

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Laboratory for Computer Science  
(Formerly Project MAC)

Computer Systems Research Division

Proposal for

Continuation of Research on

Engineering of a Computer System for which Security  
can be Certified by Auditing

Submitted to

Honeywell Information Systems Inc.

United States Air Force Electronic Systems Division  
and  
Defense Advanced Research Projects Agency

July, 1976

---

Jerome H. Saltzer  
Principal Investigator

---

Michael L. Dertouzos  
Director, Laboratory for  
Computer Science

---

George H. Dummer  
Director,  
Office of Sponsored Programs

## SUMMARY OF PROPOSAL

The Computer Systems Research Division of the M.I.T. Laboratory for Computer Science (formerly Project MAC) proposes the continuation of its research program to make possible the certification, by auditing, of the centrally protected core of a general-purpose, remote-accessed, multi-user computer system. The current contract between HISI and M.I.T. (actually a directed subcontract of a cost-sharing contract between AF/ESD and HISI) provides funding through June, 1976. Progress during the last 12 months in file system simplification has been largely as expected, so it is proposed here that the original work planned for the next 15 months be continued. A specific budget for July, 1976 through September, 1977 is provided.

## Overview

The Computer Systems Research Division of the M.I.T. Laboratory for Computer Science (formerly Project MAC) proposes to continue its work on the engineering of a computer system for which security can be certified. The overall project has made considerable progress in understanding how to construct simpler, better-structured file management systems during the current contract period, with several important experimental projects and conceptual tasks either completed or nearing completion. During the next contract period, the Division proposes to complete the experimental work on the file system area, and to continue explanation and technology transfer of what has been learned there. The remainder of this proposal briefly reviews the recently completed activities, and provides a proposed budget. This document is not self-contained; it presumes that the reader is familiar with the original proposal (October, 1973) and the first continuation proposal (April, 1975). The technical content of these two proposals is also summarized in a paper by M. D. Schroeder, which for reference is attached as Appendix I.

### Accomplishments so far: System Simplification

Several activities have been completed or are almost ready for a final writeup. Together, these activities represent significant progress on the problem of designing simpler, more easily certifiable, computer systems. These activities are:

1. Accomplishing dynamic inter-program binding without privilege. This activity, reported in MAC-TR-132 by Janson, demonstrated that efficient interprogram binding could be accomplished without special supervisor privileges. By this single change the supervisor interface was simplified and reduced in scope by more than 10%.
2. Accomplishing user name-space management without privilege. This activity, reported in MAC-TR-156 by Bratt, demonstrated that user name

space management can be efficiently handled without special supervisor privilege. In addition to reducing the size of the protected supervisor, this change also demonstrated a way of eliminating the need for the protected file system to manage variable-length character string names (path names) for objects.

3. Demonstration of a two-layer process implementation. This activity, to be reported in an S.M. thesis by Reed (and later in a technical report,) is establishing the feasibility of the following idea: by doing process implementation twice, first producing a fixed number of virtual processes, then later the dynamically varying number needed by users, it appears possible to disentangle the process implementation from the virtual memory implementation. The primary question remaining is whether this separation can be accomplished without excessive overhead; the demonstration will help answer that question.
4. Demonstration of a multiprocess virtual memory implementation. This activity, to be reported in an S. M. thesis by Huber (and later in a technical report,) has shown that separation of the management of virtual memory into a process responsible for removing primary memory pages and another responsible for removing backup memory pages does considerably simplify the structure of that area of the system. Further work is underway to verify the hypothesis that the initial, fairly clumsy, PL/I implementation can be tuned up to have performance equivalent to the older, machine language version.
5. Demonstration of an alternative system initialization strategy. This activity, to be reported in an S. M. thesis by Luniewski (and later in a technical report,) intends to demonstrate that pre-initializing the supervisor before generating a system tape rather than initializing it on the fly at each system start is a simpler, easier-to-comprehend strategy.
6. Rearrangement of process initialization/domain crossing strategy. This activity, to be reported in an S. M. thesis by Montgomery (and later in a technical report,) intends to demonstrate that a significant reduction in the quantity of protected, privileged code can be accomplished by making

process creation unprivileged. The standard mechanism to control entry to a domain can be used to force new processes to start at a controlled location. This strategy, in turn, allows login of users to be accomplished by an unprivileged creation of a process in the potential user's domain, followed by authentication of the user by an unprivileged initial program in that domain.

7. A new model of process coordination. This activity, to be reported in a journal paper by Kanodia and Reed, represents a breakthrough in thinking about coordination of cooperating sequential processes. By use of eventcounts (rather than semaphores) problems that require only that things be done in a predictable sequence can be handled without using more powerful mutual exclusion machinery. One result of this work is that it is possible for a reader of a database to synchronize its use with writers, while being completely unable to signal any writer; this ability has significant implications for multilevel security situations that require confinement. Another result of interest is that the eventcount model allows correct synchronization of processes separated geographically; this ability is important for network environments.
8. Partially ordered systems. This activity, to be reported in a Ph.D. thesis by Feiertag (and by a technical report,) is a study of the advantages of allowing partial ordering of modules when proving correctness of large systems. (The approaches suggested by other workers, such as those of SRI, allow partial ordering but do not take advantage of it.) The basic strategy required to take advantage of partial ordering is to provide a way to specify, for each module, exactly what environment the compiler should assume while compiling programs of that module. The goal of this thesis is to reduce the otherwise combinatorial build-up of effort required to prove correctness of multi-module systems.
9. Type abstraction in the design of the virtual memory manager. This activity, to be reported in a Ph.D. thesis by Janson (and a technical report,) demonstrates that the concept of type abstraction can be

successfully applied to the internal structure of the virtual memory manager that underlies the usual mechanisms for providing type abstraction. Most systems employing type abstraction as a simplifying concept have not attempted to utilize the concept in their own implementation, because the correct form of separation is not obvious. This thesis demonstrates that at least one self-consistent abstract layering of a virtual memory implementation exists. The implication is profound, since the virtual memory subsystem is the single most complex subsystem in Multics and it can now be broken apart into independently verifiable modules.

#### Proposed Continuation

The project to engineer a simpler kernel for Multics is close to completion, and we propose to finish the remaining research tasks in the next twelve months. Many of the tasks will be completed well before the end of that period, and the proposed budget reflects an anticipated tapering down of this activity. Two new tasks have been added, one that covers the writing of a final report on the project and the second to provide for continued "technology transfer" support of HISI and others who may wish to make use of our findings. All of the proposed tasks are described in detail below:

#### 1. Continuing Research Tasks

##### 1.1) Restructuring of Traffic Control.

This continuing project was described earlier, in the progress report. D. Reed, as his master's thesis research, has designed a restructured traffic controller that partitions its functions into two levels. The lower level multiplexes the real processors of the system among a fixed number of so called virtual processors. By fixing in advance the number of such virtual processors, this low level processor multiplexer need make no use of the system's virtual memory facilities. Thus there is a strict isolation and ordering between the multiplexer and the virtual memory. A higher level scheduler multiplexes some of the virtual processors among all of the currently operating real Multics processes. This higher level scheduler can use all of the facilities of the Multics virtual memory, since they are implemented at a lower level.

In addition, R. Kanodia and D. Reed have prepared a paper describing a new mechanism that they have developed for the coordination of processes. This new coordination mechanism, which uses monotonically increasing counters to record event occurrences, has several nice features not found in other known schemes, for example, it does not require that those waiting for an event have write access on the coordination element, which simplifies correctness proofs, and it seems to be especially appropriate for synchronization of distributed processes. It is a simple mechanism that seems to model many actual coordination situations very naturally.

The eventcount process synchronization mechanism seems to hold the potential for significant simplification in the area of interprocess coordination. The paper being prepared, a draft of which was published as RFC 102, discusses in some detail the theoretical benefits to be derived from eventcounts. The practical advantage of eventcounts to the Multics system can best be seen by an examination of the way they are used in the proposed implementation of the two level traffic controller, where it will be seen that the use of eventcounts has encouraged the exact sort of structure for which we are striving.

Two activities remain to be done, the completion of the detailed design and implementation study currently under way, and the completion of a paper describing eventcounts.

The design and implementation study, which was started in the beginning of 1976, is required in order to demonstrate that the two-level traffic controller has no adverse performance effects on the system. This performance issue is not addressed in the thesis, which discusses the theoretical justification for this structure and demonstrates the simplification which will result from its use. However, it is very important that the efficiency issue be addressed, since multi-level schedulers have been proposed before (although not with the goal of simplification in mind) and these two level schedulers have not operated with reasonable delay properties. We feel that the problems causing these delay properties have been side-stepped in the current design, but it is appropriate and important to confirm this.

By the start of this contract, the technical report describing the theoretical advantages of the structure should be finished, as should the

draft of the detailed design specification. The next step will be to determine exactly what sort of implementation of this design we propose to do. An implementation schedule will be constructed after this decision is made. Following the implementation, a document will be prepared which discusses the conclusions we draw from this implementation.

The publishable version of the paper on eventcounts will be prepared.

#### 1.2) Separating of Page Control and Segment Control Functions within the Active Segment Table.

Currently, the Active Segment Table is not properly organized for functional modularity, in that it contains variables used both by page control and segment control, causing an unstructured interaction between these two systems. This task is to propose a particular interface between these two systems that will result in a system that is simpler, more structured, and easier to understand. Currently, we are completing a detailed investigation of all the current sorts of interactions between the two systems, particularly as manifested by the variables in the Active Segment Table. The next step is to try to determine what the correct sorts of interactions should be. In other words, we must specify some appropriate set of dependencies that can exist between segment control and page control. It is likely, as we propose and justify such a specification, that the work now being completed by Phil Janson on a theoretical organization of the virtual memory mechanism will be applicable. An example of the sort of problem that must be solved is the question of whether the functions of quota management should be assigned to segment control or page control.

The current project is to identify all the modes of interaction currently existing between Segment Control and Page Control. Using the understanding and insight derived from this study, the next step is to propose what the correct sorts of interaction should be. The result of this should be a specification for the interface between Segment Control and Page Control. Having this specification, it will be possible to determine what sort of implementation is appropriate in order to demonstrate the validity of the specification. It may be necessary to iterate on the interface specification, if it is not possible to discover all modes of interaction except by



experiment. The final part of the task is a preparation of the technical report (thesis).

### 1.3) Study of Multics System Initialization.

A. Luniewski is continuing, as his master's thesis, the definition of a methodology for the initialization of the Multics system which is simpler, and more structured than the current strategy. The development of this methodology is considered very important, not only to help systematize a traditionally ad-hoc part of most systems, but also because all techniques for validating system security depend on being able to validate a secure initial state.

The work in this task so far has resulted in the description of an alternative mechanism for system initialization. We have performed certain selected implementations, to demonstrate that the ideas proposed in this thesis are indeed viable. The current project is to determine, by analysis and further implementation, to what extent this alternative structure is easier to understand than the one currently being used. It is not yet clear whether the alternative structure being proposed in this research will be of sufficient benefit to Project Guardian that the reimplementing of the system ought to be performed, but we feel that the research is valuable in any case, for the following reason. The programs that now perform initialization are not structured around any particular model of initialization, but are instead just put together in a way that works. As might be expected, current programs do display, a posteriori, some underlying structure, and this structure must be understood by anyone attempting to certify these programs. This study, in comparing the current structure with an alternative, may well yield insight to the certifier that will assist him in understanding initialization, whichever strategy he is called upon to certify.

One or two trial implementations of a hardware reconfiguration will be performed in order to demonstrate the validity of certain proposals made in this study. Certain other implementations, or at least preliminary coding, may be proposed in order to demonstrate the superiority of this strategy over the one currently used. The final part of this project is the preparation of the technical report.

#### 1.4) Provision of "Breakproof" Environment for User Programming.

As various parts of the operating environment are removed from the kernel, the question arises as to where they should be put. If they are placed in the same ring as the executing programs of the user, than they can be destroyed by a programming error of the user. It would be very nice if the removal of programs from the kernel did not lead to a reduced robustness of the programming environment. H. Goldberg is attempting to define a consistent collection of user support programs that can be protected from destruction by casual error of the user. This might include, for example, the linker and its tables, the linkage sections of the LOT. His thesis will provide a practical solution to the problem, will demonstrate the utility of rings for user self-protection, and will contribute to our understanding of the correct functionality of the user programming environment.

The implementation of the proposed environment now under way must be completed. The completed user environment must be evaluated by use, in order to ascertain that the proposals in this study are indeed correct. Following this, the final report must be prepared.

#### 1.5) Restructuring of Page Control.

As described in last year's annual report, the goal of this research was to utilize several asynchronous parallel processes to perform the various functions of page control. In particular, separate processes were used to remove pages from memory and from the paging device so that a free storage pool would always exist to be used for the servicing of page faults. This research has demonstrated that the use of parallel processes in this context provides an intrinsic simplification to the algorithm. The Master's thesis by A. Huber describing this implementation has been completed and will be available soon as a technical report.

The only part of this task remaining is the preparation of a report discussing the performance implications of the multi-process Page Control. It is expected that all experiments required in order to write this report will be completed early in the contracting period.

#### 1.6) Support of User Defined Object Types.

D. Hunt is continuing work on his EE thesis entitled "Building Blocks for

an Object-Oriented Virtual Memory System". The thesis deals with modular decomposition of large systems as a technique to make verification possible. In particular, the thesis focuses on the inter-module relationships which may exist. It is claimed that just a few relationships, each with simple semantics, are sufficient to express inter-module dependencies in a large number of cases. To support this claim, a virtual memory subsystem--similar to a Multics virtual memory enhanced to support type extension--is considered as a case study. A modularization for the virtual memory is suggested and justified, and the inter-module relationships are specified. There will be a report describing the work done on this project to date.

#### 1.7) Study of System Reliability and Recovery from Errors.

This past quarter H. Forsdick has developed a model of the occurrence and handling of errors in a computer system. In terms of this model, he has started to look at the way errors are handled in the current implementation of Multics. He has traced down the missing segment fault path in an attempt to develop a characterization of the structure of Multics with respect to errors. Several insights on the placement of error checks have come out of this study. A report will be written describing the work done on this project to date, but further work is not anticipated.

#### 1.8) Organization of the Virtual Memory Mechanism of a Computer System.

P. Janson has completed the writing of the first draft of his Ph.D. thesis on a method for producing modular, structured software to support the virtual memory mechanism of a computer system. The first part of the thesis describes the type extension concept that is recommended as the basis for organizing a virtual memory mechanism. A virtual memory mechanism should be regarded as implementing abstract information containers (e.g. segments) out of physical information containers (e.g. core blocks and disk records). The second part of the thesis explains how that type extension concept can be exploited to organize a virtual memory mechanism. In particular, it shows how one can implement the programs and the address space of the mechanism itself without violating modularity and structure. The last part of the thesis illustrates the use of the method by applying it to the redesign of the

virtual memory mechanism of Multics. The modularity and the structure of the resulting system are evaluated.

The work remaining on this task consists of publishing the final draft of the thesis and publishing the corresponding technical report.

## 2. New Tasks.

One new task area, not previously explored, is proposed. This task area involves analysis of input and output operations, especially to multiplexed channels such as a channel to a local or long haul network or a channel to a terminal-concentrating front-end computer system. The purpose of this analysis is to understand the impact of multiplexed channels on security requirements of a host, and on the complexity of the security kernel. In particular, those parts of the host software that perform connection setup, channel buffering, and multiplexing fan-in of output and fan-out of input apparently are security relevant. These functions are complicated by error handling, use of encryption, and host-to-host or concentrator to host messages. The ARPANET and the Multics to front-end processor interfaces are expected to be the vehicle for this study. Two specific tasks are currently anticipated.

2.1 Experimentally investigate alternative software protocols for the connections between the Multics host and multiplexed input/output channels (for example to the proposed secure front-end processor and to the ARPA network) to identify techniques that minimize the complexity of the security kernel of the host.

2.2 Experimentally investigate the division of function between the Multics host and a front-end processor for the case where the front-end processor acts not only as an I/O stream concentrator but also as an interface to a local or long-haul data communications network. The division of function is to explicitly take into account the impact on the complexity of security kernels not only in the host but also in the front-end processor. It is intended that this task develop techniques that can be transferred into the environment of the hardware and kernel of the proposed secure front-end processor.

### 3. Technology Transfer Tasks.

#### 3.1) Continuing Support Activity.

This task covers the activities needed to communicate the insights and understanding gained over the course of the project to Honeywell and to other organizations creating certifiably secure systems. It involves writing of internal memoranda and publishable papers discussing results from different points of view. It also includes attendance at meetings and discussion and review of ways in which the results of the research projects can be applied to the Multics kernel design.

#### 3.2) Final Report.

A retrospective final report will review all of the engineering proposals made or tested in the course of the kernel design project, to assess their cumulative impact on the size, performance, structure, and simplicity of the kernel required to support security requirements. It is intended that a version of this report will be a publishable paper.

#### Coordination with HISI Project Guardian

Over the past twelve months the style of interaction and coordination with Honeywell has evolved, under a variety of pressures. While early projects in the file system area were carried by M.I.T. to the stage that they could be installed in a production system, more recently the Honeywell teams have had more time and personnel available to review the conceptual issues involved in securing a file management system. In addition, Honeywell has contracted with a team at SRI, which has begun to advise Honeywell on the reimplementaion work required to achieve verifiability. For these reasons, it seems neither appropriate nor desirable that current engineering experiments within the file system area be carried to the same state of production readiness attained by earlier experiments. The proposed work force levels and budget reflect this change in style.

Budget

During the last twelve months, one new faculty member (Liba Svobodova) has joined the Computer Systems Research Division and two current faculty members (Michael Schroeder and David Redell) have announced their plans to leave the Division. Since the rate of work achievable depends on the amount of faculty supervision available, the proposed budget level for salaries and other expenses for the coming twelve months is somewhat lower than that of the previous eighteen months.

The proposed budget includes a plan to taper down to a level that supports cleanup work and technology transfer.

Because computer time may be directly furnished by the government or by Honeywell, it is separately identified.

Indirect costs are calculated at the following rates:

Faculty, Research, and Support Staff:

benefits = 24.5% of salary

overhead = 68% of salary and benefits

Students

overhead = 68% of 50% of salary

---

Budget Summary, Estimates  
(K = \$1000)  
See Attached 633-4

	Salaries and Expenses	Interactive Multics Service	CISL Multics time
7/76 - 9/76	\$ 72 K	\$ 31 K	\$ 13 K
10/76 - 12/76	76 K	34 K	14 K
1/77 - 3/77	41 K	19 K	0
4/77 - 6/77	32 K	13 K	5 K
7/77 - 9/77	32 K	13 K	5 K
Totals	\$253 K	\$110 K	\$37 K
		Total request:	\$400 K

---

Technical Personnel

## Faculty and Research Associates

Jerome H. Saltzer, Professor (Principal Investigator)  
 Liba Svobodova, Assistant Professor  
 David D. Clark, Research Associate

## Research Staff

Nancy Federman  
 Rajendra Kanodia  
 Robert Mabee  
 Kenneth Pogran  
 Douglas Wells

## Graduate Students

Richard Feiertag	Steven Kent
Eugene Ciccarelli	Allen Luniewski
Harry Forsdick	Andrew Mason
Robert Frankston	Warren Montgomery
Harold Goldberg	David Reed
Douglas Hunt	+ 2 new graduate students
Philippe Janson	

## Undergraduate Students

to be assigned

## APPENDIX A

### ENGINEERING A SECURITY KERNEL FOR MULTICS

Michael D. Schroeder

Massachusetts Institute of Technology  
Project MAC and Department of Electrical Engineering and Computer Science  
Cambridge, Massachusetts 02139

#### ABSTRACT

This paper describes a research project to engineer a security kernel for Multics, a general-purpose, remotely accessed, multiuser computer system. The goals are to identify the minimum mechanism that must be correct to guarantee computer enforcement of desired constraints on information access, to simplify the structure of that minimum mechanism to make verification of correctness by auditing possible, and to demonstrate by test implementation that the security kernel so developed is capable of supporting the functionality of Multics completely and efficiently. The paper presents the overall viewpoint and plan for the project and discusses initial strategies being employed to define and structure the security kernel.

Key words: protection, security, privacy, security kernel, operating system structure, verification; Multics.

CR Categories: 4.3

#### Introduction

This paper describes a research project of the Computer Systems Research Division of Project MAC at M.I.T. The objective of the project is to engineer a security kernel for Multics. This project is part of an effort to develop a secure version of Multics that implements the information release constraints of the military security system. Involved in that effort are the Electronic Systems Division of the United States Air Force, the MITRE Corporation, and Honeywell Information Systems Inc.

The paper starts with a brief introduction to the computer security problem and the role that verification of correctness plays in producing a secure system. Next, the plan is outlined for evolving Multics into a system, based on a security kernel, whose security properties are easier to verify. The third section presents a general discussion of a security kernel as the structural basis for a secure system. Finally, to make the nature of the Multics kernel project more precise, a sample of specific activities underway or planned is presented.

---

The research reported in this paper was supported in part by Honeywell Information Systems Inc.; and in part by the Air Force Information Systems Technology Applications Office (ISTAO) and by the Advanced Research Project Agency (ARPA) of the Department of Defense under ARPA order No. 2641 which was monitored by ISTAO under contract No. F19623-74-C-0193.

The issue of security arises when a single computer system provides computation and information storage service to a community of users. As the functional advantages of such shared systems have been recognized, so has the need to include facilities for controlling the access of the various users to the contained information. Many systems now include protection mechanisms providing enforcement of specific patterns of externally specified constraints on the access of executing programs to the contained information [1]. For applications involving sensitive information, the usefulness of a shared system can depend upon the ability of its protection mechanisms to prevent unauthorized release, modification, and sometimes denial of use of the information it contains.

A system is secure if it is known to prevent all actions defined as unauthorized by the specification of its security properties. Penetration exercises involving several different systems have made it apparent that existing shared, general-purpose systems are not secure. In all such systems confronted, a wily user can construct a program that can defeat the access constraints supposedly enforced by the system. Design and implementation flaws exist that provide paths by which the protection mechanisms in the system can be circumvented, thus violating the security of the contained information.

Building a secure system is hard because security places negative requirements on a system. To be secure all possible ways to perform unauthorized actions must be blocked; no way to

This paper was presented at the  
ACM 5th Symposium on Operating  
Systems Principles, University  
of Texas, Austin, Texas  
November, 1975.



circumvent the protection mechanisms can exist. A single flaw in the design or implementation is sufficient to allow a violation of security, and the absence of such flaws cannot be demonstrated by testing. For a system to be considered secure, a convincing logical verification that the implemented system is a correct realization of its security specification is required.

The operating systems of shared, general-purpose computers have a well-known tendency to be extraordinarily large and complex. This size and complexity interacts badly with the negative nature of security requirements. It generates many possible ways to perform unauthorized actions, some of which will go unnoticed by system designers and therefore remain unblocked by the protection mechanisms provided, and increases the probability of exploitable errors in the implementation of the protection mechanisms that are provided. It also makes the required verification of correctness impossible to perform. Available formal program verification techniques are overwhelmed and even verification by manual auditing is thwarted by the inability of one person to comprehend in detail the relevant software.

The negative nature of security properties is an intrinsic problem. The size and complexity of existing, shared, general-purpose systems, however, is not intrinsic. This research project attacks the problem of producing a secure system by exploring ways to reduce the size and complexity of the software that must be correct for claimed constraints on access to information to be enforced.

#### Method of Attack

The problem of constructing a secure system has attracted considerable interest recently and is being attacked with a variety of different strategies [2]. One approach being explored involves constructing a formal specification for the desired security (and other) properties of a system, and then, through a methodical, top-down design and implementation process, building a matching, operational system. The correspondence of the implemented system with the formal security specification is to be proved using formal program verification techniques. The hope is to achieve a level of confidence in the match of the system to the specification similar to the confidence that a mathematician has in the result of a well-wrought proof. A step necessarily left to human intuition is determining how appropriate the security properties expressed in the formal specification are to a particular real-world application. Three current projects [3,4,5] are trying to produce simple, experimental examples of secure systems using different versions of this approach.

At the other end of the spectrum are efforts [6,7] to find, catalog, and repair security flaws in existing systems. The goals are to convince skeptics that the computer security problem is real, to understand the sort of flaws that can be exploited, and to try to reduce the ease with which available systems can be penetrated.

Our project to reduce the size and complexity of the software that must be verified correct to

produce a secure system falls somewhere between these two extremes. Our plan is to evolve Multics [8], a general-purpose, remotely accessed, multiuser system, into a prototype system with the same essential features, but with a small and simple protected central core of software. This core will be a security kernel embodying all mechanisms necessary to enforce the claimed constraints on access to information in the system. The goal is a kernel sufficiently small, well-structured, and easy to understand that feasibly an expert could verify its correctness through manual auditing. Such a kernel also may be susceptible to verification through formal techniques, although the programs may have to be rewritten in a more appropriate style. The kernel will enforce access constraints that combine nondiscretionary controls reflecting the information release policies of the military security system and the discretionary controls on information sharing that were part of the original Multics design. (1)

The choice to evolve an existing system rather than produce a new one follows from our interest in problems of system structure, rather than in techniques for formal specification of security properties, structured programming, or formal verification of program correctness. Our goal is developing a security kernel that can be demonstrated to support the full set of functions that are desirable in a shared, general-purpose system. With the current understanding of computer systems, it is hard to have confidence that the full implications of a system structure are understood without complete implementation. Designing a new system without building it or building a simple, experimental system (2) would not allow the completeness of our kernel design to be tested adequately. Thus, to start fresh would mean undertaking the entire job of building a new, general-purpose system. By developing a security kernel for an existing, general-purpose system, we avoid this enormous effort.

A combination of factors makes Multics well suited as a base from which to engineer a security kernel. To start with, Multics provides a full set of functional capabilities, including high-bandwidth direct sharing of information among computations [11]. In addition, Multics has been developed from the ground up to protect the information it contains from unauthorized access. It already includes general protection mechanisms to control information sharing among users [10] and provides direct hardware support for some of these mechanisms [12]. Minimal features to support the information release constraints of the military security system recently have been added [13]. Thus, the system exhibits a set of security

---

(1) A formal specification [9] of the nondiscretionary controls is being developed by a group at MITRE. An informal specification of the discretionary controls is available in [10].

(2) Two examples of features usually left out of experimental systems that can complicate the kernel of an complete, general-purpose system are the storage quota and backup mechanisms mentioned in a later section.

properties that would be interesting to implement correctly and provides hardware support that will make the job easier. Also, the system is well organized for evolution and modification because it is relatively modular, is largely written in PL/I [14], and was originally constructed with evolution as a primary objective. Finally, Multics provides two unique opportunities to test and export the results. First, because Multics is a commercially available product, ideas developed in the course of this research that simplify the system's structure without changing its functionality or reducing its efficiency can be added to the standard system. Second, the security kernel being produced is serving as the structural basis for the secure version of Multics being developed by the Air Force, MITRE, and Honeywell in an effort of which our project is part.

#### The Security Kernel

A security kernel, a minimal, protected core of software whose correct operation is sufficient to guarantee enforcement of the claimed constraints on access, is the structural basis for organizing a secure system. Rather than being dispersed throughout the system software, all protection mechanisms are collected in the kernel, so that only this kernel need be considered in order to verify that the specified security properties are implemented correctly.

The security specification that a particular system must match limits how small and simple (3) the kernel can be. The patterns of access constraints to be enforced is an obvious factor. The set of abstract objects and operations to be controlled may be even more important. In Multics, it appears that most of the mechanism in the kernel will implement the abstract objects protected by the system, for example, processes, segments, directories, and I/O streams, and will be relatively independent of the specific patterns of access constraints enforced by the system.

A characterization of the mechanisms that should be included in a security kernel can be obtained by viewing the security specification as a set of constraints on the interaction of the various computations that occur in a computer system. The protection mechanisms of the system prevent one computation from exerting an unauthorized influence on the input, progress, or output of another. Permanently stored data is one form of the input and output of computations. This view suggests that the security kernel should embody all system-provided mechanisms that are common to more than one computation (domain), because a common mechanism is required if one computation is to influence another. A mechanism is common to two computations if it uses some set of data items whose value one computation can influence and the other can notice. The influence and notice may be direct--one writes into a data item and the other reads it--or indirect--the invocation of a procedure by one somehow alters the procedure's internal state so that the outcome of

---

(3) Unfortunately, no objective measure of overall complexity is known. The degree of complexity must be estimated subjectively.

an invocation by the other is affected. Common mechanisms are required to implement any explicit or implicit communication among computations. Thus, mechanisms implementing information sharing, interprocess communication, and physical resource multiplexing must be common. If no communication is involved, however, then a common mechanism is not required to implement a function. Common mechanisms carry a built-in risk--they make it possible for the computation of one user to exert unauthorized influence over the computations or data of another. Malicious users must exploit flaws in common mechanisms to work their will. To thwart such malicious activity the system designers must ensure that the common mechanisms have no exploitable design or implementation flaws, and must protect the common mechanisms against tampering. Thus, a security kernel should be the least amount of common mechanism necessary to implement the patterns of information sharing, interprocess communication, and physical resource multiplexing that are desired in the system. (4)

Although a security kernel contains all the mechanisms that must be verified as correct to produce a secure system, a correct kernel does not guarantee the integrity of all computations or stored data in a system. Nonkernel software still can cause undesired release of information, modification of information, or denial of its use. But if the kernel is correct, then these undesired results will not be unauthorized. To understand the meaning of this distinction, consider the nonkernel software as grouped in four categories.

First, there are the system-provided programs that execute as part of user computations. These include the library subroutines available in most systems and all the programs usually part of a supervisor that are not included in a security kernel. These system-provided programs are not common mechanisms, even though in many systems all computations may share the same nonwriteable code that embodies their algorithms. This is so because a private copy of the alterable part of these programs, the variable data, is provided for each computation. Because they are private mechanisms, no interuser interaction can occur through them. Private mechanisms may contain errors, but these errors can be triggered only by the actions of the computation that they might damage as a result. If one assumes that the system programmers who constructed them are not malicious and did not willfully plant "trojan horses," then the mistakes caused by these system-provided programs will decrease in time as all normally used functions are exercised. Under these circumstances the threat posed by a potential random error causing undesired release, modification, or denial of a user's data is acceptable for most applications. Unlike the common mechanisms of the security kernel, the nonkernel system-provided programs are not

---

(4) According to this characterization of a security kernel, a usual reason for including a mechanism in a supervisor, to protect it from accidental breakage caused by errors in user code, is not in itself sufficient to include a mechanism in a security kernel. Nonkernel mechanisms can be protected by placing them in other domains that are private to each user's computation.

susceptible to willful exploitation by other users. In any case, a user unsatisfied with their trustworthiness may choose not to use them and substitute his own programs.

The second category of nonkernel software is programs constructed by a user and executed in that user's computations. Any undesired result caused by errors in these is the user's own problem. The only possible help to the user would be providing tools to aid verifying the correctness of his own programs.

The third category, possible in many systems, is programs borrowed from other users. These are a real danger to the borrower's computations. Borrowed programs can contain trojan horse code maliciously constructed to cause results undesired by the borrower. (5) A user should borrow programs from another only when the borrower has reason to trust the lender. The inclusion of security kernel facilities to support user-constructed protected subsystems can reduce the potential damage such a borrowed trojan horse can do by isolating it in a separate domain of the borrower's computation. Due to the confinement problem [15], however, and also to the possibility of a borrowed program disrupting the borrower's computation simply by calculating incorrect results, a user-initiated verification of the borrowed program is the only complete protection.

The fourth category is common mechanisms that a group of users sets up to implement some function involving interuser communication or coordination. For example, a team producing a new compiler might set up a program development subsystem with a common mechanism to control installation of new modules into the evolving compiler. Such a mechanism makes the group susceptible to undesired interaction in the same way that an insecure supervisor does for the whole user community. If a user agrees to become party to such a common mechanism, then he must satisfy himself of its trustworthiness.

From considering these four categories of nonkernel software it is apparent that the essential mechanisms to verify correct are the common mechanism of the security kernel. The security kernel has initial control of all paths for the interaction of computations and every user of the system is forced to rely upon it. A correct kernel provides the tools with which a user may protect his computations and data against unwanted interference from the computations of other users. In a system providing for direct sharing of programs and data, however, users can agree to cooperate in ways that the security kernel cannot control. The kernel can prevent such sharing unless it is explicitly authorized, but cannot completely control the interaction between users that agree to share. The security kernel prevents

---

(5) This is a special case of a common mechanism. The data item whose value the lender can cause to change (and thereby influence the computation of the borrower) is the code of the borrowed program itself. Even if the program is nonwritable when borrowed, it was written by the lender when constructed.

activities that the security specification for the system defines as unauthorized, but not all undesired results are caused by unauthorized activities.

A fifth category of nonkernel software also needs to be considered. One important technique for simplifying the structure of the security kernel is writing it with a high-level programming language. Using a high-level language to generate the kernel seems to require that the compiler also be verified correct, a troubling thought since the compiler may well be larger than the kernel. Verification of correct function may be less of a problem for the compiler, however, than for the kernel. The kernel must work correctly for all possible inputs; the compiler must compile correctly only the specific programs of the kernel--not all possible programs. Thus, the compiler's effect on the kernel can be determined by comparing the source code specifications for each kernel module with the compiler-produced object code implementation, a task much simpler than verifying the compiler correct for all possible source programs. (6)

#### A Kernel for Multics

Engineering a security kernel for a system requires isolating a minimum set of functions capable of supporting the system and finding a way to structure the kernel to facilitate verifying its correctness. Our plan is to produce a security kernel for Multics by removing nonkernel mechanisms from the supervisor, and restructuring the remaining kernel and partitioning it into multiple protection domains. This section describes these three interrelated categories of activities and provides specific examples of work underway or planned in each. The intention of the section is to communicate the spirit of the work rather than to discuss thoroughly the various activities. The detailed results of individual activities are being communicated in other reports [16,17,18].

The first category of activities is taking functions not requiring implementation as common mechanisms out of the supervisor to be implemented in the user domains of a process. In many cases this transfer involves undoing a pattern caused by a performance characteristic of the original Multics implementation for the Honeywell 645 computer. For that machine, the multiple protection domains of a process, the so-called protection rings, were simulated in software. Cross-ring calls were quite expensive: a call that went from a user ring in a process to the supervisor ring cost much more than a call that did not change protection rings. The effect on system structure can be seen by considering two procedures, A and B. If a single invocation of A can result in a flurry of calls from A to B, then there is a performance benefit in placing both A and B in the supervisor, even if only B needs to be part of the protected, common mechanism. As a result of this performance characteristic of the

---

(6) Use of a high-level language for kernel construction can generate a dilemma. An optimizing compiler can increase system efficiency, but may make impossible matching object code with source code to verify correct compiler function.

(45 implementation, many functions that did not need to be implemented as common mechanism were included in the supervisor.

The present hardware base for Multics, the Honeywell Level 68, implements protection rings in hardware. Calls from one ring to another cost no more than calls inside a ring. With the performance penalty associated with supervisor calls removed, many modules included in the supervisor for performance reasons rather than protection reasons now can be removed. (7)

Actually, removing a module from the supervisor is more difficult than the example suggests. In most cases, the common and private parts of a facility are not neatly packaged in separate procedures but are intricately intertwined in the same supervisor procedures and data bases. The key problem is decomposing the supervisor into common and private primitive functions.

Most removal activities have centered on the file system. In one activity, now completed, the functions of dynamic intersegment linking and direction of file system searches to satisfy symbolic references have been removed from the supervisor [16,17]. This removed a vulnerable and complex mechanism from the supervisor. The vulnerability is a result of the linker having to accept user-constructed code segments as input data. Numerous accidents had shown that such a complex "argument", if maliciously misstructured, could cause the linker to malfunction while executing in the supervisor. Removing the linker eliminated 10% of the gate entry points into the supervisor and 6% of the total object code. The linker's removal also demonstrated that linking procedures together across protection boundaries (rings) could be done without resorting to a mechanism common to both domains.

A second completed activity removed from the supervisor the facilities for managing the association between reference names and the segments in the address space of a process [18]. Taking this naming mechanism out of the supervisor required that a data base central to the management of the address space, the Known Segment Table, be split into a private part that maintains the binding between reference names and segment numbers and a common part that maintains the binding between segment numbers and segments. Removal reduced fivefold the size of the protected code needed to manage the address space of a process. It also provided a new, simpler interface to the file system portion of the supervisor. Instead of identifying a directory by character string tree name locating it in the file system hierarchy, a segment number now is used. The algorithms for following a tree name through the file system hierarchy to locate the named element are now implemented by procedures executing in the user ring. (The actual file system hierarchy remains

---

(7) There may still exist other performance penalties associated with removing functions from the supervisor that will inhibit production of the smallest possible kernel. One goal of the research is to understand better the performance cost of security.

protected inside the security kernel.) Because tree names are now resolved one component at a time, the kernel had to learn to lie on occasion about the existence of file system directories. This deception keeps the kernel from divulging a directory's existence before an accessible element in the subtree rooted by the directory is found.

An activity under investigation involves making most of system initialization execute once, in a user environment of a previous system version, instead of executing inside the supervisor each time the system is started. The change is to produce a system tape with a bit pattern that, when loaded into memory, manifests a fully initialized system. At present the system bootstraps itself in a complex way each time it is loaded from a tape containing the separate pieces. The new pattern of operation removes most initialization software from the kernel. The correct initialization of the kernel also should be easier to verify, for most of the work of initialization will occur when the tape is made, in the stable environment of a fully initialized system.

Another activity is exploiting the equivalence between entering a protected subsystem and creating a new process in response to a user's log-in. The goal is to make a single mechanism do both tasks, so that the privileged, protected code used to authenticate and log-in users can be executed in the user code protection environment. The authentication algorithm still must be verified not to malfunction if the user trying to log-in behaves unexpectedly. Such verification should be easy, however, since the user/system interface severely limits user behavior.

The second category of activities is restructuring mechanisms that must remain in the kernel. Such activities can reduce both the size and the complexity of the kernel. In some cases a piece of the kernel can be eliminated and its function assumed by another kernel mechanism. For example, one activity is exploring replacement of all external I/O mechanisms (to terminals, tape drives, card readers, card punches, and printers) with the ARPA Network attachment. This would eliminate many special mechanisms for managing I/O devices and leave a single mechanism for managing the network attachment. Internal I/O functions (for managing the virtual memory, performing backup, and loading the system) would still be managed in the kernel.

A proposed buffering strategy for network input uses the virtual memory to provide a core resident buffer that appears to be infinite in length. With the present circular buffer, which has to be used over and over, complex mechanisms are required to cope with messages not removed before a complete circuit of the buffer is made. The circular buffer scheme is really providing a special-purpose storage management facility. The proposed infinite buffer uses instead the standard storage management facility of the system--the virtual memory.

Several restructuring activities involve the implementation and use of processes. One activity, now nearing the end of the design phase, is a reimplementing of processes using two layers of

mechanism. (8) This new design simplifies the interaction of the process implementation with the virtual memory management mechanisms. It also simplifies the base-level interprocess communication mechanisms of the system. The first level of mechanism multiplexes the processors into a larger fixed number of virtual processors. Because the number of virtual processes is fixed, this layer need not depend on the mechanisms for managing the virtual memory. Several of the virtual processors are permanently assigned to implement processes for the dedicated use of other kernel mechanisms, including the virtual memory management mechanism. The remaining virtual processors are multiplexed by the second layer of the process implementation into any desired number of full Multics processes that execute in the virtual memory. Use of the proposed base-level interprocess communication facility can be controlled with the standard memory protection mechanisms.

The implementation (implied above) of certain kernel mechanisms as asynchronous parallel processes also simplifies system structure, which now forces many supervisor mechanisms into sequential algorithms. The virtual memory mechanisms for moving pages among the three levels of the memory hierarchy are a good example. Whenever a missing page fault occurs in a process, the fault handler attempts to initiate the transfer of the desired page from bulk store or disk to primary memory. This can be done only if a free primary memory block is available. If none is available the fault handler must move a page from primary memory to the bulk store to make room. This, in turn, is possible only if a free block of bulk store is available. If not, a page must be moved from the bulk store, via primary memory, to a disk. At present, this series of steps occurs sequentially with page control executing in the process that took the page fault, and then in various other user processes that happen to receive the subsequent I/O interrupts. The new scheme involving multiple dedicated processes is much simpler. One process makes sure that some small number of free primary memory blocks always exist. Whenever the number of free primary memory blocks drops below that number, this process transfers pages to bulk store. Another process keeps space free on the bulk store by moving pages to disk when required. Signals from processes that have taken a page fault and not found free primary memory blocks activate the process that frees primary memory. The process that frees bulk store blocks is driven in a similar manner by signals from the process that frees primary memory blocks. The path taken by a user process on a page fault is greatly simplified. This process can just wait until a primary memory block is free and then initiate the transfer of the desired page into primary memory.

Interrupt handling is another possible application of processes in the kernel. Each interrupt handler would be assigned its own process in which to execute, rather than being forced to inhabit whatever user process was running when the interrupt occurred. As a result, the system

---

(8) This idea is being explored by others as well [3,19].

interrupt interceptor could turn each interrupt into a signal to the corresponding process. Being processes, the interrupt handlers could use the normal system interprocess communication mechanisms to coordinate their activities with one another and user processes, greatly simplifying their structure. The problem to solve here is implementing the interrupt processes so that system performance is not degraded.

A major activity is restructuring the portion of the file system that must remain in the kernel. This software implements the directory hierarchy and manages the virtual memory at the level of segments. Work in this area is just beginning, but three changes with a potential, significant cumulative effect are promising first steps. One change is removing physical attributes of segments from directory entries. The physical attributes will be recorded in data bases associated with the various secondary storage devices. (This modification is being implemented as part of a file system overhaul done by Honeywell for other reasons.) The second change is removing storage quotas from directory entries, recording them instead in a separate (possibly hierarchical) data structure. The third is eliminating the dependence of the backup mechanism on the date-time modified information recorded in directory entries and reflected up the hierarchy toward the root. Eliminating this dependence will allow the date-time modified item to be removed from directory entries as well. Backup will be driven by a queue of requests from the machinery that controls deactivation of segments from primary memory. As a result of these three changes, it appears that the management strategy for the Active Segment Table can be modified to eliminate the need for holding active the superior directories of an active segment [11].

The various restructuring activities eventually will extend to all parts of the kernel, and to its overall structure.

The third category of activities is partitioning the kernel into differently protected pieces to modularize the job of matching the kernel to the system security specification. (9) The specific projects in this category are not as well developed as the others. There appear to be several different design principles with which to generate the kernel partitions, and it is not yet clear which produces the kernel that is easier to verify. To illustrate two possible approaches to partitioning a kernel into multiple domains, imagine that the security specification is expressed as a set of security properties, each of which must be met. One design principle is to divide the kernel into domains arranged so that each property is implied by a subset of the domains. Then, to verify that the kernel implements the security specification, an independent verification of each property is required, but each involves only a subset of the

---

(9) Partitioning is really the same problem as dividing the kernel into separate procedures and data bases, with the extra property that the modularity is enforced by the system's protection mechanisms

domains in the kernel. Another design principle is to ignore any structure suggested by the security properties and divide the kernel into domains according to some other principle of structured programming (for example, Parnas' notion of information hiding [20]) so that each domain has a simple interface behavior specification. Verification of each of the security properties may involve all the kernel domains, but once each domain has been verified to match its interface specification, then only these specifications need be considered to verify each security property. Which of these two approaches is preferable--or indeed, whether they really are different approaches--remains to be seen.

Two specific methods for partitioning the Multics kernel are available. The first is dividing the part of the kernel that is in the address space of each process into multiple layers in different rings of protection. The second is placing some of the kernel processes in separate address spaces and also using the protection rings to layer them. Several suggestions have been made for layering the part of the kernel that is in each user process. One is that the bottom layer implement a file system in which all segments were named by system-generated unique identifiers. The next layer would implement a naming hierarchy on top of the primitive first-layer file system. Another suggestion is that mechanisms to provide the nondiscretionary controls on the flow of information among processes be implemented at the bottom and mechanisms to control discretionary sharing within the constraints of the nondiscretionary controls be implemented in the next layer. This last suggestion is particularly intriguing, because if correctly done, the notion of minimizing common mechanisms would be well supported. The second-layer mechanisms would be common only within the access constraints enforced by the first layer.

Partitioning through use of separate address spaces for kernel processes is being considered in the case of the processes that manage the various system resources. The protection rings in these processes then could be used to separate the policy and mechanism components of the resource managers. (10) For example, the process described earlier that removed pages from primary memory could be given its own address space with multiple rings. Programs in the most privileged ring would implement the mechanics of page removal, providing gate entry points for requesting the movement of a particular page from primary memory to a particular free block on the bulk store, and for obtaining usage information about pages in primary memory. The policy algorithm that decides which page to remove when another free primary memory block needs to be generated would execute in a less privileged ring, calling the gate entry points to collect the necessary usage statistics and to do the actual moving, once a decision was made. The policy algorithm, however, could never read or write the contents of pages, learn the segment to which each page belonged, or cause one page to overwrite

---

(10) Separation of policy from mechanism is a structural principle that has been explored by many others [21,22,23].

another. Such operations would not be available in its ring of execution. The result is that the policy algorithm could never cause unauthorized use or modification of the information stored in the pages. It could only cause denial of use. Under the circumstance that denial of use was deemed less serious than the other security violations, the policy algorithm need not be as carefully verified as the rest of the kernel. It appears that the idea of separating policy from mechanisms applies to all resource management processes.

### Conclusion

This paper has presented the plan for a research project to evolve the Multics supervisor into a security kernel capable of supporting the functionality of Multics completely and efficiently. It has described a sample of the specific strategies being employed. The broad objective is finding ways to reduce the size and complexity of the software that must be correct for a shared general-purpose system to be secure. Reduced size and complexity of security-relevant software is a prerequisite to performing a convincing logical verification that a system correctly implements the claimed access constraints, no matter what verification techniques are used. Without such verification of correctness, a system cannot be considered secure.

At the time this paper is being written, the project has run for about half of its intended four year span, and most of the initial tasks are nearing completion. So far, the expected reductions in size and simplifications of structure of the security-relevant software seem to be occurring. It is too soon to tell, however, whether the security kernel for Multics that will result will be sufficiently small and simple to be understood in detail by one person.

### Acknowledgements

In describing a group project of the Computer Systems Research Division of Project MAC at M.I.T., this paper discusses the work of several faculty members, graduate students, and staff members in the Division. Rather than list all here, they will receive credit for their contribution as specific activities are completed and reported separately. Preparation of this paper was aided by written commentaries on various drafts provided by E. Cohen, F. Corbató, R. Fabry, D. Hunt, P. Janson, D. Reed, J. Saltzer, and R. Schell, and by comments from D. Clark, A. Jones, and D. Redell.

### References

- [1] J. H. Saltzer and M. D. Schroeder, "The Protection of Information in Computer Systems," Proc. IEEE 63, 9 (Sept. 1975), pp. 1278-1308.
- [2] J. H. Saltzer, "Ongoing Research and Development on Information Protection," ACM Operating Sys. Review 8, 3 (July 1974), pp. 8-24.

- [3] L. Robinson, et. al., "On Attaining Reliable Software for a Secure Operating System," Int. Conf. on Reliable Software, Apr. 1975, pp. 267-284.
- [4] G. J. Popek and C. S. Kline, "A Verifiable Protection System," Int. Conf. on Reliable Software, Apr. 1975, pp. 294-304.
- [5] W. L. Schiller, "Design of a Security Kernel for the PDP-11/45," The MITRE Corp. Tech. Rep. ESD-TR-73-294, Dec. 1973.
- [6] J. Carlstedt, R. L. Bisbey II, and G. J. Popek, "Pattern-Directed Protection Evaluation," U. of So. Calif. Inf. Sci. Institute Tech. Rep. ISI/RR-75-31, June 1975.
- [7] P. A. Karger and R. R. Schell, "Multics Security Evaluation: Vulnerability Analysis," Air Force Elec. Sys. Div. Tech. Rep. ESD-TR-74-193, Vol. II, June 1974.
- [8] F. J. Corbató, J. H. Saltzer, and C. T. Clingen, "Multics - the First Seven Years," AFIPS Conf. Proc. 40 (SJCC 1972), pp. 571-583.
- [9] D. E. Bell and L. J. LaPadula, "Secure Computer Systems," The MITRE Corp. Tech. Rep. ESD-TR-73-278, Nov. 1973.
- [10] J. H. Saltzer, "Protection and the Control of Information Sharing in Multics," CACM 17, 7 (July 1974), pp. 388-402.
- [11] A. Bensoussan, C. T. Clingen, and R. C. Daléy, "The Multics Virtual Memory: Concepts and Design," CACM 15, 5 (May 1972), pp. 308-318.
- [12] M. D. Schroeder and J. H. Saltzer, "A Hardware Architecture for Implementing Protection Rings," CACM 15, 3 (Mar. 1972), pp. 157-170.
- [13] Honeywell Information Systems Inc., "Design for Multics Security Enhancements," Air Force Elec. Sys. Div. Tech. Rep. ESD-TR-74-176, 1974.
- [14] F. J. Corbató, "PL/I as a Tool for System Programming," Datamation 15, 6 (May 1969), pp. 68-76.
- [15] B. W. Lampson, "A Note on the Confinement Problem," CACM 16, 10 (Oct. 1973), pp. 613-615.
- [16] P. A. Janson, "Removing the Dynamic Linker from the Security Kernel of a Computer Utility," S.M. Thesis, Dept. of Elec. Eng. and Comp. Sci., M.I.T., June 1974. (Also available as Project MAC Tech. Rep. MAC-TR-132, June 1974.)
- [17] P. A. Janson, "Dynamic Linking and Environment Initialization in a Multi-Domain Computation," ACM 5th Symp. on Operating Sys. Principles, Austin, Texas, Nov. 1975.
- [18] R. G. Bratt, "Minimal Protected Naming Facilities for a Computer Utility," S.M. Thesis, Dept. of Elec. Eng. and Comp. Sci., M.I.T., July, 1975. (Also available as Project MAC Tech. Rep. MAC-TR-156, Sept. 1975.)
- [19] A. R. Saxena and T. H. Bredt, "A Structured Specification of a Hierarchical Operating System," Int. Conf. on Reliable Software, Apr. 1975, pp. 310-318.
- [20] D. L. Parnas, "On the Criteria to be Used in Decomposing Systems Into Modules," CACM 15, 12 (Dec. 1972), pp. 1053-1058.
- [21] M. J. Spier, T. N. Hastings, and D. N. Cutler, "An Experimental Implementation of the Kernel/Domain Architecture," ACM Operating Sys. Review 7, 4 (Oct. 1973), pp. 8-21.
- [22] G. R. Andrews, "COPS - A Protection Mechanism for Computer Systems," Computer Sci. Teaching Lab., Univ. of Wash., Tech. Rep. 74-07-12, July 1974.
- [23] W. Wulf, et. al., "HYDRA: The Kernel of a Multiprocessor Operating System," CACM 17, 6 (June 1974), pp. 337-345.