Navy Collaborative Integrated Information Technology Initiative

(NAVCIITI)

ONR Grant N00014-00-1-0549

Report No. 12

Quarterly Report

Year 2

Submitted to:

Mr. Paul Quinn
ONR 311
Office of Naval Research
Ballston Centre Tower One
800 North Quincy Street
Arlington VA 22217-5660

Submitted by:

Virginia Tech
Blacksburg VA 24061

1 May 2000 – 31 July 2000

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)* 30-08-2000	**2. REPORT DATE** Technical	**3. DATES COVERED** *(From - To)* 1-5-000-31July2000

4. TITLE AND SUBTITLE Navy Collaborative Integrated Information Technology Initiative	**5a. CONTRACT NUMBER** N00014-00-1-0549 **5b. GRANT NUMBER** N00014-00-1-0549 **5c. PROGRAM ELEMENT NUMBER**
6. AUTHOR(S) Kenneth Reifsnider Rick Habayeb	**5d. PROJECT NUMBER** **5e. TASK NUMBER** **5f. WORK UNIT NUMBER**
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Virginia Polytechnic Institute & State University Office of Interdisciplinary Programs 306 Burruss (0244) Blacksburg, VA 24061	**8. PERFORMING ORGANIZATION REPORT NUMBER** 12
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research± Ballston Centre Tower One 800 North Quincy Street Arlington, VA 22217-5660	**10. SPONSOR/MONITOR'S ACRONYM(S)** **11. SPONSORING/MONITORING AGENCY REPORT NUMBER**

12. DISTRIBUTION AVAILABILITY STATEMENT

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Quarterly progress report on NAVCIITI hardware and software configurations of smart antennas, multi-function antennas, secure configurable platform, command & control visualization, visualization and HCI, collaboration workspace, digital ships, mechanically flexible displays, fiberoptic LAN, network protocol interoperability, network system interoperability, and scheduling and resource management.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			K Reifsnider
U	U	U		75	**19b. TELEPONE NUMBER** *(Include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI-Std Z39-18

1.1 Wireless communications

Quarterly Report for NAVCITTI 7/30/00

Smart Antenna Research

Jeffrey H. Reed and Bill Tranter
Mobile and Portable Radio Research Group

Project Background

The goals of our smart antenna effort are to

1. Create a Smart Transmit and Receive Antennas Test Beds and Demonstrate
   - Anti-jam capability
   - Increased capacity
   - Extended range
   - Low probability of intercept (LPI)
   - Reduced fading
   - Lower transmit power (increased battery life)
2. Create a Space-Time Coding Test Bed and Demonstrate
   - Improved throughput
   - More reliable communications
3. Perform Vector Channel Modeling to create
   - Design & deployment tools
   - Speed network simulations

Efforts for the Quarter and Accomplishments

During the first quarter we did an initial investigation of the hardware necessary to create the test beds. This included the creation of a two-element transmit smart antenna test bed based on two Modular Radio Components (MRCs) produced by Rockwell. These components are the fundamental building block of the transmit smart antenna and space-time coding test beds.

We also tested our 4-element vector measurement system (a smart receive antenna array), a key element in our space-time coding test bed. Some initial measurements have been obtained for an indoor environment and the data is currently being analyzed. We are in the process of determining the parameters of the measurements that are necessary to support creating the Markov channel models. These models will greatly speed simulations involving smart antennas placed in a network.

Figure 1 provides a description of an example indoor environment in which a test was performed. The whole system was set up in room 618 of Whittemore Hall at Virginia Polytechnic Institute and State University. The transmitter was set so that there was a line of sight with the receiver. The channel provided a direct path as well as different multipaths bouncing off from the walls, metallic objects and other sources of scattering. The transmitter-receiver separation was about six feet.

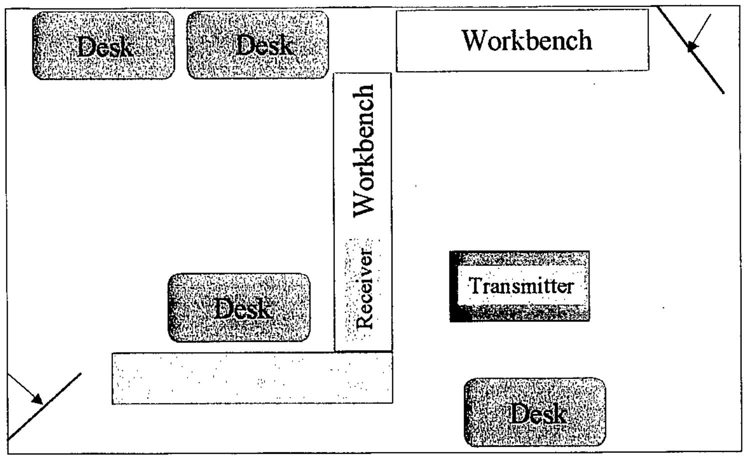

Figure 1. Demonstration environment

Power measurements were done on the received signal for both the diversity system and single-antenna systems. The received data was processed to generate the cumulative distribution function (CDF) plots that show the probability that an SNR will be above a certain threshold. Figure 2 shows the instantaneous received power plots and the corresponding CDF plots for a particular setting of the measurement system. These plots conservatively represent the gains provided by the smart antenna since the line-of-sight component is strong in this instance. We expect to see about 9 dB of gain in non-line of sight conditions.

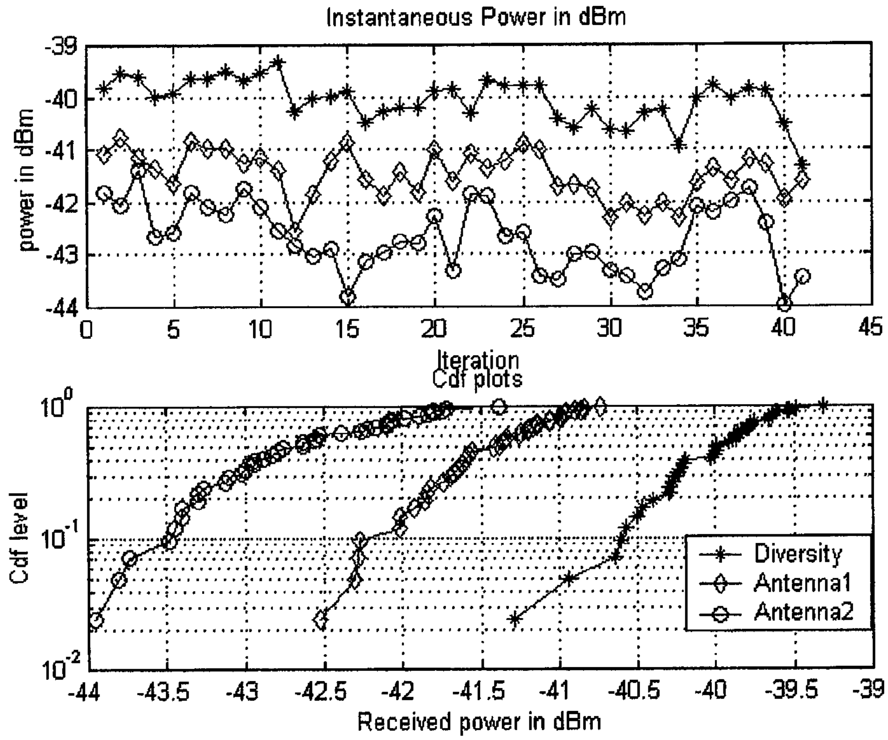

Figure 2. Instantaneous power plots and CDF plot for one data collection

Activities during the quarter: presentations, publications, workshops, visits, attending / sponsoring conferences, symposia.

Poster presentations were made on the vector propagation measurement system and the smart antenna test bed at the June MPRG Symposium on Wireless and Personal Communications. In July we demonstrate a rudimentary two-element smart transmit antenna to DARPA. Furthermore, we have also been in contact with Rich North of SPAWAR, San Diego, to transition some of the material from our research into educational material that would benefit the Navy.

Plans for next quarter.

Next quarter we plan on refining our transmit smart antenna demonstration. If Rockwell is able to fix the non-linear problems with the MRCs we will proceed to add two additional MRCs to the test bed and include amplitude adjustments in the adaptive algorithms. Additional test will be conducted to gain more experience with the performance of smart transmit antennas at handsets.

Additional propagation data will be collected using VIPER, the 4-element receive smart antenna. This data will be processed and used to help create the Markov vector channel models. Markov models use measurement data to create an on or off channel state for a system and can greatly reduce the complexity of the physical layer when doing network level simulations. The vector channel Markov model will provide this on/off state for modeling an adaptive antenna array within network simulations. The measurement data is necessary to create this Markov model and will also be used in our analysis of the trade-offs between time processing and space processing for exploiting multipath.

Our algorithm development efforts will focus on wideband techniques for space-time processing. We plan on having an extensive set of simulations demonstrating the performance of wideband smart transmit antenna algorithms. Likewise, we will begin investigating space-time coding algorithms in simulation for eventual deployment on our hardware test bed.

We plan to present papers at a workshop in August sponsored by TI on our smart antenna work. We plan on following up this presentation with the submission of a journal paper on smart transmit antennas.

Issues

We are concern that the non-linear characteristics of the MRC transmitters will limit the capabilities of our test beds. We are able to adjust the phase accurately, but amplitude characteristics are non-linear. Rockwell has been informed of the problems with the MRCs and has committed to fixing the problems. Their first attempt at this has met with limited success, but we will continue to push Rockwell to fix these problems. Contingence plans are being made to produce our own transmitters if Rockwell is unable to fix the problems with the MRCs.

1.2 Multifunction Antennas

Wordnavy/nav7-00.doc 8/7/00

RESEARCH PROGRESS FOR THE FIRST QUATER OF YEAR 2 FOR NAVCIITI PROGRAM ELEMENT 1.2, Wideband Antenna Research

Warren Stutzman, William Davis, Randall Nealy, Carey Buxton, Ko Takamizawa, and Seong-Youp Suh

Antenna Group
Electrical and Computer Engineering Department
www.ee.vt.edu/antenna

This report describes progress during the first quarter (May 1 - July 31, 2000) of Year 2 on the wideband antenna research portion of the NAVCIITI program. Our work is in three areas: development of measurement and analysis capability for wideband antennas, research on wideband element antennas, and research on wideband antenna arrays. This report discusses those areas. In addition, we have expanded our industrial involvement. Harris Corp. is building and testing wideband array antennas in cooperation with VTAG. Also, Northrop Grumman became a VTAG Industrial Affiliates and is very interested in wideband arrays for their programs that include Navy activities.

Task 1.2.1 Completion of the Test Facility

A few hardware and software anomalies remain in the near field scanner. The equipment is under warranty and the vendor is coming to correct these. Calibration efforts have produced good results. A paper was presented at a national meeting describing the initial results from anechoic chamber measurements.[1] Measurements were made for a microstrip antenna using both near field and far field anechoic chamber pattern data as well as pattern data measured with NASA Langley's anechoic chamber. Good agreement was obtained.

Task 1.2.2 Wideband Element Antenna Design

Efforts to improve the performance and bandwidth of planar dual polarized, compact radiating elements continued by investigating variations of the Foursquare antenna. VTAG simulated many structures to understand geometric parameter influences. One form, the Fourpoint antenna, that showed excellent results in simulation was constructed for 6-12 GHz operation. The reflection coefficient was measured using a 180° hybrid to generate a balanced feed to the antenna. The measured data and simulation data are compared in Fig. 1.3.2-1.

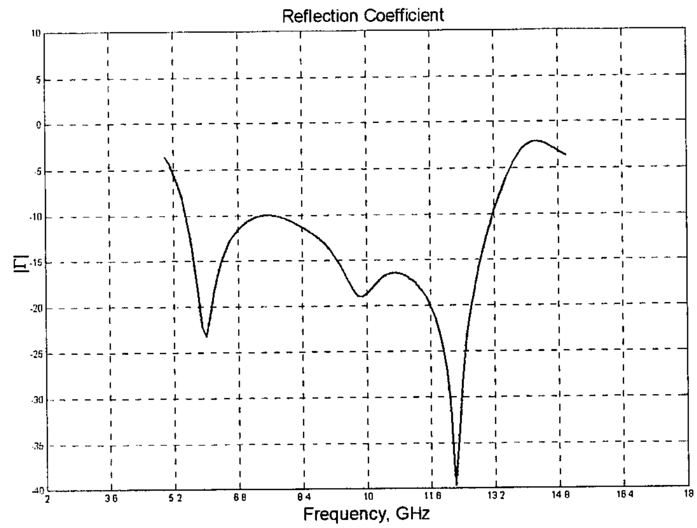

(a) Calculated reflection coefficient.

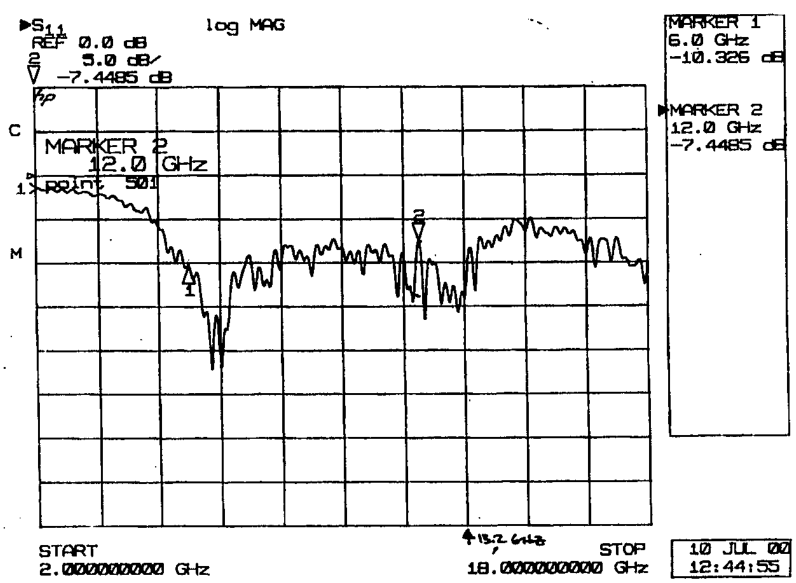

(b) Measured reflection coefficient.

Figure 1.2.2-1 Reflection coefficient of the Fourpoint Antenna operating at 6-12 GHz

The radiation pattern and gain of the Fourpoint Antenna operating at 6-12 GHz were computed using a commercial code, Fidelity. We found through comparisons to measured data that the antenna pattern and gain obtained from Fidelity are not accurate. However, the results are acceptable for estimating the patterns and gain. Fig. 1.2.2-2 shows the E-plane and H-plane patterns of the Fourpoint antenna at selected frequencies (6, 8 10, and 12 GHz) in the band. The E-plane pattern represents the pattern computed in a plane through the radiating portion of the element and the H-plane pattern corresponds to the plane through parasitic portion of the element. There is no significant pattern variation over the band in both planes.

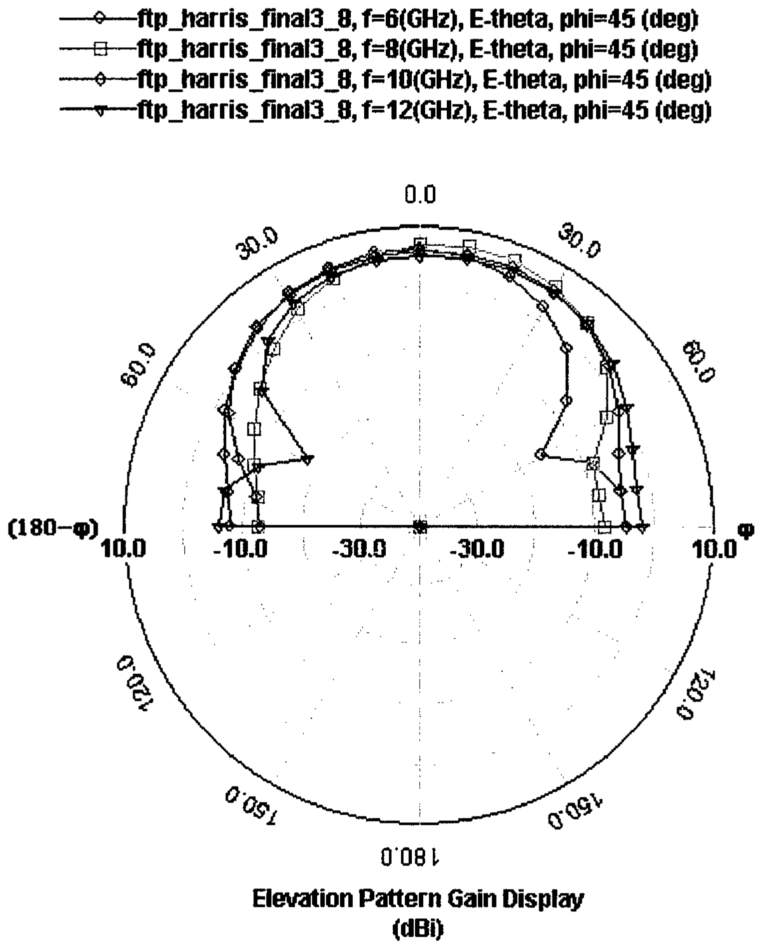

(a) Computed E-plane pattern of the Fourpoint Antenna

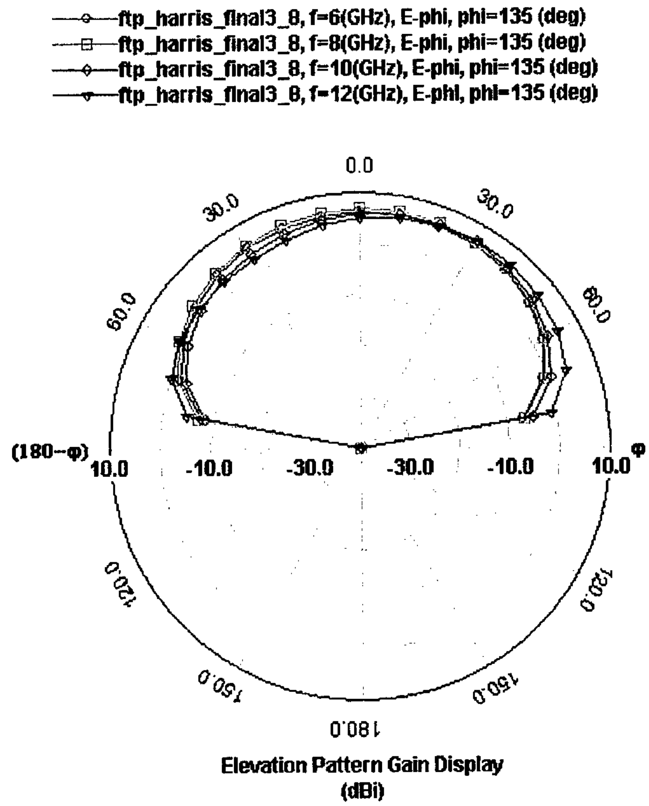

(b) Computed H-plane pattern of the Fourpoint Antenna

Figure 12.2-2 Antenna patterns of the Fourpoint Antenna operating at 6-12 GHz

Fig. 1.2.2-3 shows the calculated gain in E, H-planes; the values are very close each other, as they should, and are constant over the band. The curve with diamonds represents gain in the E-plane and the curve with squares represents H-plane gain.

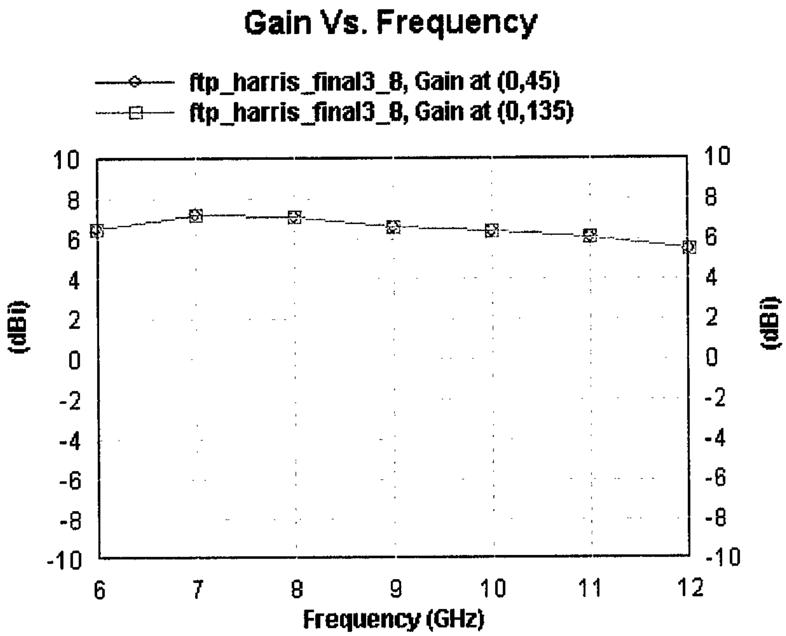

Figure 1.2.2-3 Antenna gain of the Fourpoint Antenna operating at 6-12 GHz

We also improved the performance of the dual-band Fourpoint antenna this quarter. The dual-band Fourpoint antenna reported in last quarterly report was operated in two bands: AMPS (824-894 MHz) and PCS (1820-1990 MHz). The antenna geometry was modified to obtain multi-band operation to include the AMPS, GSM (890-960 MHz), DCS (1710-1880 MHz), and PCS bands. The measured impedance and VSWR of the multi-band Fourpoint antenna are shown in Fig.1.2.2-4. The impedance bandwidth of the modified Fourpoint antenna covers both 810 MHz-1030 MHz and 1440 MHz-2200MHz at a VSWR value of 2. However, the VSWR over the four bands does not exceed 1.7. The multi-band Fourpoint antenna has several applications where multi-band operation is required. There are several such cases of interest to the Navy. Our research shows that the Fourpoint antenna can be designed to cover bands of interest to the Navy. Applications include both shipboard and airborne antennas, and terrestrial base station antenna supporting quad-band operation as well as orthogonal linear polarizations for polarization diversity.

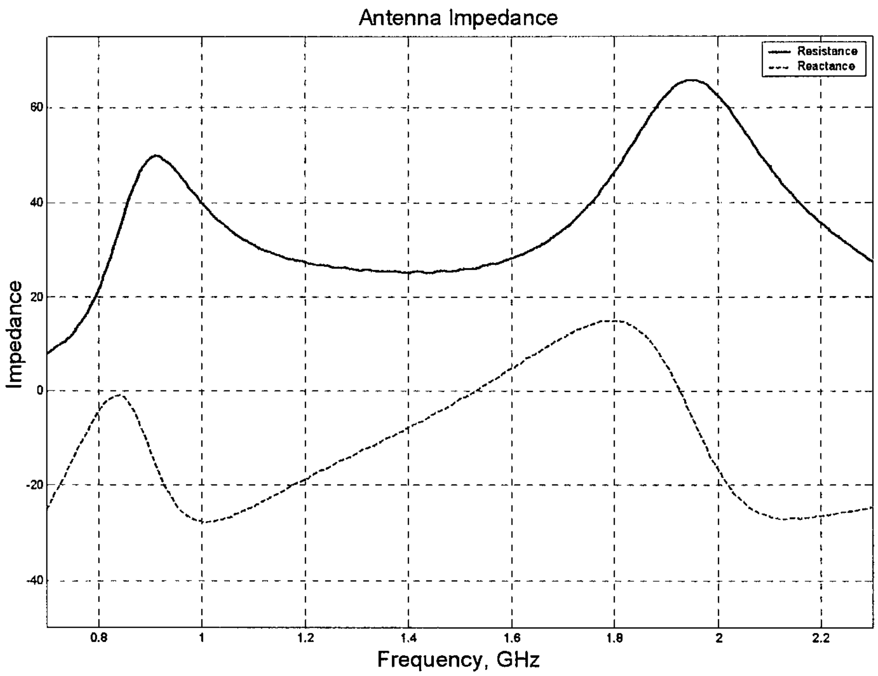

(a) Impedance of the multi-band Fourpoint Antenna

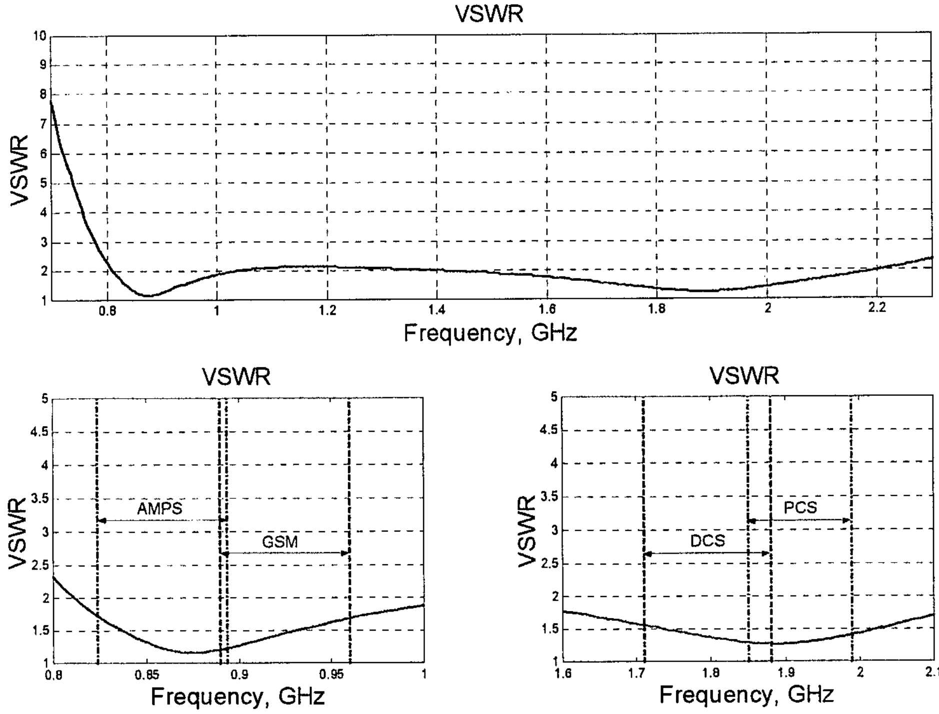

(b) VSWR of the multi-band Fourpoint Antenna

Figure 1.2.2-6 Impedance and VSWR of the multi-band Fourpoint Antenna

Task 1.2.3 Wideband Element Antenna Measurements

The multi-band Fourpoint antenna is being constructed and will be measured during the next quarter. The scanner will be used to measure the antenna pattern so that the measurement results can be compared with simulation results.

Task 1.2.4 New Approaches to Wideband Antennas

The 3x3 Foursquare Array is currently being analyzed using the in-house FDTD code. The tight spacing of the array creates strong mutual coupling between the elements. The effect of the mutual coupling on the far-field pattern of the array was calculated using the FDTD code. The E- and H-plane patterns from 3.52-7.04 GHz are shown in Figures 1.2.4-1 and -2 for the fully excited 3x3 Foursquare array. The H-plane half-power beamwidth is between 25° and 40° over the specified frequency band. The E-plane half-power beamwidth is between 18° and 30° over the specified frequency band.

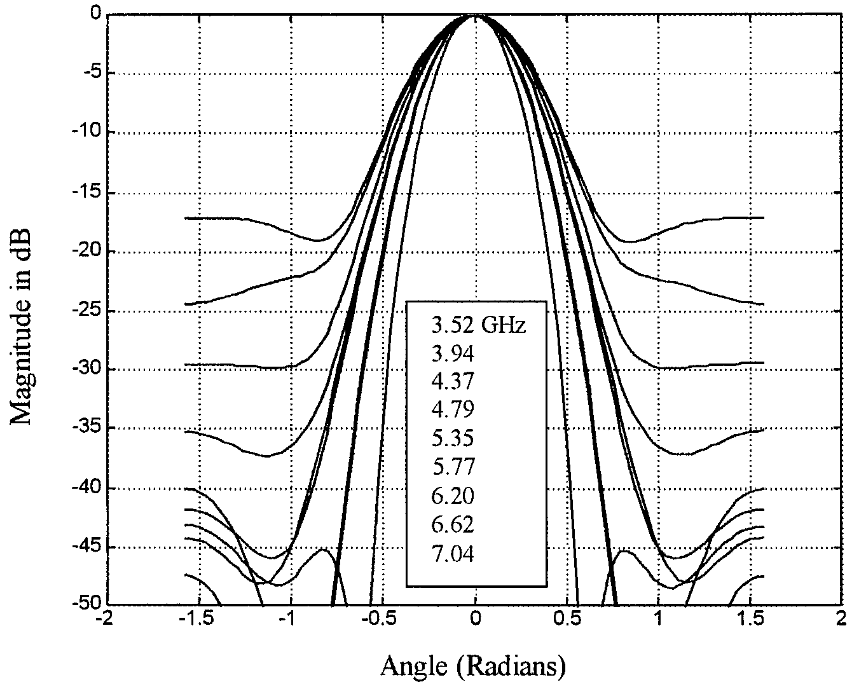

Figure 1.2.4-1 E-plane pattern of fully excited 3x3 Foursquare array.

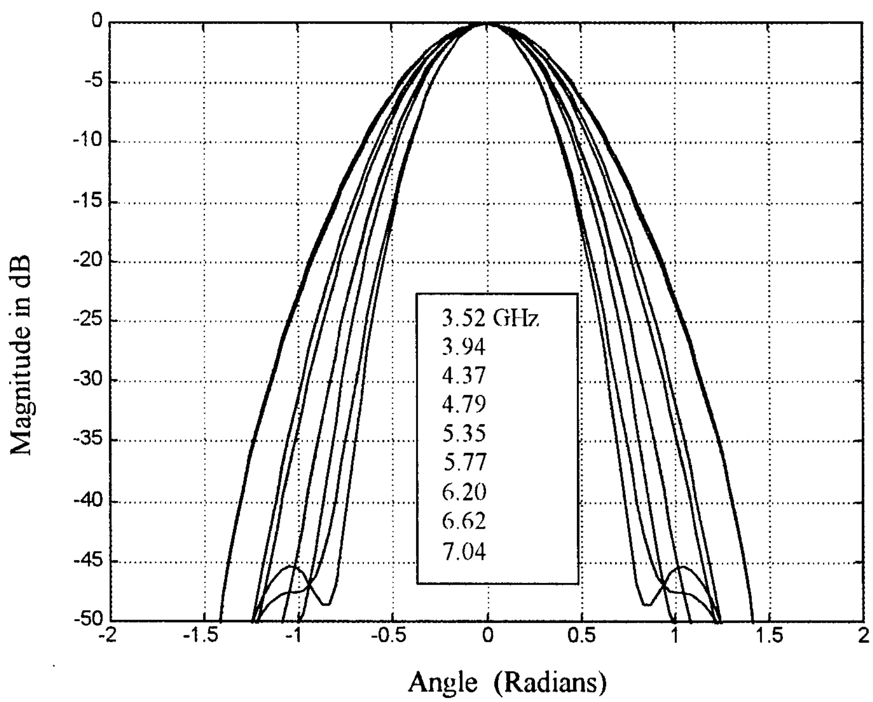

Figure 1.2.4-2 H-plane pattern of fully excited 3x3 Foursquare array.

Task 1.2.5 Array Antenna Investigations

The Foursquare antenna is an ideal choice as an element in a planar array. The Foursquare elements are typically fed with a balanced feed using two coaxial cables. Our investigation showed that there is some interaction between the feed and the antenna elements in the array of balance fed elements. We investigated the effects of feed-element interaction for a 3x3 array of Foursquare antennas using electromagnetic modeling and network modeling. The antennas were simulated using the commercial code IE3D, which is based on 2.5D Method of Moment. The simulations were conducted using infinite dielectric layers and infinite ground plane. Two types of feeding conditions were investigated: 1) all 9 elements fed with coaxial cables (see Fig. 1.2.5-1a) and 2) the center element is fed with coaxial cables and the surrounding 8 elements are terminated with 100 ohm chip resistors (see Fig. 1.2.5-1b).

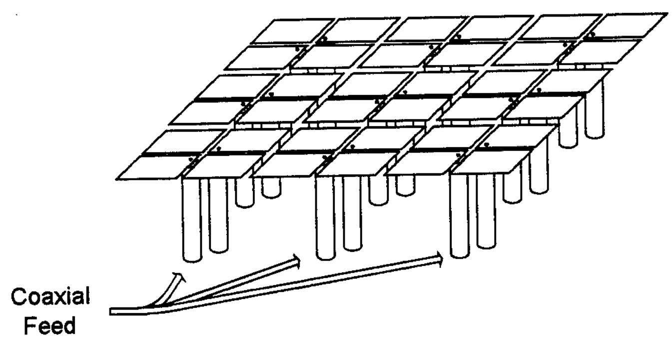

(a)

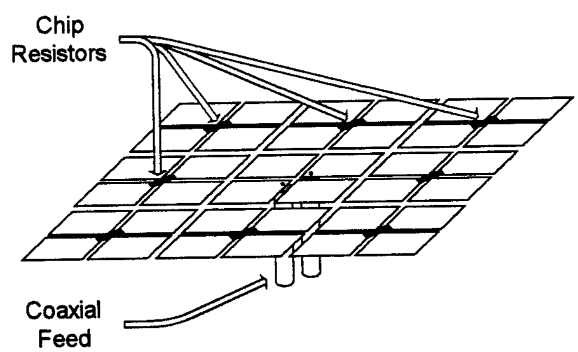

(b)

Figure 1.2.5-1 Geometry and feed configurations of 3x3 array of Foursquares simulated.

Fig. 1.2.5-2 shows the radiation patterns for the two feed configurations of Foursquare array mid band in the operational band and the radiation pattern of isolated Foursquare element at the same frequency. It shows that when all antenna elements are connected with coaxial cables, there is a significant degradation in radiation pattern shape and in gain due to coupling. The pattern degradation is a function of frequency. These results are consistent with experimental observations made at Harris. Since the radiation pattern degradation is minimum for the chip resistor terminated array as show in Fig. 1.2.5-2(b), it is hypothesized that the outer shields of the coaxial feed support the common mode currents on the elements surrounding the center Foursquare.

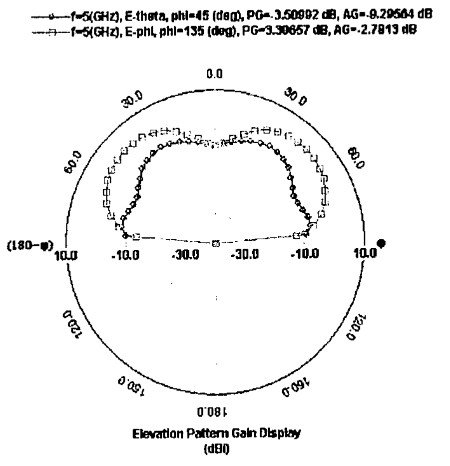
f=5(GHz), E-theta, phi=45 (deg), PG=-3.50992 dB, AG=-9.29564 dB
f=5(GHz), E-phi, phi=135 (deg), PG=3.30657 dB, AG=-2.7813 dB
0.0
30.0
30.0
60.0
60.0
(180-φ)
10.0
-10.0
-30.0
-30.0
-10.0
10.0
φ
120.0
120.0
150.0
150.0
180.0
Elevation Pattern Gain Display
(dBi)

(a)

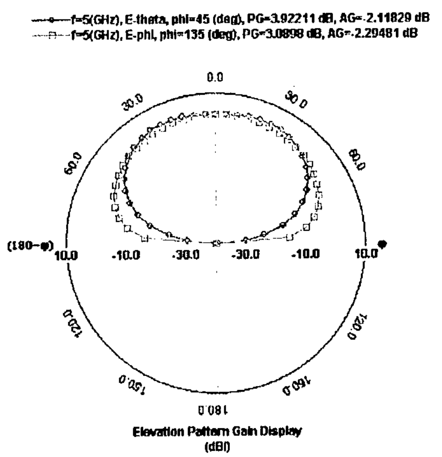
f=5(GHz), E-theta, phi=45 (deg), PG=3.92211 dB, AG=-2.11829 dB
f=5(GHz), E-phi, phi=135 (deg), PG=3.0898 dB, AG=-2.29481 dB
0.0
30.0
30.0
60.0
60.0
(180-φ)
10.0
-10.0
-30.0
-30.0
-10.0
10.0
φ
120.0
120.0
150.0
150.0
180.0
Elevation Pattern Gain Display
(dBi)

(c)

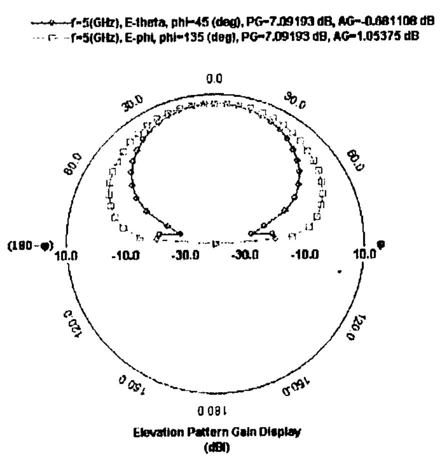
f=5(GHz), E-theta, phi=45 (deg), PG=7.09193 dB, AG=-0.681108 dB
f=5(GHz), E-phi, phi=135 (deg), PG=7.09193 dB, AG=1.05375 dB
0.0
30.0
30.0
60.0
60.0
(180-φ)
10.0
-10.0
-30.0
-30.0
-10.0
10.0
φ
120.0
120.0
150.0
150.0
180.0
Elevation Pattern Gain Display
(dBi)

(b)

Figure 1.2.5-2 Active element patterns calculated with IE3D for a Foursquare array and an isolated Foursquare element. a) 3x3 array of Foursquares with all nine elements connected with coaxial cables, b) 3x3 array of Foursquares with center element connected with coaxial cables.

Future work is necessary to determine the detailed mechanism of the pattern and gain degradations. Specifically, 3x3 array of Foursquare elements implemented at Harris for their test bed will be the simulated. Then detailed studies of currents on the outer shields of the coaxial feeds will be conducted for various feeding and feed termination conditions using the simulation. Then, assessments based on the simulation results will be made on the pattern degradation effects in the fully active array of Foursquare elements.

References

[1] K. Takamizawa, N. Cummings, W. Stutzman, and W. Davis, "Comparative Study of Analysis Techniques and Measurement Methods for a Canonical Microstrip Antenna," URSI National Radio Science Meeting (Salt Lake City, UT), July 2000.

1.3 Secure Configurable Platform

Peter Athanas
Associate Professor
Virginia Tech, Dept. of Electrical and Computer Engineering
Blacksburg, Virginia
http://www.ee.vt.edu/athanas

(1) Background

The purpose of this project is to demonstrate the viability of a potentially highly adaptable low-cost technology for naval information management and communications systems. It is the intent of this project to demonstrate an alternative means of collecting, processing, storing, and transmitting information securely throughout the next-generation seaborne naval platforms.

The technology featured in this task has been chosen based upon the following criteria:
(1) The focus of this project is on information management. The processing power of the underlying platform should be sufficient to address a variety of demanding computational tasks.
(2) The processing platform should utilize commodity (high production volume) components. The intent here is to reduce per-unit costs, manufacturing cost, and time-to-deployment as new devices become available.
(3) The hardware technology should be multifunctional, and be capable of fulfilling the operations of perhaps a variety of tasks with little or no modifications. The intent of this is to derive economy-in-scale by having one deliverable satisfy possibly several requirements. The system should handle data in a secure manner. System resources should be protected from unintentional (friendly misuse of system resources) or intentional (protection of software and hardware against theft and reverse engineering) threats.

(2) Planned Effort

The stated objective for this quarter was to acquire two configurable radio front-ends as well as a new generation configurable computing board. These are to be integrated into a typical PC platform.
In the proposal, we had stated our plan of using the Rockwell MRC for performing this function. We acquired one of these modules for integration, but unfortunately, the price of these modules has increased ten-fold. In response to this, we assigned one student to review alternatives before we commit to this platform.

(3) Accomplishments

Our approach to this project requires a hardware system that is capable of adapting to data traffic while running. This system should be able to examine traffic flowing through it and modify itself to provide optimal support for that information. Field-Programmable Gate Arrays (FPGAs) provide the type of reconfigurability desired, as they allow for both task-specific pipelines and traditional RISC processors to be embedded and modified. As a result, it was decided to use a currently available FPGA-based platform to support this project.

Several FPGA-based reconfigurable platforms were investigated, including Annapolis MicroSystem's Wildstar and Wildforce boards, and the University of Southern California's Information Sciences Institute's SLAAC1V system. While all three platforms are host-based, the goal of this research is to develop a system that does not depend on the host, and is instead entirely self-contained within the chosen platform. The

Wildforce board was the first platform examined. It contains FPGAs that do not support partial run-time reconfiguration, however, and was rejected. The Wildstar board was examined next. Unfortunately, it has some host interface problems that might significantly delay the project and was also rejected.

Finally, the SLAAC1V platform was investigated and chosen for development. It contains Xilinx Virtex FPGAs, allowing for partial reconfiguration of the system while running. It also provides onboard hardware for configuration control, allowing for operation independent of the host. Finally, the SLAAC1V platform comes with full access to the existing API and hardware schematics.

Once the hardware was chosen, processor mappings were examined to determine a design that can be implemented on the SLAAC1V platform. The processing system proposed classifies and processes continuous 32-bit data traffic in a pipelined fashion. The processing portion of the pipeline is contained in a single FPGA where it can be reconfigured while running to provide an optimal processing solution. The classification portion is contained in a second FPGA and can identify new data streams and adapt the system to handle them accordingly.

A basic version of this processor was developed using Java Hardware Description Language (JHDL). The JHDL language provides for rapid development of structural models, allowing a simple version of the processing system to be quickly developed and simulated. This model was not embedded in the SLAAC1V, however, due to problems interfacing with currently existing system support code. The existing code was written in VHDL and already embedded in some of the FPGAs. Currently, this existing system support code is being modified to provide sockets for additional features. Once the sockets are available, the processor will be ported to VHDL for a more straightforward implementation.

At the present time, the configuration control interface (CCI) is being developed. An existing CCI allowed the system to be controlled by a host via addressable registers. This interface has been extended to allow for on-board initiation of reconfiguration. The on-board controller uses the existing addressable registers to perform the operations previously available only from the host system.

In the final system, the configuration controller will govern system adaptation. Configuration information will be stored in on-board memory associated with the controller, and additional addresses have been added to the CCI to allow for reconfiguration requests and configuration store update requests. When a system processing element determines that a new configuration is needed, it will use the CCI to send a request to the configuration controller. The controller will then perform the tasks necessary to reconfigure the system as requested.

Once the configuration controller has been fully realized and embedded in the system, attention will be turned to the classification portion of the processor. This is the part of the processor that will identify data traffic, send it to an appropriate processing channel, and initiate configuration requests. The classifier will share an FPGA with the host interface, and once again require modification of existing code.

Finally, the data processing channels can be developed. They will be contained within their own FPGA, and do not require adaptation of existing SLAAC1V code.

(4) Objective Importance

This objective must be resolve before the detailed physical design phase begins.

(5) Activities

We had one meeting in June at TECSEC headquarters in Vienna, VA with Frank Adamouski (TECSEC) and Frank Deckelman (ONR) which served as a kick-off meeting.

(6) Plans for Next Quarter.

We are currently implementing a MIPS RISC core onto our platform which will serve as the means of executing the TECSEC Cryptographic Key Management software. We will work with TECSEC to cross-compile their x86 code to our custom processor. The mechanism of transfer will be object files. In addition, we will connect the radio front-end to the SLAAC1V processing engine, and prototype the stream controller.

(7) Issues

There are no issues.

2.1 Command & control visualization

Dr.Ronald D. Kriz
NAVCIITI Quarterly Report
Project 2.0: Visualization HCI and Collaboration
Task 2.1: Command and Control Visualization
SOW: 2.1.1 Identify working simulation models used by NUWC

Significant progress has been made on improving the CAVE Collaborative Console (CCC) with new HCI (Human Computer Interaction) features. Specifically a common desktop to CAVE menu interface was designed and implemented to encourage collaboration across platforms. The upgraded CCC was installed and tested on the Immersive-Desk at Dam Neck NAVSEA HPC (High Performance Computer Center) as part of the Collaborative Engineering Environment project. The CCC was used to link DAM Neck's I-Desk with Virginia Tech's CAVE. This upgrade was made possible by incorporating the "X-Wand" software so that CAVE users could function similar to the desktop user.

http://www.sv.vt.edu/future/cave/software/ccc

Significant progress has also been made on the development of the DIVERSE (Device Independent Virtual Environment: Reconfigurable, Scalable and Extensible) toolkit and graphics interface to Performer which will be used for both the NUWC (year-2) and Digital Ship (year-3) NAVCIITI project. Presently DIVERSE is being merged with VRPN (a tool used extensively at the NRL VE Lab) to link an array of devices in a distributed-shared collboartive environment. Future work will include CCC HCI features, such as participant lists, 2D- and 3D- radars, and teleportation, into DIVERSE. An improved HCI desktop CAVE simulation interface called "Persona" is designed to facilitate desktop users of shared virtual environments.

http://www.diverse.vt.edu

Under the direction of Dr. Larry Rosenblum we are waiting clarification of exact sonar simulation models from NUWC (Naval Undersea Weapons Center: Contact Ann Silva) to be incorporated into the DIVERSE collaborative interface that will connect desktops, Immersive- Workbenches, Head Mounted Display (HMD), and I-Desks with simulations running on remote-site HPC (High Performance Computers). A working Web-based prototype of and wave propagation simulation that could be extended to sonar simulation has been developed into a working prototype: section "simulation results". This particular prototype is based on a three-tier network architecture that allows remote-site users ("client") to access simulation models on a centralized computer ("server") that submits simulation jobs to supercomputers ("distributed- HPC-Resource") across the internet.

http://www.jwave.vt.edu/crcd/kriz/examples/example7

2.2 Visualization and HCI

Hix/Gabbard

Quarter 1 Year 2 – Progress report

Background

Virginia Tech has a long-term collaboration with NRL (Dr. E. Swan, Dr. J. Templeman, et al.) to research human-computer interaction (HCI) issues in virtual environments (VEs). Our recent work, funded by ONR and NRL and presented at the *VR'99 Conference* (the premiere international VR conference), was awarded "Best Technical Paper" of the conference by popular vote of conference attendees. Our work is conducting empirical studies to produce generalizable guidelines for designing VE user interfaces. The studies are using a map-based battlefield visualization VE called "Dragon", and include comparative evaluations of parameters for such critical VE design issues as stereopsis vs. non-stereopsis, ego-centric vs. exo-centric navigation, and platform (Immersive Workbench vs. desktop vs. CAVE™). This prior and on-going work places us in an especially advantageous position to uniquely contribute to the proposed Digital Ships integrated testbed that is the "glue" for this NAVCIITI research.

We are using the Dragon software to explore use of eye tracking in VEs, producing and expanding guidelines for VE user interface design. To our knowledge, eye tracking has not been incorporated into a CAVE or an IWB elsewhere, much less systematically studied within these VEs. These issues are especially critical for (human-centric) design and evaluation of our Digital Ship concept. The objective of our two-year proposed research is to concurrently explore two different facets of eye tracking:

(1) use of eye tracking for data collection for usability analyses of VEs and
(2) use of eye tracking as a multi-modal interaction technique in VEs.

Accomplishments for this quarter.

We have developed software to integrate our ISCAN eye tracking system with almost any combination of VE display and input devices. We accomplished this by integrating previously developed eye tracking software with DIVERSE (also under development as part of NAVCIITI). The current version of the software supports low level data acquisition and logging, calibration on desktop, IWB and CAVE™ systems, as well as identification and logging of eye-based events such as fixations and saccades.

Developing eye tracking software within the DIVERSE paradigm supports our main NAVCIITI goals and tasks, both short- and long-term. Additional details are given below.

<u>Task 1: Integrate eye tracker into Dragon software on multiple hardware platforms</u>

The generic nature of DIVERSE allows the eye tracking software to be incorporated with any SGI Performer-based application (currently on IRIX and LINUX). From a low-level perspective, integration with Dragon software is mostly complete. We are now focusing on the higher-level implementation issues that will specify exactly *how* eye tracking information is used within Dragon (e.g., what type of multi-modal interactions are most appropriate).

Task 2: Explore use of eye tracking for data collection for usability analyses

DIVERSE's support for abstracted VE displays allows us to create a single piece of visualization software that supports presentation of eye tracking data and analysis on a number of display devices. Currently, an explicit line-of-sight is displayed to the VE user. A separate evaluator's view is under construction that employs DIVERSE's seamless network integration. This software will allow evaluators to view the user's scene augmented with eye tracking analysis in real-time on a separate desktop, IWB, ImmersaDesk or even CAVE environment. Moreover, the evaluator's view can be displayed anywhere a network connection is available, either locally or remote. We expect these features will be extremely valuable in the coming months as we begin to scope the evaluation of Digital Ships Testbed components (currently slated for NAVCIITI year 3).

Task 3: Explore use of eye tracking as multi-modal interaction technique in VEs

DIVERSE supports a number of networked VE input devices, and thus is well-suited to support VE multi-modal interaction. Currently, the eye tracking software is coupled with DIVERSE motion tracking, allowing exploration of gesture and eye-based interaction techniques.

In addition to our software accomlishments, our hardware, the Oynx2, has undergone refinement and configuration to support a more flexible and extensible research and development environment. This task included developing a rich set of stereoscopic display support for both the desktop display and IWB.

We have been working with ISCAN to upgrade our eye tracking at no additional cost. The original infrared filters (constructed out of plastic) will be replaced by glass filters to improve the accuracy and reliability of the eye tracking systems. The systems will be upgraded during the move to the ACITC building (see below).

Activities during this quarter.

31 May – 2 June 2000: Presentation in Washington DC, 6.2 Program Review.

22 June 2000: Presentation in Blacksburg VA, ONR NAVCIITI Review.

April 2000: Attended VR2000 Conference, the premiere international conference on virtual environments, New Brunswick NJ. Was co-leader for heavily attended workshop on Perceptual and Multi-Modal Interfaces.

Plans for next quarter.

These will include getting the IWB taken apart, moved, and set back up in the new ACITC building on-campus. Plans have been finalized to move the IWB and Onyx2 to the ACITC building. Estimated moving date is the week of 21 August.

2.31 Collaboration Work Space

Activity Report: Center for Human-Computer Interaction
Navy Collaborative Integrated Information Technology Initiative (NAVCIITI)

14 August 2000

During June and July, the Center for Human-Computer Interaction continued to extend on-going work on software support for collaborative multimedia conferencing, integrating shared notebook, whiteboard, chat, visualizations, simulations, video conferencing, etc.

We focused on Task 2.3.1.1:

(1) Develop Java classes that allow a client-side application to query an SQL database via a middleware component. Communication between client-side classes and middle tier will be implemented with sockets, RMI, or CORBA, depending on platform constraints, and will be abstracted to allow evolution to new communication layers. Communication between middle tier and database will be implemented using JDBC for SQL access.

We participated in a review on June 21. Sample collaborative database access applications based on the classes developed for Task 2.3.1.1 were demonstrated at this meeting. The review was very useful in clarifying and confirming our delivered results for Task 2.3.1, but the convoluted misunderstandings, the preparation for the meeting, and carrying out the requested rewriting of our tasks was also somewhat of a drain on our NAVCIITI time and effort. It set us back about 2 weeks.

We have also begun analysis and design work on Task 2.3.1.2:

2. Develop Java classes to support archiving and printing of ONR Collaborator whiteboard content by creating images from arbitrary JDK 1.1 (or higher) user interface components, as supported by the JDK (Java Development Kit). Deliverables consist of classes that support generation of a Java Image object containing a snapshot of an arbitrary Java user interface component and conversion of the snapshot to a GIF-encoded byte stream. Documentation describing structure and use of the classes will be included.

Preliminary work on this task is focused on examining the limitations of image capture from standard and custom Java user interface components, and on reviewing available options for flexible image encoding.

Publications related to NAVCIITI currently submitted and in review

Carroll, J.M., Rosson, M.B., Isenhour, P.L., Van Metre, C., Schafer, W.A. & Ganoe, C.H. Submitted. MOOsburg: Multi-user domain support for a community network. *Internet Research*.

Publications and presentations

Carroll, J.M. 2000. Making use: Scenarios and scenario-based design. Invited plenary talk *ACM DIS'2000 Conference: Designing Interactive Systems* (Brooklyn, New York, August 17-19).

Carroll, J.M. 2000. Introduction to the Special Issue on "Scenario-Based System Development". *Interacting with Computers.* in press

Carroll, J.M. 2000. Five reasons for scenario-based design. *Interacting with Computers*. in press. (earlier version appeared in R.H. Sprague (Ed.), *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences* (Maui, January 5-8). Los Alamitos, CA: IEEE Computer Society, 1999).

Carroll, J.M., Chin, G., Rosson, M.B. & Neale, D.C. 2000. The Development of Cooperation: Five years of participatory design in the virtual school. *ACM DIS'2000 Conference: Designing Interactive Systems* (Brooklyn, New York, August 17-19).

Carroll, J.M. & Rosson, M.B. 2000. School's Out: Supporting authentic learning in a community network. Presented at *IFIP Conference on Information Technology at Home* (Wolverhampton, United Kingdom, June 28-30). Published in *Home Informatics and Telematics: Information, Technology and Society*, Kluwer Academic Publishers.

Carroll, J.M., Rosson, M.B., Isenhour, P.L., Van Metre, C., Schafer, W.A. & Ganoe, C.H. 2000. MOOsburg: Supplementing a real community with a virtual community. In S. Furnell (Ed.), *Proceedings of the Second International Network Conference: INC 2000.* (Plymouth, United Kingdom, July 3-6). Plymouth, UK: University of Plymouth, pp. 307-316.

Chin, G. & Carroll, J.M. In press. Articulating collaboration in a learning environment. *Behaviour and Information Technology.*

Dunlap, D.R., Neale, D.C. & Carroll, J.M. Teacher collaboration in a networked community. 2000. *Educational Technology and Society, 3(3),* 442-454.

Go, K., Carroll, J.M. & Imamiya, A. 2000. Bringing User's View to Design: Roles of Scenarios in System Design, *IPSJ Magazine, 41(1),* 82-87 (Information Processing Society of Japan).

Helms, J., Neale, D.C. & Carroll, J.M. 2000. Data logging: higher-level capture and multi-level abstraction of user activities. In *Proceedings of the 40th annual meeting of the Human Factors and Ergonomics Society* (pp. to appear). Santa Monica, CA: Human Factors and Ergonomics Society.

Isenhour, P.L., Rosson, M.B. & Carroll, J.M. in press. Supporting interactive collaboration on the Web with CORK. *Interacting with Computers.*

Isenhour, P.L., Carroll, J.M., Neale, D.C., Rosson, M.B. & Dunlap, D.R. 2000. The Virtual School: An integrated collaborative environment for the classroom. *Educational Technology and Society, 3(3),* 74-86.

Neale, D.C., Dunlap, D.R., Isenhour, P.L. & Carroll, J.M. 2000. Collaborative critical incident development. In *Proceedings of the 40th annual meeting of the Human Factors and Ergonomics Society* (pp. to appear). Santa Monica, CA: Human Factors and Ergonomics Society.

Briefings and meetings

Meeting with NAVSEA, June 5
Philip Isenhour, Dennis Neale and Dan Dunlap of the Center for HCI met with John Rice, Connie Krause, and Robert McLain to discuss possible research support for computer-based engineering collaboration. We were asked to sketch ideas for a white paper on critical factors in computer-supported cooperative work.

NAVCIITI project review, June 21-22

Planning meeting at Dam Neck July 31 – August 1
Meeting participants included Philip Isenhour from the Center for HCI, Ron Kriz and Fernando das Neves from the CAVE group at Virginia Tech, Glen Wheless from Old Dominion University, Connie Krause and John Rice from the Center for Experimentation and Collaborative Exchange at Dam Neck, and several

members of the technical staff at Dam Neck. Discussions focused on establishing a consortium of groups from VT, ODU, and elsewhere (industry, academia, and other Navy groups) involved in collaboration and collaborative visualization work to do requirements analysis, proof-of-concept, and demonstration systems for the Navy's collaborative engineering effort. The goal is to involve and coordinate distributed groups from different organizations in this work in order to address some of the focus, adaptability, and interoperability issues that the current collaborative engineering efforts are facing.

2.3.0 Ship on Fire

Ship on Fire Simulation

Y2Q1 Report
Kenneth L. Reifsnider and Jason A. Burdette
Department of Engineering Science and Mechanics – Materials Response Group

Objectives

The purpose of this work is to provide a simulation of command and control of a ship when part of the ship is exposed to fire. This will require the coordination of several tasks. First, the evolution of conditions on the ship in the presence of fire will be predicted. Second, the influence of the fire on the ship's structures, personnel, and equipment will be identified. This information will then be linked with the Navy's command and control models to appropriately guide the response of the ship's systems and personnel to the fire. Ultimately, the entire process will be simulated in real-time on a virtual network.

Completed Activities

In addition to a general search for background on current and past fire research (especially relating to composite materials and Naval applications), several specific tasks have been accomplished during the past quarter.

A memorandum of agreement has been signed to facilitate the exchange of information between personnel at Virginia Tech and the Naval Surface Warfare Center – Carderock Division (NSWCCD). The first meeting between these groups was held in late July at Carderock. The main purpose of the meeting was to identify specific avenues for future collaboration and discuss ways to relate the study of materials on fire to the Navy's command and control methodologies and models. Following an introduction and general description of Carderock's activities by the primary point-of-contact Gene Camponeschi (Structures and Marine Composites Department), Rick Habayeb (NAVCIITI Program Manager) gave an overview of NAVCIITI's efforts to create a real-time digital simulation of ship operation. Harry Gray (Protection and Weapons Effects Department) then discussed some specific issues relating to ship damage control and provided points of contact to obtain more general survivability information. Jason Burdette (VA Tech - Materials Response Group) provided a layout of NAVCIITI's "ship on fire" tasking and discussed the proposed fire module testing to be performed at Virginia Tech. Usman Sorathia (Head: Fire Protection and Sea Survival Department) then discussed the Navy's current activities in fire research. A tour of fire testing facilities (small and intermediate scale) was given and efforts to model the response of composites in the presence of fire were described. Ken Reifsnider (NAVCIITI PI, Virginia Tech) then gave a description of the MRLife life prediction methodology and emphasized its intended application to material damage due to fire. Plans for an upcoming meeting (late August) at Virginia Tech with additional Carderock personnel were discussed.

To address the lack of experimental work involving materials under combined fire and mechanical loading (a critical issue in the simulation of a ship on fire), some preliminary tests have already been performed at Virginia Tech. Under funding from the Center for Adhesives and Sealant Science (CASS), material samples were made by bonding either aluminum or composite (poly phenylene sulfide w/ AS-4 carbon fibers) strips with an epoxy adhesive to form a lap joint. The joints were stressed by subjecting the structures to end-loaded compression and flames from a propane burner were applied to various points on the structure. Figure 1 shows a schematic of the test set-up.

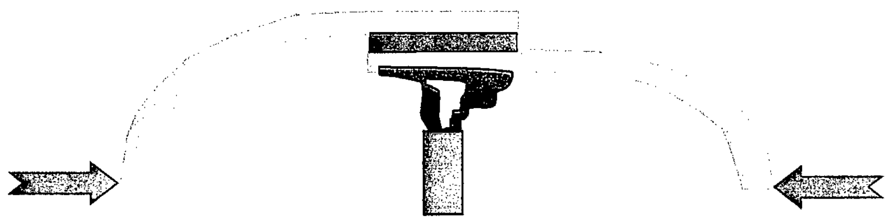

Figure 1. Set-up for fire test of mechanically loaded structures

The spread of heat, flame, and damage from the point of flame impingement were monitored and the modes of failure for the two types of structures were compared. It was found that damage to the composite was restricted to the region near flame contact and heat spread throughout the structure was inhibited by the composite's low thermal conductivity. Despite local property degradation (which was largely recovered after the flame was extinguished) the structure remained intact unless the flame was applied directly to the joint. The joints in metallic structures, on the other hand, failed regardless of where the flame was applied because heat from the flame was quickly conducted to the joint. These simple experiments demonstrate some important but inobvious differences between the behavior of metallic and composite structures on fire. Also, this was the first demonstration of a method to evaluate loaded materials while they are actually "on fire". This is a significant improvement upon existing fire experiments that focus on materials "near a fire" or "after a fire".

Planned Activities

As was stated earlier, the first step in simulating command and control of a ship on fire is the defining how the conditions on the ship evolve due to a fire. A powerful fire model called Fire Dynamics Simulator (FDS) has been obtained from the National Institute for Standards and Technology (NIST) and will be used for the first stage of the simulation. This model is capable of computing the temperature distribution, particle flow patterns, toxic gas concentrations, and smoke density within a chamber of complex geometry and can be used to represent a variety of situations encountered on a Navy ship.

This fire model can be integrated with the Navy's command and control models in various ways. Considering the effects of fire on structures aboard the ship, the model can be used to compute the environmental conditions (especially temperature) to which a loaded structure is exposed. Mechanics-based material behavior models will then be used to describe the degradation and failure of the loaded structure exposed to these conditions. Considering the effects of fire on people or equipment aboard the ship, this fire model can be used to determine the time until habitability or operability standards are exceeded. In either case, the output from these models can be integrated with the Navy's command and control models to appropriately guide the response of the ship's systems and personnel to the fire.

A special fire-testing facility will be designed and constructed at Virginia Tech to support the fire simulation effort. This facility will allow samples to be simultaneously loaded and exposed to fire. These experiments will not only add to the limited collection of fire test results, but will also be used to validate both the fire model and the material response models.

The fire models and data will be used to construct simulation models in the Digital Ships Lab. That simulation will be used as a device input to the CAVE which will be the next-generation network for collaboration, the pirmary objective of NAVCIITI.

2.32 Digital Ship / Force Multiplier

NAVCIITI Task 2.3.2 , 2.3.3, & 2.5
Digital Ships™ Laboratory

Dr. A. R. Habayeb
NAVCIITI Program Manager
Research and Graduate Studies
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061
Phone: 540-231-4353, Fax 540-231-9805, E-mail: ahabayeb@vt.edu

Introduction

This quarterly progress report describes the objectives, proposed research, current status, and planned efforts for Tasks 2.3.2, 2.3.3, and 2.5. The goal of Task 2.3.2 is to define the command and control metrics and relate them to the network centric interacting grids. Task 2.3.3 goal is to design, implement and optimize the ship information management system. The goal of Task 2.5 is to combine the ongoing research efforts in advanced technologies with modeling and simulating the components of the ship weapon system to build a common test bed at Virginia Tech. The Digital Ships Laboratory (DSL) initiative will establish a digital test bed to develop, demonstrate, evaluate, and test advanced shipboard command & control technology. The test bed will be used to integrate intra-ship, and inter-ship Information and provide a network centric focus for the NAVCIITI technology thrusts. A three-tier architecture composed of client, enterprise logic middleware applications, and database tiers will provide a flexible framework for NAVCIITI. The DSL efforts will capitalize and supplement the other NAVCIITI tasks by providing a platform for integrating the research technologies under development.

Objectives

The command and control metrics are implied and assumed ideal in the laws of combat. The formulation of the laws of combat will be examined to derive the command and control metrics. The DSL provides the mechanism to implement and optimize the ship information management system. DSL is the test bed to integrate command and control technologies. The integration initiative requires sufficient hardware and software for implementation of a flexible, reconfigurable three-tier architecture and the underlying network. The Digital Ships™ infrastructure will be designed with sufficient capacity and flexibility to provide the capability for integration and evaluation of the operational utility of NAVCIITI technologies in a system context. Because simulating the engagement of a ship weapon system with the threat involves huge blocks of data and many nodes with many components; it will be difficult to evaluate the system performance under a variety of conditions. This problem will be overcome by integrating the DSL with the CAVE.

The simulation will be projected in stereo onto three walls, the floor, and the ceiling of the CAVE and viewed with stereo glasses.

Efforts for the Quarter and Accomplishments

The preliminary investigation of the Command & Control metrics indicates that the C & C attributes form the basis of the Force Multiplier in combat. C&C functionalities provide the necessary and timely information from the sensors to the shooters. Relevant collaboration within the network centric environment- system of systems- is a fundamental requirement to achieve combat effectiveness.

The DSL concept is new to the NAVCIITI program and the research team needed guidance to focus their technologies toward the unifying theme Digital Ships. During this quarter several meetings and presentations on the Digital Ships concept were made to the NAVCIITI team. The three–tier architecture as shown in Figure 1 was chosen for the test bed to provide flexibility, scalability, connectivity and reliability. An investigation of the organization of a typical ship was performed. The investigation established that ships are organized into five basic departments. The ship organization was mapped over the Three-tier architecture of DSL, as shown in Figure 2. Equipment was ordered to build the DSL.

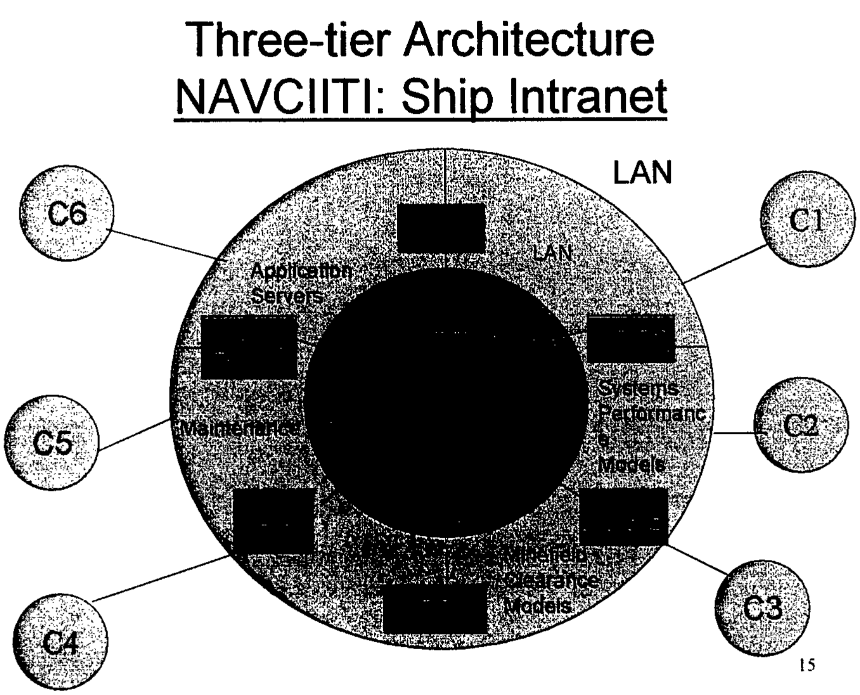

Figure 1. Three-tier architecture

A set of metrics were defined to allow evaluation of the technologies incorporated into the DSL system. The DSL will have robust simulation capability for modeling of intra and inter ship information system behavior. The simulation capability will be sufficient to generate real-world network traffic that models actual shipboard behaviors and functions. The DSL will have a set of visualization tools to allow for evaluation of the integrated technologies. These visualization tools will be developed to support shipboard Human Computer Interaction investigations for the individual shipboard action stations and Command and Control visualization systems. The required behavior of shipboard systems will be investigated and this knowledge will be used to guide the integration of the NAVCIITI technologies.

The integration of the DSL with the CAVE created the need for innovative software tools like the Virginia Tech. DIVERSE (Device Independent Virtual Environments Reconfigurable Scalable Extensible) toolkit, DTK, http://www.diverse.vt.edu/dtk/ , and its sister project the DIVERSE graphics interface, http://www.diverse.vt.edu/dgipf/. This new software has novel constructs that will aide in the integration of any distributed CAVE based simulation. DIVERSE applications can be seamlessly ported to desk top work stations and other platforms without modification.

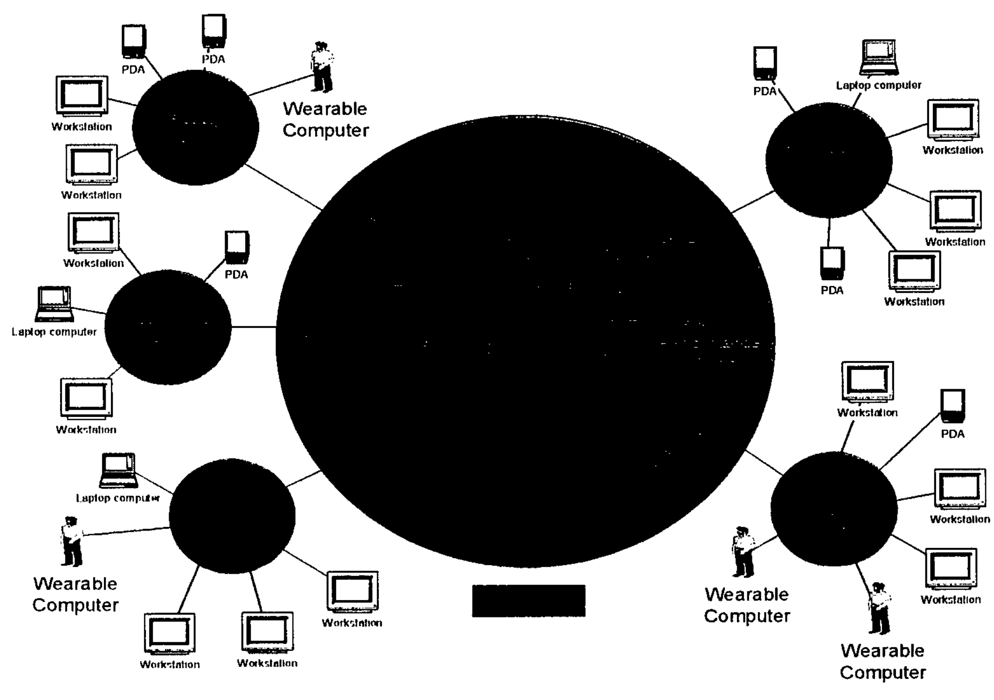

Figure 2. Three-tier shipboard organization.

Accomplishments

The required hardware and software for development of a baseline three-tier network has been defined and purchased. Currently a Digital Ship Laboratory has been established at Norris Hall. This system consists of a database server, a middleware server, and five client workstations. The complete baseline system, when delivered will consist of at least seven client workstations three middleware servers, and two database servers. Three of the client workstations are of sufficient capacity to be also used as middleware or database servers as required. The VA Tech network system will be used for initial configuration and testing of the hardware and software system. A DSL network independent of the VA Tech system will be designed. The baseline software will include both commercial and open/public-domain applications, application program interfaces, and software compilers will be used for development in the DSL. MS-SQL Server 7.0, MySQL (open-software relational database system) have been selected to serve as the initial database systems. It is expected that other relational and object oriented databases systems will be used as the research test bed evolves.

A Common Object Request Broker (CORBA) application ORBacus developed by OOC, Inc., which meets current Object Management Group CORBA standards, has been selected for development of system components. This application is free for non-commercial use. JAVA, C++, and the object and data management services provided by these software language systems will also be used to develop required client, middleware, and data management applications.

The DSL staff developed a simple three-tier system in the JAVA language. The system is used to demonstrate three-tier functionality and to aid in the understanding of the network communications between the layers of the three-tier architecture. The set of applications consists of a Helm Client, a Middlware Server, and a Database component. Changes in the course and speed of the ship are used as inputs to the Helm Client, the Middleware Server component echoes receipt of the changes and sends it to the database component for recording. This basic system will also be implemented using C++ and CORBA to increase the understanding of how these tools are behaving in three-tier implementation. These basic applications, as they mature, will be used for testing the system. The codes developed will be modified as necessary and combined with other commercial and non-commercial applications for implementation monitoring and data collection systems.

Discussions between the DSL staff and the other NAVCIITI team members have had very positive results. The researchers of Task 4.1 provided the DSL with a Combat Simulator that uses distributed processing. Currently, the simulator is used to evaluate system scheduling and resource management. This interaction has helped the team gain a better understanding of three-tier architecture and shipboard system requirements. Also, discussions between the researchers of Task 2.31 (Collaboration Work Space), and the DSL staff has resulted in a growing mutual understanding of the others accomplishments and goals. The Collaboration Work Space has a persistent data object software component, that, if combined with Task 4.1 distributed resource management research may prove valuable in development of survivable three-tier shipboard components.

The researcher of Task 3.1, (Network Protocol and Interoperability) had worked with Newport News Shipyard and modeled and evaluated the CVN-76 backbone network. He has provided valuable insight into present and emerging shipboard network designs and has offered to help in designing a DSL network to support the integration efforts. The DSL effort will continue to be the catalyst for communication between the NAVCIITI participants. These preliminary activities have been positive and show promise in being beneficial to the total effort.

The metrics used to evaluate the DSL integration effort is important to the success of the project. Initial high-level requirements have been defined and are summarized below. These metrics, their associated variables, and others recognized as the project matures will provide a means to evaluate the research efforts and provide an effective viewpoint into the problem.

Metrics	Variables
Affordability	Cost, effectiveness, requirements
Connectivity	Timeliness, robustness, channel capacity, merging of functions, error rate, quality of service (QoS), security
Interoperability	Diversity, collaboration, channel matching, data rate, security, QoS
Scalability	Growth in number, band-width, security, QoS
Flexibility	Reconfigurability, upgrades
Availability	Readiness, supportability, maintenance
Capability	Performance, data rate, security, multi-media, merging of functions
Reliability	Hardware and software, failure rate, errors

Importance of objectives

The realization of the goals of the digital ship effort will result in a cost-effective approach for gaining insights into the many vexing operational challenges facing the fleet. The DSL will provide a platform that will allow advance command and control technologies to be developed, demonstrated, and tested. The DSL will serve as a halfway house between NAVCIITI command and control technologies and the fleet. It will serve as a comprehensive testing tool for technology such as, models of antennas, wireless reconfigurable radios, interoperability, and software. The mature system will provide the infrastructure to integrate and evaluate the operational utility of technology in a system context. Also realized will be the ability to perform experiments/runs at a very low cost compared to building physical hardware. NAVCIITI researchers and sponsors will gain further insight into the value of new and brave ideas or concepts by providing a means to examine the interfaces to other technologies and the value of individual research efforts and how they fit with and can be enhanced by other efforts. Also, the information system and components developed and tested in the DSL can provide a training platform for Navy personnel.

Activities

The DSL was visited by representatives from Newport News Shipbuilding and NSWC Dam Neck. On June 22, 00, the NAVCIITI strategy and vision was presented at the kickoff briefing.

Plans for Next Quarter.

During the next quarter the DSL team will further develop the projects described above. The baseline equipment will be installed, configured, and tested. We will move into the new space in Torgersen Hall. We will also identify initial requirements for an independent, flexible, reconfigurable network and purchase the necessary hardware. The characteristics of the network will be defined by the NAVCIITI integration requirements. Emphasis will be placed on gaining a deeper understanding of shipboard organization so that these behaviors can be mapped into the three-tier architecture. The DSL team will begin documentation and visualization of the DSL three-tier architecture and the evaluation data and metrics. Integration of the other NAVCIITI technologies will be a major milestone.

Issues

none

2.4 Mechanically Flexible displays

Fiber & Electro-Optics Research Center
College of Engineering
Richard O. Claus

Our tasks as part of the NAVCIITI program have involved two primary tasks. One task has been the research of electro-optic thin film display devices that may be used as visualization tools. The second is the development of new methods for high speed reliable data transfer in optical fiber-based local area networks.

Prior reports have focused largely on the electro-optic and display results of our part of the NAVCIITI project. This report instead discusses a new approach that we have considered for some time but have only now had enough modeled results to make its presentation logical. It concerns possible ways in which near real-time processing may be achieved using physical layer devices in high speed optical fiber-based communication networks.

1. Results: Modeling Multilayered and Multipath Thin Film Phase Control Materials for Near Real-Time Information Coding and Data Compression

Introduction

In recent years the world has witnessed exploding demands on communication networks for both military and commercial use. In addition to more conventional communication devices such as the telephone and video, the military and commercial development of the Internet has provided for interconnectivity of personal computers and related devices. This interconnectivity enabled a transition of the personal computer from computing device into a new kind of personal communication device, ultimately creating an enormous worldwide network of these PCs. The booming trend of the Internet has overwhelmed the existing communication infrastructure and has placed demands non-existent just a decade ago, and similar expansion and demands on military data transmission networks has placed identical demands on military systems. To meet some of these demands, there has been a global effort to merge the segregated networks into one integrated single network capable of merging long-standing foes -- data, voice and video. This effort has in turn resulted in the birth of numerous new services such as e-commerce and video-conferencing, which only further expand the requirements placed on the communication infrastructure. The infrastructure will have to cope with the increase in the sheer amount of data traversing the infrastructure as well as with the increase in the importance of information carried by the network. The ever-increasing amount of information calls for faster and more efficient ways to compress data before it reaches the Internet. On the other hand the ever-increasing importance of information calls for much more reliable transmission of data after it reaches the Internet.

Unfortunately, both of these functions, compression and coding, require a significant amount of additional processing power, currently performed by bandwidth-limited silicon-based devices. Even with the processing power increasing tremendously in recent years, computing speeds achieved by electronic devices are generally still the main bottleneck in the end-to-end information transfer and processing. It is thus our goal to alleviate these bottlenecks by either distributing some of the functions away from the processors and into the optical physical layer, or by completely replacing the electronics. Also, it is our belief that we can achieve this goal in a cost effective manner.

<u>Technical Approach as Part of NAVCIITI Program</u>

Unlike the current implementations of compression and coding, which are done in software and silicon-based hardware, we propose and have been analyzing a new technology that offers a passive optical implementation of these functions with minimal, if any, number of active electro-optical components. By migrating some of the functionality into the physical layer of the Open System Integration (OSI) Reference Model, our goal is to tremendously increase the speed and reduce the cost and complexity of implementation of the compression and coding system.

We have considered how to optically implement the basic building blocks of digital filter systems, first and second order filter units, and methods of implementing cascade and parallel combinations of those units to realize systems with higher complexity. There are two different approaches to implementing the proposed optical signal processing: 1) multilayered optical materials model approach and 2) integrated optical waveguide model approach. Both of these are explained below.

Multilayered Optical Materials (MOM) Model

The multilayered optical materials model represents one way of implementing the basic building block of the optical signal processor. It concerns an approach for the near real-time signal processing of optical signals by controlling the way they form multiple reflections through their propagation inside multilayered and/or integrated optical structures.

Multilayered optical materials may be formed in which the thickness and index of refraction of each individual layer are controlled by the materials incorporated into each layer, along with the processing conditions used to synthesize each layer. Optical signals incident on such multilayered materials are partially reflected and partially transmitted at each layer-to-layer interface.

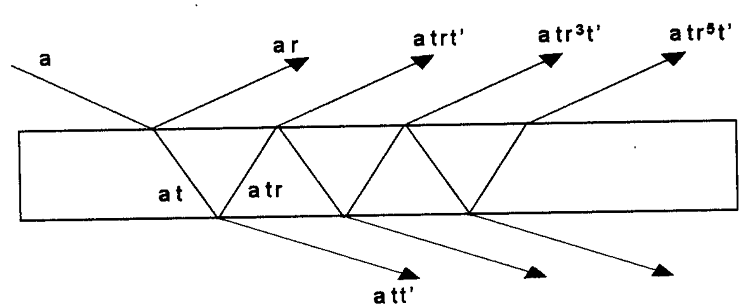

Figure 1. Basic Building Block of an MOM Model

The phases and amplitudes of each of the reflected and transmitted signals may be determined by the thicknesses and indices of the layers. Both reflected and transmitted signals are created at specific angles with respect to the plane of the substrate. By detecting the multiple reflected and/or transmitted signals at those different angles, different combinations of amplitudes and phase delays of the input signal may be obtained. The time required to obtain this set of output signals is equal to the propagation time required for the optical signals to pass through the thin film materials, or, typically, well less than several femtoseconds. The signal processing thus is effectively performed in near real-time. The basic building block of the multilayered model is a single layer of an optical material. A layer showing multiple output signals is presented in Figure 1. A functional analog of the MOM model basic building block is the all-pole IIR lattice filter, which is shown in Figure 2.

Given many layers having different indices of refraction, the output signals can be formed as sums of previous input and output signals with different and controlled phase and amplitude characteristics. The output reflection components are sums of multiple phase-delayed and amplitude-scaled signals, as shown in Figure 3. These multi-term components may be directly related to IIR and FIR signal processing components that also are time-delayed and amplitude-scaled forms of an input or output signal.

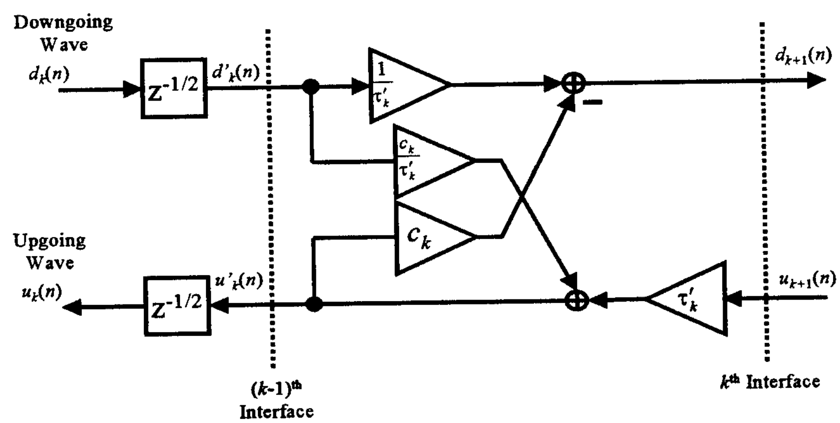

Figure 2. Basic Building Block of an All-Pole IIR Lattice Model for Layer k

Models providing the values for layer thickness, angle of reflection and refraction in each layer, and reflection coefficients at each layer to provide a desired output response are already well understood. These models were initially developed for modeling of the earth's layers for exploration seismology, where the method of reflection seismology, combined with intensive signal processing, is capable of providing a two- or three-dimensional representation of earth's subsurface down to about 20,000 to 30,000 ft. In this layered earth modeling approach, the known values are the input, which is generally a source of seismic energy, and the output, which is the collected reflections of the seismic waves from the interfaces between earth's geological layers. Then, based on the information these reflected waves carry, a representation of the earth's layers is created with high enough accuracy and resolution. The resulting outcome is that unknown garbled output is interpreted to deduce the earth's structure. The optical analog is "Rugate" filters that are based on similar multilayer structures. Also, in acoustics, there are similar analogies.

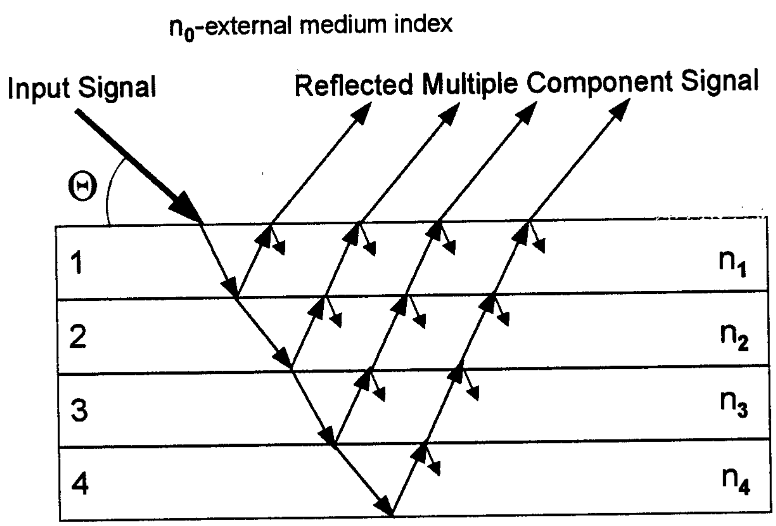

Figure 3. MOM System Implementation

The analog system implementation in DSP is shown in Figure 4.

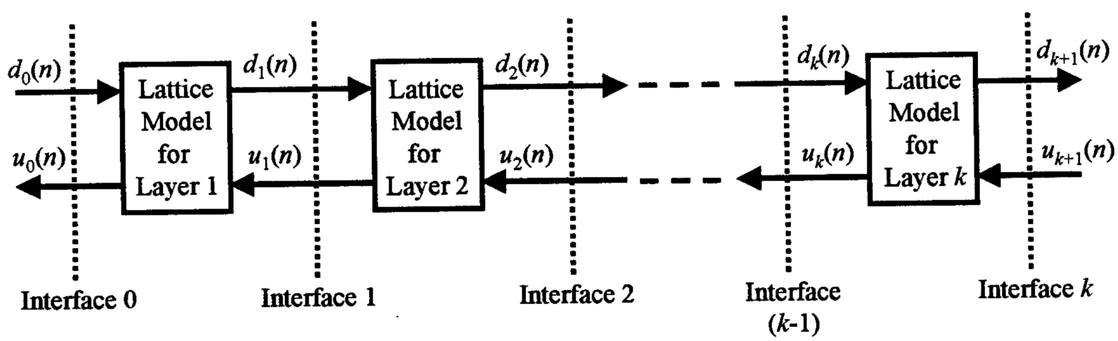

Figure 4. Lattice System Model Representation

In the approach of an optical layering model for signal processing, the reverse is applied. Based on the desired and known transfer function of our system, i.e. a particular IIR response, it is possible to derive the important parameters for each optical layer needed to produce the desired output for a given input. Given the ability of thin film forming methods to create thicknesses and indices of refraction over a wide range of values, the values required to create the necessary sequence of propagation parameters are possible. In particular, we have demonstrated electrostatic self-assembly (ESA) and sol-gel methods for forming multilayered materials having a large range of indices of refraction and controlled thickness. Such methods, and others, may be used to form the films required for the implementation of real-time signal processing transmission and reflection materials.

Integrated Optical Waveguide (IOW) Model

We have also considered another approach for implementing near real-time signal processing of optical signals involving integrated optical waveguides, where ESA or sol-gel techniques in combination with conventional optical photolithography may be employed for optical thin film assembly can be used for precise control of lengths and angles of optical waveguides, which in turn define propagation delays and attenuation or scaling of the signal. One possible basic building block configuration is presented in Figure 5.

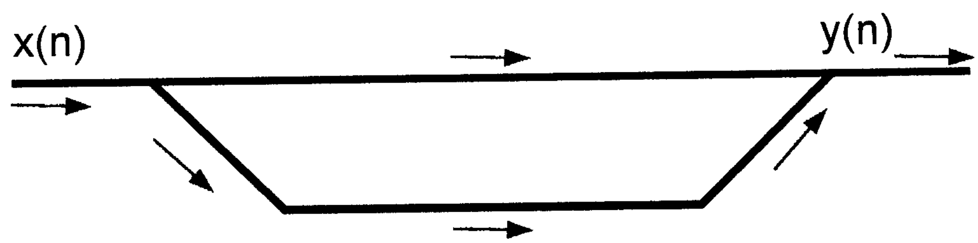

Figure 5. Type I Basic Building Block for an IOW Model

The output y can be written in the following form:

$$y(n) = b_0 x(n) + b_1 x(n-1) \qquad \text{Equation (1)}$$

where x(n-1) is a delayed component of the original signal and b_0 and b_1 are scaling factors. The delay, created by a part of the input signal taking a longer route, can be tuned very finely by increasing or decreasing the overall difference in the length between the two paths the light takes. The scaling factor can be controlled in many ways, one being the angle the branched path contains along its path, where sharper angles cause larger scatter and loss of light, and thus would represent a smaller scaling coefficient. Ultimately, there are many permutations between the propagation properties of the optical materials used for manufacturing of the waveguides and numerous geometries of the waveguides, which can provide the desired delay and scaling parameters, and thus the desired output characteristics. Changes in the length of the secondary light path relative to the first allows both constructive and destructive interference at the recombination point, and thus allows the realization of both positive and negative coefficients.

It is easy to conceive that by adding multiple branches in the waveguide, many different delays and scales of the original input can be achieved, allowing for tailoring the output according to a wide variety of needs. For example, a standard Moving Average (MA) FIR filter could be implemented by the integrated waveguide model as presented in Figure 6.

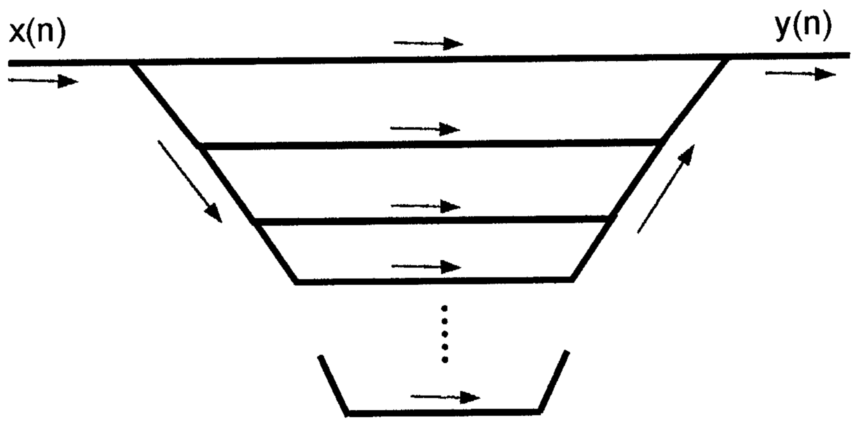

Figure 6. Direct Implementation of MA(M) Filter Equation- IOW Implementation

The output y of the general system can be written in the following form:

$$y(n) = \sum_{m=0}^{M} b_m x(n-m) \qquad \text{Equation (2)}$$

To be able to fully emulate the functions performed by digital signal processors, the IOW model needs to be able to emulate the basic building blocks of standard filter models, such as the first and second order filter sections. Doing so requires the introduction of another type of an IOW building block, as shown in Figure 7.

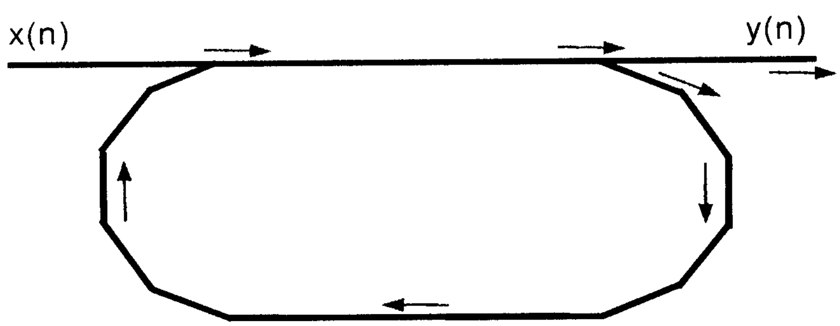

Figure 7. Type II Basic Building Block for an IOW Model

The output y can be written in the following form:

$$y(n) = a_1 y(n-1) \qquad \text{Equation (3)}$$

where the output is a delayed and scaled version of itself, and a_1 is a scaling factor.

2. Problems

Our support through the NAVCIITI program was mimimized during the transition between initial year and Year 2 funding, and, as a result, our work during this time was minimal and confined primarily to the analytical modeling studies above.

3. Plans for Next Quarter

Here at the beginning of the Year 2 effort, our plans are to further expand the analytical modeling of both the multilayered and segmented branch optical waveguide geometries to obtain near real-time signal processing of optical signals. We will also participate in the early occupation of the ACITC building on campus and design the local area network testbed to be installed for analysis in that facility during the program year.

4. Note

A patent disclosure concerning this work was submitted to Virginia Tech Intellectual Properties Inc. and has been approved for submission as a provisional patent.

3.1 Network Protocol Interoperability

NAVCIITI Task 3.1: Network Protocol Interoperability
Year 2, Quarter 1 Status Report

Scott F. Midkiff
August 4, 2000

1. Background

Task 3.1, Network Protocol Interoperability, focuses on enabling interoperability between heterogeneous networks that may belong to and be managed by different organizations, including allies and coalition partners. The proposed research addresses four specific technical areas related to the network infrastructure needed for the Navy's Virtual Operations Network (VON) initiative: (i) wireless-wireline interoperability, including mobility; (ii) quality of service (QoS) across heterogeneous networks; (iii) network management in shared heterogeneous networks; and (iv) security, especially in wireless systems. We are researching these problem areas, are investigating and will propose solutions based on commercial equipment and standards, and will demonstrate and evaluate solutions with respect to both functionality and performance using a distributed network interoperability test bed. We emphasize use of the TCP/IP protocol suite and network infrastructure support for delay sensitive applications, bandwidth-on-demand, and security. During Year 2 (the first year of this particular task), the focus is on deploying the test bed and developing and evaluating specific interoperability solutions.

2. Planned Effort

During the quarter, we had planned to continue small-scale experiments with protocols for mobility (Mobile IP and TORA routing for mobile ad hoc networks), to begin investigation of routing in mobile networks, and to begin small-scale experiments with security (IPSec) and network management (SNMPv1, SNMPv2c, and SNMPv3). The work proceeded as planned.

3. Accomplishments

In research addressing both network management and security issues, we completed a preliminary qualitative investigation of secure Simple Network Management Protocol (SNMP) deployment strategies in networks with low data rate links and we completed a small-scale experiment into the performance of various configurations of SNMPv2c and SNMPv3 with and without IP Security (IPSec). The experimental test bed consisted of two subnets connected by a one-link "backbone" network. Gateways at each subnet used Linux FreeS/WAN (version 1.4), a freeware implementation of IPSec. The qualitative investigation and quantitative experimental investigation will be documented in two technical reports which are now in near final draft form (see Section 6 below).

In the area of mobility, we have begun investigation of routing approaches for mobile networks. In this context, a mobile network refers to an entire subnet where nodes within the network are static with respect to each other, but the entire subnet is dynamic with respect to other subnets. This model is applicable to a multi-ship network where the ship-to-ship link topology is dynamic. Specifically, we are exploring the suitability of routing in mobile network using Mobile IP versus a tuned version of the industry-standard Open Shortest Path First (OSPF) routing protocol. Work also continued at a low level of effort on experiments using Mobile IP and the University of Maryland's Temporally-Ordered Routing Algorithm (TORA) mobile ad hoc network (MANET) routing protocol.

In the area of quality of service, we continue to monitor ongoing activity in the Integrated Services (IntServ) and Differentiated Services (DiffServ) working groups of the Internet Engineering Task Force.

We have also begun to plan test bed experiments to examine the use of the Reservation Protocol (RSVP) for bandwidth reservation and the integration of DiffServ and RSVP as a potential compromise between scalability and per-flow quality guarantees.

In the area of security, we have continued an ongoing study of general security vulnerabilities to develop a unified taxonomy. And, as noted above, we have deployed IPSec in a small-scale test network and explored security schemes for SNMP.

We continue to plan for the interoperability test bed and have ordered five low-end Linux personal computers that will be used to expand our current test bed prior to our move to the Advanced Communications and Information Technology Center (ACITC) in late August/early September.

4. Importance

The accomplishments above all relate to the general problems encountered in a Virtual Operations Network. Network management will be critical to the rapid deployment and operational maintenance of shared network resources. This shared network will likely utilize at least some low data rate links and security, especially authentication and, perhaps, privacy, will also be important. The work reported is providing a better understanding of alternative architectures to meet these goals and of the relative benefits and costs of each approach. Routing over point-to-point links instead of using satellite links, bandwidth allocation, and security in wireless networks are also critical to the VON.

5. Activities

We attended the following meetings directly associated with this task during the past quarter.

- Scott Midkiff and Erik Hia (graduate research assistant) met with LCDR David Jakubek of ONR, David Marlow of NSWC-DD, and other Navy personnel at the Virginia Tech's Alexandria Research Institute (Alexandria, VA) on May 26, 2000 to discuss the status of the Navy's Virtual Operations Network project and our NAVCIITI task. (This meeting was also cited in the June 2000 "bridge funding" status report.)
- Scott Midkiff and Nat Davis attended a meeting between personnel from TECSEC (Vienna, VA) and Mark Jones and Peter Athanas of Virginia Tech on June 20, 2000 to learn about TECSEC's Constructive Key Management (CKM) scheme. CKM may have applicability for key management to support protocols for privacy and authentication that we are investigating.
- Scott Midkiff briefed Navy personnel and others on the status of this task at a Year Two kick-off meeting at Virginia Tech (Blacksburg, VA) on June 22, 2000.
- Scott Midkiff and Luiz DaSilva attended a briefing by Howie Marsh of ONR on the Extending the Littoral Battlespace (ELB) demonstration project. Experiences with IEEE 802.11 and other wireless networking schemes from that program were discussed.

Other related activities during the quarter were as follows.

- Luiz DaSilva attended the International Communications Conference (New Orleans, LA) on June 18-22, 2000 to present the paper "QoS Mapping Along the Protocol Stack: Discussion and Preliminary Results."
- As part of a different, but related project, Scott Midkiff attended the DARPA Global Mobile Information Systems (GloMo) principal investigator's meeting (Eatontown, NJ) on July 11-13, 2000. He gave a presentation on hardware/software co-simulation for high-fidelity simulation of wireless networks.

6. Plans for the Next Quarter

By the end of August 2000, we plan to provide two technical white papers and one planning document to our points of contact:

1. Qualitative comparison of security schemes for SNMP considering use of low data rate wireless links,
2. Quantitative results of small-scale experiments with security schemes for SNMP, and
3. Demonstration plans for Year Two (2000-2001).

Other activities planned for the next quarter are as follows.

- Occupy space in Virginia Tech's new Advanced Communications and Information Technology Center beginning August 31, 2000 and begin to establish experimental network test bed in the ACITC. Begin to deploy equipment at the ARI and connect laboratories at the ARI and in the ACITC, initially via IP tunneling.
- Expand testing of IPSec, Mobile IP, MANET protocols as the test bed expands.
- Deploy QoS mechanisms (DiffServ and RSVP) in the test network and develop experiments.
- Continue investigation of SNMP and begin to research methods for managing resource utilization for network management tasks.
- Continue investigation of routing schemes for mobile networks and develop a plan for experiments.

7. Issues

There are no known issues at this time.

3.2 Network system Interoperability

**COMMUNICATIONS MODELING
IN SUPPORT OF THE
VIRTUAL OPERATIONS NETWORK
(NAVCIITI Task 3.2 Quarterly Report)**

Personnel

Richard E. Nance, Systems Research Center and Department of Computer Science
James D. Arthur, Department of Computer Science

Background

The Virtual Operations Network (VON) is a communications network among coalition members with varying and quite different organic bandwidth and processing capabilities. The network is to support common operating pictures to the extent needed among coalition participants. The major challenge is to understand how participants with differing capabilities can work together to accomplish a given mission overcoming the vast disparity that can exist among the members' communications capability.

Planned Effort for the Quarter

The objective for the given quarter was to obtain a definition of the VT project responsibilities in modeling the VON architecture so that the scope, objectives, requirement sources and study expectations could be completed.

Accomplishments

Please see the attached meeting reports for the description of accomplishments and work completed during the quarter.

Activities

Two meetings were held for discussions between VT personnel concerned with NAVCIITI Task 3 and Navy representatives. The first was held at the VT Alexandria Research Institute, involving all of Task 3, and the second, at SSC SD, involving only Task 3.2. The two meeting minutes attached give details of the meetings.

Meeting: Alexandria Research Institute

Alexandria, Virginia
26 April 2000

Attendees:
Dave Marlow
LCDR Dave Jacubek
Roger Merk (SPAWAR)
Karen O'Donoghue
Jim Oblinger (NUWC)
Tom Conrad (NUWC)
Scott Midkiff (VT)
Sean Arthur (VT)

The purpose to the Alexandria meeting was to discuss the progress and current status of the VON effort. LCDR Jacubek started the meeting by providing and overview of the current status of VON. His presentation focused on (a) existing components, (b) desired components and capabilities, and (c) related work. His presentation provided the basis on which later presentations and responses were formulated.

In listening to LCDR Jacubek's presentation and expectations in terms of PI responses, we realized that the presentation we prepared was lacking a few desired components. In particular, the presentation lacked detailed information as to what we needed from the VON sponsors to perform our tasks. We modified the presentation to include the expected components (these were the handwritten slides)

The presentation describing our 3.2 task focused on several items. More specifically it

(1) outlined the view of the total VON effort and then identified our task within that effort,

(2) stated our assumption regarding the VON architecture, and in particular, those assumptions related to our task,

(3) provided an overview of the steps we needed to take to achieve our task objectives,

(4) outlined, from a very high-level perspective, the classes of data we needed to support our research effort and introduced why we needed that data,

(5) outlined both short- and long-term goals, and the anticipated steps to achieving those goals.

During the final question and answer session, we reiterated the task objectives, described our approach to achieving those objectives and discussed potential information sources.

Plans

During the next quarter the following is planned:

(2) Data sources and access to them are to be determined so that the Static Model can be defined.

(3) Formatting and record definition issues are to be resolved prior to the definition of a prototype database.

(4) A prototype database for the Structured Model is to be defined.

(5) Definition of the Structured Model is to be done, leading to an initial design.

Issues

None at the moment, but access could pose a problem.

4.1 Scheduling and Resource Management

Adaptive Resource Management of Asynchronous, Decentralized Real-Time Systems

Binoy Ravindran
The Bradley Department of Electrical and Computer Engineering
Virginia Polytechnic Institute and State University, Blacksburg, VA 24061
Phone: 540-231-3777, Fax: 540-231-3362, E-mail: binoy@vt.edu
WWW: www.ecpe.vt.edu/faulty/ravindran.html

1. Introduction

This first quarterly (Q1) report of NAVCIITI Year Two describes the objectives, proposed research, current status, related activities, and planned efforts of Task 4.0 "Real-Time Resource Management" that was performed during the period May 1, 2000 through July 31, 2000. The organization of the report is as follows: We describe the background and objectives of our proposed research in Section 2. A description of the proposed research and our planned efforts is given in Section 3. We report the current status of the work and accomplishments in Section 4. Section 5 discusses the publication, presentation, and other related activities during this quarter. Finally, we conclude the report by describing the planned efforts for the next quarter in Section 6.

2. Background and Objectives of Proposed Research

Real-time computer systems that are emerging for the purpose of strategic mission management such as coordination of multiple entities that are manufacturing a vehicle, repairing a damaged reactor, or conducting combat are subject to great uncertainties at the mission and system levels. The computations in the system are "asynchronous" in the sense that processing and communication latencies do not have known upper bounds and event arrivals have non-deterministic distributions. Computational parameters such as task execution times, communication delays, periods, and event arrivals are non-deterministic as they are dependent upon the conditions of the external environment where the system is deployed and the operational scenario of the system that is situation-specific.

Such real-time mission management applications require decentralization because of the physical distribution of application resources and for achieving survivability in the sense of continued availability of application functionality that is situation-specific. Because of their physical dispersal, most real-time distributed computing systems are "loosely" coupled using communication paradigms that employ links, buses, rings, etc., resulting in additional uncertainties e.g., variable communication latencies, regardless of the bandwidth. The characteristics of the application combined with the physical laws involved in distribution thus contribute to the non-deterministic and asynchronous behavior of the system [Jen92, Jen99].

New advances in real-time distributed systems research [Quo97] have produced quality-of-service (QoS) technologies that allow asynchronous, decentralized real-time systems to specify and negotiate service expectations such as timeliness and survivability, which was not previously possible under classical real-time theory. However, the QoS management literature or the classical real-time literature does not address three fundamental problems in asynchronous, decentralized real-time systems.

First, the literature does not provide solutions to the problem of achieving the real-time and survivability requirements of *aperiodic* computations in asynchronous, decentralized real-time systems. Classical real-time theory provide scheduling and resource management solutions for aperiodic computations primarily in static, centralized systems where the computational parameters such as execution times, communication delays, and (distributions on) event arrivals are known with absolute certainty. On the other hand, the QoS management technologies focus primarily on periodic computations, where adaptation of the periodics to the "right level" of their desired QoS is dynamically performed to accommodate run-time uncertainties and workload changes. It is very difficult, if not impossible to apply the QoS techniques for aperiodics. This is because it is impractical to model the real-time and survivability requirements of aperiodic computations as multiple levels of service that vary in their quality of computational results and requirements of resources. Furthermore, allocation of resources (by adaptation or negotiation) for aperiodic computations *at event arrival* is practically infeasible due to their high cost.

Second, the classical real-time communication literature focuses primarily on static systems, where the communication parameters of the application such as message latencies, message periods, and message arrivals are deterministically known. Further, the QoS literature is sparse with works on dynamic adaptation mechanisms at the trans-node *message-level*. Rather, the QoS technologies primarily focus on process-level adaptation mechanisms such as process replication and imprecise computations. We believe (and demonstrate in Section 4.2) that dynamic adaptation mechanisms at the message-level can produce significant benefits. Furthermore, adaptive communication technologies assume great practical significance with the emerging communication standards such as IEEE 802.1p [IEEE802] that supports message prioritization on Ethernet networks that are cheap and widely available.

Third, the performance of the real-time QoS management algorithms that appear in the literature is only *empirically* studied for the most part. The performance of the QoS mechanisms is sometimes studied through simulation or they are implemented as part of system software layers (e.g., middleware) and performance characterization is done using real-time benchmarks such as [MV+96, SWR99]. This certainly validates the effectiveness of the techniques. However, the *general feasibility conditions* of the QoS techniques for satisfying the real-time and survivability requirements are unknown. The fundamental limitations of the QoS management techniques and the boundaries of satisfiability of the requirements are not known. This is a difficult, but fundamental problem since there does not exist a theory for designing QoS management algorithms and for reasoning, evaluating, and validating their behavior under all possible operating conditions of the system.

We propose fundamental technology that solves each of the three problems outlined above. We propose *proactive resource management strategies* that achieve the real-time and survivability requirements of aperiodic computations. The proactive techniques will forecast the resource needs of the aperiodic computations using *regression theory*, given the current internal resource situation and the desired external load. The forecasted resource needs will drive resource allocation strategies that achieve the desired real-time and survivability requirements. The resource allocation strategies, if feasible, will determine optimal resource allocations that satisfy the collective requirements of the application—real-time and survivability requirements of periodics and aperiodics—through adaptation mechanisms such as process replication. If optimal resource allocations that achieve the collective application requirements are not feasible, the algorithms will perform negotiation of requirements across computations for optimal satisfaction of requirements.

Further, we propose *adaptive communication mechanisms* that achieve the application real-time requirements for IEEE 802.5 and 802.1p networks through adaptation mechanisms such as *dynamic message reprioritization* and *run-time alteration of token holding times at processor stations*. We propose to develop distributed algorithms that identify (at run-time) inter-process, trans-node messages of the application that are contributing to low timeliness of end-to-end computations due to changes in application workloads. The algorithms will adapt the "ailing" trans-node messages to workload changes through mechanisms such as message reprioritization for IEEE 802.5 and IEEE 802.1p network standards. The algorithms will determine current optimal dynamic priorities for the messages that will alter their communication bandwidths such that the real-time requirements of end-to-end computations are achieved. Further, we propose to develop adaptive communication algorithms for the Time Token Protocol where the token holding times of processor stations will be adapted to workload changes. The algorithms will determine the current optimal dynamic token holding times for processor nodes that have application processes generating "ailing" trans-node messages. The algorithms will alter the token holding times of the processor nodes to change the communication bandwidths of the nodes and the messages such that the end-to-end real-time requirements are achieved.

Finally, we propose *hybrid feedback control theory* that will enable the mathematical design and evaluation of adaptive resource management mechanisms. We propose to develop the theory by modeling asynchronous, decentralized real-time systems in a feedback control-theoretic setting using *non-linear hybrid models*.[1] The hybrid models will contain elements of different modeling paradigms such as difference equations and

[1] The objectives of feedback control theory are remarkably similar to adaptive resource management techniques – achieve the desired behavior of a physical system that operates under uncertain conditions through dynamic monitoring, feedback, and adaptation.

finite state machines. Thus, the models will be inherently non-linear. Using the hybrid models, we propose to formulate adaptive resource management problems that require satisfaction of the real-time and survivability requirements of the system as optimal control problems. The control problems will be analytically solved using game-theoretic approaches to counter the disturbances and run-time uncertainties of the system. The proposed theory will thus allow the mathematical design of adaptive resource management strategies and will enable their behavior to be studied analytically. This will provide guarantees of optimality or desired transient and steady state performance even in the presence of disturbances and uncertainties.

3. Description of Proposed Research and Planned Efforts

In this section, we describe each of the research problems outlined in Section 2 in greater depth; we discuss the state-of-the-art to illustrate the significance of the problems and our proposed solution approaches. Section 3.1 discusses proactive resource management of aperiodic computations. Adaptive communication mechanisms are described in Section 3.2. Section 3.3 discusses hybrid feedback control theory-based adaptive resource management.

3.1 Proactive Resource **Management of Aperiodic Computations**

Traditionally, scheduling and resource management of aperiodic computations is performed through reservation of resources and on-line admission control schemes [CSR86, DTB93, KS92, LRT92, LSS87, Mok83, NS94, RCF97, RSZ89, RTL93, TLS96, Spru90, Sta89, SB96, SLS88, SR91, SRC85, SSL89, ZRS87a, ZRS87b]. Admission control strategies are employed to ensure the adequacy of resources for achieving the real-time requirements of aperiodics before accepting the events that trigger their execution. The fundamental premise of the classical real-time theory that provide hard guarantees for aperiodics is that the (worst-case) parameters of the computations such as execution times, communication delays, and distributions of event-arrivals are deterministically known, with absolute certainty in advance. On the other hand, asynchronous decentralized real-time systems are characterized by load patterns that are impossible to deterministically characterize and obtain upper bounds. Thus, computational parameters such as execution times and communication delays have unknown upper bounds and event arrivals have non-deterministic distributions in such systems. Hence, determining the worst-case resource needs and performing an *a-priori* reservation of resources is not a viable approach.

Also, it is practically infeasible to perform admission control for aperiodic events. Aperiodic computations in asynchronous, decentralized real-time systems are often triggered by (asynchronous) events that occur in the application environment such as detection of hostile threats that need to be engaged and destroyed in an air-engagement combat system. Furthermore, the computations perform mission-critical calculations such as determining the launch velocities and flight trajectories of weapons (e.g., interceptor missiles) that are detonated to engage the threats. The ability to meet the real-time and survivability requirements of the computations are thus critical to the success of the mission. Therefore, "rejection" of aperiodic events (as a result of admission control) that are asynchronously triggered to engage and intercept threats is impossible as every aperiodic event must be "accepted" and the resulting computation must be timely and survivable. This means that aperiodic scheduling and resource management algorithms that emphasize average-case performance are practically infeasible for asynchronous real-time systems since meeting the real-time requirements of each computation is as critical as every other.

The QoS literature also does not address the problem of achieving the real-time and survivability requirements of aperiodic computations in asynchronous, decentralized real-time systems. The fundamental premise of the QoS technologies is that the application can tolerate multiple levels of service that vary in their quality of computational results and requirements of resources. Therefore, given a set of QoS levels at which the application can operate, an adaptation mechanism can determine the "right level" of QoS depending upon the availability of internal resources and external load. Such QoS adaptation models and algorithms are presented for periodic computations in [AAS97, AS98, CSST97, HSNL97, HBNB97, RLLS98, RLLS97, RS99, RSY98, RSYJ97, RWS99].[2] However, the proposed QoS adaptation models cannot be applied for achieving the requirements of aperiodic computations. This is because it is difficult to determine the multiple levels of service that the aperiodic computations can tolerate that vary in their quality of computational results and resource requirements. Quality and accuracies of the aperiodic computations are in fact, very critical to mission success. Furthermore, the QoS adaptation algorithms that perform the adaptation—selection of the "right level" of service—*during event arrival* are impractical for aperiodics due to the high cost incurred in performing the adaptation.

To address this problem, we propose *proactive* resource management algorithms for aperiodics that forecast their resource needs, given the current internal resource availability and desired external load. We propose to develop forecasting techniques that predict the end-to-end resource requirements of aperiodic computations using *regression theory*. Further, we propose to model the desired number of aperiodic events that needs to be satisfied by the system in a timely and survivable manner as the desired QoS, as opposed to the accuracy or quality of the computations. Thus, the application will desire a certain number of aperiodic events that it would prefer to be "ready for" at all times. The proposed proactive adaptation mechanism will forecast the resource requirements that are needed to satisfy the desired number of aperiodic events, given the current resource availability, workloads, and real-time and survivability requirements of other, possibly, periodic tasks.

The forecasted requirements will be used to drive resource allocation strategies. For example, if the forecasted resource requirements are not currently available, then resource allocation is performed through adaptation mechanisms such as replication of processes of the aperiodic computation. Resource allocation strategies will determine (1)

[2] Adaptive QoS mechanisms are also described in the multimedia literature [ACH96]. Examples include the QoS-A framework [CCH94], the Heidelberg QoS model [VW+96], V-net [FZM95], NetWorld [CC+96], XRM [LBL95], the OMEGA end-point architecture [NS96], and the QoS Broker [NS95].

candidate processes of the aperiodic computation that needs to be replicated, (2) optimal number of replicas that are needed that can satisfy the desired real-time and survivability requirements, and (3) optimal processor and network resources that are required for the computational and communication needs of the replicas—processors for executing the replicas and networks that will allow the replicas to communicate with the rest of the application—that will satisfy the desired requirements. The proactive resource allocation algorithms will determine whether it is possible to satisfy the real-time and survivability requirements of the aperiodics through dynamic adaptation and without jeopardizing the timeliness and availability of other computations, given the current situation. Furthermore, if a feasible resource allocation is possible, the algorithms will determine the optimal resource allocation. If the algorithms determine that it is not feasible to achieve the collective requirements of the application, then negotiation strategies will be employed to achieve the optimal satisfaction of requirements (real-time and survivability) across computations. Thus, proactive resource management algorithms perform resource allocation and negotiation based on forecasted application performance to achieve the desired requirements.

We propose to use regression theory for forecasting the resource requirements of aperiodics. The resource usage characteristics of the application such as execution times and communication delays will be profiled at a set of internal resource utilizations (e.g., CPU utilization, memory utilization, network utilization) and external loads (e.g., sensor reports, aperiodic events). The application profile data will then be used to determine regression equations that compute resource usage characteristics at arbitrary internal resource utilization levels and external loads. The regression equations will be used to forecast resource needs, determine feasibility of the resource allocations in achieving the requirements, optimal resource allocations, and perform negotiations.

We have conducted preliminary investigations to determine the viability of regression theory for proactive resource management. The investigations (discussed in the subsection that follow) reveal that we can predict the resource needs of application processes such as execution times using second order non-linear regression equations at arbitrary resource utilizations and external loads with very high degrees of predictability. However, the degree of predictability with which we can forecast the resource needs of inter-process trans-node message links such as communication latencies depends upon the underlying network. For Ethernet networks, where randomness is built into the network layers to resolve packet collisions through mechanisms such as exponential back offs, we propose probabilistic prediction strategies. However, on IEEE 802.5 token ring networks, we can forecast the communication latencies with much higher degrees of predictability.

3.2 Adaptive Communication Mechanisms

Past efforts in real-time communication research have focused primarily on static message scheduling [ACZD92, BM91, CAZ92, Fer92, FV90, KZ93, KMZ93, MZ95, MKZ96, Sath93, SLS88a, TBW95]. The fundamental thesis of these efforts is that the communication behavior of the application can be made to be deterministic through extensive *a-priori* knowledge about communication load parameters, message latencies, and precedence relationships. The a-priori information is then used to perform end-to-end static schedulability analysis that determines the ability of the application to meet the end-to-end timing requirements. It is very difficult if not impossible, to employ such techniques for achieving the end-to-end real-time requirements in asynchronous real-time systems because of two fundamental reasons.

First, message latencies in asynchronous decentralized real-time systems, like task latencies, are significantly dependent upon the conditions of the external environment of the application and the current operational scenario of the system. For example, message latencies are functions of environment-dependent application parameters such as number of sensor reports generated by sensor devices (e.g., radars) and are transmitted for trans-node processing.

Secondly, the geographically distributed nature of the system (e.g., radars, missile launchers, and computing nodes of a shipboard combat system that performs air-engagement) constitutes a highly loosely coupled environment where the communication paradigm used often employs buses, links, and rings that produces significant run-time uncertainties in the message latencies (unlike communication using private shared memory that is more predictable).

Thus, adaptive resource management that involves run-time monitoring, feedback, and adaptation at the trans-node message-level to encounter changing workloads and uncertainties becomes crucial in achieving the end-to-end real-time requirements of the application. The real-time communication literature is very sparse in this area. Interestingly, the new efforts on QoS management technologies such as [Quo97] have also primarily focused upon dynamic adaptation at the process-level, using techniques such as process replication for exploiting concurrency and load sharing [RWS99, RSY98, RSYJ97], and the imprecise computation model that trades-off computational accuracy of processes against their resource utilizations [HSNL97, LSTS99]. To the best of our knowledge, the QoS literature does not contain any works on adaptive resource management for asynchronous real-time systems where adaptation is performed at the trans-node message level.

To address this problem, we propose adaptive communication strategies that achieve the real-time requirements through mechanisms such as *dynamic message reprioritization* and *dynamic alteration of token holding times at processor stations*. We propose to develop distributed, adaptive communication algorithms for IEEE 802.5 and 802.1p networks that support priority-based message scheduling. The adaptive communication algorithms will assign *initial* priorities to application messages based on end-to-end task deadlines and relative priorities of tasks. For example, deadlines can be assigned to messages from end-to-end task deadlines using optimal strategies such as EQF [Kao95]. The message deadlines can then be used to determine their priorities using optimal techniques such as Earliest Deadline First (EDF) [SSRB98]. Such an initial assignment of message priorities is possible for asynchronous real-time systems only by estimating the initial operating conditions of the system and deriving the initial values of computational parameters such as execution times and communication delays. This assignment will clearly not suffice as the computational parameters change at run-time.

Therefore, the proposed adaptive communication algorithms will monitor the application at run-time to determine low timeliness situations that are caused by workload changes such as high number of sensor reports. The algorithms will adapt the application to changing workloads by first identifying messages that are experiencing fluctuations in their communication latencies (due to variations in data size loads) and then changing their priorities. For example, we can assign high priorities to those messages that are experiencing increasing communication latencies and thereby increase the communication bandwidth of the messages and reduce their communication latencies. We propose to develop algorithms that will determine the *current optimal dynamic priorities* of the messages that will achieve the real-time requirements, given the current resource availability, workloads, and real-time and survivability requirements of other tasks. Thus, our proposed algorithms will dynamically adapt the application to changing workloads and achieve the real-time requirements through dynamic message reprioritization.

We propose to develop the message reprioritization algorithms for IEEE 802.5 (the priority-driven protocol) and IEEE 802.1p standards. The 802.1p standard [IEEE802] is a very interesting development as it allows message priorities on Ethernet networks that are cheap and widely available. Thus, we believe that our approach will have a strong impact. Further, we propose to develop adaptive communication strategies for the Timed Token Protocol (TTP) that is employed in network standards such as FDDI [ANSI87], HSDB/HSRB [SAE88a, SAE88b, Uhl91], and SAFENET [GM89, DoD94]. Since TTP does not support priorities for messages, we propose to develop adaptive communication strategies that perform adaptation by changing the token holding times at processor stations. We propose to develop algorithms that will determine the initial target token rotation times and initial token holding times at stations based on estimates of initial operating conditions of the system. The adaptive communication algorithms will then monitor the application at run-time to determine low timeliness situations caused by changes in workloads. The algorithms will identify processor stations with application programs that are generating messages that are experiencing variations in their communication latencies (due to variations in data size loads). The algorithms will perform adaptation by changing the token holding times of the stations to change their communication bandwidths. For instance, we can increase the token holding times of those stations that are generating messages that are experiencing increasing communication latencies. By increasing the token holding times of such "bottleneck" stations, we can increase the communication bandwidth of the stations and thus the bandwidth of the messages that are generated from the stations. This will reduce the communication latencies of the messages.

We propose to develop algorithms that will determine the *current optimal token holding times* of the processor stations that will achieve the real-time requirements, given the current resource availability, workloads, and real-time and survivability requirements of other tasks. Thus, our proposed algorithms will dynamically adapt the application to changing workloads and achieve their real-time requirements on TTP networks through dynamic alteration of token holding times at processor stations.

3.3 Hybrid Feedback Control Theory-based Adaptive Resource Management

As discussed in Section 1, recent advances in real-time distributed systems research [Quo97] have produced a suite of QoS technologies that allow asynchronous, decentralized real-time systems to achieve their real-time and survivability requirements, which was not previously possible under classical real-time theory. However, the effectiveness of the QoS techniques in achieving the desired requirements is only *empirically* studied in the literature. This only validates the viability of the techniques for the set of operating conditions under which they have been studied. For example, efforts such as [BN+98, HBNB97, HSNL97, RWS99, RSY98, RSYJ97] present techniques to monitor the application for adherence to the desired requirements, detect situations when the application fails to meet the requirements or exhibit trends for failures, and recover from the failures. The effectiveness of the techniques is studied either through simulation or by implementation (e.g., middleware) and evaluation using real-time benchmarks. From such experimental data, the effectiveness of the techniques for performing a recovery from a timing or survivability failure is known, given the nature of the failure, conditions of external load, and internal resource situation. However, the general feasibility conditions of the techniques for achieving the desired real-time and survivability requirements are not known. For instance, it is not known under what general conditions of resource availability (processors and networks) and external load (number of sensor reports and aperiodic events) do the QoS techniques achieve the desired requirements.

We believe that it is impossible to determine the boundaries of feasibility and the limitations of the techniques through empirical study. To determine the feasibility conditions, a formal, mathematical study of the behavior of the techniques is essential. Knowledge of the feasibility conditions will allow the behavior of the techniques to be reasoned under all possible operating conditions of the system.

This is a very difficult problem, as there does not exist any analytical frameworks for studying the behavior of resource management techniques for systems that have uncertain operating conditions. More significantly, there does not exist a theory and hence, tools and experimental methods for designing adaptive resource management techniques and for reasoning, evaluating, and validating their behavior.

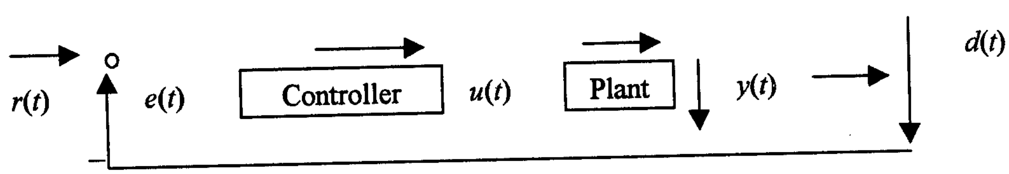

Figure 1. Classical Feedback Control

To develop a mathematical foundation for designing adaptive resource management techniques and to understand and evaluate their behavior, we propose resource management techniques that are based on *feedback control theory*. In the discipline of feedback control theory, control theorists design control functions that optimize variables of physical systems based on continuous or discrete time feedback (Figure 1). The functions are analytically designed and their behavior is also studied analytically [PH96, SP96].

Feedback control, therefore, is remarkably similar to adaptive resource management techniques. However, control functions that are analytically designed and studied in control theory for the most part, are linear-time invariant systems in continuous and discrete times. Asynchronous real-time systems are inherently non-linear. Therefore, we require a new paradigm to design control functions for such systems. We propose adaptive resource management control functions that are based on *hybrid* models that contain elements of different modeling paradigms [MS00 and references therein]. For instance, a hybrid control system can have a finite state machine (FSM) model with each state represented by continuous dynamics in terms of controlled vector fields or difference equations. The controller can be explicitly part of the continuous dynamics or be implied, so that the task of the controller is to determine the destination state that is reached after a transition either synchronously or asynchronously. We believe that the class of hybrid system model that contains asynchronous finite state machines with continuous ordinary differential equations is the most appropriate model for asynchronous decentralized real-time systems.

4. Current Status and Accomplishments

In this section, we describe the current status and our accomplishments toward solving each of the problems. Section 4.1 presents predictive resource management algorithms that forecast the resource needs of periodic computations and achieve their real-time requirements through regression theory-based prediction. Section 4.2 discusses an adaptive communication heuristic that adapts periodic computations to changing workloads through dynamic message reprioritization and achieves their real-time requirements. Finally, Section 4.3 presents an adaptive resource management control function that is based on the classical proportional integral derivative (PID) feedback control function. We study the performance of the techniques through a combination of simulation and benchmarking.

4.1 Validation of the Effectiveness of Regression Theory for Adaptive Resource Management

To determine the effectiveness of regression theory for forecasting the resource needs of computations and perform proactive resource management, we have developed predictive algorithms for *periodic* computations that satisfy their real-time requirements. We consider an example asynchronous, decentralized real-time system with trans-node tasks that are required to process data (sensor reports) that arrives periodically, within a specified end-to-end deadline. The upper bound on the size of the data that arrives during each period is assumed to be unknown a-priori. Hence, the upper bound on the end-to-end execution latencies of the tasks during each period is not deterministically known.[3] The sub-tasks of the tasks (executable application programs) are assumed to be replicable so that the replicas can be dynamically executed on multiple computing nodes to exploit concurrency and achieve load sharing. Such application "scaling" is performed to reduce end-to-end task latencies when the data size increases at run-time and causes unacceptable task timeliness. The application hardware is assumed to consist of a set of distributed processors with private memory (e.g., host machines) that share a common communication medium (e.g., Ethernet).

We consider two predictive resource allocation algorithms, called Predictive I and Predictive II that are based on regression theory for achieving the real-time requirements of the tasks. The algorithms determine the number of replicas of application programs of the tasks that are needed to adapt the tasks to changes in workloads through regression theory-based prediction. The resource usage characteristics of application programs of the tasks are determined for a number of resource utilization conditions and external load situations. We use the real-time benchmark application that has resulted from our past effort [Rav98, SWR99] as the example application.

Figures 2 and 3 show sample plots of execution latency measurements of two application programs of the benchmark called *Filter* and *EvalDecide*. The latency measurements are made for increasing number of sensor reports processed by the programs at a number of different CPU utilization levels (blue lines called "y" in the figures). For the latency measurements made at a CPU-utilization, we first determine a second order non-linear regression equation that computes execution latency as a function of data size (red lines called "Y" in the figures). Finally, we combine the equations for latency measurements made at the different CPU utilizations into a single regression equation that computes execution latency as a function of data size and CPU utilization (green lines called "Y-" in the figures). The regression equation thus obtained for the benchmark is given by $e = f(u,d) = (a_1u^2 + a_2u + a_3)d^2 + (b_1u^2 + b_2u + b_3)d$, where e is the execution latency (in milliseconds), u is the CPU utilization (%), d is the data size (in hundreds of data items), and a_i 's and b_i 's are constants that are dependent upon the application program. Table 1 illustrates the values that were obtained through measurements for a_i 's and b_i 's for the two programs, Filter and EvalDecide.

[3] Examples of such tasks can be found in distributed real-time combat systems such as the Navy's Anti-Air Warfare (AAW) system [WRSB98, WR+96]. The tasks process sensor reports, which are periodically generated by radar systems. Upper bounds on the number of sensor reports ("radar tracks") are not deterministically known, as they depend upon the operational scenarios of the system and vary significantly between different scenarios.

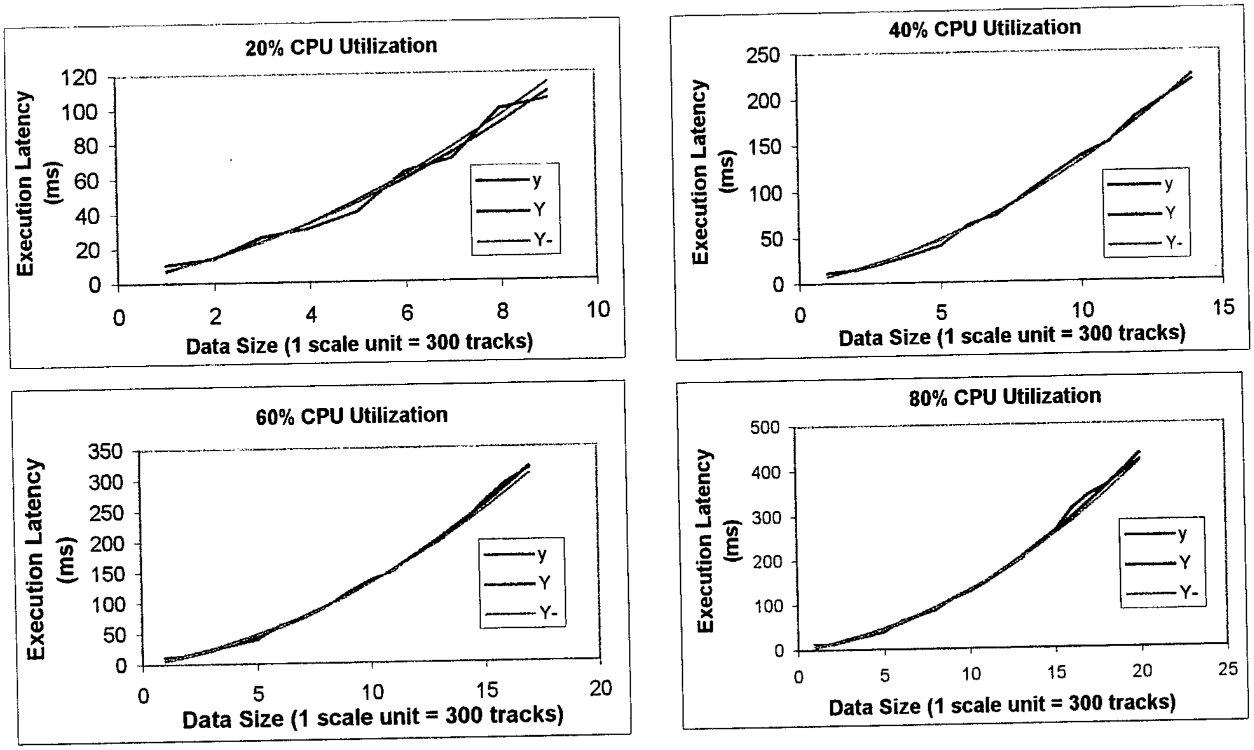

Figure 2. Execution Latencies of Filter at Different CPU Utilizations and External Loads

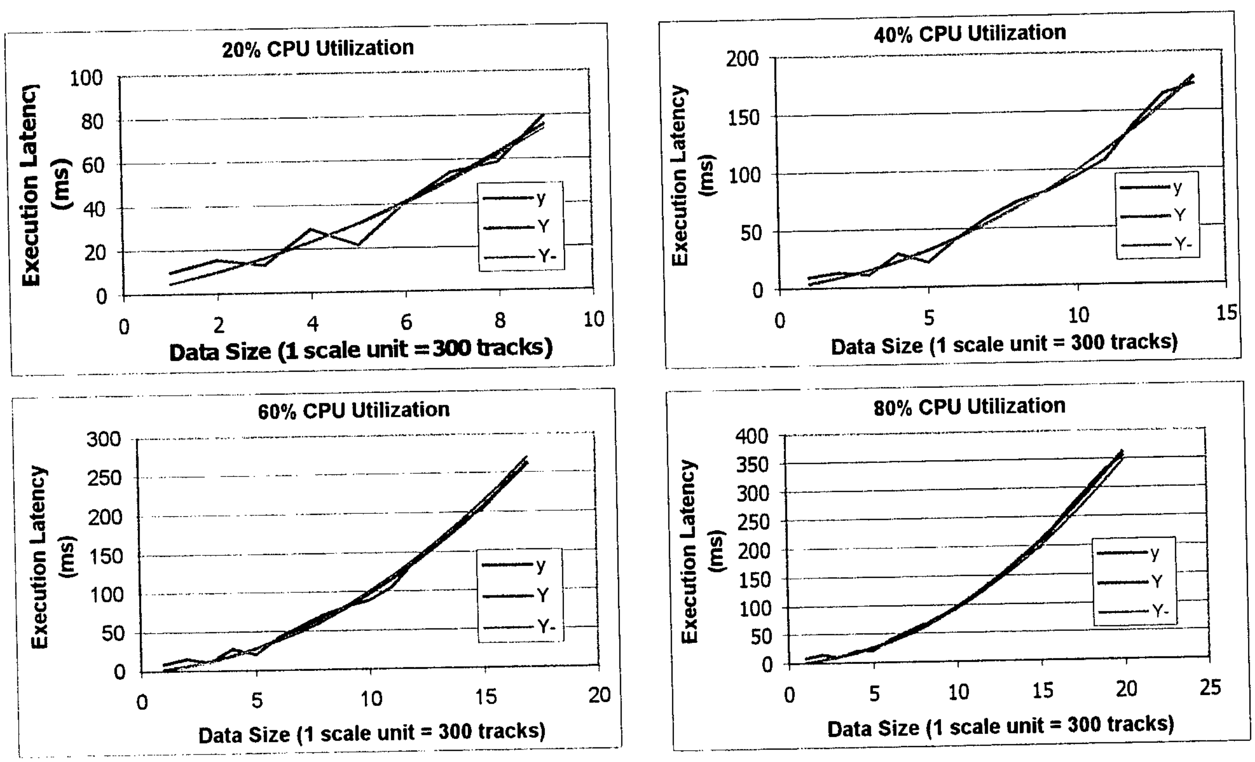

Figure 3. Execution Latencies of EvalDecide at Different CPU Utilizations and External Loads

At run-time, when the application workload increases and causes a low timeliness situation, the algorithms detect the situation and perform "diagnosis" to identify candidate programs of the task for replication. The two algorithms differ in the way they perform diagnosis.

Predictive algorithm I performs diagnosis by assigning deadlines to individual programs of the task from the end-to-end task deadline using a variant of the Equal Flexibility Strategy (EQF) proposed in [Kao95]. The algorithm identifies those application programs for replication that have missed their individual deadlines.

Predictive algorithm II, on the other hand, analyzes the trend of execution latencies of the application programs during a window of past task periods and identifies those programs for replication that are exhibiting higher execution time growth rates. Once the candidate programs are identified, the algorithms determine the optimal number of replicas that are needed to achieve the real-time requirement by iteratively forecasting the program latency using the regression equation. The regression equation is used to predict the execution latency of a replica on the set of candidate processors on which it can be executed. The observed CPU utilizations of the processors and the data size that the replica will need to process are used in the equation. The forecasted execution latency is then compared with the program deadline. The process is repeated in an iterative fashion, increasing the number of replicas with each iterative step until the program deadline is satisfied. Observe that with each additional replica, the data size load of each of the existing replicas is reduced and thus their forecasted execution time is also reduced.

Table 1. Coefficients of Regression Equation for Predicting Program Execution Latencies

Filter	EvalDecide
a_1 = 1.53554E-05	a_1 = -1.59635E-05
a_2 = -0.001553929	a_2 = 0.002123039
a_3 = 0.118161741	a_3 = 0.022324378
b_1 = -0.000285037	b_1 = 0.000107844
b_2 = 0.02982759	b_2 = -0.023927144
b_3 = 0.983699036	b_3 = 1.443762034

We compare the performance of the predictive algorithms with two non-predictive algorithms that determine the number of replicas using heuristic strategies. Our first non-predictive algorithm, called "external load-based" (or EL-based) determines the number of replicas based on the rate of change in the data size load of the programs. The algorithm decides on the number of replicas for sub-tasks of the tasks as a linear function of the rate of change in the data size load of the sub-tasks during a window of past task periods.

The second non-predictive algorithm, called "resource availability-based" (or RL-based) determines the number of replicas based on resource availability. The algorithm identifies processors that are exhibiting utilization levels below a certain threshold value and replicates all the (replicable) sub-tasks of the tasks on such processors. We study the performance of the algorithms using the metric of missed deadline ratios during increasing workloads caused by high data sizes.

Figure 4 shows the missed deadline ratios of the four algorithms. Figures 5, 6, and 7 illustrate the average CPU utilization, average network utilization, and average number of replicas used by the four algorithms, respectively. From the figures, we observe that both the predictive algorithms perform as effectively or better than the non-predictive algorithms. Also, we note that there is no significant performance difference between the two predictive algorithms. This is due to the ability of the diagnostic techniques employed by the algorithms to discriminate against poorly performing sub-task programs equally well. We also note that the external load-based algorithm misses more number of deadlines than all other algorithms almost at all times.

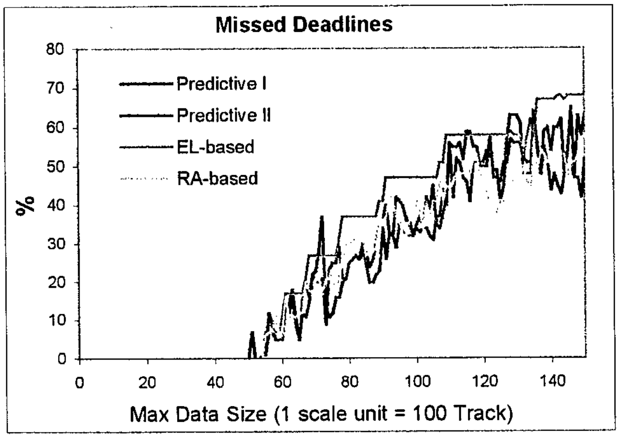

Figure 4. Missed Deadline Ratios of Predictive and Non-predictive Algorithms

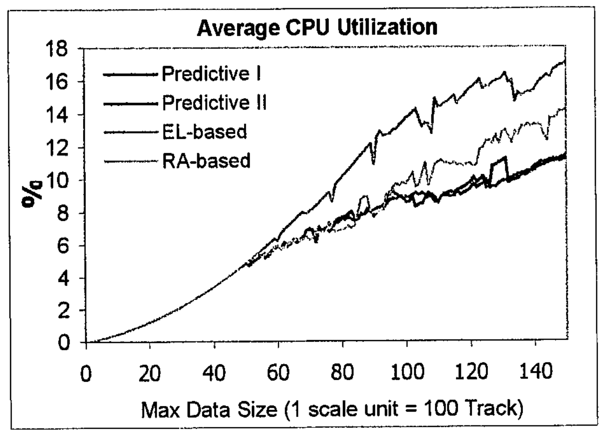

Figure 5. Average CPU Utilization of Predictive and Non-predictive Algorithms

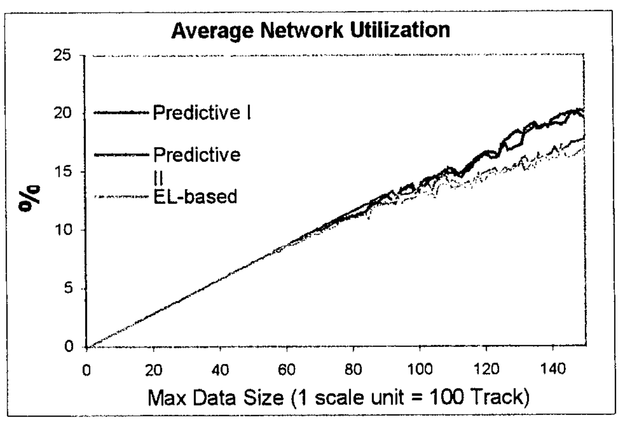

Figure 6. Average Network Utilization of Predictive and Non-predictive Algorithms

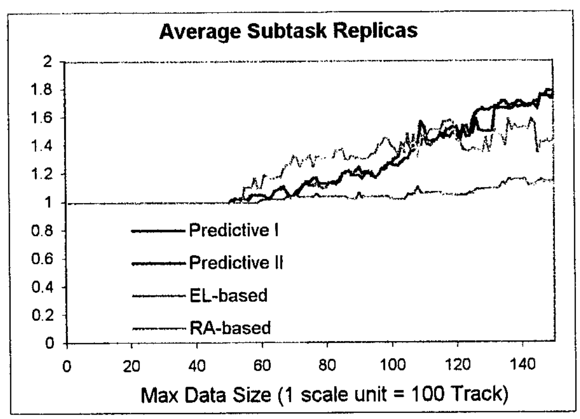

Figure 7. Average Number of Replicas of Predictive and Non-predictive Algorithms

Thus, the experimental results establish the effectiveness of regression theory for performing adaptive resource management.

4.2 Validation of the Effectiveness of Adaptive Communication Mechanisms for IEEE 802.5 Standard

To determine the effectiveness of adaptive communication mechanisms, we have developed a heuristic technique for performing message reprioritization on IEEE 802.5

networks. We consider the same example asynchronous, decentralized real-time system discussed in Section 4.1. We assign initial deadlines to the sub-tasks (or application programs) and messages of the periodic tasks from their end-to-end deadlines using a variant of the EQF strategy [Kao95]. We perform the initial deadline assignment based on estimates of the initial operating conditions of the system. The message deadlines are then used to determine their (initial) priorities based on EDF.

We have developed a centralized, adaptive communication algorithm called local message reprioritization (LMR) that monitors the performance of the end-to-end tasks at run-time and determines low timeliness situations by detecting missed deadlines of messages. Once a low timeliness situation is detected, LMR performs message reprioritization using a *heuristic* strategy. The algorithm first determines the set of messages of the task that have missed deadlines and orders the set according to their laxities (difference between deadline and message communication latency). Priorities are given to the messages of this set according to the Least Laxity First (LLF) strategy, starting from the highest priority that has been previously assigned to the sub-tasks or messages of the task. The priorities of the messages of the task that did not miss their deadlines are now reassigned according to EDF, starting from the next lower priority that has been given to the LLF-ordered message set. The intuition behind such a priority assignment is that we can exploit the communication bandwidth of the messages that are doing well and have higher laxities. This bandwidth can be used by the messages that are experiencing low laxities. LMR thus partitions the total message set of the task into a set of failed messages and a set of successful messages. Priorities are reassigned to the set of failed messages according to LLF and to the set of successful messages according to EDF within the priority range of the task (and hence the heuristic is "local").

We study the performance of the LMR heuristic strategy using the metric of missed deadline ratios and compare its performance with an adaptive resource management algorithm called PR that performs adaptation through process replication. The PR algorithm is a combination of the Predictive I and EL-based algorithms presented in Section 4.1. The algorithm monitors end-to-end tasks for low timeliness situations and identifies candidate sub-tasks of the tasks for replication by assigning individual deadlines to the sub-tasks using a variant of EQF [Kao95]. The algorithm identifies those sub-tasks for replication that have missed their individual deadlines and determines the number of replicas that are needed as a linear function of the rate of change in the data size load of the sub-tasks during a window of past task periods. Further, the algorithm selects the processors for executing the replicas using a "first-fit" load balancing strategy where the capacity of the processors is characterized using their CPU utilizations.

The performance of the LMR and PR algorithms is studied during high workload situations caused by high number of sensor reports to be processed by the tasks, under three specific load scenarios: (1) communication latencies of the messages grow significantly with respect to process execution latencies as the workload increases (load scenario called LS-1), (2) execution latencies of processes grow significantly with respect to communication latencies as the workload increases (load scenario called LS-2), and (3) both message communication latencies and process execution latencies grow significantly as the workload increases (load scenario called LS-3).[4]

Figures 8, 9, and 10 show the missed deadline ratios, average CPU utilization, and average network utilization for the two algorithms during the three load conditions, respectively. From Figure 8, we observe that LMR outperforms PR during LS-1 when message latencies increase significantly with respect to process latencies. This is intuitive, since LMR performs adaptation at the message-level. In contrast, PR performs poorly during such situations when increase in process latencies is not substantial, thereby diminishing the prospects of the algorithm to make an impact through process-level adaptation.

However, note that LMR performs as good as PR during LS-2 when increase in process latencies is significantly higher with respect to increase in message latencies! This is due to the fact that increase in process latencies also results in an increase in message latencies. Further, the ability of LMR to improve task timeliness for such indirect smaller increases in message latencies is significant. In fact, this is significantly better than that of PR for smaller direct increases in process latencies (observe the curves of the algorithms in Figure 8 for the data size range 1 through 16 for LS-2).

However, as the message latencies indirectly increases and become substantial, LMR's ability to improve task timeliness also decreases and the performance becomes almost the same as that of PR. This intuition is clearly substantiated when we compare the performance of LMR during LS-1 and LS-2. During LS-2 when message latencies increase indirectly and less significantly, the algorithm is able to make a greater impact than during LS-1 when the increase in message latencies is direct and substantial. Finally, we observe that LMR performs almost as good as PR during LS-3 when both process and message latencies increase.

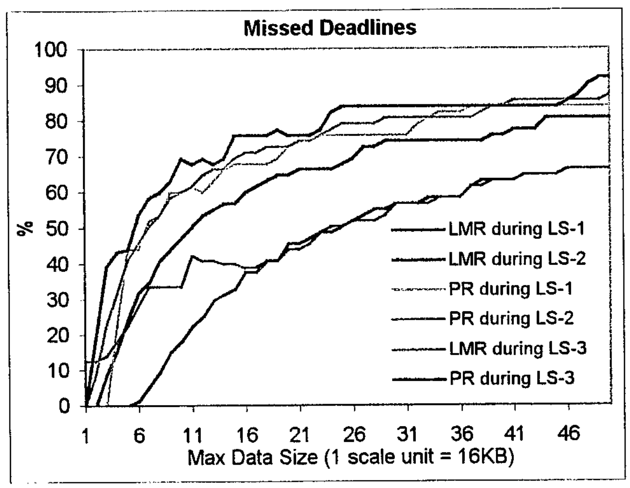

Figure 8. Missed Deadline Ratios of LMR and PR

[4] Such load patterns are typical in radar track processing systems such as the AAW system, which contain components for radar track filtering, track classification, and threat evaluation. The radar may generate very high data loads that contain only small number of valid tracks, but significant amount of noise. This causes execution and communication latencies "up to" the filtering component to increase significantly during the data period. However, the latencies that follow the filtering component will be significantly lower.

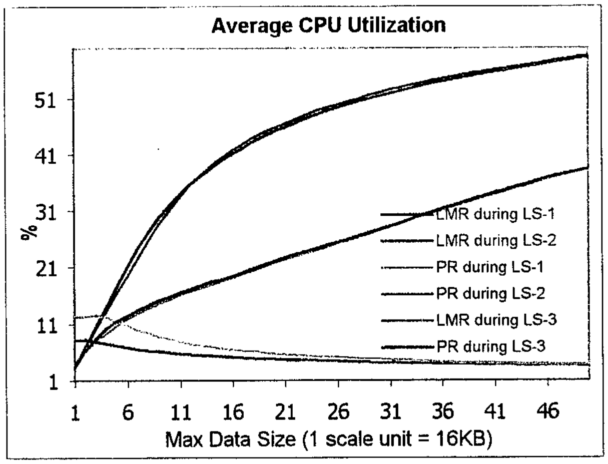

Figure 9. Average CPU Utilization of LMR and PR

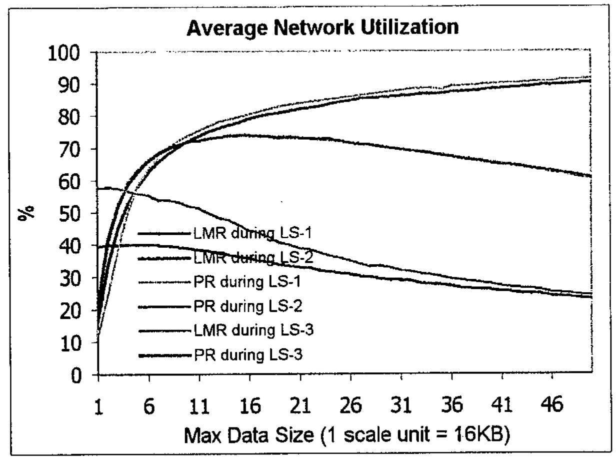

Figure 10. Average Network Utilization of LMR and PR

Thus, our experimental results clearly validate the effectiveness and superiority of adaptive communication mechanisms through message reprioritization.

4.3 Validation of the Effectiveness of Feedback Control Theory-based Adaptive Resource Management

To illustrate the flexibility and correctness of the hybrid control paradigm, we describe a hybrid model for an example problem. We consider the same example asynchronous, decentralized real-time system discussed in Section 4.1. We can design an adaptive resource management controller for this system that determines the optimal number of replicas that are needed for each sub-task of the periodic tasks during the task periods. The hybrid system for the problem can be represented by a strongly connected directed graph where the graph nodes represent the system states and the control law has to decide when to perform a state transition and the destination state that is reached after a transition. For example, if we have two processors namely P and Q, that can execute two sub-tasks u and v, then the total number of different states that the system can be in is given by:

$$\text{Total number of states with two processors and two sub-tasks} = \left(\sum_{j=1}^{2} C_j^2\right)^2$$

This turns out to be nine different states. The states correspond to [(u,P), (v,P)], [(u,P), (v,Q)], [(u,P), (v,P;v,Q)], [(u,Q), (v,P)], [(u,Q), (v,Q)], [(u,Q), (v,P;v,Q)], [(u,P;u,Q), (v,P)], [(u,P;u,Q), (v,Q)], and [(u,P;u,Q), (v,P;v,Q)]. The states are shown in Figure 11. Here, [(u,P), (v,Q)] indicates that the sub-task u is being executed in processor P and the sub-task v is being executed in processor Q. [(u,P;u,Q), (v,P)] indicates that the sub-task u is being executed in processor P and processor Q, and the sub-task v is being executed in processor P. Each of the finite number of states contains continuous dynamics to model the sub-task execution. The continuous differential equations within each state are simply $\dot{x} = -1 + d$ that model the execution of each sub-task in any processor, where the variable d represents the uncertainties in the execution time of a sub-task. At each instant in time, the controller has to decide whether to transition from the current state or to continue in the current state. If the decision is to transition from the current state, then the controller has to decide the destination state of the transition.

Figure 11. Strongly Connected Graph of the Sample Problem in Hybrid Design Paradigm

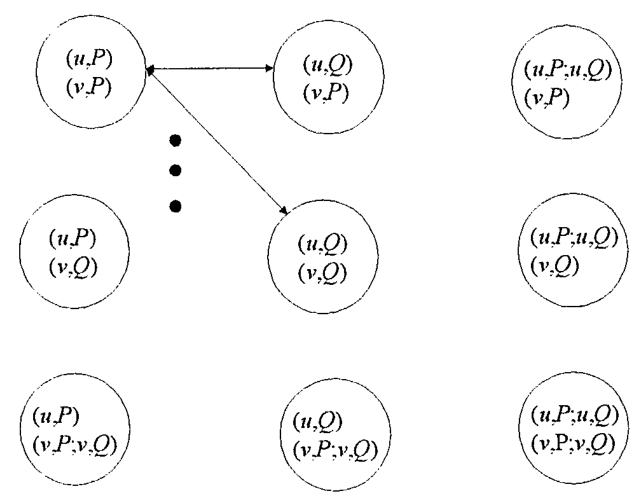

The hybrid model thus indicates the execution time for each sub-task and the set of states that the system can be transitioned into, as a result of the decision of the controller. The objectives of the controller can include minimization of task missed deadline ratios and processor and network utilizations. We can formulate an optimal control problem using this model and can use game-theoretic approaches to counter the disturbances and uncertainties. Also, we can add projection dynamics to enable the system to be "viable" at all times. To model situations such as processor failures, the topology of the FSM itself can be made to be time varying.

To determine the effectiveness of adaptive resource management techniques that are based on feedback control theory, we designed an adaptive resource management controller for the above problem that determines the number of sub-task replicas using the classical proportional integral derivative (PID) feedback control function. The function uses the sum of weighted error, integral of error, and derivative of error terms as the control variable. We define the error term as

$e(k) = w_1(x(k) - X(k)) + w_2(u(k) - U(k)) + w_3 n(k) + w_4 m(k)$, where k is the sampling instant, $x(k)$ is the actual task execution time for processing the data that arrived in the previous period, $X(k)$ is the desired task execution time, $u(k)$ is the actual average utilization of the processors during the previous period, $U(k)$ is the desired utilization, $n(k)$ is the actual average network utilization during the previous period, and $m(k)$ is the missed deadline ratio. Based on the error term, a PID control function that computes the change in the number of replicas for each sub-task of the task is given by:

$$\Delta st_i(k) = k_{1i}e(k) + k_{2i}\sum_{j=0}^{k} e(j) + k_{3i}e(k-1); \qquad i \in [0,1,..n]$$

Through a simulation study, we evaluated the performance of the PID-based technique by comparing it with two "static" controllers that always use a constant number of sub-task replicas. One static controller called Static-3, uses half of the available number of processors as the (constant) number of sub-task replicas and the other controller called Static-6, uses the maximum available number of processors as the number of sub-task replicas, to exploit maximum concurrency. We study the performance of the resource management controller functions during situations of high workloads caused due to increasing number of sensor reports.

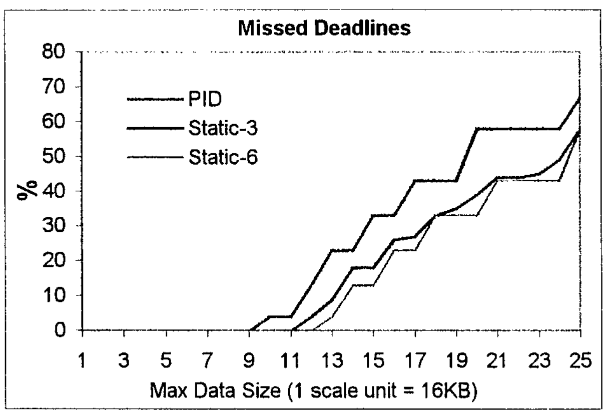

Figure 12. Missed Deadline Ratios of Feedback and Static Controllers

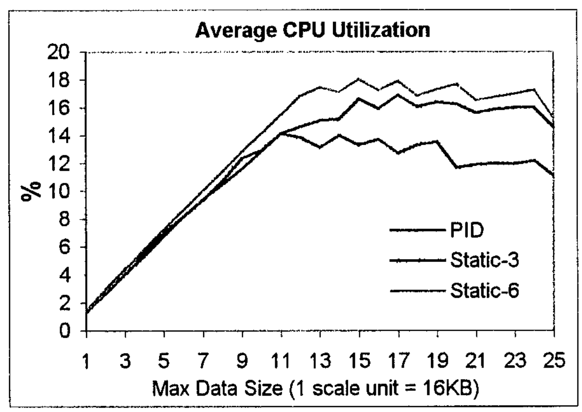

Figure 13. Average CPU Utilization of Feedback and Static Controllers

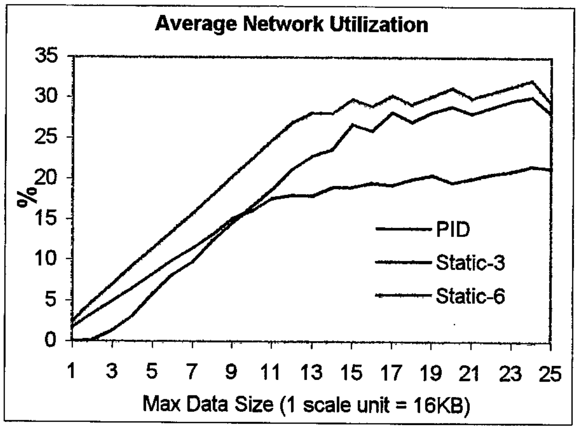

Figure 14. Average Network Utilization of Feedback and Static Controllers

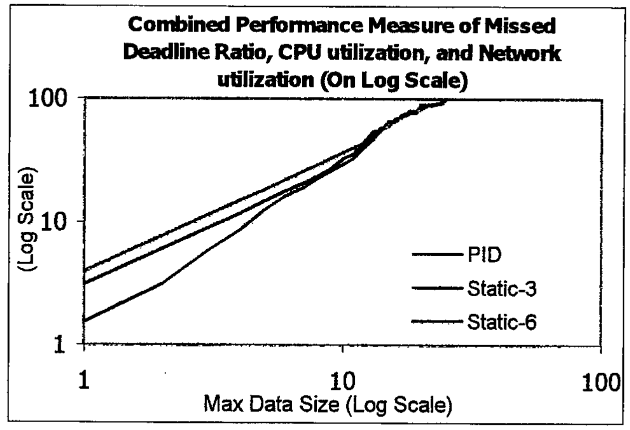

Figure 15. Combined Performance of Feedback and Static Controllers

Figures 12, 13, 14, and 15 show the missed deadline ratios, average CPU utilization, average network utilization, and a performance metric that combines missed deadline ratios and CPU and network utilizations for the three resource management controller functions, respectively. From Figure 12, we observe that the missed deadline ratios of the Static-6 controller is the best, that of the Static-3 controller is the worst, and that of the PID function lies in between. This is because the Static-6 controller always produces the maximum number of sub-task replicas and thereby exploits the maximum available concurrency in the system. Hence, the performance of the controller is the best. However, this significantly increases the CPU and network utilization (Figures 13 and 14). The Static-3 controller on the other hand, utilizes only half the available concurrency. This may not be sufficient for adapting the application to increasing data sizes. Hence, the controller performs poorly. However, the CPU and network utilization caused by the controller is significantly low.

The PID controller adapts the application to increasing data size situations by using the minimal number of replicas that is determined through feedback control. The performance of the controller thus falls between the best and the worst for all three metrics. The PID function thus produces the lowest value for the combined performance metric that aggregates the three metrics of missed deadlines, average CPU utilization, and network utilization at all times (Figure 15).

The experiments thus illustrate the effectiveness of feedback control theory for designing adaptive resource management techniques. However, note that the parameters of the control function are determined here using the experience of the designer. This might work in many practical applications, but does not give a generic model and analytical performance guarantees for most of the actual problems. It is clear that a complete theory needs to be developed to enable the design of feedback controllers for asynchronous decentralized real-time systems that provide the desired specified behavior. We propose the use of the hybrid system philosophy as a framework and methodology for designing feedback controllers for such systems. The traditional control designs lack in providing this required framework due to their inherent difficulties in dealing with non-linear models. We believe that this would provide guarantees of optimality or desired transient and steady state performance even in the presence of disturbances and uncertainties.

5. Publication, Presentation, and Related Activities

Our publication, presentation, and other related activities during this quarter includes the following:

Publications

1. B. Ravindran, "On Predictive Resource Management Algorithms for Asynchronous, Decentralized Real-Time Systems," *Journal of Systems Integration*, Submitted April 2000 (under review).
2. B. Ravindran and P. Kachroo, "Adaptive Resource Management in Dynamic, Distributed Real-Time Systems Using Feedback Control Functions," *Twentieth Conference on the Foundations of Software Technology and Theoretical Computer science*, Submitted June 2000 (under review).
3. B. Ravindran, P. Kachroo, and T. Hegazy, "Intelligent Feedback Control–based Adaptive Resource Management for Asynchronous, Decentralized Real-Time Systems," *IEEE Transactions on Systems, Man, And Cybernetics Part C: Applications And Reviews*, Submitted July 2000 (under review).
4. B. Ravindran and S. Edwards, "Palette: A Reuse-Oriented Specification Language for Real-Time Systems," *Software Reuse: Advances in Software Reusability*, William Frakes (Editor), *Proceedings of the Sixth International Conference on Software Reuse, Lecture Notes in Computer Science*, Springer-Verlag, Volume 1844, pages 20-40, June 2000, Available at: http://www.ee.vt.edu/~binoy/papers.html

Presentations

1. B. Ravindran, "Palette: A Reuse-Oriented Specification Language for Real-Time Systems," *The Sixth International Conference on Software Reuse*, Vienna, Austria, *Lecture Notes in Computer Science*, Springer-Verlag, Volume 1844, pages 20-40, June 2000.

Other:

Guest Editor (with E. Douglas Jensen), Special Issue on "Asynchronous Real-Time Distributed Systems," *IEEE Transactions on Computers,* Scheduled for publication in 2001.

6. Planned Efforts for the Next Quarter

Our planned activities for the next quarter includes the following:

1. Develop proactive resource management algorithms that achieve the real-time requirements of aperiodic computations without affecting the timeliness of other computations (negotiation mechanisms will be developed in the subsequent quarter)
2. Design metrics for evaluating the effectiveness of proactive resource management algorithms

3. Develop adaptive communication (centralized) heuristic strategies for Timed Token Protocol
4. Develop decentralized adaptive communication heuristic algorithms for (1) IEEE 802.5 priority-driven protocol and (2) Timed Token Protocol
5. Design hybrid models for benchmark problems in asynchronous decentralized real-time systems (e.g., determining the number of current optimal sub-task replicas for satisfying the real-time requirements)

References

[ACH96] C. Aurrecoechea, A. Cambell, and L. Hauw, "A Survey of QoS Architectures," *Proceedings of The Fourth IFIP International Conference on Quality of Service*, March 1996.

[ACZD92] G. Agrawal, B. Chen, W. Zhao, and S. Davari, "Guaranteeing Synchronous Message Deadlines with Timed Token Protocol," *Proceedings of The Twelfth IEEE International Conference on Distributed Computing Systems*, June 1992.

[ANSI87] ANSI Standard X3T9.5, "Fiber Distributed Data Interface (FDDI): Token Ring Medium Access Control (MAC), May 1987.

[AS98] T. Abdelzaher and K. G. Shin, "End-host Architecture for QoS-Adaptive Communication," *Proceedings of The Fourth IEEE Real-Time Technology and Applications Symposium*, pages 121-130, June 1998.

[AAS97] T. F. Abdelzaher, E. M. Atkins, and K. G. Shin, "QoS Negotiation in Real-Time Systems and its Application to Automated Flight Control," *Proceedings of The IEEE Real-Time Technology and Applications Symposium*, June 1997.

[BM91] A. Banerjee and B. Mah, "The Real-Time Channel Administration Protocol," *Proceedings of The Second International Workshop on Network and Operating System Support for Digital Audio and Video*, pages 160-170, November 1991.

[BN+98] S. Brandt, G. Nutt, et. al., "A Dynamic Quality of Service Middleware Agent for Mediating Application Resource Usage," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 307-317, December, 1998.

[CAZ92] B. Chen, G. Agrawal, and W. Zhao, "Optimal Synchronous Capacity Allocation for Hard Real-Time Communications with the Timed Token Media Access Control Protocol," *Proceedings of The IEEE Real-Time Systems Symposium*, December 1992.

[CC+96] D. Chen, R. Colwell, H. Gelman, P. K. Chrysanthis, and D. Mosse, "A Framework for Experimenting with QoS for Multimedia Services," *Proceedings of the International Conference on Multimedia Computing and Networking*, 1996.

[CCH94] A. Cambell, G. Coulson, and D. Hutchison, "A Quality of Service Architecture," *ACM Computer Communications Review*, April 1994.

[CSR86] S. Cheng, J. Stankovic, and K. Ramamritham, "Dynamic Scheduling of Groups of Tasks with Precedence Constraints in Distributed Hard Real-time Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, 1986.

[CSST97] S. Chatterjee, J. Snydir, B. Sabata, and T. Lawrence, "Modeling Applications for Adaptive QoS-based Resource Management," *Proceedings of The Second IEEE High-Assurance System Engineering Workhsop*, August 1997.

[DoD94] U.S. Department of Defense, "Survivable Adaptable Fiber Optic Embedded Network," MIL-STD-2204A, January 1994.

[DTB93] R. I. Davis, K. W. Tindell, and A. Burns, "Scheduling Slack Time in Fixed-Priority Preemptive Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 222-231, December 1993.

[Fer92] D. Ferrari, "Real-Time Communication in an Internetwork," *Journal of High Speed Networks*, Volume 1, Number 1, pages 79-103, 1992.

[FV90] D. Ferrari and D. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks," *IEEE Journal on Selected Areas of Communication*, Volume 8, Number 4, pages 368-379, April 1990.

[FZM95] B. Field, T. Znati, and D. Mosse, "V-net: A Framework for a Versatile Network Architecture to Support Real-Time Communication Performance Guarantees," *InfoComm*, 1995.

[GM89] D. T. Green and D. T. Marlow, "SAFENET: A LAN for Navy Mission Critical Systems," *Proceedings of The Conference on Local Computer Networks*, October 1989.

[HBNB97] M. Humphrey, S. Brandt, G. Nutt, and T. Berk, "The DQM Architecture: Middleware for Application Centered QoS Resource Management," *Proceedings of The IEEE Workshop on Middleware for Distributed Real-Time Systems and Services, The Eighteenth IEEE Real-Time Systems Symposium*, December 1997.

[HSNL97] D. Hull, A. Shankar, K. Nahrstedt, and J. W. S. Liu, "An End-to-End QoS Model and Management Architecture," *Proceedings of The IEEE Workshop on Middleware for Distributed Real-Time Systems and Services, The Eighteenth IEEE Real-Time Systems Symposium*, pages 82-89, December 1997.

[IEEE802] IEEE's 802.1p Standard (discussion), Available at: http://www.nwfusion.com/news/tech/0907tech.html

[Jen92] E. D. Jensen, "Asynchronous Decentralized Real-Time Computer Systems," Book Chapter, *Real-Time Computing, Proceedings of the NATO Advanced Study Institute*, St. Martin, Springer-Verlag, October, 1992.

[Jen99] E. D. Jensen, "Adaptive Real-Time Distributed Computer Systems," Available at: http://www.realtime-os.com/

[Kao95] B. C. Kao, "Scheduling in Distributed Soft Real-Time Systems With Autonomous Components," *PhD Thesis*, Princeton University, November 1995.

[KS92] G. Koren and D. Shasha, "D-Over: An Optimal On-line Scheduling Algorithm for Overloaded Real-Time Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 290-299, December 1992.

[KZ93] S. Kamat and W. Zhao, "Real-Time Schedulability of Two Token Ring Protocols," *Proceedings of The Thirteenth IEEE International Conference on Distributed Computing Systems*, June 1993.

[KMZ93] S. Kamat, N. Malcolm, and W. Zhao, "The Probability of Guaranteeing Synchronous Real-Time Messages with Arbitrary Deadlines in an FDDI Network," *Proceedings of The IEEE Real-Time Systems Symposium*, December 1993.

[LBL95] A. Lazar, S. Bhonsle, and K. Lim, "A Binding Architecture for Multimedia Networks," *Journal of Parallel and Distributed Computing*, Volume 30, pages 204-216, November 1995.

[LRT92] J. P. Lehoczky and S. Ramos-Thuel, "An Optimal Algorithm for Scheduling Soft-Aperiodic Tasks in Fixed-Priority Preemptive Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 110-123, 1992.

[LSS87] J. P. Lehoczky, L. Sha, and J. K. Strosnider, "Enhanced Aperiodic Responsiveness in Hard Real-time Environments," *Proceedings of The IEEE Real-Time Systems Symposium*, 1987.

[LSTS99] C. Lu, J. Stankovic, G. Tao, and S. Son, "Design and Evaluation of a Feedback Control EDF Algorithm," *Proceedings of The Real-Time Systems Symposium*, December 1999.

[Mok83] A. K. Mok, "Fundamental Design Problems of Distributed Systems for the Hard-Real-Time Environment," *PhD Thesis*, M.I.T, Cambridge, Massachusetts 02139, 1983.

[MS00] A. S. Matveev and A. V. Savkin, "Qualitative Theory of Hybrid Dynamical Systems," Birkhauser, 2000.

[MV+96] R. C. Metzger, B. VanVoorst, L. S. Pires, R. Jha, W. Au, M. Amin, D. A. Castanon, and V. Kumar, "C3I Parallel Benchmark Suite – Introduction and Preliminary Results," *Supercomputing*, 1996.

[MZ95] N. Malcolm and W. Zhao, "Hard Real-Time Communication in Multiple-Access Networks," *Journal of Real-Time Systems*, Volume 8, Number 1, pages 35-77, January 1995

[MKZ96] N. Malcolm, S. Kamat, and W. Zhao, "Real-Time Communication in FDDI Networks," *Journal of Real-Time Systems*, Volume 10, Number 1, pages 75-107, January 1996.

[NS94] M. D. Natale and J. A. Stankovic, "Dynamic End-to-End Guarantees in Distributed Real-Time Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 216-227, December 1994.

[NS95] K. Nahrstedt and J. Smith, "The QoS Broker," *IEEE Multimedia*, Volume 2, Number 1, pages 53-67, 1995.

[NS96] K. Nahrstedt and J. Smith, "Design, Implementation, and Experiences with the OMEGA End-point Architecture," *IEEE JSAC*, September 1996.

[PH96] C. L. Phillips and R. D. Harbor, "Feedback Control Systems," Prentice Hall, 1996.

[Quo97] DARPA ITO, "Quorum," Available at http://www.ito.darpa.mil/research/quorum/projlist.html, August 1997.

[RCF97] I. Ripoll, A. Crespo, and A. G. Fornes, "An Optimal Algorithm for Scheduling Soft Aperiodic Tasks in Dynamic Priority Preemptive Systems," *IEEE Transactions on Software Engineering*, Volume 23, Number 6, pages 388-400, June 1997.

[RLLS97] R. Rajkumar, C. Lee, J. Lehoczky and D. Siewiorek, "A Resource Allocation Model for QoS Management," *Proceedings of The 18th IEEE Real-Time Systems Symposium*, pages 298-307, 1997.

[RLLS98] R. Rajkumar, C. Lee, J. Lehoczky and D. Siewiorek, "Practical Solutions for QoS-based Resource Allocation Problems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 296-306, 1998.

[RS99] D. Rosu and K. Schwan, "Faracost: An Adaptation Cost Model Aware of Pending Constraints," *Proceedings of The IEEE Real-Time Systems Symposium*, December 1999.

[RSY98] D. Rosu, K. Schwan, and S. Yalamanchili, "FARA – A Framework for Adaptive Resource Allocation in Complex Real-Time Systems," *Proceedings of The IEEE Real-Time Technology and Applications Symposium*, pages 79-84, June 1998.

[RSYJ97] D. Rosu, K. Schwan, S. Yalamanchili, and R. Jha, "On Adaptive Resource Allocation for Complex Real-time Applications," *Proceedings of The 18th IEEE Real-Time Systems Symposium*, pages 320-329, December 1997.

[RSZ89] K. Ramamritham, J. A. Stankovic, and W. Zhao, "Distributed Scheduling of Tasks with Deadlines and Resource Requirements," *IEEE Transactions on Computers*, Volume 38, Number 8, pages 1110-1123, August 1989.

[RTL93] S.Ramos-Thuel and J.P. Lehoczky, "On-line Scheduling of Hard Deadline Aperiodic tasks in Fixed-Priority Systems," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 160-171, 1993.

[RWS99] B. Ravindran, L. R. Welch, B. Shirazi, "Resource Management Middleware for Dynamic, Dependable Real-Time Systems," *Journal of Real-Time Systems*, 1999, Accepted for publication, To appear.

[Sath93] S. Sathaye, "Scheduling Real-Time Traffic in Packet-Switched Networks, *PhD Thesis*, Carnegie Mellon University, Pittsburgh, PA, June 1993.

[Spru90] B. Sprunt, "Aperiodic Task Scheduling for Real-Time Systems," *PhD Thesis*, Carnegie Mellon University, Pittsburgh, PA, August 1990.

[Sta89] J. A. Stankovic, "Decentralized Decision Making for Task Reallocation in a Hard Real-Time System, *IEEE Transactions on Computers*, Volume 38, Number 3, pages 341-355, March 1989.

[SAE88a] SAE, Aerospace Systems Division, Committee AS-2, "Linear Token-passing Multiple Data Bus," AS4074.1, Version 4.0, January 1988.

[SAE88b] SAE, Aerospace Systems Division, Committee AS-2, "High Speed Ring Bus," AS4074.2, January 1988.

[SB96] M.Spuri and G.Buttazzo, "Scheduling Aperiodic Tasks in Dynamic Priority Systems," *Journal of Real-Time Systems*, Volume 10, pages 179-210, March 1996.

[SLS88] B. Sprunt, J. Lehoczky, and L. Sha, "Exploiting Unused Periodic time for Aperiodic Service Using the Extended Priority Exchange Algorithm," *Proceedings of The IEEE Real-Time Systems Symposium*, pages 251-258, 1988.

[SLS88a] J. K. Strosnider, J. Lehoczky, and L. Sha, "Advanced Real-Time Scheduling Using the IEEE 802.5 Token Ring," *Proceedings of The IEEE Real-Time Systems Symposium*, December 1988.

[SP96] S. Skogestad and I. Postlethwaite, "Multivariable Feedback Control: Analysis and Design," John Wiley & Sons, 1996.

[SR91] J.A. Stankovic and K.Ramamritham, "The Spring Kernel: A New Paradigm for Real-time Systems," *IEEE Software*, Volume 8, Number 3, pages 62-72, May 1991.

[SRC85] J.A. Stankovic, K.Ramamritham, and S.Cheng, "Evaluation of a Flexible Task Scheduling Algorithm for Distributed Hard Real-time Systems," *IEEE Transactions on Computers*, Volume C-34, Number 12, pages 1130-1141, December 1985.

[SSL89] B.Sprunt, L.Sha, and J.Lehoczky, "Aperiodic Task Scheduling in Hard Real-time Systems," *Journal of Real-Time Systems*, Volume 1, Number 1, pages 27-60, 1989.

[SSRB98] J. A. Stankovic, M. Spuri, K. Ramamritham, and G. C. Buttazzo, "Deadline Scheduling for Real-Time Systems," Kluwer Academic Publishers, 1998.

[SWR99] B. Shirazi, L. R. Welch, and B. Ravindran, "DynBench: A Benchmark Suite for Dynamic Real-Time Systems," *Journal of Parallel and Distributed Computing Practices*, 1999, Accepted for publication, To appear.

[TBW95] K. Tindell, A. Burns, and A. J. Wellings, "Analysis of Hard Real-Time Communications," *Journal of Real-Time Systems*, Volume 9, Number 2, pages 147-171, September 1995.

[TLS96] T-S. Tia, J.W.-S. Liu, and M.Shankar, "Algorithms and Optimality of Scheduling Soft Aperiodic Requests in Fixed-Priority Preemptive Systems," *Journal of Real-Time Systems*, Volume 10, pages 23-43, January 1996.

[Uhl91] R. W. Uhlhorn, "The Fiber-Optic High-Speed Data Bus for a New Generation of Military Aircraft," *IEEE LCS*, Volume 2, Number 1, February 1991.

[VW+96] C. Volg, L. Wolf, R. Herrwich, and H. Wittig, "HeiRAT – Quality of Service Management for Distributed Multimedia Systems," *Multimedia Systems Journal*, 1996.

[WRSB98] L. R. Welch, B. Ravindran, B. A. Shirazi, and C. Bruggeman, "Specification and Modeling of Dynamic, Distributed Real-Time Systems," *Proceedings of The 19th IEEE Real-Time Systems Symposium*, pages 72 - 81, December 1998.

[WR+96] L. R. Welch, B. Ravindran, R. D. Harrison, L. Madden, M. W.Masters, and W. Mills, "Challenges in Engineering Distributed Shipboard Control Systems," *Proceedings of The Work-In-Progress Session, The Seventeenth IEEE Real-Time Systems Symposium*, pages 19-22, December 1996.

[ZRS87a] W. Zhao, K. Ramamritham, and J. A. Stankovic, "Preemptive Scheduling Under Time and Resource Constraints," *IEEE Transactions on Computers*, Volume 36, Number 8, pages 949-960, August 1987.

[ZRS87b] W. Zhao, K. Ramamritham, and J. A. Stankovic, "Scheduling Tasks With Resource Requirements in Hard Real-Time Systems," *IEEE Transactions on Software Engineering*, Volume 13, Number 5, pages 564-577, May 1987.

Financial Status

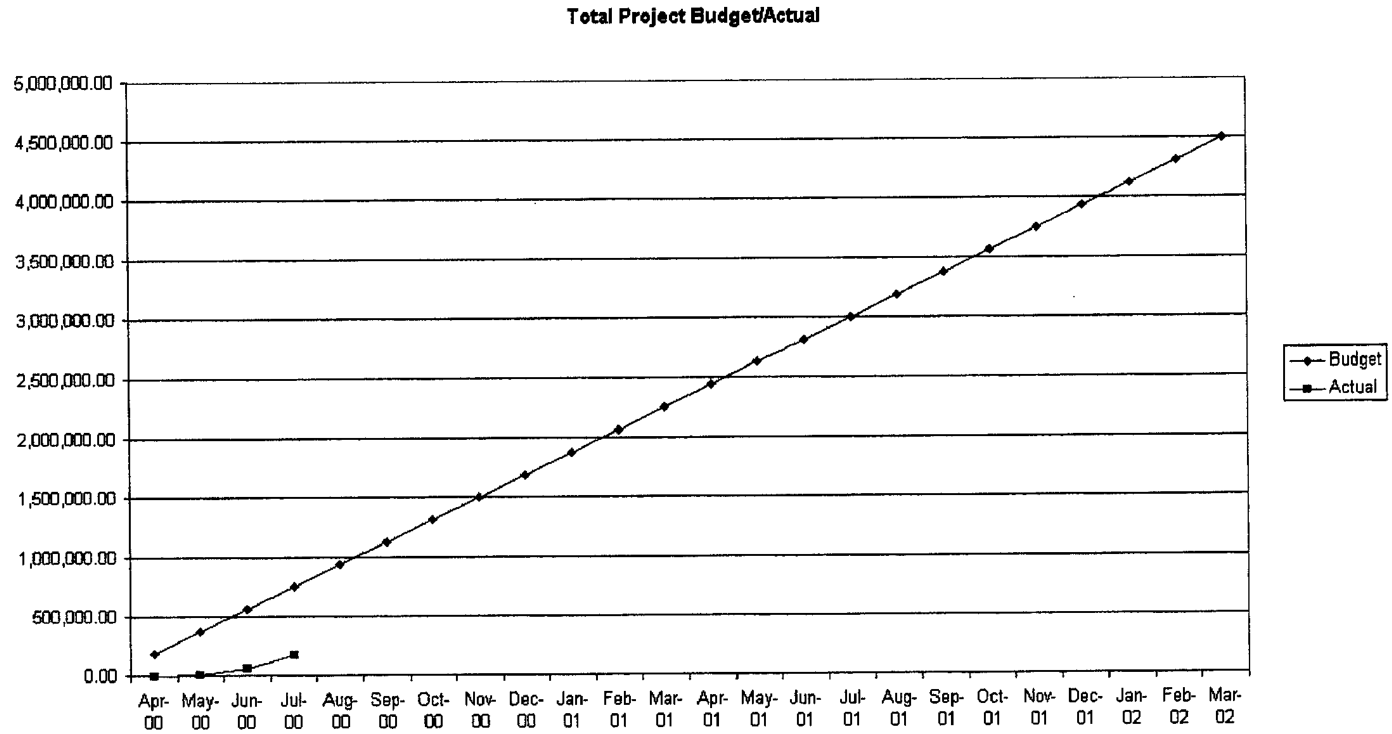

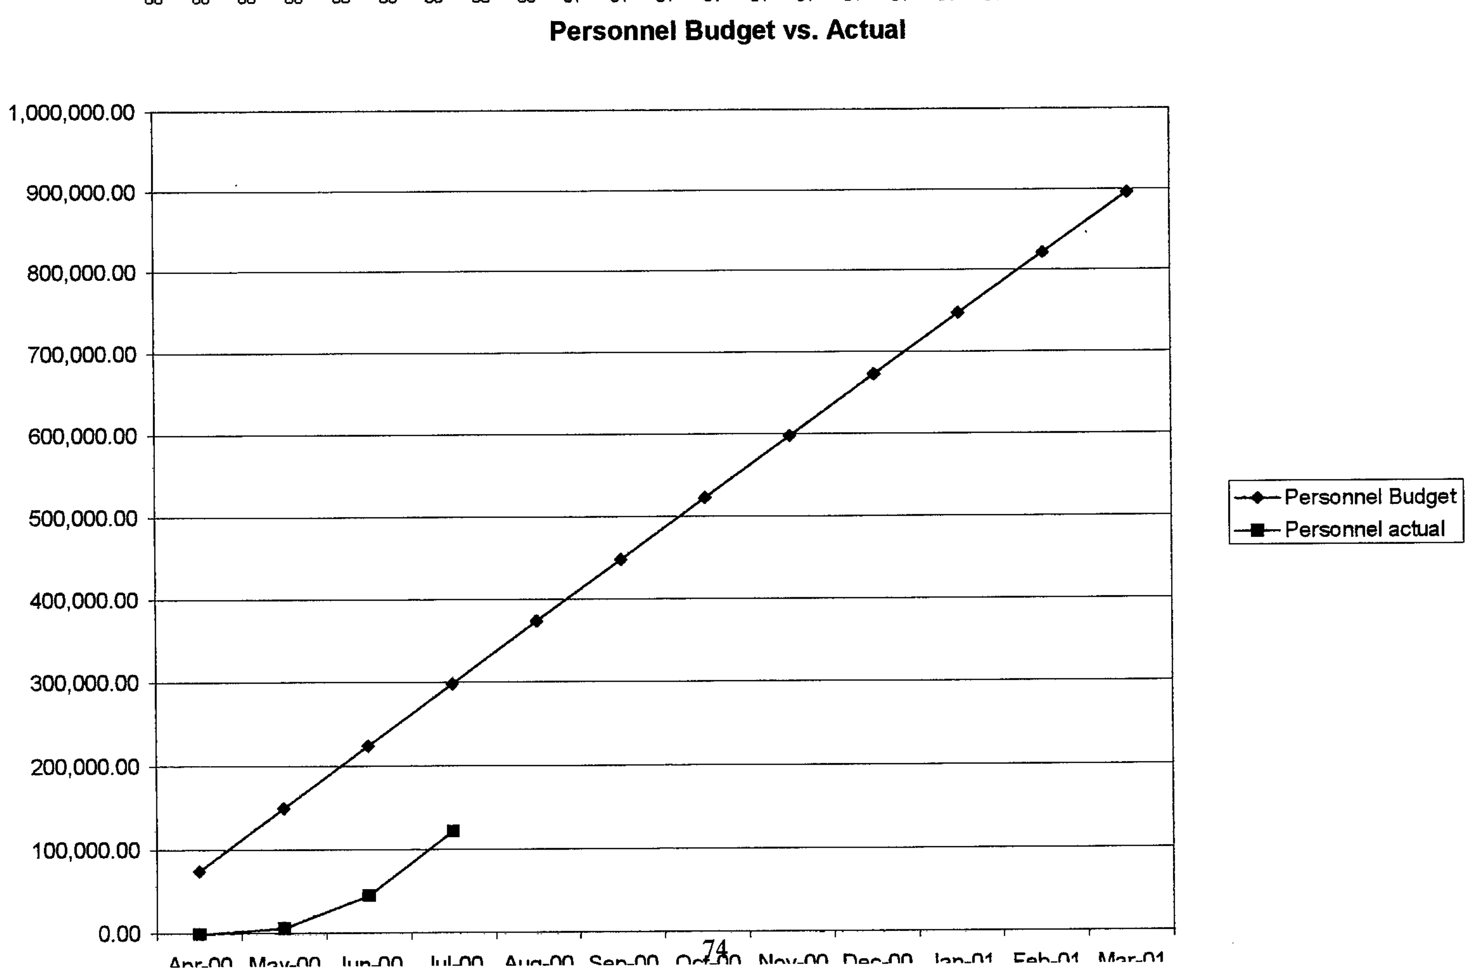

Equipment Budget vs. Actual

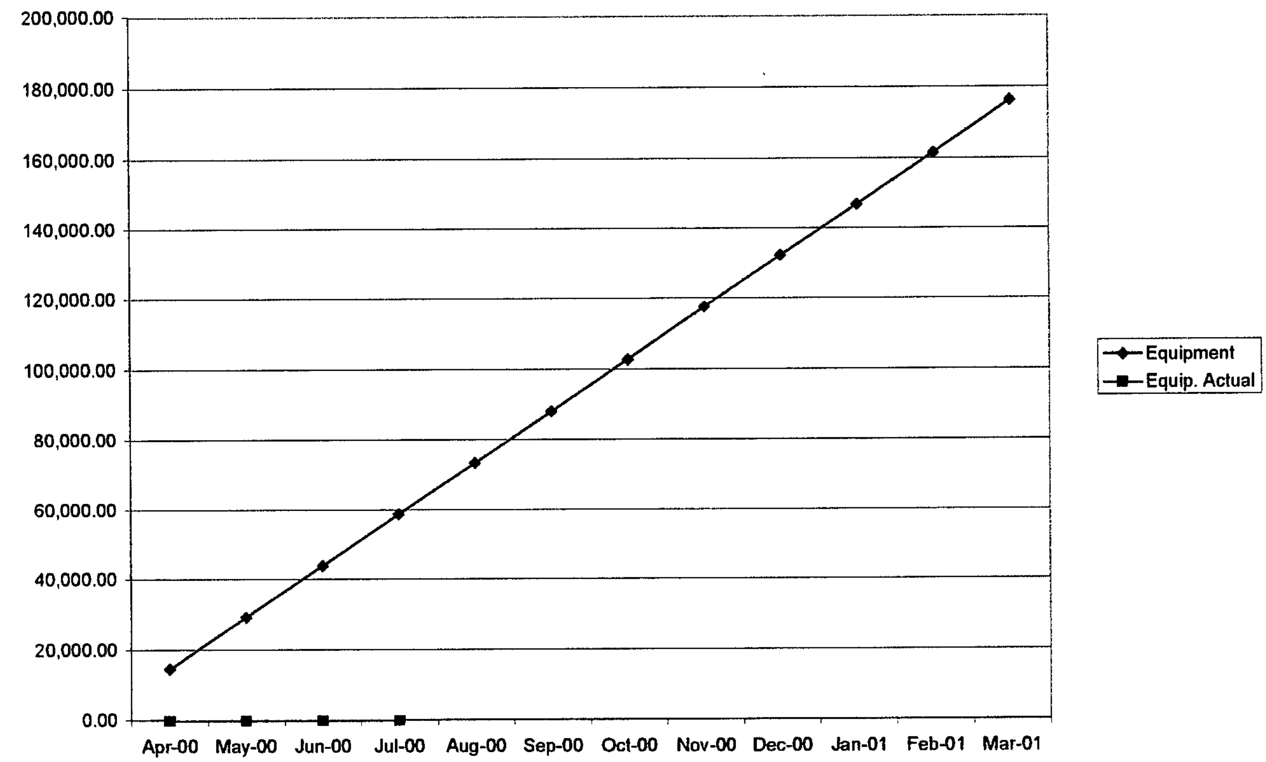

Other Direct Costs vs. Budget