REPORT DOCUMENTATION PAGE

OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
		R&D Status Report 10/1/97-12/31/9

4. TITLE AND SUBTITLE
Applications of the Theory of Distributed and Real Time Systems to the Development of Large-Scale Timing Based Systems

5. FUNDING NUMBERS
C-
F19628-95-C-0118

6. AUTHOR(S)
Nancy Lynch

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02138

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Department of the Airforce
Electronic Systems Center (AFMC)
Hanscom Air Force Base, MA 01731

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES
N/A

12a. DISTRIBUTION/AVAILABILITY STATEMENT
No limits on disclosure

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

19980205 092

14. SUBJECT TERMS

15. NUMBER OF PAGES
9

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298
Prescribed by ANSI
298-102

DTIC QUALITY INSPECTED 3

January 15, 1998

Mr. Harry Koch
ESC/ENS
5 Eglin Street, Building 1704
Hanscom Airforce Base, MA 01731-2116

Dear Mr. Koch:

This letter contains our R & D Status Report covering the period from October 1, 1997 to December 31, 1997 for Contract F19628-95-C-0118, entitled "Applications of the Theory of Distributed and Real-Time Systems to the Development of Large-Scale Timing-Based Systems".

Technical Progress

The group this year consists of Prof. Lynch and graduate students Victor Luchangco, Roberto DePrisco, Mandana Vaziri, Henrik Jensen, Josh Tauber (who joined in November), Roger Khazan, Carl Livadas, and Kate Dolginova. Dr. Steve Garland and graduate student Anna Chefter of Guttag's group are also working closely with us. Information about these people can be found at URL http://theory.lcs.mit.edu/tds/people.html.

I. Proposal preparation

A large part of Prof. Lynch's time was occupied on writing proposals for renewal of funds for 1998-2001. She prepared proposals for DARPA (4 BAAs) and NSF.

II. Modelling and verification tools

We continued our project on the IOA language and toolset, which are designed to support our formal approach to distributed system design and analysis. The design of the IOA language is substantially complete, and appears in a language manual on the web.

This quarter, work continued on the development of tools for the IOA language; our toolset will include a parser and static semantic checker, composition routine, support for levels of abstraction, interfaces with theorem provers and model checkers, a simulator, and a code generator for real distributed code.

- Garland worked on the static semantic checker. He enhanced the checker to recognize the built-in types, to check invariants and simulation relations involving both primitive and composite automata, and to annotate IOA programs so that they can be printed using LaTeX. Preliminary versions of a parser and static semantic checker have now been completed, and are available for the use of other researchers who want to develop tools for IOA. The parser has already been given to Devillers and Vaandrager at Nijmegen, who are interested in connecting IOA to the PVS theorem prover.

- We begin work on connecting IOA to a model checker, for validation of small instances of designs. Specifically, Vaziri started translating IOA to PROMELA, the input language of the SPIN model checker. She devised a translation scheme for a subclass of IOA programs, and implemented a translator based on this scheme, using the existing parser.
- We began work on the ambitious project of translating distributed IOA programs to running Java or C++ code. Tauber began by evaluating various communication services that might be used in the implementation; he selected the Message Passing Interface (MPI). He defined an interface between local IOA programs and the MPI service, and designed a set of I/O automata that formally describe the behavior of an implemented MPI system. Tauber also devised a translation scheme for a subclass of (non-distributed) IOA programs and implemented a translator into Java, based on this scheme, using the existing parser. The next step is to compose the modelled communication interface with the existing translation to extend the scheme to distributed systems.
- Garland and Lynch wrote an extended abstract for submission to PODC, on the IOA language and the design of the toolset.

III. Applications

A. Distributed system building blocks

We continued our work on building-blocks for fault-tolerant distributed systems. Much of our progress on this topic this quarter involved dynamic view-oriented group-communication services.

- Khazan continued his work on modeling a load-balancing replicated data server. His implementation relies on the underlying group-communication service to achieve fault-tolerance and efficiency. During this reporting period, Khazan completed the assertional proof of correctness and has begun work on the performance and fault-tolerance property.
- DePrisco, Fekete, Lynch, and Shvartsman have made great progress on the work started last quarter on *dynamic quorum* versions of view-synchronous group communication services, in which view changes are restricted so that clients in new views can always obtain information propagated from old views. Requiring that each view's membership comprise a majority of the processes would suffice, since then all pairs of views would have nonempty intersections. However, the service can allow more freedom than this if it learns that the clients in some views have completed exchanges of information. We have provided a formal specification of a dynamic quorum service. To show the significance of such a specification we have provided both an implementation of the service and an application that runs on top of the service. An extended abstract describing the work done has been completed and submitted for publication. The work done has focused on the safety properties of the service. Future work include performance analysis and other applications.

- In response to comments received from the journal reviewers, Luchangco revised the ESDS journal submission, including a significant extension of the analysis of the performance and fault-tolerance of the algorithm proposed, and an outline of further possibilities for fault recovery, and the performance under such conditions. He also reworked several of the proofs to be simpler and more complete than before.

- Shvartsman, with Shlomi Dolev and Roberto Segala are working on new results on performing work in reconfigurable networks using group communication services. A report is being prepared.

B. Multiprocessor shared memory models

We finally completed our work on the RAID case study. We also continued our work on developing the theory needed to understand how to program using weakly coherent memory models. Some of that work involved reading and reporting on the literature and outlining directions, but there has been some more specific written progress, described below:

- Luchangco and Frigo resumed work on computation-centric memory models, which characterize memories from the point of view of the programmer. A computation is a generalization of an instruction stream. Memory models are expressed in terms of these computations, allowing the programmer to reason about what a program specifies rather than low-level system details. They define sequential consistency in this framework, along with several weak consistency models, and show some characteristics of these models, as well as relationships amongst them. They also define properties that characterize "reasonable" memory models, i.e., they argue that memory models not having these properties are undesirable from the programmer's point of view. On this basis, they suggest a candidate for the weakest "reasonable" memory model.

- Vaziri finished writing a conference and journal version of her work on proving correctness of a controller algorithm for the RAID level 5 system. The conference paper was submitted to FTCS-28.

C. Automated Transportation Systems

This quarter, we wrote and submitted two papers on our automated vehicle work to the 1998 Hybrid Systems workshop, and both were accepted. We also worked hard on our new project on modelling/verification of controlled aircraft systems. All our work is based on our hybrid I/O automaton (HIOA) model.

- Livadas' work on using hybrid I/O automata to model safety-critical hybrid systems and in particular automated vehicle protection subsystems was submitted and accepted to the workshop on Hybrid Systems: Computation and Control to be held at Berkeley California in April 1998.

- Lygeros and Lynch used hybrid I/O automata to model the problem of emergency deceleration of a string of vehicles. They established conditions under which such a maneuver will be safe, in the sense that any collisions are guaranteed to be at low relative velocities. In this quarter, this work was submitted and accepted to the Berkeley workshop on Hybrid Systems: Computation and Control.

- Livadas, Lygeros and Lynch have been working on the analysis and verification of simplified models of the Traffic Alert and Collision Avoidance System II (TCAS II). The goal is to use analysis techniques developed for automated vehicle transportation systems to the analysis of the complex algorithms used in detecting and resolving collision threats among aircraft in flight. We have completed a model of the system, have a preliminary sketch of a safety proof for normal operation, and a list of some "less normal" cases to consider. In early December this work was presented to the TCAS development/analysis group at Lincoln Laboratories. The feedback we got was very encouraging. The work was also presented in the 36^{th} Conference on Decision and Control, San Diego, California, December 10-12, 1997. An extension of our approach has been applied by Lygeros and researchers at U.C. Berkeley to the verification of the Center TRACON Automation System (CTAS). It has been accepted and will be presented in the workshop on Hybrid Systems: Computation and Control, Berkeley, California, April 13-15, 1997. Future work will involve the completion of the analysis of the TCAS II system and the development of automated tools to help in the verification of complex systems modeled and analyzed using hybrid I/O automata.

- Lynch submitted a full version of her paper, [10], to appear as a book chapter.

IV. Algorithms and impossibility results

- Lynch and Rajsbaum completed a journal version of their paper with Borowsky and Gafni, on "The BG Distributed Simulation Algorithm", and submitted it for publication.

- Jensen has continued work on using abstraction techniques to reduce large/infinite state verification problems to small finite ones amenable to model-checking. Having successfully applied his technique to the Burns n-process mutual exclusion algorithm and the Bakery n-process mutual exclusion algorithm, Jensen is now attempting to apply the technique to the Bounded Concurrent Timestamp algorithm of Dolev and Shavit.

A paper reporting the work on Burns n-process mutual exclusion algorithm has been accepted for the TACAS'98 conference on tools and algorithms for the construction and analysis of systems.

- M.S. student Gunnar Hoest and visiting faculty member Prof. Nir Shavit from Tel Aviv University have been working on a full journal version of their paper on a mathematical complexity framework for fault-tolerant asynchronous systems; they are near completion. Their work uses topological models and methods to analyze time complexity in the *iterated immediate snapshot* model, a restricted type of atomic snapshot shared memory model. They obtained tight bounds for the approximate agreement problem, and a fundamental time vs. number of names tradeoff for the renaming problem. A paper appeared in PODC'97, and Hoest completed his M.S. thesis.

Special Programs and Major Items of Equipment

None.

Changes in Key Personnel

This year, the group has been reduced in size with the departure of research associates Alex Shvartsman (now a professor at the University of Connecticut), John Lygeros (now a postdoc at Berkeley), and Nir Shavit (a Professor at Tel Aviv University). All retain ties with the group as visitors and research collaborators.

Tentative plans have been made for Dr. Idit Keidar of the Hebrew University to join the group in summer, 1998.

Prof. Lynch is on sabbatical in Spring, 1998, but is travelling only sporadically. She will continue to supervise the project closely.

Trips, Talks and Conferences

1. Nancy Lynch. "Practical Specs for Practical Communication Services." DIMACS Workshop on Communication, New Brunswick, NJ, October, 1997.
2. John Lygeros. "On the Formal Verification of the TCAS Conflict Resolution Algorithms." 36th IEEE Conference on Decision and Control, San Diego, CA, December 1997.
3. Nancy Lynch and John Lygeros. "On the Formal Verification of the TCAS Conflict Resolution Algorithms." Lincoln Labs, Bedford, MA, December 1997.
4. Alex Shvartsman. "Networked Resource Management and Distributed System Building Blocks." DIMACS Workshop on Communication, New Brunswick, NJ, October, 1997.

Areas of Concern

Funding needs renewal at the end of this semester.

Statement of Sufficiency

The contractually prescribed effort appears to be sufficient to achieve the objectives of this contract.

Degrees Awarded

None this quarter.

Related Accomplishments

During this reporting period the following papers were submitted for publication, accepted for publication, or published:

Submitted for publication:

[1] Elizabeth Borowsky, Eli Gafni, Nancy Lynch, and Sergio Rajsbaum. The BG Distributed Simulation Algorithm. Submitted for journal publication, December, 1997.

[2] Mandana Vaziri and Nancy Lynch. Proving Correctness of a Controller Algorithm for the RAID Level 5 System. Submitted for publication, December 1997.

[3] Stephen J. Garland and Nancy A. Lynch. The IOA Language and Toolset: Support for Mathematics-Based Distributed Programming. Submitted for publication, January 1998.

[4] Roberto De Prisco, Alan Fekete, Nancy Lynch and Alex Shvartsman. A Dynamic View-Oriented Group Communication Service. Submitted for publication, January 1998.

Accepted:

[5] Roberto Segala. Compositional Verification of Randomized Distributed Algorithms. *Proceedings of Compositionality - the Significant Difference (COMPOS)*, Malente (Holstein), Germany, September 1997. Invited talk. To appear published by Springer-Verlag.

[6] Henrik Jensen and Nancy Lynch. A Proof of Burns N-Process Mutual Exclusion Algorithm using Abstraction. *TACAS'98 (Conference on Tools and Algorithms for the Construction and Analysis of Systems)*, Gulbenkian Foundation, Lisbon, Portugal, March 31-April 2, 1998. To appear.

[7] Carolos Livadas and Nancy A. Lynch. Formal Verification of Safety-Critical Hybrid Systems. *Hybrid Systems: Computation and Control*, Berkeley, California, April 13-15, 1998. To appear.

[8] John Lygeros and Nancy Lynch. Conditions for Safe Deceleration of Strings of Vehicles. *Hybrid Systems: Computation and Control*, Berkeley, California, April 13-15, 1998. To appear.

[9] John Lygeros, George Pappas, and Shankar Sastry. An Approach to the Verification of the Center-TRACON Automation System. *Hybrid Systems: Computation and Control*, Berkeley, California, April 13-15, 1998. To appear.

[10] Nancy Lynch. A Three-Level Analysis of a Simple Acceleration Maneuver, with Uncertainties. *AMAST Workshops on Real-Time Systems 95 & 96*, World Scientific Publishing Company, AMAST Series in Computing, book chapter. To appear.

Appearing:

[11] Stephen J. Garland, Nancy A. Lynch and Mandana Vaziri. IOA: A Language for Specifying, Programming, and Validating Distributed Systems. Draft, September 1997. URl `http://larch-www.lcs.mit.edu:8001/~garland/ioaLanguage.ps`.

[12] Carolos Livadas. Formal Verification of Safety-Critical Hybrid Systems. Technical Report MIT/LCS/TR-730, Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, September 1997 (Technical Report printed November, 1997). Masters thesis.

[13] Roberto Segala. Quiescence, Fairness, Testing and the notion of Implementation. *Information and Computation*, 130(2):194–210, November 1997.

[14] Elizabeth Borowsky, Eli Gafni, Nancy Lynch, and Sergio Rajsbaum. The BG Distributed Simulation Algorithm. Technical Memo MIT/LCS/TM-573, Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, December, 1997.

[15] John Lygeros and Nancy Lynch. On the Formal Verification of the TCAS Conflict Resolution Algorithms. *36th IEEE Conference on Decision and Control (CDC)*, San Diego, California, pages 1829-1834, December 10–12, 1997. Extended abstract.

[16] John Lygeros, Claire Tomlin and Shankar Sastry. Multi-objective hybrid controller synthesis. *36th IEEE Conference on Decision and Control*, San Diego, California, December 10–12, 1997.

[17] George Pappas, Claire Tomlin, John Lygeros, Datta Godbole and Shankar Sastry. A next generation architecture for air traffic management systems. *36th IEEE Conference on Decision and Control*, San Diego, California, December 10–12, 1997.

Papers in progress

[18] Nancy Lynch, Roberto Segala, Frits Vaandrager, and H. B. Weinberg. Hybrid I/O Automata." Journal version. In progress.

[19] Nancy Lynch and Alex Shvartsman. Robust Emulation of Shared Memory Using Dynamic Quorum-acknowledged broadcasts, Journal version. In progress.

[20] Alan Fekete, David Gupta, Victor Luchangco, Nancy Lynch, and Alex Shvartsman. Eventually-Serializable Data Services. Journal version. In progress.

[21] John Lygeros and Nancy Lynch. On the Formal Verification of the TCAS Conflict Resolution Algorithms. Technical Report in progress.

[22] Matteo Frigo and Victor Luchangco. Computation-Centric Memory Models. In progress.

[23] Alex Shvartsman, Shlomi Dolev and Roberto Segala. Using Group Communication Service for Dynamic Load Balancing. In progress.

[24] Alex Shvartsman and Oleg Cheiner. Implementing and Evaluating an Eventually-Serializable Data Service. In progress.

[25] Henrik Jensen. An Abstract Interpretation of the Bounded Concurrent Timestamp Algorithm. In progress.

[26] Mandana Vaziri and Nancy Lynch. Translating IOA to Promela. In progress.

Theses in progress

Roberto Deprisco. PhD thesis (Untitled). Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Kate Dolginova. MEng thesis. "Safety Verification of Automated Car Maneuvers." Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Henrik Jensen. PhD Thesis. "Integration of Deductive and Algorithmic Methods for Verification of Reactive Systems." Aalborg University, Denmark. Visiting MIT.

Roger Khazan. Masters Thesis. "Group Communication as a Base for a Load-Balancing, Replicated Data Service." Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Carl Livadas. PhD thesis (Untitled). Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Victor Luchangco. PhD Thesis. "Consistency Models for Distributed Memories." Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Josh Tauber. PhD Thesis (Untitled). Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Mandana Vaziri. PhD thesis (Untitled). Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

Awards:

None

Sincerely,

Nancy Lynch

Nancy Lynch
NEC Professor of Software Science and Engineering

Electrical Engineering and Computer Science
(617)253-7225
lynch@theory.lcs.mit.edu

MIT Laboratory for Computer Science

Applications of the Theory of Distributed Real-Time Systems To the Development of Large-Scale Timing-Based Systems

Prof. Nancy Lynch, Principal Investigator

R & D Status Report
Program Financial Status
ARPA Contract # F19628-95-C-0118
CLIN # 0002
Quarterly Report (10/97 - 12/97)

	Planned Expenditures	Actual Expenditures at Report Date	% Completion	Budget At Completion	Latest Revised Estimate	Remarks
Total Base Contract	858,443	518,314	60.38%	858,443	858,443	
Current Funding Profile	574,447	518,314	90.23%		518,314	*
Equipment	35,308	23,554	66.71%			

* Data reflects all received fundingthrough 12/97. Current funding is sufficient through 12/97.