAL CHART

At SC'95, Classical IP over ATM AAL5 was generated from a Router with an ATM Interface to the sitewide AAI ATM Switch. The ONYX was attached to the Router through an FDDI connection. The FDDI Ring and the Router formed the Legacy LAN portion and coupled with the Router ATM connection to the AAI Switch, represented the Hybrid LAN/Local OC-3 based ATM Network.

The AAI portion of the Network was interconnected to IWAY in San Diego approximately six weeks before SC'95 was scheduled to begin. All AAI participating sites had submitted Application Proposals and Network Requirements months in advance. The WAN Switches were configured and tested for optimum performance with network updates provided by two Address Resolution Protocol (ARP) servers located in Kansas City, Missouri, and in Washington, D.C.

Local network configurations and the WAN configuration took weeks of preparation and planning. Network testing was only allowed between the SSC BBG and the SC'95 AAI Switch prior to the day of the demonstration.

End-to-end testing between the C916 and the ONYX was approved ten minutes before the demonstration and continued for the next fifteen minutes. The actual demonstration started five minutes behind schedule. This was due to the unstable nature of the router software and its ATM Interface Code in particular. For one week prior to the demonstration, the router had crashed repeatedly on a daily basis when the ATM Interface was configured.

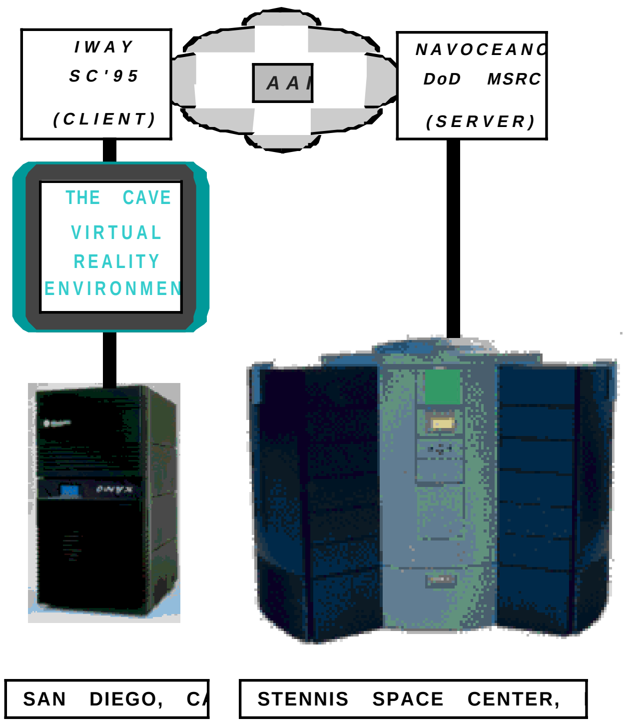

Figure 2: ATM NETWORK LINKING CLIENT AND SERVER

4 Results Obtained

During the Sea Of Japan Ocean Current Model Demonstration at SC'95, valuable statistics were gathered. The demo required approximately 36 Mbps bandwidth in order to run in real time. This was very close to the theoretical maximum of the AAI ATM WAN. Since AAI sites were internetworked with DS-3 (45 Mbps) links, the highest throughput we could expect was about 38 Mbps. (The remaining bandwidth on a T-3 or DS-3 is reserved for network signalling and overhead.)

Available bandwidth for the DS-3 interfaces on the Local Site AAI ATM Switches and on AAI Intra-Network Switches was increased to 45 Mbps. The entire bandwidth was dedicated to the assigned ATM Permanent Virtual Path/Permanent Virtual Circuit (PVP/PVC) of the link. During the demonstration Network Session, all other ATM traffic between the two Site Switch DS-3 Interfaces was de-prioritized. Network traffic and statistics were monitored by AAI and by Network Engineers at both sites. Data integrity was monitored by Systems Engineers at both sites.

MONITOR STATION	SUS-TAINED THROUGH PUT	% CELL / PACKET LOSS
BBG NAVOCEANO (SERVER)	80 Mbps	5 %
AAI ATM SWITCH NAVOCEANO	36 Mbps	0 %
AAI ATM TIOC NMS	36 Mbps	0 %
AAI ATM SWITCH SC'95	36 Mbps	0 %
SGI ONYX SC'95 (CLIENT)	78 Mbps	0 %

Figure 3: NETWORK STATISTICS DURING DEMONSTRATION

The application had the capability of using all available CPUs on the C916. Pre-demonstration application testing indicated that the server component of the application would perform best with 4 CPUs.

The client component of the application running on the ONYX used 1 of the 2 CPUs available.

5 Lessons Learned

The following observations were based on tests conducted during the demo:

On the C916, the distributed visualization application exhibited poor performance relative to the Client/Server Model and its utilization of hosted hardware. This was due to inefficient tuning of interprocess communications between the client and server components. We learned that with better tuning, the High-Speed ATM Network could be better utilized.

The FDDI Network Interface of the SGI ONYX for the client component was unable to keep pace with the ATM Network and the server's output. Synchronization between the server and the client was accomplished by throttling back (close to idle) the server component to allow the ONYX graphics engine to process the data flow.

A marked increase (5%) in dropped cells due to collisions occurred on the ATM interface card of the BBG. This is believed to be caused by the card's buffering scheme.

Transmission Control Protocol (TCP) windowing parameter changes could be adjusted to provide performance enhancements for better optimization of the ATM Protocol.

PVP/PVC (Static) configurations were required to ensure Constant Bit Rate (CBR) availability across the WAN. Dynamic routing in the ATM Network must undergo continued extensive testing prior to production network implementation.

Router crashes encountered during SC'95 were caused by unstable router ATM software.

Planning and testing a project of this magnitude requires extensive lead time. The participants utilized unproven hardware and software to meet the deadline.

Failure incidences and lack of pre-demonstration testing caused extreme anxiety during the week of SC'95 and during the day of the demonstration. Due to the instability of the Hybrid Network at SC'95, we could not perform end-to end testing until the demonstration was about to begin. Improved control over both ends of the network could alleviate such situations and provide ample time for disaster recovery if required.

6 Planned Projects

Projects planned for near-term implementation include TCP windowing optimization for ATM networking between CRAY computers; Open Systems Foundation (OSF) Distributed Computing Environment (DCE) Client/Server implementation and testing over ATM networks; and a remote visualization application that will run on the C916 at NAVOCEANO and be displayed on an SGI server during the Special Interest Group on Computer Graphics (SIGGRAPH'96) Conference in New Orleans.