

AD-A118 954

GEORGIA INST OF TECH ATLANTA SCHOOL OF INFORMATION A--ETC F/G 9/2
THE GEORGIA INSTITUTE OF TECHNOLOGY RESEARCH PROGRAM IN FULLY D--ETC(U)
AUG 82 N00014-79-C-0873

UNCLASSIFIED

NL

1 04 1

AD-A
18954



END
DATE
FILMED
10-82
DTIC

12

AD A118954



DTIC
ELECTE
SEP 0 1 1982

School of

E

Information and Computer Science

GEORGIA INSTITUTE OF TECHNOLOGY

A UNIT OF THE UNIVERSITY SYSTEM OF GEORGIA

82 08 18 016

This document has been approved
for public release and sale; its
distribution is unlimited.

DTIC FILE COPY

12

THE GEORGIA INSTITUTE OF TECHNOLOGY
RESEARCH PROGRAM IN
FULLY DISTRIBUTED PROCESSING SYSTEMS

Quarterly Progress Report Number 11
1 March, 1982 - 31 May, 1982

August, 1982

Supported by

Office of Naval Research (ONR)
Contract: N00014-79-C-0873
GIT Project: G36-643

U.S. Air Force Rome Air Development Center (RADC)
Contract: F30602-78-C-0120
GIT Project: G36-654

U.S. Air Force Rome Air Development Center (RADC)
Contract: F30602-81-C-0249
GIT Project: G36-659

U.S. Army Research Office (ARO)
Contract: DAAG29-79-C-0155
GIT Project: G36-638

National Science Foundation (NSF)
Contract: MCS-7924370
GIT Project: G36-652

0 1 1982

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, Georgia 30332

1. INTRODUCTION

This is the Eleventh Quarterly Progress Report prepared covering the Georgia Tech Research Program in Fully Distributed Processing Systems (FDPS).

a. Program Description.

The Georgia Tech Research Program in Fully Distributed Processing Systems is a comprehensive investigation of data processing systems in which both the physical and logical components are extremely loosely coupled while operating with a high degree of control autonomy at the component level. The definition of the specific class of multiple computer systems being investigated, and the operational characteristics and features of those systems is motivated by the desire to advance the state-of-the-art for that class of systems that will deliver a high proportion of the benefits currently being claimed for distributed processing systems. The scope of individual topics being investigated under this program ranges from formal modeling and theoretical studies to empirical examinations of prototype systems and simulation models. Also included within the scope of the program are areas such as the utilization of FDPS's and their interaction with management operations and structure.

b. Program Support.

The principle support for the program is a Selected Research Opportunity contract from the Office of Naval Research; however, there are a number of other sources of funding which also support the program. A list of the currently active contracts and grants is given below.

Title: "Research on Fully Distributed Data Processing Systems"
Funding Agency: Office of Naval Research (ONR)
Contract Number: N00014-79-C-0873
GIT Project No.: G36-643/336
Principle Investigator: Philip H. Enslow, Jr.

Title: "Evaluation of Distributed Control Models"
Funding Agency: U.S. Air Force Rome Air Development Center (RADC)
Contract Number: F30602-78-C-0120
GIT Project No.: G36-654
Principle Investigator: Philip H. Enslow, Jr.

Title: "System Support Capabilities for Fully-Distributed /
Loosely-Coupled Processing Systems"
Funding Agency: U.S. Air Force Rome Air Development Center (RADC)
Contract Number: F30602-81-C-0249
GIT Project No.: G36-659
Principle Investigator: Philip H. Enslow, Jr.

Title: "Theory of Systems of Asynchronous Parallel Processors"
Funding Agency: U.S. Army Research Office (ARO)
Contract Number: DAAG29-79-C-0155
GIT Project Number: G36-638/332
Principle Investigator: Nancy A. Lynch

Title: "Complexity and Computability for Distributed Data Bases"
Funding Agency: National Science Foundation (NSF)
Contract Number: MCS-7924370
GIT Project Number: G36-652/340
Principle Investigator: Nancy A. Lynch

2. ORGANIZATION AND STAFFING

Faculty

DeMillo, Richard A. - Professor
Enslow, Philip H. Jr. - Professor
Griffeth, Nancy A. - Assistant Professor
Jensen, Alton P. - Professor
LeBlanc, Richard J. - Assistant Professor
Livesey, Jon - Assistant Professor
Lynch, Nancy A. - Associate Professor (currently visiting at MIT)
McKendry, Martin S. - Assistant Professor
Miller, Raymond - Professor
Underwood, William - Assistant Professor

Staff

McDonell, Sharon - Administrative Secretary
Myers, Jeanette - Research Scientist
Pinion, Nancy - Part-time Secretary
Mongiovi, Roy - Research Technologist I

Students

There are approximately 30 students working on various projects in the FDPS Research Program. Of these, 12 are in the Ph.D. program, and 5 are preparing their M.S. Thesis on topics in FDPS.

3. CURRENT RESEARCH PROJECTS

The specific research projects have been organized into the major areas identified in the basic program proposal.

A. Theoretical and Formal Studies

A.3 Reliable Systems
A.4 Time Performance of Distributed Systems
A.5 Audit Algorithms
A.6 Ticket Systems
A.9 Theory of Distributed Databases
A.16 Stochastic Synchronization

- A.17 Research Allocation in a Failure-Prone Environment
- A.18 Multilevel Atomicity
- A.19 Formal Semantics and Specification of Distributed Systems
- A.20 Nested Transactions with Aborts

B. Physical Interconnection and Networking

- B.2 Local Networking in Fully Distributed Processing Systems

C. Distributed Operating Systems

- C.4 Local Operating System
- C.5 Communications Support for Distributed Systems
- C.8 Distributed Software Tools
- C.9 Command Languages in an FDPS
- C.10 Distributed Operating System Implementation

D. Distributed Data Bases

- D.1 Concurrency Control in Distributed Database Systems
- D.3 Implementation of the Audit Algorithm
- D.4 User Interfaces to Database Systems

E. Fault-Tolerance

F. Special Hardware to Support FDPS

G. Application of Distributed Processing

H. System Design Methodologies

- H.2 Coordinating Large Programming Projects

I. System Utilization

- I.1 A Language for Distributed Programming
- I.2 System Implementation Language Development
- I.3 Experiments with a Distributed Compiler

J. Security

- J.1 Process Structures
- J.2 System Security

Accession For	
NTIS	X
DDIC	
Under	
J. 1	
<i>on file</i>	
Ext.	
Dist.	
Avail.	
Dist	Special
A	

DTIC
COPY
INSPECTED
2

K. System Management

L. Evaluation and Comparison

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)

M. FDPS Testbed

M.1 Establishment of FDPS Testbed Facility

M.3 Fully Distributed Operating System Simulation Testbed

M.4 Interactive Monitoring of Distributed Programs

4. SUMMARY OF PROGRESS

A.3 Reliable Systems (Lynch, Fischer, Fowler, Merritt)

A paper, "Cryptographic Protocols", was presented at SIGACT 82.

A.4 Time Performance of Distributed Systems (Lynch, Fischer, Lazowska, Schönage)

No significant progress to report.

A.5 Audit Algorithms (Griffeth, Fischer, Lynch)

No significant progress to report.

A.6 Ticket Systems (Fischer, Griffeth, Guibas, Lynch)

A draft of a paper, "Analysis of a Network Resource Allocation Algorithm", has been prepared for presentation at the June ACM workshop on probabilistic complexity.

A.9 Theory of Distributed Databases (Lynch, Griffeth)

Initial work on a new formulation of concurrency control, providing a more basic definition of correctness than earlier work, was carried out.

A.16 Stochastic Synchronization (DeMillo, R. Miller, Lipton)

No significant progress to report.

A.17 Resource Allocation in a Failure-Prone Environment (Fischer, Lynch, Burns, Borodin)

No significant progress to report.

A.18 Multilevel Atomicity (Lynch)

No significant progress to report.

A.19 Formal Semantics and Specification of Distributed Systems (Lynch, Stark)

The model was used to specify and prove correct a simple arbiter algorithm. Both safety and fairness properties are easily expressed and proved.

A.20 Nested Transactions with Aborts (Lynch, Liskov)

A paper, "Concurrency Control for Resilient Nested Transactions", was written. It defines the semantics of resilient nested transactions, and uses the framework for proving correctness of a version of Moss's locking algorithm (the implementation of nested transactions used in the Argus system).

B.2 Local Networking in FDPSs (Enslow, Myers, Brundette, Hutchins, Arius)

The Net/One Network Configuration Facility (NCF-2) has been successfully installed and performs according to published specifications. A Whitesmith C compiler has been received for the NCF-2 to be used for local software development.

C.4 Local Operating System (Livesey, Fukuoka)

A paper by Fukuoka presenting a comprehensive taxonomy of IPC facilities based on the semantic aspects associated with the IPC has been submitted to ACM Transactions on Programming Languages and Systems (TOPLAS). This paper has been forwarded by TOPLAS for review by ACM Computing Surveys. The paper includes a survey of IPC mechanisms in fifteen existing and proposed programming languages and systems for distributed processing.

C.5 Communications Support for Distributed Systems (Enslow, Skowbo)

The formal project proposal is essentially complete and is currently being reviewed, revised, and extended. The specification and design of tools for an experimental evaluation of proposed alternatives for communications support is also in progress.

C.8 Distributed Software Tools (Myers, Livesey, Hopkins, Lee, Fox)

Hosts, which initiate additional user processes to perform functions on local and remote hosts, have been implemented but require some modification as to the "initial" state of the processes. Since we have found the overhead of "logging in" another process to be high, we are proposing a special network-server process to be used instead of an additional user process when a specific user environment is not required.

The domain-structured file system has been implemented and awaits testing and incorporation into DSWT. A paper describing the file system was presented at the ACM Southeastern Regional Conference, April, 1982, in Knoxville, Tennessee.

An IPC facility has been incorporated into the SWT I/O Subsystem. Currently, this facility requires that virtual circuit connections be made to specific ports on specific systems. This is being modified so that port numbers and system names will be completely transparent to the user program.

C.9 Command Languages in an FDPS (Badre, Myers, Greene)

During the past quarter, work on characterizing the user, his requirements, and his needs has continued. User activity was monitored throughout the quarter resulting in a data set that spans one year. This information has been gathered in an effort to understand the differences, if any, between "guru" or sophisticated users and "novice" or unsophisticated users. It is hoped that this will result in a definition of the extent of function a user will use in an interface. Specifically, the processing task has included calculating frequency counts of the commands used by both groups and sampling of the terminal sessions of members of both groups of users. Again, the purpose of the analysis of this data is to understand what is needed in an interface--to understand the difference in functional need and use of a computer by a user, whether a sophisticated or unsophisticated user.

C.10 Distributed Operating System Implementation (McKendry, Allchin, Thibault)

The CLOUDS project is constructing a distributed operating system for a group of workstations connected by a high-speed local-area network. The fundamental aims of the project are to provide a testbed for evaluation of algorithms developed within the FDPS program and to evaluate structural concepts for distributed operating systems.

The overall structure of the operating system has been defined, the interprocess communication mechanism has been designed and documented in an internal working paper, and considerable progress has been made in the study of data consistency requirements for the system.

Research is currently concentrating on data management and resource management. An implementation of the kernel for PERQ workstations is underway, and preparation of a conference paper is in progress.

D.1 Concurrency Control in Distributed Database Systems (Griffeth, Livesey, Lynch)

Correction to the March 1982 Quarterly Progress Report: The project, "Distributed Database Algorithms", reported under Project D.4, "User Interfaces", should have been reported under this heading.

Work is underway on analyzing preliminary hypotheses and developing the simulation. Currently, special attention is being paid to establishing results which will simplify the simulation (e.g., read/write mix has no effect, the same ratio of transaction load to database size has the same effect regardless of absolute database size, a database with a skewed distribution of frequencies of data-item requests behaves like a smaller database with a uniform distribution). Combinatoric and queuing-theoretic techniques are being used.

D.3 Implementation of the Audit Algorithm (Griffeth, Livesey, Lynch)

No significant progress to report.

D.4 User Interfaces to Database Systems (Griffeth)

Pilot studies have been run using a "cops-and-robbers" game and a registration problem to test the effectiveness of the relational calculus in a problem-solving situation. Preliminary indications are that each subject will required more than three hours.

H.2 Coordinating Large Programming Projects (Enslow, Smith)

The proposed metric for effectiveness of communication during large software development has been refined. A major focus has been the determination of the criteria that such a metric must meet. A series of experiments has been planned to evaluate the proposed metric.

I.1 A Language for Distributed Programming (LeBlanc, Maccabe, Mongiovi)

Intensive work on implementation and evaluation has continued this quarter. Work has also been done to further develop features for handling process and processor failures. Maccabe's Ph.D. thesis is in preparation.

I.2 System Implementation Language Development (LeBlanc, McKendry, Wilkes)

Work continues on the Pascal compiler which is intended to be the basis of our implementation.

I.3 Experiments with a Distributed Compiler (LeBlanc, J. Miller)

Further experiments have been conducted to study the effects of buffering messages. This concept appeared as a major factor in our earlier studies. A journal paper based on these experiments is in preparation.

J.1 Process Structures (DeMillo, Lipton, R. Miller, Merritt, Thomas)

No significant progress to report.

J.2 System Security (Livesey, Davida, DeMillo)

No significant progress to report.

L.1 Simulation of Distributed Algorithms (Griffeth, Lynch)

Hardware and software selection are underway. The initial simulations will be run on a PRIME 550, VAX 780, or CYBER 170/760. Available simulation languages include GPSS, SIMULA, and SIMSCRIPT. The ticket system simulation written in FORTRAN includes event list management and statistical routines.

M.1 Establishment of FDPS Testbed Facility (Myers, Mongiovi, Fox)

Test facilities are being developed in conjunction with development work in Projects C.8 and C.10.

M.3 FDOS Simulation Testbed (LeBlanc, Saponas, Myers)

No significant progress to report.

M.4 Interactive Monitoring of Distributed Programs (LeBlanc, Robbins)

With the development of distributed computing systems, it becomes necessary to provide programmers with appropriate tools to effectively utilize them. The first required tool is a programming language which supports the design and construction of distributed programs. One such language, called PRONET, has been developed as part of the FDPS Research Program (see Project I.1). The newly initiated project described here is concerned with the next required tool: a monitor which will allow programmers to examine the behavior of distributed programs, interactively. Monitoring a distributed program presents significant new challenges, since the "state" of such a program involves information about an arbitrary number of processes running on a number of machines. This problem is far more complex than monitoring a typical program on a single machine, in which case, all of the state information is in a single address space. The desired monitoring capability should be generalized, so that it can be used both for debugging and performance analysis. Such generality is a reasonable goal, since the most important aspect of monitoring distributed programs will concern the collection of data about the interactions among program parts, a task that is independent of the use intended for the data.

Initial work on the design of a monitoring system has begun and proposals have been prepared and submitted in order to obtain support for this work.

5. TRAVEL RELATED TO THE FDPS PROGRAM

Date of Trip: 18-20 May, 1982

Individual(s) Traveling: Philip Enslow

Itinerary: Rome Air Development Center

Contact: Tom Lawrence

Purpose: Participate in RADC Distributed Processing Technology Exchange Meeting

Date of Trip: May, 1982

Individual(s) Traveling: Nancy Lynch

Itinerary: Marina del Rey

Contact:

Purpose: Attended Symposium on Principles of Database Systems. Presented the paper, "Multilevel Atomicity".

6. VISITORS

No visitors to report.

7. PUBLICATIONS

Author(s): T. Allen Akin and Richard J. LeBlanc

Title: The Design and Implementation of a Code Generation Tool

Type: journal paper

Status: accepted for publication in Software - Practice and Experience

Author(s): Hirobumi Fukuoka

Title: Interprocess Communication Facilities for Distributed Systems: A Taxonomy and a Survey

Type: journal paper

Status: submitted to ACM Transactions on Programming Languages and Systems; subsequently forwarded to ACM Computing Surveys for review.

GIT Number: GIT-ICS-82/06

Author(s): N.J. Livesey

Title: Extending File Systems to Distributed Systems

Type: conference paper

Status: presented at the ACM April Southeastern Regional Conference

GIT Number: GIT-ICS-82/07

Publ. Date: April, 1982

Author(s): Richard DeMillo, Nancy Lynch, Michael Merritt

Title: Cryptographic Protocols

Type: conference paper

Status: presented

GIT Number: TBA

Publ. Date: May, 1982

