Active

Project #: E-21-Z50 Cost share #: E-21-398 Rev #: 0
Center # : 10/24-6-R8546-0A0 Center shr #: 10/22-1-F8546-0A0 OCA file #:
Work type : RES
Contract#: ECS-9411846 Mod #: Document : GRANT
Prime #: Contract entity: GTRC

Subprojects ? : N CFDA: 47.041
Main project #: PE #: N/A

Project unit: ECE Unit code: 02.010.118
Project director(s):
YALAMANCHILI S ECE (404)894-2940

Sponsor/division names: NATL SCIENCE FOUNDATION / GENERAL
Sponsor/division codes: 107 / 000

Award period: 950601 to 961130 (performance) 970228 (reports)

Sponsor amount	New this change	Total to date
Contract value	140,000.00	140,000.00
Funded	140,000.00	140,000.00
Cost sharing amount		140,000.00

Does subcontracting plan apply ?: N

Title: HIGH PERFORMANCE COMPUTING FOR ENGINEERING APPLICATIONS

PROJECT ADMINISTRATION DATA

OCA contact: Jacquelyn L. Bendall 894-4820

Sponsor technical contact	Sponsor issuing office
GEORGE LEA	GRACIELA NARCHO
(703)306-1339	(703)306-1218
NATIONAL SCIENCE FOUNDATION	NATIONAL SCIENCE FOUNDATION
4201 WILSON BLVD.	4201 WILSON BLVD.
ARLINGTON, VA 22230	ARLINGTON, VA 22230

Security class (U,C,S,TS) : U ONR resident rep. is ACO (Y/N): N
Defense priority rating : N/A NSF supplemental sheet
Equipment title vests with: Sponsor GIT X

Administrative comments -
INITIATION OF PROJECT. THIS IS AN EQUIPMENT GRANT.

GEORGIA INSTITUTE OF TECHNOLOGY
OFFICE OF CONTRACT ADMINISTRATION

U

NOTICE OF PROJECT CLOSEOUT

Closeout Notice Date 03/10/97

Project No. E-21-Z50 | Center No. 10/24-6-R8546-0A0

Project Director YALAMANCHILI S | School/Lab ECE

Sponsor NATL SCIENCE FOUNDATION/GENERAL

Contract/Grant No. ECS-9411846 | Contract Entity GTRC

Prime Contract No.

Title HIGH PERFORMANCE COMPUTING FOR ENGINEERING APPLICATIONS

Effective Completion Date 961130 (Performance) 970228 (Reports)

Closeout Actions Required:	Y/N	Date Submitted
Final Invoice or Copy of Final Invoice	N	
Final Report of Inventions and/or Subcontracts	N	
Government Property Inventory & Related Certificate	N	
Classified Material Certificate	N	
Release and Assignment	N	
Other	N	

Comments
LETTER OF CREDIT APPLIES. 98A SATISFIES PATENT REPORT.

Subproject Under Main Project No.

Continues Project No.

Distribution Required:

Project Director	Y
Administrative Network Representative	Y
GTRI Accounting/Grants and Contracts	Y
Procurement/Supply Services	Y
Research Property Managment	Y
Research Security Services	N
Reports Coordinator (OCA)	Y
GTRC	Y
Project File	Y
Other	N
	N

E-21-Z50
#1

OMB Number 345-0058

NATIONAL SCIENCE FOUNDATION
4201 Wilson Blvd.,
Arlington, VA 22230

BULK RATE
POSTAGE & FEES PAID
National Science Foundation
Permit No. G-69

PI/PD Name and Address

Sudhakar Yalamanchili
Electrical & Computer Engineering
[redacted]
Atlanta GA 30332-0250

NATIONAL SCIENCE FOUNDATION
FINAL PROJECT REPORT

PART I - PROJECT IDENTIFICATION INFORMATION

1. Program Official/Org. George Lea - ECS

2. Program Name COMMUNICATIONS & COMPUTATIONAL SYSTEMS P

3. Award Dates (MM/YY) From: 06/95 To: 11/96

4. Institution and Address
[redacted]
Administration Building
Atlanta GA 30332

5. Award Number 9411846

6. Project Title
REG: High Performance Computing for Engineering Applications

** You are encouraged to submit your Final Project Report electronically **
** through the NSF FastLane home page (www.fastlane.nsf.gov). **

This Packet Contains
NSF Form 98A
And 1 Return Envelope

!SF Grant Conditions (Article 17, GC-1, and Article 9, FDP-11) require submission of a Final Project eport (NSF Form 98A) to the NSF program officer no later than 90 days alter the expiration of the ward. Final Project Reports for expired awards must be received before new awards can be made \SF Grants Policy Manual Section 677).

elow, or on a separate page attached to this form, provide a summary of the completed projects and technical information. Be ure to include your name and award number on each separate page. See below for more instructions.

PART II - SUMMARY OF COMPLETED PROJECT (for public use)

---

ne summary (about 200 words) must be self-contained and intelligible to a scientifically literate reader. Without restating the oject title, it should begin with a topic sentence stating the project's major thesis. The summary should include, if pertinent the project being described, the following items:

The primary objectives and scope of the project
The techniques or approaches used only to the degree necessary for comprehension
The findings and implications stated as concisely and informatively as possible

Please see attached

PART III - TECHNICAL INFORMATION (for program management use)

---

st references to publications resulting from this award and briefly describe primary data, samples, physical collections, ventions, software, etc. created or gathered in the course of the research and, if appropriate, how they are being made available the research community. Provide the NSF Invention Disclosure number for any invention.

:ertify to the best of my knowledge (1) the statements herein (excluding scientific hypotheses and scientific opinion) are true and complete, and) the text and graphics in this report as well as any accompanying publications or other documents, unless otherwise indicated, are the original ork of the signatories or of individuals working under their supervision. I understand that willfully making a false statement or concealing a aterial fact in this report or any other communication submitted to NSF is a criminal offense (U.S. Code, Title 18, Section 1001).

Principal Investigator/Project Director Signature	**Date**

IMPORTANT:
MAILING INSTRUCTIONS
Return this *entire* packet plus all attachments in the envelope attached to the back of this form. Please copy the information from Part I, Block I to the *Attention block* on the envelope.

Project Summary

The explosive growth in inter-processor communication technologies and the continuing improvement in the cost performance of processor and memory technologies is constantly re-defining architectural solutions to computationally intensive engineering applications. The objective if this project is to create a computing testbed with a cluster of commodity multiprocessor and uniprocessor graphics workstations interconnected by state of the art high speed communications networks to existing parallel supercomputers. This testbed will provide a feasible hardware basis to pursue the development and application of new computational paradigms for large scale engineering applications. In particular, we foresee the ability to harness such high speed networks of specialized graphics and parallel machines such that components of applications can utilize distinct parallel machines while results from these distinct platforms can be visualized, integrated and manipulated in real-time on workstations. This integration opens new avenues for solutions to solving many large engineering problems by enabling integrated, interactive, computing environments for engineering applications.

Summary of NSF Research Equipment Grant Support

Title: High Performance Computing for Engineering Applications
Award: $280,000 (includes $140,000 matching from Georgia Tech) Grant Number ECS-9411846
Duration: June 1st, 1995 through November 30th, 1996

Project Summary:

The NSF support has been used to establish a state of the art computational facility. This facility provides for the integration of parallel supercomputers and state of the art workstations using commodity high speed interconnects. This emerging technology base is supporting a wide range of research projects. This NSF Research Engineering equipment grant has impacted Georgia Tech far beyond its actual funding level for two reasons: (1) due to the simultaneous awards of a substantial NSF Research Infrastructure grant and a NSF CISE research equipment grant and (2) due to new research awards from other agencies, including DARPA. The equipment purchased specifically under this grant is shown below.

Qty	Description	Cost
5	Sun Ultra Enterprise 2 Model 1170(167 Mhz, 64 Mb, 2.1 Gbyte, CDROM)	62,985.00
5	Sun Ultra CPU Model for 167 MHZ Model 1170)	15,000.00
2	32 MB memory expansion	1,620
4	Sun 64 MB expansion	5,900
8	Sun Fast Ethernet SBus Adaptors	3752.40
1	Sun UltraServer Model 170E (167 Mhz, 64 MB, 2.1 Gbyte)	14,157.05
1	SCSI II Bus Adaptor	705.05
1	Sun 64 MB memory expansion	2271.50
1	Sun CDROM Drive	236.00
1	Sun Internal Floppy Drive	88.50
1	High Speed Serial Interface	1,342.25
8	Sun UltraServer Model 170	96736.40
1	Sun 4.2 Gbyte SCSI II disk	1224.25
4	Sun Ultra 1 Model 170E)	66,028.20
	Towards the ATM infrastructure for the test bed	11,533.00
	Total	**283,579.55**

The total funded amount was $280,000 (including $140,00 in matching funds from Georgia Tech). The difference above was funded from other sources. The projects of the PIs that are actively utilizing the cluster are described in the following. This does not include a host of other projects by faculty and graduate students in engineering and computer science that also utilize the facilities.

Adaptive Resource Management

Many applications are characterized by heterogeneous HPC computing platforms, sharing the computing resources among application subsystems, and dynamic variability in the performance demanded of these various subsystems. Moreover, these applications are reactive in nature, since they may have to respond to changes in the external environment. Often, the run-time behavior of these applications is heavily data-dependent, governed by input data parameters. Consequently their computing resource requirements may vary considerably during execution, and are for the most part statically unpredictable. The proposed solution is the dynamic adaptation of resources to changing application requirements.

Dynamic adaptation of application programs to changing resource requirements is predicated on the ability to obtain reliable and accurate information about resource usage at run-time, i.e., some mechanism for on-line monitoring. The monitored information must be maintained on-line for analysis to determine if and when it is beneficial to re-allocate resources. Re-allocation decisions must be based on models of application and system behavior and prevent thrashing - a state where the system is constantly re-allocating resources. Finally, once the decision to re-allocate is made, fast algorithms are necessary to compute a new assignment on-line. Solutions to these problems rapidly become intractable as we increase the number of distinct types of resources to be managed, and/or the number of applications to be simultaneously managed. Our long term strategy is to first address these complex issues in the context of the execution of a single parallel/distributed application. Our goal is to build a framework of a distributed environment which contains all of the major components of any dynamically adaptive system, and to demonstrate this framework individually on sample, dynamically adaptive applications. Successful completion provides the insight, algorithms, and software components that we can build on to address the more complex issues of dynamic resource management in multi-resource, multi-application environments. A first version of this environment was built under ARPA funding, and the development is continuing under DARPA funding for two additional years.

High Performance Applications

This NASA-funded project is developing a high performance, parallel atmospheric modeling code, in conjunction with researchers from the School of Earth and Atmospheric Sciences at Georgia Tech, and funded by a joint NASA grant. This code exhibits the following novel attributes:

- i) It is interactively executable, where researchers can 'steer' the ongoing computation by manipulation of selected parameters. Currently, such on-line steering may be performed with respect to the dynamic computations (i.e., constituent transport) being performed, specifically addressing alternative settings for vertical constituent motion.
- It has high computational and communication requirements, where the on-line monitoring and visualization of selected output data (and corresponding observational data) must be performed such that little or no performance penalties arise. This requires significant infrastructure support in terms of the transport and processing of monitoring, steering, and visualization information, especially when multiple machines are involved in these processes.
- The steering interface utilizes data visualizations, rather than the typical code module or program visualizations employed by other efforts addressing program steering (or by debuggers). As a result, end users may inspect and steer programs in terms with which they are familiar.

The project has been focussing on interactivity of the application with a moderate number of distributed end users, as well as with porting the application from a shared memory platform to mixed shared and distributed memory platforms:

Interactive and Distributed Simulation

This project is developing high performance parallel and distributed simulation tools to support parallel and potentially interactive simulations of large-scale telecommunication networks. Over the past nine months, our parallel simulator (Georgia Tech Time Warp, or GTW) that was originally designed to execute on shared memory multiprocessor platforms has been adapted to execute on the equipment purchased with NSF funds. Experiments have been performed executing simulations of wireless networks on the Power Challenge, Sun Ultrasparc workstations, and SGI Indy workstations.

Using the NSF equipment, experimental research completed over the first year of the grant focused on (1) partitioning ATM network simulations for parallel execution, (2) developing new dynamic load balancing algorithms to enable Time Warp simulations to execute ``in background'' on the NOW platform (a system that performs dynamic load management is now operational), (3) developing techniques to efficiently perform incremental state saving in Time Warp, transparent to the simulation application, and (4) developing a variety of techniques to enable efficient execution of Time Warp programs.

This research is already beginning to have a substantial impact outside the Georgia Tech community. Within the last year MITRE has developed a commercial air traffic simulation on top of GTW that allows them to perform simulations that previously took 1.5 hours using a SIMSCRIPT-based simulator in only one or two minutes using GTW on a multiprocessor Sparc, while producing comparable simulation results. MITRE is planning to provide this tool to the FAA for use in air traffic management. Our research has also heavily influenced the *High Level Architecture* effort in DoD in defining a common distributed simulation infrastructure for all simulations in the DoD.

Technology Developments

Several research projects are developing underlying, enabling technologies to support high performance applications on NOW machines and on parallel supercomputers. Three are primarily concerned with ``middleware'' software and technologies that are directly utilized by user applications. The first project is studying the dynamic monitoring, adaptation, and interactive steering of high performance computations for on-line control of ``virtual laboratory instruments'' and for ``what-if?'' experimentation with complex simulation models by distributed laboratory users. The second project is exploring the efficient execution of simulation programs, especially discrete-event simulations of telecommunication network models, on multi-granular compute servers (servers containing both tightly-coupled multiprocessors and loosely-coupled workstations). The third project is exploring the collaboration technologies required in distributed laboratories.

The fourth project is concerned with the distributed systems and telecommunication networking technologies underlying distributed laboratories. The emphasis in distributed systems research is on support for shared state in multi-granular parallel and distributed computing environments. Networking research is concerned with the communications in such environments, including multicast protocols and protocol configuration.

Collaborative Systems Support

This project has concentrated on three areas. First, work has been done on generalized techniques for shared visualizations on the Open Inventor platform in conjunction with an Atmospheric Science application. This software is designed to provide a base capability to support shared views of visualization between distributed participants.

Second, work has been done on new algorithmic techniques to support multiple coupled views in a distributed setting. This work has been aimed at extending existing ``constraint-based" techniques for user interface implementation, which have been very successful in the single user domain, to a distributed setting. Although constraint techniques provide very powerful, high-level abstractions for implementation of several aspects of user interfaces, until now the efficient update algorithms needed to support them have been limited to single user sequential systems. Work done this year has developed a new distributed algorithm for efficient (incremental and lazy) update of constraints distributed across a network. This result extends a prior optimal algorithm in the sequential domain and preserves its performance characteristics (and optimality when used sequentially). In addition, the new algorithm has been proven to provide an optimal level of concurrency (in keeping with preserving causal semantics of updates).

The final area of work in collaborative systems has been in systems to support awareness in distributed work groups. There are a large number of factors that make working in a distributed setting more awkward than working in a co-located group. Many of these reasons are informal and relate to awareness of one's colleagues. Awareness serves as a backdrop and catalyst upon which more explicit communications is built. This work has looked at several specific techniques for promoting awareness in distributed work groups. It has been particularly concerned with tradeoffs that occur between transmission of awareness information and privacy (for the sender), and between awareness and disruption (for the receiver).

Work next year will continue to concentrate on the same three areas: specific tools to be applied to collaboration in the atmospheric sciences application, use of constraint systems to support distributed user interfaces, and new techniques to support awareness in distributed work groups.

Indigo

The Indigo project explores user level support for building distributed state sharing techniques in workstation cluster environments. In particular, Indigo enables the implementations of both distributed shared memory as well as distributed shared abstractions that are defined and maintained in an application specific way. A paper describing the Indigo system was selected as the best one in the Conf. of High Performance Distributed Computing and an extended version of this paper was chosen for publication in Concurrency: Practice and Experience.

Flex

The flex project explores scalable implementations of distributed object systems. In particular, it investigates object caching issues and the performance benefits made possible by them. Several papers describing this work have been published. One appeared at the Conference on Object-oriented Technologies and Systems and was selected as one of four papers that were recommended for publication in the Computing Systems Journal.

Optimizing Inter-Processor Communication

The goal of the proposed research is the development, implementation, and evaluation of primitives to minimize, and/or hide, the inter-processor communication latency in parallel/distributed programs. We are interested in user level and compiler level access to a set of primitives that can be utilized to exploit locality in inter-processor communication to fine tune the communication performance of these programs. These primitives are general enough to be applicable across a wide range of communication substrates and enable optimizations for low latency, real-time guarantees, and reliable message delivery.

The proposed program has a strong experimental systems component based on the newly established computing facilities. As a result, we will focus on quantitative evaluation of candidate approaches using a suite of programs drawn distinct application domains. We expect to produce application specific solutions as well as identify optimizations that are portable across applications and interconnection networks. This work complements existing approaches to the development of low latency message layers, e.g., Active Messages, Fast Messages, PORTS, etc. Furthermore, the techniques can be used across the range of commodity interconnects that are becoming available for cluster computing. Thus, these efforts seek to leverage foreseeable advances in cluster communication networks, and focuses on providing effective implementations of communication services for parallel/distributed programs.

Summary Data on Project Personnel

Since this is a research equipment grant, no personnel were supported on grant funds.