

AD-A116 955

GEORGIA INST OF TECH ATLANTA SCHOOL OF INFORMATION A--ETC F/6 9/2
THE GEORGIA INSTITUTE OF TECHNOLOGY RESEARCH PROGRAM IN FULLY D--ETC(U)
MAR 82 N00014-79-C-0873

UNCLASSIFIED

NL

1 of 1

00000000



END
DATE
FILMED
8-82
DTIC

AD A116955



School of
Information and Computer Science

**GEORGIA INSTITUTE
OF TECHNOLOGY**

A UNIT OF THE UNIVERSITY SYSTEM OF GEORGIA

DISTRIBUTION STATEMENT A
Approved for public
Distribution

82 05 25 195

1377

12

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

Quarterly Progress Report Number 10
1 December, 1981 - 28 February, 1982

March, 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

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

DTIC
SELECTED
JUL 12 1982
H D

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

1. INTRODUCTION

This is the Tenth 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

Accession For	
NTIS	GR&I
DTIC	TDS
Unannounced	
Justification	
By <i>EL-112/Jan</i>	
Distribution/	
Availability Codes	
Dist	Avail and/or
	Special
A	



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
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.1 Decentralized and Distributed Control
 - C.2 Resource Allocation and Work Distribution in an FDPS
 - 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

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

4. SUMMARY OF PROGRESS

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

A new Byzantine Generals algorithm was devised, with better performance than previously known algorithms, in terms of amount of message traffic. The number of rounds is also considerably smaller than in previous algorithms with good message performance.

The lower bound proof for number of rounds in an environment allowing arbitrary authentication capabilities, has been considerably refined and clarified.

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

No significant progress to report.

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

The paper, "Global States of a Distributed System", was accepted for a special issue of Transactions on Software Engineering (May, 1982). Also, work is in progress on analysis of some significant special-case systems.

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

The ticket system simulation has been modified to reduce the variance in results and generalized to allow returns and reallocation of tickets. Initial tests have been run. Further modifications are planned to test limiting cases.

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

No significant progress to report.

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)

Some work is being done to integrate the multilevel atomicity concept with related work on nested transactions.

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

Equivalence has been proved for three natural definitions for the class of possible behaviors of distributed systems.

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

Formal correctness conditions for nested transaction systems, and correctness proofs for implementations using locking, are in process of development.

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

A 12-1/2 hour video Net/One Programming class created by Ungermann-Bass was taken by members of this project. This course described in detail Net/One operation and the recently released software development support package.

We are waiting for the delivery of a new component for the local network, an NCF-2, which replaces the MCZ. (The MCZ is no longer supported by Ungermann-Bass). The NCF-2 includes the software development support package, support for the C programming language and, due to an IEEE-448/79 interface to the local network, reduces downloading time from approximately 90 seconds to 19 seconds per board.

C.1 Decentralized and Distributed Control (Enslow, LeBlanc, Saponas)

A technical report concluding Tim Saponas's research in this area is ready for distribution. This report analyzes in detail the factors contributing to the particular evaluations of each of the control models covered in a previously distributed technical report.

C.2 Resource Allocation and Work Distribution in an FDPS (Enslow, Sharp)

A technical report concluding Don Sharp's research in this area is ready for distribution. Results of simulation experiments indicated that the best criterion for work distribution is to minimize communication between processes represented as nodes of a task graph.

C.4 Local Operating System (Livesey, LeBlanc, McKendry, Myers, Allchin, Fukuoka, Maccabe, Pitts, Spafford)

Work continues on the identification and research of those design concepts related to distributed operating systems. This includes atomicity, consistency, recovery, fault tolerance, interprocess communication, file systems, etc. Of special importance is the area of distributed databases. Due to the amount and distribution of information required by a distributed operating system, we are looking into building a distributed database as an intrinsic part of the operating system.

Project C.10 has been created to attempt a "first try" implementation of this combination of data base and operating system.

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, Fox)

Work continues on the implementation of Distributed Software Tools and preparation of a technical report describing our experiences.

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

During this quarter, the main effort has been directed toward characterizing the user, his requirements, and his needs. To accomplish this in part, user activity has been monitored and these data processed to determine how people are using the current facilities. We are studying what commands are being used and in what general class of processing these commands are used (e.g. text manipulation, program development, etc.). Also, the commands available in the same system have been examined to determine what these commands functionally provide to a user. The information gathered from these approaches is being used as a strong basis for the friendly user interface that is the goal of this work.

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

Realizing that there remain many unsolved problems and open areas of research related to distributed operating systems, we believe that it

would be beneficial to begin an implementation based on what we currently know that we can do. We have named our prototype operating system "CLOUDS", an acronym for "Coalescing Local Operating Systems Under Decentralized Supervision." The descriptor "Coalescing" is an important one. CLOUDS will be capable of stand-alone operation, but when networked with other CLOUDS, will naturally come together or "coalesce" into one distributed operating system.

Current efforts are centered around designing a scaled down version of such an operating system for the first implementation.

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

No significant progress to report.

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

No significant progress to report.

D.4 User Interfaces to Database Systems (Griffeth)

Several interfaces have been completed for the PRIME relational database. They are: (1) a relational algebra interface, (2) a relational calculus interface, and (3) a network interface. Preliminary tests on the effectiveness of these interfaces will begin next quarter.

A project, "Distributed Database Algorithms", will be funded by NSF for the period July 1982 - June 1984.

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

A number of additional data sources were identified. A tentative metric for the quality of communication activities during large software development was proposed and is currently under investigation.

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

Design work has been completed; implementation and evaluation are in progress. A paper was submitted to the Third International Conference on Distributed Computing Systems.

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

Implementation of a Pascal compiler, using the code generator previously developed under this project, is nearly complete. Design of extensions for system implementation support is in progress. Implementation should begin during next quarter.

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

Work by Miller has confirmed and considerably improved earlier results. Design for balanced message flows was found to be crucial for best performance. A paper was submitted to the Third International Conference on Distributed Computing Systems.

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)

No significant progress to report.

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

No significant progress to report.

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

No significant progress to report.

5. TRAVEL RELATED TO THE FDPS PROGRAM

Date of Trip: 25-27 January, 1982
Individual(s) Traveling: Richard LeBlanc
Itinerary: Albuquerque, New Mexico
Purpose: Attend ACM Principles of Programming Languages Symposium

Date of Trip: February, 1982
Individual(s) Traveling: Nancy A. Lynch
Itinerary: Brandeis University, Northeastern University, Boston University
Purpose: Speak about new Byzantine Generals algorithm.

6. VISITORS

Dates of Visit: February, 1982

Visitor: Bharat Bhargava

Contact: Nancy A. Lynch (at MIT)

Purpose: Discussions about correctness proofs of concurrency control algorithms, and about reliability properties of distributed algorithms.

7. PUBLICATIONS

Author(s): John A. Miller and Richard J. LeBlanc

Title: Distributed Compilation: A Case Study

Type: conference paper

Status: submitted

Author(s): Arthur B. Maccabe and Richard J. LeBlanc

Title: The Design of a Programming Language Based on Communication Networks

Type: conference paper

Status: submitted

Author(s): Michael J. Fischer, Nancy D. Griffeth, and Nancy A. Lynch

Title: Global States of a Distributed System

Type: journal paper

Status: accepted for publication in Transactions on Software Engineering

Publ. Date: May, 1982

Author(s): Nancy A. Lynch

Title: Multilevel Atomicity

Type: conference paper

Status: accepted by Principles of Database Systems Conference

Publ. Date: March, 1982

Author(s): Nancy A. Lynch, Michael J. Fischer, and Robert Fowler

Title: A Simple and Efficient Byzantine Generals Algorithm

Type: conference paper

Status: submitted to IEEE Symposium on Reliability in Distributed Software and Database Systems.

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

Title: Cryptographic Protocols

Type: conference paper

Status: accepted by SIGACT

DATE
FILMED
8-8