AD-A114 199 GEORGIA INST OF TECH ATLANTA SCHOOL OF INFORMATION A--ETC F/G 9/2
RESEARCH PROGRAM IN FULLY DISTRIBUTED PROCESSING SYSTEMS.(U)
JAN 82 P H ENSLOW, N A LYNCH, A P JENSEN N00014-79-C-0873

UNCLASSIFIED NL

END
DATE FILMED
-82
DTIC

(12)

AD A114113

DTIC
ELECTE
MAY 07 1982
S
E

DTIC FILE COPY

School of Information and Computer Science

GEORGIA INSTITUTE OF TECHNOLOGY

A UNIT OF THE UNIVERSITY SYSTEM OF GEORGIA

82 04 20 079

(12)

THE GEORGIA INSTITUTE OF TECHNOLOGY

RESEARCH PROGRAM IN

FULLY DISTRIBUTED PROCESSING SYSTEMS

Quarterly Progress Report Number 9
1 September, 1981 - 30 November, 1981

January, 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

U.S. Army Institute for Research in
Management Information and Computer Science (AIRMICS)
Contract: DAAK70-79-D-0087
GIT Project: G36-647

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

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

1. INTRODUCTION

This is the Ninth 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: "Support of MILPERCEN Data Storage Concept"
Funding Agency: U.S. Army Institute for Research in Management Information and Computer Science (AIRMICS)
Contract Number: DAAK70-79-D-0087
GIT Project Number: G36-647
Principle Investigator: Alton P. Jensen

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 - Assistant Professor
Jensen, Alton P. - Professor
LeBlanc, Richard - Assistant Professor
Livesey, Jon - Assistant Professor
Lynch, Nancy A. - Associate 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.2 Decomposition of Parallel Systems
A.3 Reliable Systems
A.4 Time Performance of Distributed Systems
A.5 Audit Algorithms
A.6 Ticket Systems
A.7 Synchronous Simulation
A.9 Theory of Distributed Databases
A.10 Arbiter Design
A.14 Using Complementary Distributed System Models
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

D. Distributed Data Bases

D.1 Concurrency Control in Distributed Database Systems
D.2 Support of MILPERCEN Data Storage Concept
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.2 Decomposition of Parallel Systems (Lynch, Fischer)

This project was completed with the submission of a report for journal publication.

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

A lower bound on time requirements for solution to the Byzantine Generals problem, in an environment allowing authentication, has been undergoing revision for presentation at SIGACT 1982. Work is also proceeding on development of better Byzantine Generals algorithms, requiring less communication than currently known algorithms.

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

Project Complete.

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

The paper "Global States of a Distributed System" is being rewritten for submission to TODS.

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

Algorithms for sequential resource allocation are being developed, to prove the upper bound on ticket allocation time when tickets are allocated one after the other. The algorithms have been written in such a way as to allow generalization to the case in which allocation of different tickets may proceed concurrently.

A.7 Synchronous Simulation (Lynch, Fischer, Arjomandi)

This project was completed with the submission of a report for journal publication.

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

No significant progress to report.

A.10 Arbiter Design (Lynch, Schönhage, Fischer)

Project Complete.

A.14 Using Complementary Distributed System Models (Lynch, Rounds, R. Miller)

Project Complete.

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

We have begun work under the new NSF contract to extend our previous results in this area.

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

No significant progress to report.

A.18 Multilevel Atomicity (Lynch)

Connections between the multilevel atomicity model and the nested transaction model (as used by Leskov) are being studied.

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

Dr. Lynch is acting as major Ph.D. advisor to Eugene Stark, an MIT Ph.D. student. His thesis work involves formal models for describing distributed system behavior. The model has some similarities to previous work of Lynch and Fischer (there is no implicit synchronization, and finite delay is assumed). However, his model is at a much more abstract level than the previous work.

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

Dr. Lynch is studying the nested transaction structure used as a basis for Argus, Dr. Leskov's new distributed computing language. The language includes failure of subactions explicitly in its semantics. Dr. Lynch is formulating precise abstract semantics for the language, thereby attempting to prove correctness of some of the algorithms used in the implementation of Argus. A key difficulty is guaranteeing the preservation of consistency in the presence of "orphans": active processes with ancestors which have aborted.

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

A technical report "Initial Experiences with a Local Network --- Net/One by Ungermann-Bass, Inc." was completed and distributed.

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

Further simulation studies were carried out to evaluate the distributed control models. The report on the evaluation of these models has been published. A draft of Saponas's Ph.D. Dissertation was prepared for the final report. Saponas presented a defense of his thesis on this subject.

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

Algorithms developed for work distribution continue to be tested and evaluated using the simulator developed in project M.3.

C.4 Local Operating System (Livesey, LeBlanc, Saponas, Maccabe, Alchin, Fukuoka)

A project has been initiated to study high level inter-process communication based on a database description of a distributed system. A design of a domain-structured file system for a distributed system has been completed, and incorporated in the technical report being prepared for project C.8.

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

Work continues on identifying and describing the requirements for communications support in detail. Several network technologies are being considered to evaluate and compare their applicability to the fulfillment of these requirements. The formal project proposal is nearing completion and is currently under review.

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

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

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

Work continues on the literature survey for command languages and the design of a "user-friendly" command language. Two abstracts have been submitted for consideration as papers for the IFIP Technical Comittee-2, Working Group 2.7 Conference on Operating Systems Interfaces to be held in September, 1982.

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

A system for simulating concurrency control in distributed database systems has been proposed, as outlined in technical report GIT-ICS-81/15.

D.2 Support of MILPERCEN Data Storage Concept (Jensen, Doyle, Gehl, Bingham)

This project has been completed. The research conducted under this contract has identified and summarized many issues which affect the manner in which human resource information is managed for the Army. Several recommendations concerning how the Army should address itself in upgrading its Automated Manpower and Personnel Resources Management Informations Systems are made in the Final Report, "Automating the Exchange of Military Personnel Data Among Selected Army Organizations." A second report is a comprehensive set of appendices containing data collected during the course of study.

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

No significant progress to report.

D.4 User Interfaces to Database Systems (Griffeth)

An experimental system to test the usability and power of various database user interfaces is being developed. The actual storage structure will be based on the relational model, because it is easier to present a variety of interfaces when the underlying model is relational. However, the user data model may include such structures as CODASYL sets, repeating groups, vectors, semantic nets, or even higher level objects such as entities, relationships, aggregations, generalizations, etc. Furthermore, the user language may include navigational features as well as the ability to reference rows by value.

Currently, the underlying database is one operation short of implementation (this is the division operation). A pretest to determine the subject's background and database sophistication has been developed and is being refined. The language interpreters are in the design stage. A game for the subjects to play has been designed and is close to implementation.

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

During this quarter several sources for data on large software development projets were investigated. Of the sources identified, the RADC Productivity Database was considered to be most suitable. A copy of the data was obtained and preliminary analysis begun. Noting deficiencies in the coverage of the database, a search for additional data sources will proceed in parallel with continued analysis of the data so far obtained.

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

The PRONET language features (described in technical report GIT/ICS-81/03) are currently being implemented. Most of the work done so far has been concerned with the required run-time support software.

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

A compiler for an extended version of Pascal is currently under development, using the code generation tool previously implemented as part of this project. This compiler will be extended to implement PRONET (project I.1) and will be used in future work in project I.3.

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

No significant progress to report.

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

We have continued to develop new cryptographic protocols. This work will be summarized in a forthcoming paper.

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

A proposal for a secure operating system based on the paper "Secure Architectures - the Master/Slave Model" has been completed, and awaits submission to NSF.

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

See D.1; otherwise, no significant progress due to loss of student working on simulation program.

M.1 Establishment of FDPS Testbed Facility (Myers, Mongiovi, Pitts, 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: 30 August - 4 September, 1981
Individual(s) Traveling: N. Griffeth
Itinerary: Boston, Massachusetts (MIT)
Contact: N. Lynch
Purpose: Develop distributed algorithms for ticket systems and discuss further work on performance studies of ticket systems.

Date of Trip: 14-16 October, 1981
Individual(s) Traveling: J. Livesey (for P. Enslow)
Itinerary: Rome, New York
Contact: Tom Lawrence
Purpose: RADC Technical Exchange Meeting

Date of Trip: 21 October, 1981
Individual(s) Traveling: P. Enslow (accompanied by Dr. Robert Grafton, ONR)
Itinerary: New London, Connecticut
Contact: Naval Underwater Systems Center
Purpose: Orientation on Georgia Tech FDPS Research Program

Date of Trip: 22 October, 1981
Individual(s) Traveling: P. Enslow (accompanied by Dr. Robert Grafton, ONR)
Itinerary: Newport, Rhode Island
Contact: Naval Underwater Systems Center
Purpose: Orientation on Georgia Tech FDPS Research Program

Date of Trip: October, 1981
Individual(s) Traveling: N. Lynch
Itinerary: Nashville, Tennessee
Purpose: Attend FOCS Conference

Date of Trip: November, 1981
Individual(s) Traveling: N. Lynch
Itinerary: New Haven, Connecticut
Contact: Mike Fischer
Purpose: Work on projects and give presentation on latest progress.

6. VISITORS

Dates of Visit: September, 1981
Visitor: M. Fischer
Contact: N. Lynch (at MIT)
Purpose: Work on various projects.

Dates of Visit: 23 October, 1981
Visitor: Herman J. Weegenaar (Centraal Beheer, The Netherlands)
Contact: P. Enslow
Purpose: Discuss distributed operating system functions

Dates of Visit: 29-30 October, 1981
Visitor: Tom Lawrence (RADC) and Rudy Nothdurft
Contact: P. Enslow
Purpose: Review and planning session on new RADC contract

7. PUBLICATIONS

Author(s): M. Fischer, N. Lynch
Title: A Lower Bound for the Time to Assure Interactive Consistency
Type: technical report
Status: published
GIT Number: GIT-ICS-81/13
Publ. Date: September, 1981

Author(s): P. Enslow, P. Manno, and J. Myers
Title: Initial Experience with a Local Network - NET/ONE by Ungermann-Bass
Type: technical report
Status: published
GIT Number: GIT-ICS-81/11
Publ. Date: October, 1981

Author(s): R. DeMillo, N. Lynch, and M. Merritt
Title: Cryptographic Protocols
Status: abstract submitted to conference
Publ. Date: Winter '82

Author(s): A. Jensen, J. Bingham, J. Doyle, and J. Gehl
Title: Automating the Exchange of Military Personnel Data Among Selected Army Organizations
Type: final report
Status: published
Publ. Date: June, 1981

Author(s): A. Jensen, J. Bingham, J. Doyle, and J. Gehl
Title: Automating the Exchange of Military Personnel Data Among Selected Army Organizations
Type: appendices
Status: published
Publ. Date: June, 1981

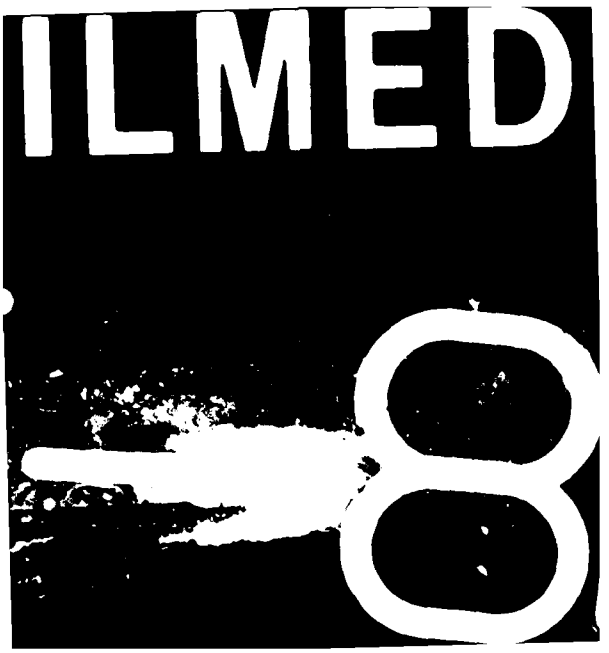
ILMED