

## INFORMATION TO USERS

This material was produced from a microfilm copy of the original document. While the most advanced technological means to photograph and reproduce this document have been used, the quality is heavily dependent upon the quality of the original submitted.

The following explanation of techniques is provided to help you understand markings or patterns which may appear on this reproduction.

1. The sign or "target" for pages apparently lacking from the document photographed is "Missing Page(s)". If it was possible to obtain the missing page(s) or section, they are spliced into the film along with adjacent pages. This may have necessitated cutting thru an image and duplicating adjacent pages to insure you complete continuity.
2. When an image on the film is obliterated with a large round black mark, it is an indication that the photographer suspected that the copy may have moved during exposure and thus cause a blurred image. You will find a good image of the page in the adjacent frame.
3. When a map, drawing or chart, etc., was part of the material being photographed the photographer followed a definite method in "sectioning" the material. It is customary to begin photoing at the upper left hand corner of a large sheet and to continue photoing from left to right in equal sections with a small overlap. If necessary, sectioning is continued again — beginning below the first row and continuing on until complete.
4. The majority of users indicate that the textual content is of greatest value, however, a somewhat higher quality reproduction could be made from "photographs" if essential to the understanding of the dissertation. Silver prints of "photographs" may be ordered at additional charge by writing the Order Department, giving the catalog number, title, author and specific pages you wish reproduced.
5. PLEASE NOTE: Some pages may have indistinct print. Filmed as received.

### **Xerox University Microfilms**

300 North Zeeb Road  
Ann Arbor, Michigan 48106

76-514

DENNING, Dorothy Elizabeth Robling, 1945-  
SECURE INFORMATION FLOW IN COMPUTER SYSTEMS.

Purdue University, Ph.D., 1975  
Computer Science

**Xerox University Microfilms**, Ann Arbor, Michigan 48106

SECURE INFORMATION FLOW IN COMPUTER SYSTEMS

A Thesis

Submitted to the Faculty

of

Purdue University

by

Dorothy Elizabeth Robling Denning

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

May 1975

PURDUE UNIVERSITY

Graduate School

This is to certify that the thesis prepared

By Dorothy Elizabeth Robling Denning

Entitled Secure Information Flow in Computer Systems

Complies with the University regulations and that it meets the accepted standards of the Graduate School with respect to originality and quality

For the degree of:

Doctor of Philosophy

Signed by the final examining committee:

Robert A. Schwetman, chairman  
Samuel Rosen  
Kenneth O'Mahen  
Cedric A. Ahlstrom

Approved by the head of school or department:

May 1 1975 [Signature]

To the librarian:

This thesis <sup>is</sup> ~~is not~~ to be regarded as confidential

Robert A. Schwetman  
Professor in charge of the thesis

## ACKNOWLEDGMENTS

I am grateful to Herbert D. Schwetman for his assistance in the preparation of this thesis, and to Kenneth J. Omahen, Saul Rosen, and Andrew B. Whinston for their advice and support. I wish to thank Carl Ellison and Butler Lampson for their suggestions on handling program interrupts. I am indebted to my husband, Peter J. Denning, for his criticisms, suggestions, patience, and constant encouragement.

I wish also to acknowledge IBM for awarding me a fellowship, and the National Science Foundation for their support through grants GJ-43176 and GJ-41289.

## TABLE OF CONTENTS

	Page
LIST OF FIGURES . . . . .	v
LIST OF SYMBOLS . . . . .	vii
ABSTRACT . . . . .	x
CHAPTER	
1 INTRODUCTION . . . . .	1
1.1 Problem Statement . . . . .	2
1.2 Previous Research . . . . .	7
1.3 Organization and Results of Thesis . . . . .	10
2 A MODEL OF INFORMATION FLOW . . . . .	13
2.1 Description of the Model . . . . .	16
2.2 Security Requirements . . . . .	27
2.3 The Lattice of Information Flow . . . . .	37
2.3.1 Derivation . . . . .	37
2.3.2 Examples . . . . .	46
2.3.3 Additional Lattice Properties . . . . .	49
2.4 Classification of Systems . . . . .	50
2.4.1 Single-Class, Static Binding . . . . .	54
2.4.2 Single-Class, Dynamic Binding . . . . .	55
2.4.3 Single-Class, Static and Dynamic Binding . . . . .	57
2.4.4 Multi-Class, Static Binding . . . . .	60
2.4.5 Multi-Class, Dynamic Binding . . . . .	67
2.5 Determinacy Requirements . . . . .	68
2.5.1 Single-Class, Dynamic Binding . . . . .	69
2.5.2 Single-Class, Static and Dynamic Binding . . . . .	75
2.5.3 Multi-Class, Dynamic Binding . . . . .	76
3 PROGRAM CERTIFICATION UNDER STATIC BINDING . . . . .	78
3.1 Expressions, Assignment, and I/O of Elementary Data Structures . . . . .	83
3.2 Conditional Statements . . . . .	88
3.3 Gotos . . . . .	92
3.4 Interrupts . . . . .	99

CHAPTER	Page
3 (cont.)	
3.5 Complex Data Structures . . . . .	105
3.5.1 Arrays . . . . .	105
3.5.2 Records . . . . .	107
3.5.3 List Structures . . . . .	108
3.6 Procedure Calls . . . . .	110
3.7 Applications . . . . .	114
3.7.1 Covert Channels . . . . .	114
3.7.2 Confinement . . . . .	115
3.8 3.7.3 Data Bases and Data Banks . . . . .	117
4 PROGRAM CERTIFICATION UNDER DYNAMIC BINDING . . . . .	119
4.1 Hardware Mechanism . . . . .	122
4.2 Software Mechanism . . . . .	129
4.3 Software Mechanism for Selective Updating . . . . .	136
4.4 Extensions . . . . .	148
4.4.1 Files and I/O . . . . .	148
4.4.2 Arrays . . . . .	149
4.4.3 Procedure Calls . . . . .	150
4.4.4 Interrupts . . . . .	151
4.4.5 Static Binding . . . . .	153
4.4.6 Verify Statement . . . . .	155
4.5 Applications . . . . .	158
5 CONCLUSIONS . . . . .	159
5.1 Summary . . . . .	159
5.2 Areas of Future Research . . . . .	161
LIST OF REFERENCES . . . . .	163
VITA . . . . .	168

## LIST OF FIGURES

Figure	Page
1.1-1. Communication Channel from a to b . . . . .	4
2.1-1. Examples of Partially Ordered Information Flow Structures . .	18
2.1-2. Transition Operators . . . . .	24
2.2-1. Information Flow from A to B by an Operation t . . . . .	28
2.2-2. Operations that Cause Information Flow from A to B . . . . .	31
2.2-3. Examples of Operation Sequences that Cause Information Flow from A to B . . . . .	33
2.3-1. Example of a Function Dependent Flow Relation . . . . .	41
2.3-2. Lattice of Subsets of $X = \{x, y, z\}$ . . . . .	48
2.4-1. Information Flow in a Single-Class, Static Binding System . .	56
2.4-2. Information Flow in a Single-Class System with Static and Dynamic Binding . . . . .	58
2.4-3. Example of an Information Flow Structure Satisfying Fenton's Model . . . . .	66
2.4-4. Example of Figure 2.4-3 Extended to a Lattice . . . . .	66
2.5-1. Example of a Process System . . . . .	70
2.5-2. Example of a Non-Deterministic System . . . . .	73
3.1-1. Syntax and Certification Semantics for Assignment and I/O . .	85
3.1-2. Certification Tree of an Assignment Statement . . . . .	87
3.1-3. Certification Tree of an Input Statement . . . . .	89
3.2-1. Syntax and Certification Semantics for Conditional Statements . . . . .	91



Figure	Page
3.2-2. Certification Trees of Conditional Statements . . . . .	91
3.3-1. Example of a Program with Gotos . . . . .	95
3.3-2. Control Flow Graph of Program of Figure 3.3-1 . . . . .	98
3.5-1. Syntax and Certification Semantics for Array References . .	106
3.5-2. Syntax and Certification Semantics for Records . . . . .	109
3.6-1. Syntax and Certification Semantics for Procedure Calls . .	112
3.6-2. Syntax and Certification Semantics for Restricted Procedure Calls . . . . .	112
4.1-1. Execution Traces of an Insecure Program . . . . .	128
4.2-1. Transformation Algorithm T1 . . . . .	131
4.2-2. Transformation of Insecure Programs . . . . .	132
4.3-1. Transformation Algorithm T2 . . . . .	138
4.3-2. Example of a Program Transformed by T1 . . . . .	145
4.3-3. Example of a Program Transformed by T2 . . . . .	146
4.4-1. Modifications to T2 to Support Static Binding . . . . .	154
4.4-2. Syntax of Verify Statement . . . . .	156

## LIST OF SYMBOLS

Symbol	Meaning	Page
$\rightarrow$	Flow relation and partial ordering relation on the security classes SC . . . . .	16
$\Rightarrow$	Derives relation on storage objects, security classes, and items of information . . . . .	29
$\oplus$	Class-combining operator and least upper bound operator on pairs of security classes of SC . . . . .	16
$\oplus$	Least upper bound operator on sets of security classes of SC . . . . .	42
$\otimes$	Greatest lower bound operator on pairs of security classes of SC . . . . .	43
$\otimes$	Greatest lower bound operator on sets of security classes of SC . . . . .	44
$-$	Operator denoting the security class of an element . . . . .	38
$\leftarrow$	Partial ordering relation on the processes P . . . . .	69
A,B	Security classes of SC . . . . .	16
a,b	Logical storage objects of N . . . . .	19
$\underline{a}$	Security class of object a . . . . .	38
$(\underline{a})_i$	Security class of object a in state $s_i$ . . . . .	38
AM	Access control matrix . . . . .	21
c	Class sequence . . . . .	71
$D_p$	Domain (input set) of process p . . . . .	69
d	Structure index of a program . . . . .	133
F	Set of files . . . . .	50
FM	Information flow model . . . . .	16

Symbol	Meaning	Page
FS	Information flow structure . . . . .	16
f	Function or operation . . . . .	17
H	Highest security class of SC . . . . .	45
h	Table of upper bounds on Security Classes . . . . .	137
HS	Hardware stack of security classes . . . . .	123
<u>HS</u>	Security class on top of HS . . . . .	123
IFD	Immediate forward dominator block . . . . .	97
L	Lowest security class of SC . . . . .	39
l	Table of lower bounds on security classes . . . . .	137
M	Set of physical memory locations . . . . .	19
N	Set of logical storage objects ( $N = F \cup W$ ) . . . . .	19
P	Set of processes . . . . .	19
p,q	Processes of P . . . . .	19
[p	Initiation of process p . . . . .	69
p]	Termination of process p . . . . .	69
$R_p$	Range (output set) of process p . . . . .	69
r	Process initiation or termination event . . . . .	71
S	Set of states . . . . .	16
s	State of S . . . . .	19
SS	Software stack of (low,high) pairs of security classes . . .	137
<u>SS.l</u>	Low component of security class pair on top of SS . . . . .	137
<u>SS.h</u>	High component of security class pair on top of SS . . . . .	137
T	Set of transition operators . . . . .	16
t	Transition operator of T . . . . .	27
u,v	Physical storage locations of M . . . . .	19

Symbol	Meaning	Page
V	Set of objects assigned values in a program . . . . .	62
W	Working store or computational store . . . . .	50
$W_p$	Working store of process p . . . . .	51
$\alpha, \beta$	Items of information . . . . .	27
$\delta$	Maximum path length in partial order graph . . . . .	144
$\lambda$	Empty sequence . . . . .	71
$\nu$	Final value of a class sequence . . . . .	72
$\xi$	Expression . . . . .	61
$\Pi$	Process system . . . . .	69
$\rho$	Execution sequence of a process system . . . . .	71
$\sigma$	State sequence corresponding to an execution sequence . . . .	71
$\tau$	Transition operator sequence . . . . .	27
$\Phi$	Program or abstract program structure . . . . .	61
$\Psi$	Set of basic blocks . . . . .	97
$\omega$	String of terminal and non-terminal symbols . . . . .	86

## ABSTRACT

Denning, Dorothy Elizabeth Robling. Ph.D., Purdue University, May 1975.  
Secure Information Flow in Computer Systems. Major Professor: Herbert D.  
Schwetman.

This thesis investigates mechanisms that guarantee the secure flow of information between a set of security classes in a computing system. These mechanisms are examined within the framework of a mathematical model suitable for concisely formulating the requirements of secure information flow. The model provides a unifying view of all systems that restrict information flow and enables a classification of them according to security objectives.

The central component of the model is a lattice structure derived from the security classes and justified by the semantics of information flow. The lattice properties are used to develop two mechanisms that verify the secure execution of a program addressing storage locations associated with different security classes. The first operates in an environment in which the security class associated with each location is invariant. The second operates in a more complex environment in which the security class associated with each location is to be determined by the information contained in it.

## INTRODUCTION

## Chapter 1

This thesis is a theoretical investigation of mechanisms that guarantee secure information flow in a computer system. By information flow we mean that information in one object (e.g., a variable or file) is used to derive information in another. By secure information flow, or simple security, we mean that no unauthorized flow of information is possible. The mechanisms that guarantee the security of a system are referred to as protection mechanisms.

We restrict our study to the special case where the flow of information is to be constrained by a "flow relation" on a set of "security class". We develop a mathematical model of information flow suitable for formulating the security requirements of such a system. The properties of the model are used to construct automatic program certification mechanisms that verify the secure execution of a program.

In the remainder of this introduction we discuss informally the security problem we wish specifically to solve and the approach followed thereafter to its solution. (See [Po74,Hf69] for a survey of other security problems.)

### 1.1 Problem Statement

The problem under investigation may be informally stated as follows: given a set of "security classes" corresponding to classes of information, and a specification of allowable paths by which information can flow among them, construct a mechanism which guarantees that the flows which arise during program executions do not violate the specification. The security classes are intended to model the important concepts of "security classifications", "security categories", and "need to know" [We69,Ga72]; as will be seen, however, they are more general in their scope. Information is considered to flow from one class to another if it flows from an object associated with the former to one associated with the latter.

Consider for example a system, in which all information is partitioned into four security levels:  $A_1, A_2, A_3, A_4$ , where  $A_1 < A_2 < A_3 < A_4$ . The security requirement of this system is that information may flow from  $A_i$  to  $A_j$  only if  $i \leq j$ . Interpreting  $A_i < A_j$  to mean that " $A_i$  is at a lower security level than  $A_j$ ," this requirement states simply that high security information cannot flow into low security classes. In a government or military system, the classes  $A_1, \dots, A_4$  could correspond to unclassified, classified, secret, and top secret information respectively. In this case, the security requirement is to prevent events such as top secret data's from entering an unclassified file or reaching a user with only a secret clearance. This partially describes the security requirements of ADEPT-50, a time sharing system designed to handle sensitive information in government and military systems [We69].

To guarantee security in such a system, it is necessary to block all communication channels or "flow paths" between objects in class  $A_i$  and class  $A_j$  whenever  $A_i > A_j$ . For instance, consider two files  $a$  and  $b$  belonging to classes  $A_i$  and  $A_j$  respectively. The security requirement of the system implies that data may be transferred from  $a$  to  $b$  if and only if  $A_i \leq A_j$ . According to this requirement, transferring data from a top secret file  $a$  to an unclassified file  $b$  is not permitted.


A major implementation problem with this type of security requirement is detecting and monitoring all possible flow paths from  $a$  to  $b$ ; most direct paths are simply detected, but indirect paths — involving one or more other processes and files — are considerably more difficult to detect. The problem can be stated precisely as follows. For a file  $a$  and process  $p$ , let  $a \Rightarrow p$  denote the proposition " $p$  can read  $a$  into its local computational memory," and  $p \Rightarrow a$  denote the proposition " $p$  can write from its local computational memory into  $a$ ." Then let  $a \Rightarrow^* b$  mean there exists a process  $p$  such that  $a \Rightarrow p \Rightarrow b$ , and  $a \Rightarrow^* b$  mean there exist processes  $p_1, \dots, p_n$  and files  $c_0, \dots, c_n$  ( $n \geq 1$ ) such that

$$a = c_0 \Rightarrow p_1 \Rightarrow c_1 \Rightarrow p_2 \Rightarrow \dots \Rightarrow p_{n-1} \Rightarrow c_{n-1} \Rightarrow p_n \Rightarrow c_n = b.$$

When  $n = 1$  the flow  $a \Rightarrow^* b$  is "direct", and when  $n > 1$  it is "indirect". The foregoing security requirement is now stated simply: no flow  $a \Rightarrow^* b$  for  $a \in A_i$  and  $b \in A_j$  is permissible unless  $A_i \leq A_j$ .

Figure 1.1-1 illustrates two representations of the flow  $a \Rightarrow^* b$ : a graphical form and an "access matrix" form. The latter form was first introduced by Lampson [La71] and developed by Graham and





Graphical Representation

Files

	a	c <sub>1</sub>	c <sub>2</sub>	...	c <sub>n-2</sub>	c <sub>n-1</sub>	b
p <sub>1</sub>	r	w					
p <sub>2</sub>		r	w				
.							
.							
p <sub>n-1</sub>					r	w	
p <sub>n</sub>						r	w

Processes

Access Matrix Representation

Figure 1.1-1. Communication Channel from a to b.

Denning [GD72]. It has proved to be a useful conceptual basis within which to study problems of access control [Jo73]. However, since it does not include notions of security classes and information flow, one of our tasks will be extending it so that, for example, granting a read or write access to a process is not permitted if it introduces an insecure flow path.

The example above is easily extended to allow for processes that communicate directly between their computational memories (for example, using "send message" and "get message" operations [Ha70]) by regarding some of the  $c_i$  as message buffers rather than files.

The example is of course a rather simple instance of a more general problem. The complications which we shall study include, but are not limited to:

1. The security classes  $A_i$  do not conform to a strict linear ordering relation.
2. The processes involved in effecting a flow  $a \xrightarrow{*} b$  are themselves associated with security classes — e.g., according to the classes of their owning users — and these classes may vary during execution. Process security classes must be accounted for in deciding whether a flow  $a \xrightarrow{*} b$  is permissible.
3. The classes with which files are associated may be allowed to change during a computation. Similarly, the security level of a process's local computational memory may change according to the highest class of information residing within it.

4. Programs, or the systems on which they operate, must be certified as to whether they meet the stated security requirements before they are permitted to execute.

Even if adequate solutions for controlling information flow among specified security classes are instituted, there remain many subtle ways information can flow in violation of the stated security requirements. Some of these troublesome flows utilize the mechanisms put into place to control the direct and indirect flows cited above! Indeed, were it not for the so-called "covert channels" along which information can flow, and were it not for the difficulty in identifying such channels, protection might not be the enormous problem we find today [La73]. A famous example of a covert path is a process that alters its demands on some system resource in a pattern (e.g., Morse Code) that can be observed by an accomplice process; though the capacity of such a channel may be at best a few bits per second, this may be quite sufficient for a patient perpetrator. Another example is a process that discloses a confidential value  $x$  by opening  $x$  files, again under the observation of an accomplice. Other examples can be constructed around the principle of using error-checking mechanisms to divulge information by the very existence of an error; we shall see specific cases of this later. These problems have been perplexing researchers for many years (see, for example, Lampson [La73], Rotenberg [Ro74, and Fenton [Fe74a]). The mechanisms developed in this thesis provide solutions to many of them.

It is important to note that the protection requirements we have specified here are stronger than the read-write access control

requirements found in many contemporary systems, such as MULTICS [Or72,SS72] and HYDRA [Wu74], in that indirect flow paths are not controlled in these latter systems. For example, in MULTICS, if a process  $p$  is given read-access to a file  $a$ , it may be able to communicate, via a second file  $b$ , the entire contents of  $a$  to another process  $p'$  that was denied read-access to  $a$ . More precisely, existence of a path  $a \Rightarrow p$  may imply the existence of a path  $a \stackrel{*}{\Rightarrow} p'$  (namely  $a \Rightarrow p \Rightarrow b \Rightarrow p'$ ), even though  $p'$  was not permitted to establish a direct path  $a \Rightarrow p'$ . Similarly, in HYDRA, if a process is given a "capability" to read a file, it can pass this capability to another process (provided the capability has the "copy" option). This does not imply that the protection mechanisms of MULTICS or HYDRA are insecure; it implies only that they are inadequate for the problem under investigation in this thesis. It is also important to realize that the access control mechanisms of these systems handle additional security requirements that we are not explicitly considering here. For example, the mechanisms of both HYDRA and MULTICS permit the controlled execution of code segments; the mechanism for creating new types of objects in HYDRA permits greater control over the use of files. Here, for instance, the "owner" of a file can require that access to it be only through a specified set of procedures.

## 1.2 Previous Research

The problem under investigation has also been studied by Weissman [We69], Rotenberg [Ro74], Walter et. al. [Wa74], Bell and LaPadula [BL73a,BL73b], and Fenton [Fe73,Fe74a,Fe74b]. Weissman describes the ADEPT-50 system (mentioned earlier), in which control over information

flow is based on a set theoretic model of access rights. The security classes (or profiles, as they are called) are sets derived from two properties, Authority and Category, corresponding respectively to government and military Classification and Compartments. Each user, terminal, program, and file is assigned a profile. The basic set operations (e.g., union, intersection) are used to control access to objects and the flow of information. Rotenberg proposes the use of "privacy restrictions" for controlling information flow in a large time sharing system. Here, the security classes are sets of restrictions. A Privacy Restriction Processor records the flow of information in the system by propagating restrictions among the restriction sets associated with segments and processes. The restriction sets are then used to prevent unauthorized releases of information. Walter et. al. report on an investigation at Case Western Reserve to design a Security Kernel that monitors the flow of information between files and processes in a system like MULTICS. Their approach has been to develop a series of three models,  $M_0$ ,  $M_1$ ,  $M_2$ .  $M_0$  represents the most abstract model and contains agents and repositories which must obey four basic axioms of information flow. Structured modeling techniques are then used to arrive at consistent, but successively less abstract models  $M_1$  and  $M_2$ . They are currently investigating implementation of the final model  $M_2$ . Bell and LaPadula describe a project at MITRE to develop a set theoretic model of protection for government and military systems. Like ADEPT and the CASE system, their model is intended for systems that control information flow between files and processes. Fenton examines the problem of controlling information flow within a program; i.e.,

controlling the flow of information between the storage registers accessible to a program. He proposes a restricted hardware design in which each storage register is tagged according to its information class and the instruction execution mechanism inhibits instructions that would produce information flows not permitted by the tags.

We shall return to these works, examining them in more detail in the light of the general model of information flow to be developed in Chapter 2. The objective of this thesis is to extend and improve these research efforts. The model presented in the next chapter gives a formal basis for proving that security requirements are met. By providing a general enough context, it gives a single medium of discussion for comparing and contrasting the five systems above and others as well. The semantics of the security problem impose a well-developed structure on the permissible flows among security classes. This structure is exploited to construct efficient verification and checking algorithms.

The problem of certifying security properties has also been previously investigated by Bell and Burke [BB74], Neumann and Fabry [NF74], Popek and Kline [PK74] and Walter et. al. [Wa74]. However, their research efforts have been directed primarily toward the problem of manually certifying the correct operation and implementation of a complete system, whereas our efforts have been directed exclusively toward the problem of automatically certifying the secure execution of a program in an otherwise secure system.

### 1.3 Organization and Results of Thesis

In Chapter 2 we introduce a formal model of information flow and state a definition of security in terms of the model. An examination of the mathematical properties of the model shows that under certain assumptions the set of security classes forms a lattice. These assumptions are not arbitrarily chosen, but rather follow from the semantics of the problem. The lattice structure is exploited in later chapters to develop efficient security certification algorithms. Special cases of the model are examined in terms of their security requirements and relation to previous work. Finally, the problem of guaranteeing "determinacy" of security classes in a multiprogrammed environment is analyzed.

In Chapters 3 and 4 we investigate the special case of guaranteeing the secure flow of information within a program written in a well-structured high level language, such as PASCAL [Wr71]. In particular, we present algorithms for proving (or disproving) certain security properties for an arbitrary program. The significance of these algorithms is that the secure execution of a program can be established prior to its execution by a certification procedure. In Chapter 3 we present an algorithm for certifying the secure execution of a program in an environment in which the security class of each storage register remains constant throughout the lifetime of the program. Because it exploits the lattice structure of the security classes, the certification procedure can easily be embedded into the analysis phase of a compiler without adding substantially to the compilation time of a program. In Chapter 4 we propose a hardware design, based

on tagged architecture, for dynamically binding registers to security classes, so that the class of a register can be a function of its contents. We then present two possible algorithms for certifying the secure execution of a program in this environment.

In Chapter 5 we state the conclusions reached by this research and propose areas for future study. The important contributions of this research are summarized below:

1. A formal model of information flow and security is developed. It is used for stating and proving the security and determinacy requirements of systems that control information flow.
2. Analysis of the properties of the model, and of the security problem itself, leads to the conclusion that it is semantically reasonable to assume that the set of security classes forms a lattice. The lattice, which appears frequently in studies of programming languages and other semantic problems in computer science, is found to be important in the security problem as well. That lattices are inherent in this problem is the principal reason that verification of security requirements is not only feasible — it is understandable and efficient.
3. The model is a generalization of previous work on the security problem. It gives a unifying view of this work, enables a classification of it according to security objectives, and suggests some new approaches.
4. The model enables the specification of compile-time algorithms for certifying the secure execution of a program, one in an



environment in which the security class of each storage register is invariant, the other in the more complex environment in which the security class of some (or all) storage registers can change in accordance with the class of the information contained in them.

5. The model allows the study of run-time verification mechanisms, leading to the conclusion that such mechanisms are, in most cases, insufficient to enforce properly the security requirements of a system. Certain forms of certification must be performed prior to program execution (point 4, above) to guarantee the correct use of run-time mechanisms.

## A MODEL OF INFORMATION FLOW

## Chapter 2

In Section 1.1 we observed that the access control mechanisms of most systems do not control the flow of information in a system. Instead they monitor only a process's immediate read and write access rights (or privileges) to objects. In these systems, whether or not information will flow correctly is an undecidable question [HR74]. Consequently the security of such general systems relies greatly on trust.

The objective of this research is to eliminate this element of undecidability and reduce the need for trust by finding suitable constraints under which the security of a system can be decided and enforced. Our results have been more than encouraging: not only do such constraints exist, but they are rich in mathematical structure.

In Section 2.1 we introduce a model suitable for formulating the concept of information flow and specifying the requirements of secure information flow. The model consists of three components: an information flow structure, states, and transition operators. The information flow structure consists of a set of security classes  $A, B, \dots$  representing different classes of information; a binary relation  $\rightarrow$ , where  $A \rightarrow B$  means that information associated with class  $A$  is permitted to flow to class  $B$ ; and a binary operator  $\oplus$ , where  $A \oplus B$  represents the

class associated with information derived from classes A and B (e.g., the result of an arithmetic operation performed on information in classes A and B). The states consist of logical storage objects, an address mapping function from these objects to physical memory, processes, functions that bind objects and processes to security classes, and relations specifying processes' read and write access rights to objects. The transition operators consist of all operations that can cause a change of state.

In Section 2.2 the security requirements of the model are stated: no set of processes must be able to effect a flow of information from class A into class B unless  $A \rightarrow B$ . This flow could occur, for example, as the result of one or more processes transferring (via a sequence of read and write operations) information from a storage object assigned to class A to one assigned to class B. The analysis reveals that security can be guaranteed if each operation is independently secure when the flow relation  $\rightarrow$  is transitive.

An investigation of the mathematical properties of the model in Section 2.3 shows that under certain assumptions, justified by the semantics of information flow, the security classes form a lattice. Specifically, we show the semantics of information flow imply that we are justified in assuming that a) the security classes are partially ordered by the flow relation  $\rightarrow$ , b)  $\oplus$  is a least upper bound operator, and c) there exists a greatest lower bound operator, which we denote by  $\otimes$ . This result permits concise formulations of the security requirements of different systems and facilitates the construction of mechanisms that enforce these requirements.

In Section 2.4 special cases of the model are classified along two dimensions: static v. dynamic binding and single v. multi-class. "Static binding" means that the security class associated with an object remains constant over the lifetime of the system; at the opposite extreme, "dynamic binding" means that the security class of an object is a function of its contents and therefore varies with its contents. The difference between static and dynamic binding is illustrated with a simple example. Suppose a process  $p$  wishes to move information from its working store (i.e., computational memory)  $W_p$  to a file  $b$ . Under static binding, it is necessary to verify that the flow from  $W_p$  to  $b$  is valid exactly once, at the time  $p$  opens  $b$  for writing. Under dynamic binding, it is necessary to update the class of  $b$  by that of  $W_p$  each time  $p$  writes into  $b$ .

A "single-class" working store of a process is treated as a single object bound to a single security class; in contrast, a "multi-class" working store of a process is treated as more than one object and allows each object to be bound to a different class. Security in a single-class system depends only on file input and output operations and therefore can be verified at low cost. In contrast, security in a multi-class system may depend on every instruction executed by a process and therefore is inherently more costly to verify. All of the known mechanisms for controlling information flow are characterized by these two dimensions.

A determinacy problem arises in dynamically bound systems if the security class of an object depends on the order in which a set of processes execute. Since this may be undesirable in some systems, in

Section 2.5 we state additional constraints that guarantee determinate operation with respect to security classes.

## 2.1 Description of the Model

An information flow model FM is defined by

$$FM = \langle FS, S, T \rangle$$

where

FS is an information flow structure.

S is a set of states.

T is a set of transition operators.

The information flow structure FS models the valid flow paths of a system and is defined by

$$FS = \langle SC, \rightarrow, \oplus \rangle$$

where

SC = {A, B, ...} is a set of security classes closed under  $\oplus$ .

$\rightarrow \subseteq SC \times SC$  is a binary flow relation.

$\oplus: SC \times SC \rightarrow SC$  is a binary class-combining operator.

The security classes SC correspond to disjoint classes of information.

The flow relation  $\rightarrow$  on  $SC \times SC$  is defined by  $A \rightarrow B$  if and only if information in class A is permitted to "flow" into class B. (A full discussion of the implications of "flow" comes later). If  $A \rightarrow B$ , we say that A is a lower security class than B, or equivalently that B is a higher security class than A. If there exists a class A such that  $A \rightarrow B$  for all  $B \in SC$ , then A is said to be a lowest security class. If there exists a class B such that  $A \rightarrow B$  for all  $A \in SC$ , then B is said to be a highest security class. In Section 2.3 we shall show that the semantics of the problem imply we are justified to assume the existence of

unique lowest and highest security classes.

The class-combining operator  $\oplus$  specifies the class that corresponds to the combination of information in two classes (e.g., by addition, multiplication, or insertion into a common file); that is,  $A \oplus B$  is the class corresponding to the result of an operation  $f$  applied to information in classes  $A$  and  $B$ . In Section 2.3 we shall show that the semantics of the problem imply we are justified in assuming that the  $\oplus$  operator is a least upper bound operator on the set of classes  $SC$ .

Figure 2.1-1 shows two examples of information flow structures on four security classes, where the flow relation  $\rightarrow$  is a partial ordering of the classes; i.e.,  $\rightarrow$  is reflexive, transitive, and anti-symmetric. Example (a) shows the special case of a linear ordering, with a lowest security class  $A_1$  and a highest security class  $A_4$ ; it corresponds to the example described in Chapter 1. Example (b) shows a non-linear (partial) ordering, with a lowest security class  $A_{00}$  and a highest security class  $A_{11}$ . The graphical representations are standard precedence graphs for a partial order, showing only the non-reflexive, immediate (irredundant) relations; the remaining relations are implied by transitivity (i.e., a full partial order graph is the transitive closure of its precedence graph). In Section 2.3 we shall show that the semantics of the problem imply we are justified in assuming that the flow relation always partially orders the security classes.

Although we have defined an information flow model  $FM$  in terms of a single flow structure  $FS$ , an actual system may require more than one flow structure. For example, several disjoint flow structures may be needed to specify the valid flows between different types of objects.

$$SC = \{A_1, A_2, A_3, A_4\}$$

$$A_i \rightarrow A_j \quad \text{iff} \quad i \leq j$$

$$A_i \oplus A_j = A_{\max(i,j)}$$



Description


Representation

## a) A Linear Ordering

$$SC = \{A_{00}, A_{11}, A_{10}, A_{01}\}$$

$$A_{ij} \rightarrow A_{i',j'} \quad \text{iff} \quad i \leq i', j \leq j'$$

$$A_{ij} \oplus A_{i',j'} = A_{\max(i,i'), \max(j,j')}$$



Description

Representation

## b) A Non-Linear (Partial) Ordering

Figure 2.1-1. Examples of Partially Ordered Information Flow Structures.

In this case it is a simple matter to extend the model to include a set of flow structures  $\{FS_1, \dots, FS_n\}$ .

The states  $S$  model the state-dependent components of a protection system. These are the set of active processes, the set of logical memory objects (i.e., the receptacles of information), the binding of objects to security classes, the binding (if any) of processes to security classes, the current configuration of processes' read and write access rights (i.e., configuration of the access matrix), the binding of physical memory locations to objects and classes, and in general all information pertaining to the "security state" of the system. For each state  $s \in S$ , the following vector is sufficient for our purposes here:

$s = \langle N, M, P, \underline{loc}, \underline{class}, \underline{tag}, \underline{clearance}, \underline{read}, \underline{write} \rangle$

where

$N = \{a, b, \dots\}$  is a set of logical memory objects.

$M = \{u, v, \dots\}$  is a set of physical memory locations.

$P = \{p, q, \dots\}$  is a set of processes.

loc:  $N \rightarrow \text{powerset}(M)$  is a mapping function from logical memory objects to subsets of physical memory.

class:  $N \rightarrow SC$  is a function that binds each object in logical memory to a security class.

tag:  $M \rightarrow SC$  is a function that binds each location in physical memory to a security class.

clearance:  $P \rightarrow SC$  is a function that binds each process to a security class.

read  $\subseteq P \times N$  is a relation specifying processes' read access rights (or privileges) to objects.



write  $\subseteq P \times N$  is a relation specifying processes' write access rights (or privileges) to objects.

The logical memory  $N$  is the set of logical storage objects (i.e., the information receptacles) in the system. We refer to the elements of  $N$  simply as objects. Objects may be files, segments, or even program variables, depending on the level of detail under consideration. The physical memory  $M$  is the set of physical storage locations in the system. Some elements of  $M$  may be located in the main memory, others on secondary storage devices such as disks or drums. Each location  $u$  in  $M$  represents the smallest unit of memory that can be assigned to a logical object (see the description of loc below). For example, in a paged virtual memory system,  $u$  may be a single page in main memory or a sector of a track on drum. We assume that all physical memory locations are addressed only through logical memory objects. The processes  $P$  are the active agents responsible for all information flow in the system.

The function loc:  $N \rightarrow \text{powerset}(M)$  assigns one or more (usually contiguous) physical storage locations to each object; that is, for  $a \in N$ , loc( $a$ ) =  $U$ , where  $U \subseteq M$ . For example, if  $a$  is a segment in a paged virtual memory system,  $U$  would be the set of pages in main memory and the set of track sectors in secondary memory occupied by  $a$ .

The functions class:  $N \rightarrow SC$ , tag:  $M \rightarrow SC$ , and clearance:  $P \rightarrow SC$  respectively bind each object, physical location, and process to a security class. The tag function is so named to suggest the use of hardware tags for binding physical memory locations to classes. Although tagged architectures [Fu73] present a possible direct

implementation, tags can also be implicit, as fields in page or segment tables or in file descriptor tables. Moreover, the latter form (file tags) are easily added to most existing file systems. However, it is necessary to make explicit the requirement that the implementation of security classes must conform to the logical structure; i.e., for object  $a$ , we require that  $\text{tag}(u) = \text{class}(a)$  for all  $u \in \text{loc}(a)$ . If location  $u$  has contents  $\alpha$  and tag  $A$  (i.e.,  $\text{tag}(u) = A$ ), this will be illustrated by

A	$\alpha$
---	----------

$u$

We assume that all objects are logically and, therefore, physically disjoint; that is,  $\text{loc}(a) \cap \text{loc}(b) = \emptyset$  for all  $a, b \in N$ . Without this assumption, for any  $a, b$  not independent (i.e.,  $a \cap b \neq \emptyset$  and therefore  $\text{loc}(a) \cap \text{loc}(b) \neq \emptyset$ ), several complications arise when information flows in or out of  $\text{loc}(a) \cap \text{loc}(b)$ . For example, if information flows into object  $a$ , but is stored in  $\text{loc}(a) \cap \text{loc}(b)$ , it must be guaranteed that it is permitted to flow into  $b$  as well as  $a$ . Or, if the class of  $a$  is changed, causing a corresponding change to  $\text{tag}(u)$  for all  $u \in \text{loc}(a)$ , then the class of  $b$  should also change.

The relations read and write on  $P \times N$  specify processes' read and write access rights (or privileges) to objects. These relations are equivalent to the well known  $|P| \times |N|$  "access control matrix" AM defined by

$r \in \text{AM}[p, a]$  iff  $p$  read  $a$ ;

$w \in \text{AM}[p, a]$  iff  $p$  write  $a$ .

(See Figure 1.1-1 for an example of an access matrix and [La71, GD72] for a thorough treatment of the use of an access matrix as a model of

protection.) Of course, if a process has read access to an object a and write access to an object b (i.e.,  $p \text{ read } a$  and  $p \text{ write } b$ ), this does not necessarily mean that it can exercise these privileges and transfer information from a to b; the information transfer must also be permitted by the flow relation (i.e.,  $\text{class}(a) \rightarrow \text{class}(b)$ ).

The read and write relations need not be implemented as a matrix. Some methods that are commonly used include "capabilities", "access control lists", and "locks and keys". The capability method corresponds to storing the matrix by rows. Associated with each process p is a set of pairs

$$CL = \{(a, AM[p,a]) \mid AM[p,a] \neq \emptyset\}.$$

The elements of the set are referred to as "capabilities" and the set CL as a "capability list". Capabilities were first introduced by Dennis and VanHorn [DV66] and their utility is becoming increasingly apparent [An74,EL72,En72,Fa68,Fa71,Fa74,Il68,Jo73,Li73,Ne72,St74,Wu74]. The access control method corresponds to storing the matrix by columns. Associated with each object a is a set of pairs

$$ACL = \{(p, AM[p,a]) \mid AM[p,a] \neq \emptyset\}$$

referred to as an "access control list" [SS72]. The lock and key method is a combination of the capability and access control list methods and is commonly used to control access to physical memory [IB68]. Popek examines methods for finding optimum representations of access control relationships expressed with locks and keys [Po73]. For a comparison of the relative costs and reliability aspects of these methods see [De74,GD72,Po74,Ts73,Wi72].

When we need to identify the components of a particular state  $s_i$  to distinguish them from those of another state  $s_j$ , we shall subscript the components with  $i$ :  $N_i$ , class<sub>i</sub>, tag<sub>i</sub>, etc..

The transition operators  $T$  model all important operations that affect or effect the flow of information in the system. For example, granting a process  $p$  write access to an object  $a$  affects the flow of information if  $p$  may subsequently move information into  $a$ ; executing an assignment operation " $b := a$ " effects the flow of information by causing data to be transferred from  $a$  to  $b$  (actually from loc( $a$ ) to loc( $b$ )). The operations under consideration here are given in Figure 2.1-2. For each operation, a brief description is given together with the conditions necessary for execution of the operation (i.e., pre-conditions on the state variables) and the conditions that result from its execution (i.e., post-conditions on the state variables). We assume that no changes are made to state variables except as explicitly stated.

The operation transfer( $f, a_1, \dots, a_n, b$ ) represents any action that causes  $f(a_1, \dots, a_n)$  to flow to  $b$ . For example, execution of the statements " $b := a$ ", " $b := a_1 + a_2$ ", and "output a to  $b$ " will respectively be represented by transfer( $:=, a, b$ ), transfer( $+, a_1, a_2, b$ ), and transfer(output,  $a, b$ ). Note that the operation transfer( $f, a_1, \dots, a_n, b$ ) does not necessarily replace the entire contents of loc( $b$ ) with  $f(a_1, \dots, a_n)$ ; it may instead extend or modify it by  $f(a_1, \dots, a_n)$  (as, for example, with an output operation to a file  $b$ ). We shall assume that all transfer operations to an object  $b$  replace the contents of loc( $b$ ) when  $b$  is a single element (e.g., a scalar variable) and modify it when  $b$  is an aggregate of elements (e.g., a file or segment).

Operation	Description	Pre-Conditions	Post-Conditions
$\text{transfer}(f, a_1, \dots, a_n, b)$	An n-ary operation $f$ is performed on $a_1, \dots, a_n$ and the result stored in $b$ .	$a_1, \dots, a_n, b \in N$ ; $\exists p \in P: p \text{ read } a_i \text{ and } p \text{ write } b$	$\text{class}(b)$ and $\text{tag}(u)$ $\text{vucloc}(b)$ may change
$\text{create object}(a, A, U)$	Object $a$ is created with class $A$ and assigned to storage locations $U$ .	$a \in N$ ; $\forall b \in N: \text{loc}(b) \cap U = \emptyset$	$a \in N$ ; $\text{loc}(a) = U$ ; $\text{class}(a) = A$ ; $\text{tag}(u) = A \forall u \in U$
$\text{delete object}(a)$	Object $a$ is deleted.	$a \in N$	$a \notin N$
$\text{move}(a, U)$	Object $a$ is moved to locations $U$ .	$a \in N$ ; $\text{loc}(a) = V$ ; $\forall b \in N: \text{loc}(b) \cap U = \emptyset$	$\text{loc}(a) = U$ ; $\text{tag}(u) = \text{class}(a) \forall u \in U$
$\text{reclassify}(a, A)$	Class of $a$ is changed to $A$ .	$a \in N$	$\text{class}(a) = A$ ; $\text{tag}(u) = A \forall u \in \text{loc}(a)$
$\text{clear}(u)$	Location $u$ is cleared.	$\#b \in N: u \in \text{loc}(b)$	$\text{tag}(u)$ may change
$\text{create process}(p, A)$	Process $p$ is created with a clearance of $A$ .	$p \notin P$	$p \in P$ ; $\text{clearance}(p) = A$
$\text{delete process}(p)$	Process $p$ is deleted.	$p \in P$	$p \notin P$
$\text{grant}(p, a, \{v_w^f\})$	Process $p$ is granted read and/or write access to $a$ .	$p \in P$ ; $a \in N$	$p \text{ read } a \text{ and/or } p \text{ write } a$
$\text{ungrant}(p, a, \{v_w^f\})$	Process $p$ loses read and/or write access to $a$ .	$p \in P$ ; $a \in N$	$\neg p \text{ read } a \text{ and/or } p \text{ write } a$

Figure 2.1-2. Transition Operators.

The operation create object(a,A,U) creates a new logical object a in class A, and assigns it to (unoccupied) storage locations U. The tag of each element  $u \in U$  is also set to A, however, the contents of u is not initialized. Note that the operation create object(a,A,U) does not automatically give the process that executed it read or write access to the object; this must be acquired by subsequent execution of a grant operation (see below). This is done so that all changes to the "access matrix" are modeled by just two operations (grant and ungrant). The operation delete object(a) removes object a, releasing its physical memory locations.

The operation move(a,U) assigns new (unoccupied) storage locations U to object a, while transferring the contents of the old locations V to U. The tag of each element  $u \in U$  is also set to class(a). The move operation would be performed, for example, by a "storage compacter" or by a relocation mechanism in a virtual memory system. It is easy to see that the pre-condition " $\forall b \in N: \text{loc}(b) \cap U = \emptyset$ " on the operations create(a,A,U) and move(a,U) guarantees that no two objects will ever share physical memory, since these are the only two operations that assign storage to objects.

The operation reclassify(a,A) changes the class of object a and the tag of its corresponding storage locations, without changing its contents. This operation might be useful in a government or military system, where it is necessary to change the classification of data.

The operation clear(u) clears (unoccupied) location u (e.g., by filling with zeros or an undefined value). The tag of u is then set to the security class associated with unoccupied storage locations.

The operation create process(p,A) creates a new process p with clearance A. The class A could, for example, correspond to the security clearance of the user on whose behalf p will execute. Sometimes there will exist processes for which no clearance is specified or even necessary; in such cases the class A associated with p would be undefined. Note that the create process operation does not allocate any storage to p; this would have to be done by execution of create object and grant operations. The delete process(p) operation simply removes process p from the system.

The grant and ungrant operations add and remove elements of the read and write relations (i.e., change entries in the access matrix). For example, the operation grant(p,a,r) could be performed when a process p requests that a file a be opened for input (i.e., reading). In this case, granting p this access privilege may depend on the clearance of p, what other processes have access to a, or other state variables not explicitly accounted for in our model. In some systems the controlled granting and ungranting of access privileges may be coordinated with the security requirements -- i.e., so that all accesses granted are consistent with permissible flows. In other systems the granting and ungranting of access privileges may not be coordinated with the security requirements; in this case, a process may be denied a request to write into a file to which it has an access privilege if this would create a disallowed flow path.


## 2.2 Security Requirements

The security requirements of the model are simply stated: all information flow must satisfy the constraints of the flow structure; that is, information in class A must not flow into class B unless  $A \rightarrow B$ . A precise formulation of the meaning of "flow" can be given in terms of sequences of state transition operators. For state  $s$  and sequence  $\tau = t_1 \dots t_m$  of transition operators, we say that  $\tau$  is feasible from state  $s$  if it can be executed beginning in state  $s$  (i.e., the pre-conditions for each  $t_i$  ( $1 \leq i \leq m$ ) are satisfied); in this case we write  $s \cdot \tau$  to denote the successor state of  $s$  under  $\tau$ . Two observations are important: 1) in any state, all information in the system is stored in some physical location, and 2) new information is derived only by applying a function  $f$  to existing information (in our terms, by some process executing an operation transfer( $f, a_1, \dots, a_n, b$ )).


Now, suppose the system is in state  $s_i$  and some process executes an operation  $t$  that causes a transition to state  $s_j$  (i.e.,  $s_i \cdot t = s_j$ ). Let  $\alpha$  denote an item of information and let  $\alpha \in A$  mean that  $\alpha$  is stored in a location  $u$  such that  $\text{tag}(u) = A$ . Then information can flow from class A to class B as a result of a single operation  $t$  if and only if there exists an information item  $\alpha \in A$  in  $s_i$  that derives an information item  $\beta \in B$  in  $s_j$  ( $\alpha = \beta$  is possible). This can occur as the result of one of three possibilities (see Figure 2.2-1):

- a) In  $s_i$ ,  $\alpha$  is stored in location  $u$  for which  $\text{tag}_i(u) = A$ , and in  $s_j$ ,  $\alpha$  is stored in location  $v$  for which  $\text{tag}_j(v) = B$  ( $v = u$  is possible).
- b) There exist locations  $u_1, \dots, u_n$  ( $n \geq 1$ ) such that: the






a) Direct transfer ( $u \neq v$ ) or tag change ( $u = v$ ).



b) Operand transfer ( $v = u_j$  for some  $u_j$  is possible).



c) Function value transfer ( $v = u_j$  for some  $u_j$  is possible).

Figure 2.2-1. Information Flow from A to B by an Operation  $t$ .

contents of  $u_k$  is  $\alpha_k$  and  $\text{tag}_1(u_k) = A_k$  ( $1 \leq k \leq n$ ), and the result  $\beta = f(\alpha_1, \dots, \alpha_n)$  of some function  $f$  applied to the contents of locations  $u_1, \dots, u_n$  is stored in a location  $v$  for which  $\text{tag}_j(v) = B$ . Then  $\alpha = \alpha_k$  and  $A = A_k$  for some  $k$ .

- c) Same conditions as case (b), except that  $\alpha = \beta = f(\alpha_1, \dots, \alpha_n)$  and  $A = A_1 \oplus \dots \oplus A_n$ . (This extended use of the  $\oplus$  operator is justified in the next section.)

Note that case (a) corresponds to a transfer between locations  $u \neq v$  or a tag change in a given location  $u = v$ ; case (b) corresponds to the given value's being an input operand to a function which influences the result place in location  $v$ ; case (c) corresponds to the given value's being generated functionally, in a class to which none of the operands may belong, and then being placed in location  $v$ .

We will use the notation  $\alpha \xRightarrow{s \cdot t} \beta$  to denote the derivation of information  $\beta$  in  $s \cdot t$  from information  $\alpha$  in  $s$  as the result of executing operation  $t$  in  $s$ , and  $A \xRightarrow{s \cdot t} B$  to denote the flow of information from class  $A$  to class  $B$  as the result of executing  $t$  in  $s$ . Clearly  $A \xRightarrow{s \cdot t} B$  if and only if there exists an  $\alpha \in A$  in  $s$  and a  $\beta \in B$  in  $s \cdot t$  such that  $\alpha \xRightarrow{s \cdot t} \beta$ . Since a flow  $A \xRightarrow{s \cdot t} A$  occurs trivially whenever an operation is performed that does not affect a location  $u$  (i.e.,  $\text{tag}(u) = A$  before and after), we consider only flows  $A \xRightarrow{s \cdot t} B$  for  $A \neq B$ .

Consider two states  $s_i$  and  $s_j$  such that  $s_i \cdot t = s_j$  for some  $t$ . It is easy to see that for delete object, clear, create process, delete process, grant, or ungrant there is no flow of information. For move(a,U), information cannot change classes since the contents of  $a$ 's storage locations are moved to  $U$ , replacing the former contents of

u and changing its tag. Hence, any information item  $\alpha \in \text{loc}_i(a)$  is moved to  $U = \text{loc}_j(a)$  without changing class (the previous information  $\beta \in U$  is lost). This leaves just the three operations: transfer, create object, and reclassify to consider as flow-causing operations.

For  $t = \text{transfer}(f, a_1, \dots, a_n, b)$ ,  $A \xrightarrow{s_i \cdot t} B$  whenever  $\text{class}_i(a_k) = A$  for some  $a_k$  ( $1 \leq k \leq n$ ) and/or  $\text{class}_i(a_1) \oplus \dots \oplus \text{class}_i(a_n) = A$ , and  $\text{class}_j(b) = B$ . This is illustrated in Figure 2.2-2 (a). For

$t = \text{create object}(b, B, U)$ ,  $A \xrightarrow{s_i \cdot t} B$  whenever  $\text{tag}_i(u) = A$  for some  $u \in U$ , indicating that it was previously occupied by an object  $a$  in class  $A$  (which has now been deleted or moved elsewhere), but that its contents have not been cleared. Hence, the information in  $u$  flows from  $A$  to  $B$  since the create object operation sets  $\text{tag}_j(u) = B$ . This type of flow occurs frequently in systems that do not clear storage before assigning it to new objects. It is illustrated in Figure 2.2-2 (b). For

$t = \text{reclassify}(a, B)$ ,  $A \xrightarrow{s_i \cdot t} B$  whenever  $\text{class}_i(a) = A$ , since reclassify sets  $\text{class}_j(a) = B$  and  $\text{tag}_j(u) = B$  for all  $u \in \text{loc}_i(a)$ , thus also changing the class of information in  $\text{loc}_i(a)$ . Figure 2.2-2 (c) illustrates this type of flow.

Information can also flow from a class  $A$  to a class  $B$  as the result of execution of a sequence of operations  $\tau = t_1 \cdot t_2 \cdot \dots \cdot t_m$ . For this case, let  $\alpha \xrightarrow{s \cdot \tau} \beta$  denote the derivation of information item  $\beta$  in  $s \cdot \tau$  from information item  $\alpha$  in  $s$  as the result of execution of  $\tau$  beginning in  $s$ . Clearly  $\alpha \xrightarrow{s \cdot \tau} \beta$  implies there exists information items  $\gamma_0, \dots, \gamma_m$  (not necessarily distinct), such that

$$\alpha = \gamma_0 \xrightarrow{s \cdot t_1} \gamma_1 \xrightarrow{s_1 \cdot t_2} \gamma_2 \Rightarrow \dots \Rightarrow \gamma_{m-1} \xrightarrow{s_{m-1} \cdot t_m} \gamma_m = \beta,$$

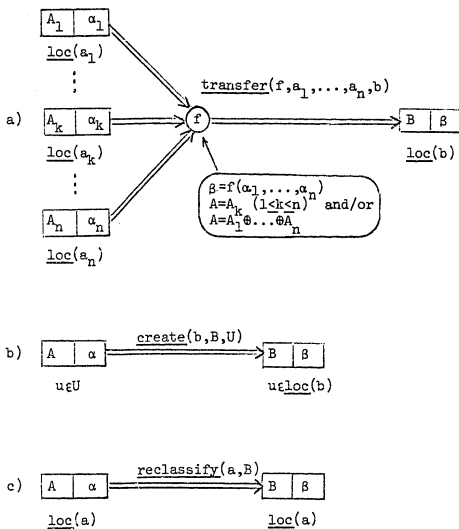


Figure 2.2-2. Operations that Cause Information Flow from A to B.

where  $s_1, \dots, s_m$  is the sequence of successor states of  $s$ . Also, let  $A \xrightarrow{s \cdot \tau} B$  denote the flow of information from  $A$  to  $B$  as the result of execution of the sequence  $\tau$  beginning in state  $s$ . Then it is also clear that  $A \xrightarrow{s_i \cdot \tau} B$  if and only if there exists an item  $\alpha$  in  $s_i$  stored in location  $u$  such that  $\text{tag}_i(u) = A$  and an item  $\beta$  in  $s_j = s_i \cdot \tau$  stored in location  $v$  such that  $\text{tag}_j(v) = B$ , and  $\alpha \xrightarrow{s_i \cdot \tau} \beta$ . Examples of operation sequences  $\tau$  that cause information flow from  $A$  to  $B$  are shown in Figure 2.2-3.

The security requirements of the model are now precisely stated: An operation sequence  $\tau$  is a secure sequence if and only if for every initial state  $s$  such that  $\tau$  is feasible (if any),  $A \xrightarrow{s \cdot \tau} B$  implies  $A \rightarrow B$ . A flow model FM is secure if and only if every possible operation sequence is secure.

To verify the security of a particular system, it is necessary first to show that the system fits the model and then to show that each possible operation sequence  $\tau = t_1 \cdot t_2 \cdot \dots \cdot t_m$  is secure. For  $m = 1$ , this involves showing that the operations transfer, create object, and reclassify cannot alone cause a flow from  $A$  to  $B$  when  $A \neq B$ . Furthermore, if the system is designed to always clear memory before allocating it to a new object, then it is sufficient to establish the secure execution of the operations transfer and reclassify. To see how this can be done, let  $s_i$  be an initial state and let  $s_j = s_i \cdot t_1$ . For  $t_1 = \text{reclassify}(a, B)$ , security can be established by verifying that  $\text{class}_i(a) \rightarrow B$ . Note that this implies that any operation that "declassifies" data (e.g., by changing it from "secret" to "unclassified") is considered to be insecure. This does not mean that such operations must




Figure 2.2-3. Examples of Operation Sequences that Cause Information Flow from A to B.

never be allowed; but it does mean that if such operations are necessary, they must be regarded as a potential threat to the security of the system and should be carefully controlled.

For  $m > 1$ , verifying the security of a sequence  $t_1 \cdot t_2 \cdot \dots \cdot t_m$  is much more difficult, since a prefix of it may be secure, while the entire sequence may not be. For example, suppose the flow relation satisfies the properties  $A \rightarrow B$  and  $B \rightarrow C$  but not  $A \rightarrow C$ , and let the system be in state  $s_i$  with  $\text{class}_i(a) = A$ . It then executes the sequence  $\text{transfer}(\text{:=}, a, b) \cdot \text{delete}(a)$ , which leaves it in state  $s_j$  with  $\text{class}_j(b) = B$ . Clearly execution of this sequence is secure. However, if the operation  $\text{transfer}(\text{:=}, b, c)$  is now executed, leaving the system in state  $s_k$  with  $\text{class}_k(c) = C$ , then the complete sequence  $\text{transfer}(\text{:=}, a, b) \cdot \text{delete}(a) \cdot \text{transfer}(\text{:=}, b, c)$  is not secure, since information in  $A$  has been transferred to  $C$  although  $A \not\rightarrow C$ . We observe that this problem would not have arisen if the flow relation  $\rightarrow$  were transitive (i.e.,  $A \rightarrow B$  and  $B \rightarrow C \Rightarrow A \rightarrow C$ ). Indeed, if we consider only those systems in which the flow relation is transitive, we can ignore the problem of verifying operation sequences of length  $m > 1$ , as the following theorem establishes.

Theorem 2.2-1. Let  $\rightarrow$  be a transitive flow relation and  $\tau'$  and  $\tau''$  operation sequences. If  $\tau'$  and  $\tau''$  are secure sequences, then  $\tau = \tau' \cdot \tau''$  is also a secure sequence.

proof. Consider an arbitrary initial state  $s$  such that  $\tau$  is feasible. If there is none, then  $\tau$  is trivially secure, so assume such an  $s$  exists. Then  $\tau'$  secure implies that for any  $\alpha \in A$  and  $\beta \in B$ , if  $\alpha \xrightarrow{s \cdot \tau'} \beta$ , then  $A \rightarrow B$ . Also,  $\tau''$  secure implies that for any  $\beta \in B$  and  $\gamma \in C$ , if  $\beta \xrightarrow{s' \cdot \tau''} \gamma$  then  $B \rightarrow C$ , where  $s' = s \cdot \tau'$ . Therefore, for

$\alpha \in A$  and  $\gamma \in C$ , if  $\alpha \xrightarrow{S \cdot \tau} \gamma$ , then  $A \rightarrow C$  by transitivity. Hence,  $\tau = \tau' \cdot \tau$  is secure.

Corollary 2.2-1. Let  $\rightarrow$  be a transitive flow relation. If every operation sequence of length one is secure, then every sequence of operations is secure.

proof. Immediate from preceding theorem.

Given this result, we will first show that the semantics of the problem imply we are justified in assuming that  $\rightarrow$  is transitive. This is done in Section 2.3.1. In the remainder of the thesis, we will then examine protection mechanisms that keep the system in a state such that execution of the next operation is guaranteed secure. Any state that satisfies this property is said to be a safe state. By Corollary 2.2-1, we know that if we can keep the system in a safe state, it will be secure. This result was also established for the MITRE model [BL73a].

Mechanisms for maintaining the system always in safe states are of two types: run-time mechanisms that verify the secure execution of each operation before it is performed (preventing those that are not secure) and certification mechanisms that guarantee that no insecure operation will be attempted. That run-time mechanisms can be used to "leak" information (i.e., violate flow paths - see below) is a compelling argument for certification mechanisms. However, a practical system may need both types of mechanisms to provide the fullest range of defenses.

As an example of using a run-time mechanism to leak information (despite the intention of its designers to prevent leaks!), consider



two classes A and B such that  $B \neq A$ , and suppose that process p has access to some information item  $\beta$  in class B, but that it is not allowed to communicate this information to its owner Smith who is in class A (since  $B \neq A$ ). However, p may attempt to transfer  $\beta$  from B to A,  $\beta$  times; and though it will fail, the record of the aborted attempts contains the value  $\beta$ . If Smith is permitted to learn of these violations, the information  $\beta$  has effectively been transferred from B to A. This example shows that under certain circumstances, "audit trails" (i.e., records of attempted violations), can be used for improper transfers of information. Rotenberg presents many examples demonstrating that such mechanisms designed to prevent leaks can be used to transmit information in an encoded form in the recorded or observable history of the security mechanism [Ro74]. This form of leakage is one example of a "covert channel".

The problem can be dealt with in systems that use run-time mechanisms. One way simply is to not "sound the alarm" when a security violation is attempted. Although this prevents leakage, it has the disadvantage of not reporting security violations either to the violator (whose intentions may have been innocent, in which case he should be informed of his error) or to the security administrator. The more comforting approach is to sound the alarm and keep an "audit trail" so that all such violations can be scrutinized and the record used as evidence in indictments of wrongdoing. However, for the reasons noted above, access to the audit records must be restricted.

This problem of covert leakage is in principle eliminated with mechanisms that certify that no insecure operations will be attempted.

Certification mechanisms are presented in Chapters 3 and 4. It is important to realize, however, that the need for run-time mechanisms is not obviated by certification mechanisms — they are still necessary, for example, to detect errors wrought by hardware malfunctions that alter the contents of memory locations or errors which may arise from incompletely debugged compilers.

## 2.3 The Lattice of Information Flow

### 2.3.1 Derivation

We now examine the mathematical properties of the information flow structure  $FS = \langle SC, \rightarrow, \oplus \rangle$  and show that under certain assumptions,  $FS$  forms a universally bounded lattice. These assumptions are not arbitrarily chosen, but rather follow from the semantics of information flow. By this we mean either that they are required for consistency or that no generality is lost by them. Consistency means that if one can write a program that effects a permissible flow, then other flows implied by that flow should also be permitted by the flow relation. Stated another way, a program should not be secure by one interpretation and insecure by another. It is important that the flow structure be consistent since it would otherwise be possible to masquerade insecure operations as secure ones.

A lattice is any structure consisting of a partially ordered set in which each pair of elements has a well-defined "least upper bound" and "greatest lower bound". A universally bounded lattice is a lattice with unique least upper and greatest lower bounds on the entire set. A general discussion of lattices and their properties can be found in

Stone [St73], Birkhoff [Bi67], or Abbott [Ab69]. The reader should note that this is not the first time a lattice has appeared in a study involving the semantics of programs or programming. Scott has shown that lattices, especially continuous functions defined on lattices, apply to data types and control flow diagrams [Sc71]. These structures are partially ordered by a relation based on the idea of "approximation".

To show that  $FS = \langle SC, \rightarrow, \oplus \rangle$  forms a universally bounded lattice, we shall show that for semantic reasons the following four assumptions are justified:

- 1)  $\langle SC, \rightarrow \rangle$  is a partially ordered set.
- 2)  $SC$  is finite.
- 3)  $SC$  has a lower bound  $L$  such that  $L \rightarrow A$  for all  $A \in SC$ .
- 4)  $\oplus$  is a least upper bound operator on  $SC$ ; i.e.,  $A \oplus B$  is the unique least upper bound of  $A$  and  $B$  for all  $A, B \in SC$ .

From these assumptions we shall then establish that there exists a greatest lower bound operator on  $SC$ , which we denote by  $\otimes$ ; i.e.,  $A \otimes B$  is the greatest lower bound of  $A$  and  $B$  for all  $A, B \in SC$ . Therefore, the structure  $\langle SC, \rightarrow, \oplus, \otimes \rangle$  is a lattice with greatest lower bound  $L$  and least upper bound, which we denote by  $H$ .

To simplify the discussion, we introduce a shorthand underbar notation to denote the class function, where  $\underline{a}$  denotes class( $a$ ) for object  $a$ . Similarly, for an information item  $\alpha$  we let  $\underline{\alpha}$  denote the class associated with  $\alpha$ . For the most part, we do not need to make the state dependency explicit; if we do, we will use the notation  $(\underline{a})_i$  to denote class <sub>$i$</sub> ( $a$ ) (i.e., the class of  $a$  in state  $s_i$ ).

To justify assumption (1), that  $\langle SC, \rightarrow \rangle$  is a partially ordered set, we must show that the relation  $\rightarrow$  should be reflexive, transitive, and anti-symmetric; that is, for all  $A, B, C, e \in SC$ :

- a)  $A \rightarrow A$  (reflexive).
- b)  $A \rightarrow B$  and  $B \rightarrow C \Rightarrow A \rightarrow C$  (transitive).
- c)  $A \rightarrow B$  and  $B \rightarrow A \Rightarrow A = B$  (anti-symmetric).

Reflexivity follows from the observation that performing an operation  $t$  in state  $s$  that does not affect an object  $a$  does not alter  $a$ , implying a flow  $a \rightarrow a$  from state  $s$  to state  $s \cdot t$ . Therefore, an inconsistency arises if  $A \neq A$  for all  $A \in SC$ . Transitivity follows from the observation that  $A \rightarrow B$  implies that it is permissible to move an information item  $\alpha$  from an object in  $A$  to one in  $B$ . Since this puts  $\alpha$  in  $B$ ,  $B \rightarrow C$  implies it is then permissible to move  $\alpha$  to an object in  $C$  and thereby effect a flow from  $A$  to  $C$ . Therefore, an inconsistency arises if  $A \neq C$ . Anti-symmetry follows from the observation that  $A \rightarrow B$  and  $B \rightarrow A$  implies anything in one class can be moved into the other. Therefore, one of the two classes is redundant.

To justify assumption (2), that the set of security classes  $SC$  is finite, we observe simply that no practical system is capable of supporting an infinite number of security classes.

To justify assumption (3), that there exists a lower bound on  $SC$ , we argue that all integers, reals, Booleans, and other data items that normally appear as "constants" in programs should be permitted to flow into all classes (or at least all classes that are bound to the same type of data). Therefore, we assume these values are bound to a lowest security class, which we denote by  $L$ .  $L$  could also be associated with

unoccupied storage locations (i.e., execution of the operation clear(u) could have the effect of setting tag(u) := L). Note also that no generality is lost by the requirement for a single lowest class: if more than one class is desired for different types of constants, no objects need be assigned to the lowest class L.

Before justifying assumption (4), it is necessary to discuss the assumption that  $A \oplus B$  is unique — that is, for  $\underline{a} = A$  and  $\underline{b} = B$ ,  $f(\underline{a}, \underline{b}) = A \oplus B$  regardless of the function  $f$ . It is possible to define different flow relations for different functions; Figure 2.3-1 suggests an example where the function  $f$  places its results in class C and function  $g$  in class D. One problem with this is the possible semantic ambiguity that arises for results in the range of both  $f$  and  $g$  — that is, a value  $\alpha = f(\underline{a}, \underline{b}) = g(\underline{a}, \underline{b})$  for  $\underline{a} = A$  and  $\underline{b} = B$ . Such values are produced from the same operands, and it is possible, under at least some interpretations of the security problem, to say that  $\alpha$  has been associated with the wrong class. At this point, we prefer to avoid this murky issue by stipulating that  $\oplus$  is independent of the function  $f$  used to combine operands. In Section 2.4.1 we show that the effect of a function dependent operator  $\oplus$  can be achieved with processes.

To justify assumption (4), that  $\oplus$  is a least upper bound operator, we must show that for all  $A, B, C \in SC$ :

- a)  $A \rightarrow A \oplus B$  and  $B \rightarrow A \oplus B$ .
- b)  $A \rightarrow C$  and  $B \rightarrow C \Rightarrow A \oplus B \rightarrow C$ .

Property (a) must hold, since it is semantically absurd to suppose that operands cannot flow into the class of a result generated from them.

Also, inconsistencies arise if an operation such as " $c := a + b$ " is




Figure 2.3-1. Example of a Function Dependent Flow Relation.

permitted whereas " $c := a$ " is not, since the latter operation can be performed by executing the former with  $b = 0$ . For part (b), consider four objects  $a, b, c$ , and  $d$  such that  $\underline{a} \rightarrow \underline{c}$ ,  $\underline{b} \rightarrow \underline{d}$ , and  $\underline{c} = \underline{d}$ ; and consider the following program and corresponding operation sequence:

```

c := a;          transfer(:=,a,c)
d := b;          transfer(:=,b,d)
d := f(c,d)      transfer(f,c,d,d)

```

Execution of this program assigns to  $d$  information derived from  $a$  and  $b$ . Therefore, the flow  $\underline{a} \oplus \underline{b} \rightarrow \underline{d}$  is implied semantically, and for consistency we require that the flow relation reflect this fact. Thus for any two classes  $A$  and  $B$ ,  $A \oplus B$  is the least upper bound, also referred to as the "join", of  $A$  and  $B$ . Note that  $A \oplus B \rightarrow C$  also implies  $A \rightarrow C$  and  $B \rightarrow C$  by transitivity and part (a).

A useful property of a least upper bound operator on a partially ordered set is that its domain can be extended to subsets containing more than two elements. This is done as follows: For a non-empty finite subset  $\{A_1, \dots, A_n\}$  of SC, define  $\bigoplus$  by

$$\bigoplus \{A_1, \dots, A_n\} = \begin{cases} A_1 & \text{if } n = 1 \\ \left[ \bigoplus \{A_1, \dots, A_{n-1}\} \right] \oplus A_n & \text{if } n > 1 \end{cases}$$

The following propositions are easily proved by induction on  $n$ , using the definition of  $\oplus$  for two classes and assumption (4).

Proposition 2.3-1.  $\bigoplus$  as defined above specifies the unique least upper bound of  $\{A_1, \dots, A_n\}$ .

Since the ordering of elements in the set  $\{A_1, \dots, A_n\}$  is immaterial - i.e.,  $A_1 \oplus A_2 = A_2 \oplus A_1$  for all  $A_1, A_2$  - the order in which

$\oplus\{A_1, \dots, A_n\}$  is evaluated is immaterial. Therefore, we may write simply

$$A_1 \oplus \dots \oplus A_n \text{ for } \oplus\{A_1, \dots, A_n\}.$$

Proposition 2.3-2.  $A_i \rightarrow B$  ( $1 \leq i \leq n$ ) iff  $A_1 \oplus \dots \oplus A_n \rightarrow B$ .

This proposition states that information in objects  $a_1, \dots, a_n$  can flow into an object  $b$  if and only if information derived from  $a_1, \dots, a_n$  (e.g.,  $f(a_1, \dots, a_n)$ ) can flow into  $b$  (since  $\underline{a_i} \rightarrow \underline{b}$  ( $1 \leq i \leq n$ ) iff  $\underline{a_1} \oplus \dots \oplus \underline{a_n} \rightarrow \underline{b}$ ). Therefore, from an initial state  $s_i$ , the safeness of state  $s_j = s_i \cdot \text{transfer}(f, a_1, \dots, a_n, b)$  is established by verifying either

$$\begin{cases} (\underline{a_k})_i \rightarrow (\underline{b})_j & (1 \leq k \leq n) \\ \text{and } (\underline{b})_i \rightarrow (\underline{b})_j & \text{if the contents of } b \text{ are modified,} \end{cases}$$

or

$$\begin{cases} (\underline{a_1})_i \oplus \dots \oplus (\underline{a_n})_i \rightarrow (\underline{b})_j & \text{if the contents of } b \text{ are replaced} \\ (\underline{a_1})_i \oplus \dots \oplus (\underline{a_n})_i \oplus (\underline{b})_i \rightarrow (\underline{b})_j & \text{if the contents of } b \text{ are modified.} \end{cases}$$

Proposition 2.3-2 is the basis of the verification procedures presented in Chapters 3 and 4.

We now prove that assumptions (1) - (4) imply the existence of a greatest lower bound operator on the set of security classes SC. For any pair of classes  $A$  and  $B$ , define

$$LB(A, B) = \{C \mid C \rightarrow A \text{ and } C \rightarrow B\}.$$

$LB(A, B)$  is the set of possible lower bounds for both classes. Observe that the lowest class  $L$  is a member of  $LB(A, B)$  by assumption (3), so that  $LB(A, B)$  is non-empty for all  $A, B$ . Define the operator  $\oplus$  to be



the least upper bound of the lower bounds:

$$A \otimes B = \bigoplus LB(A, B).$$

Our claim that  $\otimes$  is a greatest lower bound operator is established by the following proposition.

Proposition 2.3-3. For a flow structure  $FS = \langle SC, \rightarrow, \otimes \rangle$ , the operator  $\otimes$  is a greatest lower bound operator on  $SC$ ; i.e.,  $\otimes$  satisfies the following properties for all  $A, B, C \in SC$ :

- a)  $A \otimes B \rightarrow A$  and  $A \otimes B \rightarrow B$ .
- b)  $C \rightarrow A$  and  $C \rightarrow B \Rightarrow C \rightarrow A \otimes B$ .

proof. By Proposition 2.3-1 and  $L \in LB(A, B)$ ,  $\otimes$  is well-defined.

By Proposition 2.3-2,  $\bigoplus LB(A, B) \rightarrow A$  and  $\bigoplus LB(A, B) \rightarrow B$ , establishing property (a). If  $C \rightarrow A$  and  $C \rightarrow B$ , then  $C \in LB(A, B)$  and thus  $C \rightarrow \bigoplus LB(A, B)$  by Proposition 2.3-1. Therefore property (b) also holds.

Thus  $\otimes$  is shown to be the greatest lower bound operator, also called the "meet" operator, of the set of security classes  $SC$ . Note that its existence follows from the earlier assumptions and requires no new semantic assumptions. Also note that if  $C \rightarrow A \otimes B$ , then  $C \rightarrow A$  and  $C \rightarrow B$  by transitivity and property (b).

As with the least upper bound operator, the greatest lower bound operator  $\otimes$  can also be extended to operate on subsets of the security classes  $SC$ . In this case for a non-empty finite subset  $\{A_1, \dots, A_n\}$  of  $SC$ , we define  $\bigotimes$  by

$$\bigotimes \{A_1, \dots, A_n\} = \begin{cases} A_1 & \text{if } n = 1 \\ [\bigotimes \{A_1, \dots, A_{n-1}\}] \otimes A_n & \text{if } n > 1 \end{cases}$$

Using inductive arguments, it is easily shown that the following is true.

Proposition 2.3-4.  $\otimes$  as defined above specifies the unique greatest lower bound of  $\{A_1, \dots, A_n\}$ .

Therefore, we can write

$$A_1 \otimes \dots \otimes A_n \text{ for } \otimes \{A_1, \dots, A_n\}.$$

Proposition 2.3-5.  $B \rightarrow A_i$  ( $1 \leq i \leq n$ ) iff  $B \rightarrow A_1 \otimes \dots \otimes A_n$ .

This proposition allows us to verify the security of a set of transfer operations  $\{\text{transfer}(f, a, b_1), \dots, \text{transfer}(f, a, b_n)\}$  by showing that  $a \rightarrow b_1 \otimes \dots \otimes b_n$ . As with Proposition 2.3-2, this result forms the basis of the verification algorithms.

We now observe that the lowest class  $L$  is just the greatest lower bound of the entire (finite) set of security classes  $SC$  (i.e.,  $L = \otimes SC$ ). Furthermore, there exists a highest class, denoted by  $H$ , which is the least upper bound of the entire set  $SC$  (i.e.,  $H = \oplus SC$ ). Although information in  $H$  can be derived from information in any other class, it is not allowed to flow outside of  $H$  to another class. As with the requirement that there exists a lowest class  $L$ , no generality is lost by the requirement that  $H$  exist, since no objects need be assigned to it (this is also true of any class needed to complete the lattice structure).

That the information flow structure  $FS$  forms a lattice  $\langle SC, \rightarrow, \otimes, \oplus \rangle$  with universal bounds  $L$  and  $H$  now follows from assumptions (1) - (4) and Proposition 2.3-3.

### 2.3.2 Examples

An example of a structure satisfying the lattice property is derived from a linear ordering on a set of security classes SC:

- a)  $SC = \{A_1, \dots, A_n\}$
- b)  $A_i \rightarrow A_j$  iff  $i \leq j$
- c)  $A_i \oplus A_j \equiv A_{\max(i,j)}$
- d)  $A_i \ominus A_j \equiv A_{\min(i,j)}$
- e)  $L = A_1; H = A_n$

This is just a generalization of the example illustrated in Figure 2.1-1 (a) for the special case  $n = 4$ . This structure is suitable for any system in which the classes are linearly (or hierarchically) ordered. A common case is a government or military system in which the security classes are determined solely from the four security levels: unclassified, classified, secret, and top secret. Another case is found in a system that needs only two classes: unconfidential (L) and confidential (H), with the single security requirement that confidential information cannot flow into an unconfidential object. This case is considered by Denning, Denning, and Graham [DD74] and also by Fenton (using the names null and priv for the lowest class L and the highest class H, respectively) [Fe74a].

A richer structure satisfying the lattice property is derived from a non-linear partial ordering on the set of all subsets that can be constructed from a finite set X:

- a)  $SC = \text{powerset}(X)$
- b)  $A \rightarrow B$  iff  $A \subseteq B$
- c)  $A \oplus B \equiv A \cup B$

$$d) \quad A \otimes B \equiv A \cap B$$

$$e) \quad L = \emptyset; \quad H = X$$

Figure 2.3-2 illustrates this "lattice of subsets" for  $X = \{x, y, z\}$ . This structure is suitable for any system in which the classes are determined by a set of properties  $X$ , and an information item  $\alpha$  is not permitted to flow into an object  $b$  unless  $b$  has at least all of the properties that  $\alpha$  has. Consider, for instance, a system that contains medical, financial, and criminal records on individuals (i.e.,  $X = \{\text{med}, \text{fin}, \text{crim}\}$ ). Then medical information would be permitted to flow into only those objects  $a$  for which  $\text{med} \in \underline{a}$ , or a combination of medical and financial information would be permitted to flow into only those objects  $a$  for which  $\text{med} \in \underline{a}$  and  $\text{fin} \in \underline{a}$ .

The information flow structure FS of ADEPT [We69] forms a lattice  $\langle SC, \rightarrow, \otimes, \oplus \rangle$  determined by the (cartesian) product ordering of two lattices  $FS_1$  and  $FS_2$ , where  $FS_1$  is a linear ordered lattice and  $FS_2$  is a subset lattice. The lattice  $FS_1$  is derived from a linear ordering of a set of Authority Levels  $SC_1 = \{A_1, \dots, A_m\}$  corresponding to the unclassified, confidential, etc. levels of government and military security.  $FS_2$  is derived from an ordering by subsets of the set of all subsets  $SC_2$  determined by a collection of Categories (properties)  $X = \{x_1, \dots, x_n\}$ , corresponding to special control compartments used to restrict access by project and area. The information structure FS then forms the lattice  $\langle SC, \rightarrow, \otimes, \oplus \rangle$  defined by:

$$a) \quad SC = SC_1 \times SC_2$$

$$b) \quad \langle A_i, B \rangle \rightarrow \langle A_j, B' \rangle \text{ iff } i \leq j \text{ and } B \subseteq B'$$

$$c) \quad \langle A_i, B \rangle \otimes \langle A_j, B' \rangle \equiv \langle A_{\max(i,j)}, B \cup B' \rangle$$




Figure 2.3-2. Lattice of Subsets of  $X = \{x, y, z\}$ .

$$d) \langle A_i, B \rangle \otimes \langle A_j, B' \rangle = \langle A_{\min(i,j)}, B \cap B' \rangle$$

$$d) L = \langle A_1, \emptyset \rangle; H = \langle A_m, X \rangle$$

where  $B$  and  $B'$  are subsets of  $X$  (i.e., elements of  $SC_2$ ). The Franchise component of ADEPT's security profiles (corresponding to "need to know") is not considered to be a component of its information flow structure, since it is not partially ordered (transitivity is violated). However, this component can be accounted for with a mechanism controlling processes' read and write access rights.

### 2.3.3 Additional Lattice Properties

We now examine some additional well known properties of lattices and show that they are consistent, as they apply to  $\otimes$ , with the semantics of the problem:

- 1)  $A \otimes A = A$ ;  $A \otimes A = A$  (idempotent)
- 2)  $A \otimes B = B \otimes A$ ;  $A \otimes B = B \otimes A$  (commutativity)
- 3)  $A \otimes (B \otimes C) = (A \otimes B) \otimes C$ ;  
 $A \otimes (B \otimes C) = (A \otimes B) \otimes C$  (associativity)
- 4)  $A \rightarrow B \Rightarrow A \otimes B = B$  and  $A \otimes B = A$  (absorption)

The idempotent law is obviously consistent with the semantics since information derived only from information in a class  $A$  should remain in class  $A$ . The commutativity and associativity laws are consistent since the class of an information item  $\alpha$  derived from three objects  $a$ ,  $b$ , and  $c$  should be independent of the order in which the information in  $a$ ,  $b$ , and  $c$  is combined. For example, if  $\alpha$  is derived from the expression " $(a + b) + c$ " and  $\alpha'$  from the expression " $b + (c + a)$ ", then it should be the case that both  $\alpha = \alpha'$  and  $\underline{\alpha} = \underline{\alpha}'$ . Finally the absorption law is consistent since, if it is permissible

to cause a flow from A to B, then information in B in fact represents the combination of information derived from A and B. This law implies that the lowest class L is an identity element over  $\oplus$ , so that information combined with data in the lowest security class (e.g., constants) remains in its original class, and that the highest class H is an identity element over  $\oplus$ . However L is an absorbing element over  $\oplus$  since  $L \oplus A = L$ , and H is an absorbing element over  $\oplus$  since  $A \oplus H = H$ . (The last case states that once information is linked with highest security data, it remains in the highest security class.)

#### 2.4 Classification of Systems

Mechanisms for maintaining a system in secure states can be classified along two independent, structural dimensions. The first is the number of security classes permitted in each process's working store; single-class working stores are an interesting special case. The second is the extent to which objects are permanently bound to security classes, and ranges from "static binding" (no changes permitted) at one extreme to "dynamic binding" (an object changes class according to the class of information affecting its content) at the other.

The first dimension we consider is the structure imposed on the working store (also referred to as the computational memory or virtual memory) W of a system. The working store is the subset of all logical objects N in which standard arithmetic and logical operations (e.g., addition and comparison) can be performed. It is distinguished from the file store F in that information in a file must be transferred explicitly into the working store of some process (via an input instruction) before it can be operated upon or moved to another file.

We assume that  $N = W \cup F$ . Letting  $W_p$  denote the working store of process  $p$  (i.e.,  $W_p \subseteq W$ , and  $p$  read  $a$  or  $p$  write  $a$  for all  $a \in W_p$ ), we consider two cases: 1)  $|W_p| = 1$  (single-class) and 2)  $|W_p| \geq 1$  (multi-class).

The single-class case specifies the working store of a process to be a single object, and therefore bound to a single class. Hence, all operations that cause information flow within the working store of a process are necessarily secure, and this includes most of the non-I/O instructions found in programming languages. This case corresponds to the specifications of ADEPT [We69], Rotenberg's Privacy Restriction Monito. [Ro74], the CASE system [Wa74], and the MITRE model [BL73a].

The multi-class case specifies that the working store of a process may be more than one object, and therefore bound to more than one class. Here, all operations that cause flow within the working store of a process are not necessarily secure, so mechanisms are needed to guarantee the secure execution of all non-I/O as well as I/O instructions. This case corresponds to the abstract model proposed by Fenton [Fe73, Fe74a, Fe74b] and the systems described in Chapters 3 and 4 of this thesis.

The alert reader may wonder why the single-class case restricts  $W_p$  to a single object, rather than a set of objects from the same security class. The reason is that it is impossible to combine the notion of dynamic binding with a single-class, multiple object, working store, since the classes of objects bound dynamically can change over time.



The single-class model is useful for studying systems in which a process's working store is the only private store it has and files are generally subject to access by many processes. In contrast, systems which permit objects in virtual memory to be shared (i.e., the working store of a process is not completely private) must usually be treated with the multi-class model.

In the context of single-class systems, the only transfer operations that require verification are I/O operations that transfer data either from a file  $a$  into the working store  $W_p$  of a process  $p$  (i.e., transfer(input, a,  $W_p$ )) or from a working store  $W_p$  into a file  $b$  (i.e., transfer(output,  $W_p$ , b)). Therefore, the security requirements will be satisfied if these operations satisfy the conditions:

- i)  $a \oplus W_p \rightarrow W_p$ , for transfer(input, a,  $W_p$ );
- ii)  $b \oplus W_p \rightarrow b$ , for transfer(output,  $W_p$ , b).

In the context of multi-class systems, I/O operations such as the above need verification as well as all operations that cause a flow of information within a process's working store (e.g., assignment operations).

The second dimension of security system classification is the method of associating (binding) objects to security classes. Static binding means that the security class of an object is constant or invariant over all states of the system unless explicitly altered by the reclassify operation. This is the type of binding used in the CASE system, the MITRE model, and in Fenton's model. Dynamic binding means that the security class of an object varies with its contents as follows: if an operation transfer(f,  $a_1, \dots, a_n, b$ ) is performed, the security requirements are satisfied if the class of  $b$  is updated

according to the assignments:

$\underline{b} := \underline{a_1} \oplus \dots \oplus \underline{a_n}$ , if the contents of  $b$  are replaced; or

$\underline{b} := \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{b}$ , if the contents of  $b$  are modified.

Consequently, execution of a transfer operation in a dynamic binding system is always secure (provided, of course, the updating mechanism functions correctly). In the case of dynamic binding, security classes are actually associated with the information stored in an object rather than with the object itself.

A system based purely on dynamic binding may not be practical: most users, along with their private files, are usually considered to belong to fixed classes. Hence, we must also consider realistic combinations of static and dynamic binding — that is, some objects are statically bound and others are dynamically bound. Such are the specifications of ADEPT and Rotenberg's Privacy Restriction Monitor.

Recall from the results of Section 2.2 that security is guaranteed if the system can be kept in a safe state; i.e., execution of the next operation cannot cause a flow from a class  $A$  to a class  $B$  unless  $A \rightarrow B$ . Recall also that the only operations that can potentially cause a flow from  $A$  to  $B$  ( $A \neq B$ ) are create object, reclassify, and transfer. Now, the security of object creation is assured if the system always clears memory before assigning it to a new object. The security of reclassification is assured by a simple mechanism that checks  $\underline{a} \rightarrow A$  whenever an operation reclassify( $\underline{a}, A$ ) is attempted to change the class of object  $\underline{a}$ . The security of transfer operations is, therefore, the only remaining problem to consider when analyzing the security requirements of the different types of systems. This will be done in the following sections which describe five types of systems.

### 2.4.1 Single-Class, Static Binding

To guarantee the secure execution of all transfer operations in a single-class system with static binding of objects to security classes, it is necessary and sufficient to show that if a process  $p$  executes an instruction "input  $x$  from  $a$ " then  $\underline{a} @ \underline{W}_p \rightarrow \underline{W}_p$  and if it executes an instruction "output  $x$  to  $b$ " then  $\underline{b} @ \underline{W}_p \rightarrow \underline{b}$ , where  $\underline{W}_p$  is  $p$ 's working store. Also, by the results of Proposition 2.3-2 and reflexivity, these conditions respectively reduce to  $\underline{a} \rightarrow \underline{W}_p$  and  $\underline{W}_p \rightarrow \underline{b}$ . Since execution of these instructions is conditioned on process  $p$  having the necessary read or write access privileges, security is guaranteed if in every state  $s$ ,

a)  $p$  read  $a \Rightarrow \underline{a} \rightarrow \underline{W}_p$ , and

b)  $p$  write  $b \Rightarrow \underline{W}_p \rightarrow \underline{b}$ .

Since grant is the only operation that adds new read and write relations and reclassify is the only operation that changes the class of an object, a safe state can be maintained with a simple mechanism attached to these operations that verifies conditions (a) and (b) above. Hence, in this case, it is not necessary to monitor each I/O operation.

If a process  $p$  has an associated clearance (see Section 2.2), it could be used to determine the class of  $\underline{W}_p$ ; that is, if a process  $p$  is created, its working store  $\underline{W}_p$  would be created and assigned to  $p$  by execution of the operation sequence  $\underline{\text{create}}(\underline{W}_p, p) \cdot \underline{\text{grant}}(p, \underline{W}_p, rw)$ , where  $\underline{p}$  denotes  $\underline{\text{clearance}}(p)$ . In this case,  $\underline{p}$  specifies the highest security class that  $p$  can read from, and the lowest security class that it can write into. If  $a_1, \dots, a_m$  are the files  $p$  can read from, and  $b_1, \dots, b_n$  the files  $p$  can write into, then security requires that

$$\underline{a}_1 \oplus \dots \oplus \underline{a}_m \rightarrow p \rightarrow \underline{b}_1 \oplus \dots \oplus \underline{b}_n$$

by conditions (a) and (b) above and the results of Propositions 2.3-2 and 2.3-5. See Figure 2.4-1. For example, in a government or military system, a process having a "secret" clearance would not be able to read from a "top secret" file or write into an "unclassified" one (even if it had processed only unclassified information during execution). Such are the CASE system and the MITRE model. In fact, it does not make sense to design a single-class, static binding system without such a clearance function, since some mechanism is needed to determine the class of a process's computational store when it is created.

In certain cases the effect of a function dependent  $\oplus$  (i.e.,  $\underline{a} \oplus \underline{b}$  depends on the operation performed on  $\underline{a}$  and  $\underline{b}$  — see Section 2.3) can be achieved with processes. Let  $p$  and  $q$  be processes such that

- a)  $p$  and  $q$  read from common input files  $\underline{a}$  and  $\underline{b}$ , where

$$\underline{a} \oplus \underline{b} \rightarrow p \text{ and } \underline{a} \oplus \underline{b} \rightarrow q;$$

- b)  $p$  writes into output file  $\underline{c}$ , where  $\underline{p} \rightarrow \underline{c}$ ;

- c)  $q$  writes into output file  $\underline{d}$ , where  $\underline{q} \rightarrow \underline{d}$ ; and

- d)  $\underline{c} \neq \underline{d}$ .

Clearly, all security requirements are satisfied. Now, execution of  $p$  has the effect of combining information in classes  $\underline{a}$  and  $\underline{b}$  and putting the result in class  $\underline{c}$ ; execution of  $q$  has the effect of combining information in classes  $\underline{a}$  and  $\underline{b}$  and putting the result in class  $\underline{d}$ . Therefore,  $p$  implements  $\underline{a} \oplus \underline{b} = \underline{c}$ , and  $q$  implements  $\underline{a} \oplus \underline{b} = \underline{d}$ .

#### 2.4.2 Single-Class, Dynamic Binding

To guarantee the secure execution of all transfer operations in a single-class system with dynamic binding of objects to security

Input Files

Output Files




Figure 2.4-1. Information Flow in a Single-Class, Static Binding System.

classes, it is necessary and sufficient to attach to the I/O processor a mechanism that correctly updates the class of  $W_p$  whenever a process  $p$  executes an instruction "input x from a" and the class of a file  $b$  whenever  $p$  executes an instruction "output x to b", namely

$W_p := W_p \oplus a$ , for input, and

$b := W_p \oplus b$ , for output.

Such is the mechanism used by Rotenberg's Privacy Restriction Monitor for controlling flow to segments and the mechanism used by ADEPT for changing the "history" class associated with the working store of a process.

The existence of a clearance function does not seem necessary in this type of system since the automatic class updating mechanism attached to the I/O processor makes it unnecessary to constrain what files a process can access — security is always guaranteed.

#### 2.4.3 Single-Class, Static and Dynamic Binding

The concepts of static and dynamic binding can be combined in a single system, with some objects statically bound and others dynamically bound. For example, in ADEPT, all existing files are statically bound, while new files and the working stores of processes are dynamically bound. One possible specification for a secure system incorporating both types of binding and making use of a clearance function is now outlined. (See Figure 2.4-2). Each process  $p$  is initialized with a dynamically bound working store  $W_p$  in the lowest security class  $L$ ; as it inputs information from files,  $W_p$  is updated as described for the dynamic binding case. However, unlike a pure dynamic binding system,  $p$  is permitted to gain (and keep) read access to a file  $a$  only if the




Figure 2.4-2. Information Flow in a Single-Class System with Static and Dynamic Binding.

current class of file  $a$  satisfies the constraint  $\underline{a} \rightarrow p$ . It is free to acquire write access to a file  $b$  (except as constrained by determinacy requirements, "need to know", or other conditions), though it is not permitted to exercise such a right if  $b$  is statically bound and the current class of  $\underline{W}_p$  does not satisfy the constraint  $\underline{W}_p \rightarrow \underline{b}$  (if  $b$  is dynamically bound, it can, of course, exercise the right since  $\underline{b}$  will be updated to  $\underline{b} := \underline{b} \oplus \underline{W}_p$ ). It is precisely in this last requirement that ADEPT has (from our point of view) a security flaw. In ADEPT, a process with "top secret" clearance can write into any statically bound file  $b$  for which  $\underline{b} \rightarrow p$  (rather than  $\underline{W}_p \rightarrow \underline{b}$ ), making it permissible to write "top secret" data into an "unclassified" file!

In such a system, we define a routing process to be any process  $p$  with highest clearance (i.e.,  $p = H$ ), since such a process is free (insofar as the model is concerned) to read any file and thus route any information in a class  $A$  to a file in a resulting class  $B$ , provided only that  $A \rightarrow B$ . Routing processes are particularly applicable to data base systems in which the data is distributed among several files associated with different classes, is to be accessed directly only by a specified set of data manipulation processes (the routing processes), and is to be transmitted to users according to a specified flow relation.

There are two major limitations with the single-class systems outlined above. The first is that it is not possible to distinguish different classes of data in a process's working store, since everything must get lumped into a single class. Hence, a process cannot directly access and manipulate data in segments (or registers) belonging to different classes as would be required by certain data base systems



where it is not practical to physically isolate the different classes of information into separate files. The second is that it is not possible to account for the shared use of segments (or registers) between two processes  $p$  and  $p'$ . Hence, it is not possible to model inter-process communication via message buffers, semaphores, and so on.

#### 2.4.4 Multi-Class, Static Binding

The primary difficulty with guaranteeing security in a multi-class system lies in detecting (and monitoring) all transfer operations. This is because all transfer operations in a program are not explicitly specified — or indeed even executed! We showed in the previous sections that the only flows possible in a single-class system are "explicit" ones caused directly by I/O operations. When we turn to multi-class systems we find that some transfer operations applied within the working store cause "implicit" flows in addition to the explicit flows for which they were programmed; this "extra" flow complicates security checking. As an example, consider the statement "if  $a = 0$  then  $b := 0$ "; if  $b \neq 0$  initially, testing  $b = 0$  on termination of this statement is tantamount to knowing whether  $a = 0$  or not. In other words, information flows from  $a$  to  $b$  regardless of whether or not the then clause is executed.

To deal with this problem, we distinguish between two types of flow: "explicit" and "implicit". Explicit flow to an object  $b$  occurs as the result of executing any statement (e.g., assignment or I/O) that directly transfers to  $b$  information derived from some sequence of operand objects  $a_1, \dots, a_n$ . Examples of explicit flow are:

<u>Statement</u>	<u>Explicit Transfer</u>
$b := a$	<u>transfer</u> (:=,a,b)
$b := a * b$	<u>transfer</u> (*,a,b,b)
<u>input</u> b <u>from</u> a	<u>transfer</u> ( <u>input</u> ,a,b)
<u>output</u> a <u>to</u> b	<u>transfer</u> ( <u>output</u> ,a,b)

Implicit flow to b occurs as the result of executing or not executing a statement that causes an explicit flow to b when that statement is conditioned on the value of a Boolean expression. Examples of statements that cause implicit as well as explicit flow are:

<u>Statement</u>	<u>Explicit Transfer</u>	<u>Implicit Transfer</u>
<u>if</u> a = 0 <u>then</u> b := 0	<u>transfer</u> (:=,0,b)	<u>transfer</u> (=,a,0,b)
<u>if</u> a = 0 <u>then</u> <u>input</u> b <u>from</u> c	<u>transfer</u> ( <u>input</u> ,c,b)	<u>transfer</u> (=,a,0,b)
<u>while</u> a > 0 <u>do</u> b := b + 1; a := a - 1 <u>end</u>	<u>transfer</u> (+,b,1,b) <u>transfer</u> (-,a,1,a)	<u>transfer</u> (>,a,0,b) <u>transfer</u> (>,a,0,a)
<u>if</u> a = 0 <u>then</u> <u>if</u> b = 0 <u>then</u> c := d	<u>transfer</u> (:=,c,d)	<u>transfer</u> (=,a,0,c) & <u>transfer</u> (=,b,0,c)

To guarantee the secure execution of a program in the presence of implicit flows, it is convenient to consider an abstract representation of programs that preserves the flows but not necessarily all of the original structure. Define recursively an abstract program  $\Phi$  by:

1.  $\Phi$  is an elementary statement; that is, an assignment or I/O statement.
2. There exist  $\Phi_1$  and  $\Phi_2$  such that  $\Phi = \Phi_1; \Phi_2$ .
3. There exist  $\Phi_1, \dots, \Phi_m$  and an m-valued expression  $\xi$  such that  

$$\Phi = \xi; \Phi_1, \dots, \Phi_m.$$

Step (1) declares simple statements as abstract programs. Step (2) declares sequences of simpler programs as abstract programs. Step (3) declares conditional structures, in which an expression  $\xi$  selects among alternative programs, as abstract programs.

The conditional structure is used here to represent all conditional (including iterative) statements found in programming languages. For example, "if  $\xi$  then  $\phi_1$  else  $\phi_2$ " is represented by " $\xi:\phi_1,\phi_2$ ". Both "if  $\xi$  then  $\phi_1$ " and "while  $\xi$  do  $\phi_1$ " are represented by " $\xi:\phi_1$ ", and "do case  $\xi$  of  $\phi_1;\dots;\phi_m$ " is represented by " $\xi:\phi_1,\dots,\phi_m$ ". Structures arising from the use of goto statements can also be represented by the conditional structure, but to do so is more difficult (see Section 3.3).

The security requirements for any program of the above form are now stated simply. Firstly, for any elementary statement  $E$ ,  $\phi = E$  is secure if any explicit flow caused by  $E$  is secure (no implicit flow is possible). Specifically, if  $E$  represents an operation transfer( $f,a_1,\dots,a_n,b$ ), the security requirements stipulate that

$$\underline{a_1} \oplus \dots \oplus \underline{a_n} \rightarrow \underline{b}, \text{ if } E \text{ replaces the contents of } b;$$

$$\underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{b} \rightarrow \underline{b}, \text{ if } E \text{ modifies the contents of } b.$$

For example, if the unit of protection is a single register and  $E$  is the statement " $d := a + b * c$ ", security requires that  $\underline{a} \oplus \underline{b} \oplus \underline{c} \rightarrow \underline{d}$ . Secondly,  $\phi = \phi_1;\phi_2$  is secure if both  $\phi_1$  and  $\phi_2$  alone are secure (see Theorem 2.2-1). Thirdly,  $\phi = \xi:\phi_1,\dots,\phi_m$  is secure if each  $\phi_k$  alone is secure and all implicit flows from  $\xi$  are secure. To make this specific, let  $V(\phi) = \{b_1,\dots,b_n\}$  denote the set of objects that information can explicitly flow to in  $\phi$  — that is, for each  $b_i$  ( $1 \leq i \leq n$ ), there is an operation transfer( $\dots,b_i$ ) is some  $\phi_k$ ; all implicit flow

is secure if  $\underline{E} \rightarrow \underline{b}_i$  ( $1 \leq i \leq n$ ) or equivalently (by Proposition 2.3-5) if  $\underline{E} \rightarrow \underline{b}_1 \text{ @ } \dots \text{ @ } \underline{b}_n$  after any execution of the structure.

The foregoing requirements on multi-class systems with static binding can be implemented as run-time checks (e.g., on tags associated with objects in the working store) or as compiler-verification procedures (to be discussed in Chapter 3).

An interesting example of a multi-class, static-binding system has been proposed by Fenton [Fe73, Fe74a]. Fenton chose to cast his results in the context of a Minsky machine [Mi67], which is an automaton modeling a multi-register computer. Fenton extends the Minsky machine by permanently tagging each register with the class to which information residing in it belongs; this corresponds to the tag function defined earlier. The Minsky machine has only two instructions in addition to a halt instruction: the "increment" instruction for any register  $a$ , having the effect " $a := a + 1$ "; and the "decrement and jump if zero" instruction, having the effect "if  $a = 0$  then goto  $m$  else  $a := a - 1$ ". (Minsky has shown that this machine has universal computing power. Fenton used it as a formal model within which to prove a secure environment.) Now, the decrement instruction permits information to flow (implicitly) from  $a$  to any register affected by any statement reachable from the statement labelled  $m$ . To enable flow-checking, Fenton tags the program counter  $p$ , updating the class of  $p$  according to  $\underline{p} := \underline{p} \text{ @ } \underline{a}$  on execution of the above decrement instruction (in Fenton's model,  $\underline{p}$  is the only class that can be changed). The successful execution of any increment or decrement instruction applied to register  $a$  depends on the pre-condition  $\underline{p} \text{ @ } \underline{a} \rightarrow \underline{a}$ , which is verified prior to executing the

instruction. It is the nature of the above mechanism that the sequence of classes of  $p$  is nondecreasing in security. To permit the class of  $p$  to be decreased, Fenton introduces a stack of  $p$ -classes and a new instruction "if  $a = 0$  then call  $m$  else  $a := a - 1$ ", where the operation "call  $m$ " pushes  $(p, p)$  onto the stack and sets  $(p, p) := (m, p \oplus a)$ . An instruction return is also provided so that the block of instructions beginning at location  $m$  can return control, pop the stack, and continue execution following the above if-statement. Fenton shows that secure execution is guaranteed with the new instruction added.

Because he develops his  $\oplus$  operator to be consistent with the properties of the Minsky machine without considering the semantics of the information flow problem, Fenton's flow structure  $\langle SC, \rightarrow, \oplus \rangle$  is not a lattice, though it is a partial ordering of  $SC$ . The discrepancy is that his  $\oplus$  operator does not satisfy the least upper bound property, " $A \rightarrow C$  and  $B \rightarrow C \Rightarrow A \oplus B \rightarrow C$ ", and is, as a consequence, not associative. The reason for this "semantic defect" can be understood by observing that, in a Minsky machine, it is impossible to extract directly information from two registers and place the result in a third register. It is possible only to use the contents of one register to modify the contents of another. (For example, the only Minsky program to cause the sum  $a + b$  to appear in register  $c$ , assuming  $c = 0$  initially, uses the circuitous procedure of adding  $a$  to  $b$  then copying the result to  $c$ .) For this reason, Fenton defined his  $\oplus$  operator as follows:


$$A \oplus B = \begin{cases} A & \text{if } B \rightarrow A \quad (B \text{ can modify } A) \\ B & \text{if } A \rightarrow B \quad (A \text{ can modify } B) \\ H & \text{otherwise (neither } A \text{ nor } B \text{ can modify the other).} \end{cases}$$

An example of a complete structure satisfying these properties is given

in Figure 2.4-3. A requirement such as "the combination of information in classes A and B can be placed in class C" is handled by defining  $A \rightarrow C$  and  $B \rightarrow C$  (i.e.,  $A \oplus C = C$  and  $B \oplus C = C$ ), so that a program can use information in classes A and B to modify registers of class C and create the desired combination piecemeal in registers of C. Since it is possible to have  $A \rightarrow C$  and  $B \rightarrow C$  but not  $A \oplus B \rightarrow C$ ,  $\oplus$  is not a least upper bound operator.

When considering how to modify Fenton's  $\oplus$  operator so that it becomes a lattice operator, one quickly discovers that it cannot be done without adding new security classes. Figure 2.4-4 shows a modification of the structure of Figure 2.4-3 to this end, a new class AB being added. Since Fenton's structure can be extended to a lattice, one can inquire whether his structure is semantically meaningful. The answer is, of course, affirmative in the context of the Minsky machine for which it was derived.


The problem under consideration — secure information flow in multi-class systems with static binding of security classes — can be solved completely by a compile-time mechanism, without the need for tagged architecture or run-time hardware checking mechanisms (see Chapter 3). However, the run-time mechanism could be used as a redundant check in case, for example, a hardware malfunction introduced an error into a previously certified program. (See Fabry [Fa73] for a general discussion of redundancy checks to enhance system reliability.) The compile-time mechanism has the further attraction that it eliminates the possibility of "covert leakages" caused by tripping alarms built into run-time mechanisms.

Graph of  $\rightarrow$ 

	L	A	B	C	D	H
L	L	A	B	C	D	H
A	A	A	H	C	D	H
B	B	H	B	C	D	H
C	C	C	C	C	H	H
D	D	D	D	H	D	H
H	H	H	H	H	H	H

Matrix of  $\oplus$ 

Figure 2.4-3. Example of an Information Flow Structure Satisfying Fenton's Model.

Graph of  $\rightarrow$ 

	L	A	B	AB	C	D	H
L	L	A	B	AB	C	D	H
A	A	A	H	AB	C	D	H
B	B	H	B	AB	C	D	H
AB	AB	AB	AB	AB	C	D	H
C	C	C	C	C	C	H	H
D	D	D	D	D	H	D	H
H	H	H	H	H	H	H	H

Matrix of  $\oplus$ 

Figure 2.4-4. Example of Figure 2.4-3 Extended to a Lattice.

### 2.4.5 Multi-Class, Dynamic Binding

Mechanisms for guaranteeing security when all (or some) objects are bound dynamically to security classes require both compile-time and run-time support and are the most complex of all. As an example, consider a system in which the unit of protection is a single register and all objects are bound dynamically to classes. Suppose the following program (proposed by Fenton [Fe73] is executed to set *b* equal to the value of *a* (assumed to be either 0 or 1 initially):

```

b := 0;
c := 0;
if a = 0 then c := 1;
if c = 0 then b := 1

```

The "naive" run-time mechanism (see also [DD74]) for dynamically binding registers to classes operates by updating the class of a register whenever an assignment is made to it. For the above program, this proceeds as follows:

<u>Statement</u>	<u>Action</u>
b := 0	<u>b</u> := L
c := 0	<u>c</u> := L
if a = 0 then c := 1	if c := 1 is executed, <u>c</u> := <u>a</u>
if c = 0 then b := 1	if b := 1 is executed, <u>b</u> := <u>c</u>

When this program terminates, the class of *b* is the lowest class L (i.e., b = L) irrespective of the class of *a*, but the value of *b* is identically that of *a*. This is because the mechanism does not account for the implicit flow that occurs when the statements "*c* := 1" and "*b* := 1" are not executed. Therefore, a security violation results unless a = L.

The above program illustrates that a run-time mechanism is not sufficient to guarantee security since implicit flows that occur in the absence of explicit flows (from assignment or I/O operations)



go undetected. For this reason, Fenton dismisses the problem as insoluble in his framework which relies solely on a run-time mechanism. In Chapter 4 we show that the problem can be solved with a compile-time procedure that inserts special code (if need be) into a program so that all implicit flows are accounted for. In the program given above, for example, the compiler would insert code so that execution of the statement

if a = 0 then c := 1

results in the assignment

c := a

even if a ≠ 0. The point is, the problem of implicit flow can be solved if one is willing to augment the run-time mechanism with a compile-time mechanism.

## 2.5 Determinacy Requirements

The existence of appropriate security flow controls does not mean that the system is fully secure if information can still be communicated along "covert channels" [La73]. One such case arises if the system permits the class of a dynamically bound object to be observed, and the final class of the object can be made to depend on the relative speeds of parallel processes; this non-determinacy can be used to transmit information to an observer. For such systems, additional constraints may be required to guarantee determinate operation. Determinacy is, of course, usually considered desirable in its own right, because it guarantees the reproducibility of a system's actions under arbitrary timings of processes that share information.

To formulate conditions that guarantee determinate behavior, we consider first the simple case of a single-class system under dynamic binding of security classes (see Section 2.4.2). We then extend the results to a single-class system under both static and dynamic binding of security classes (see Section 2.4.3) and to a multi-class system under dynamic binding of security classes (see Section 2.4.5). In all cases, we assume that the set of all objects and processes is constant over the lifetime of the system (i.e., no create or delete operations are performed). Our treatment closely follows that of Coffman and Denning [CD73].

### 2.5.1 Single-Class, Dynamic Binding

We assume that each process  $p$  performs a computation in its private working store  $W_p$  using as input information from a domain set  $D_p \subseteq F$  and producing as output information in a range set  $R_p \subseteq F$ , where  $F$  is the set of files in the system. The domain  $D_p$  is the set of files to which  $p$  has read access; the range  $R_p$  is the set to which it has write access. We make the restriction that  $D_p$  and  $R_p$  do not change during the execution of  $p$ .

The set  $P$  of all processes in existence is partially ordered by a relation  $<$ , where  $p < p'$  means that  $p$  must terminate before  $p'$  can begin. The system  $\Pi = \langle P, < \rangle$  is called a process system. Figure 2.5-1 is an example of a process system represented by a precedence graph. Associated with each process  $p$  is an "initiation event" corresponding to the operation initiate( $p$ ), denoted by "[ $p$ ". Similarly, there is a "termination event" corresponding to terminate( $p$ ), denoted by " $p$ ".




Figure 2.5-1. Example of a Process System.

For the purposes of proving determinacy, no generality is lost by associating all input operations with the initiation of  $p$  and all output operations with the termination of  $p$ . The reason is that such a proof relies only on the assumption that  $p$  is determinate — i.e., that the class assigned to an output file depends uniquely on the classes of  $p$ 's input files, as long as no input file classes change during  $p$ 's execution; exactly when the classes of output files are updated is, therefore, irrelevant.

An execution sequence of a process system  $\Pi = \langle P, \prec \rangle$  is a string  $\rho = r_1 r_2 \dots r_{2n}$ ,  $n$  being the number of processes in  $P$ , and each  $r_i$  ( $1 \leq i \leq 2n$ ) being a process initiation or termination event such that:

- 1) For each  $p \in P$ , " $p$ " and " $p$ " appears exactly once in  $\rho$ ,
- 2) " $p$ " precedes " $p$ " in  $\rho$ , and
- 3) " $p$ " precedes " $p$ " in  $\rho$  when  $p \prec p'$ .

A valid execution sequence for the system of Figure 2.5-1 is

$[p_1 [p_2 p_1] [p_4 [p_5 p_2] [p_6 p_5] p_6] [p_7 [p_3 p_3] p_7] p_4]$ .

More than one valid execution sequence will exist for  $\Pi$  unless  $\prec$  is a linear ordering relation.

Let  $\sigma(\rho) = s_0 s_1 \dots s_{2n}$  denote the state sequence corresponding to an execution sequence  $\rho$ ; that is,  $s_0$  is the initial state, and

$s_i = s_{i-1} \cdot r_i$  ( $1 \leq i \leq 2n$ ).

Define a class sequence  $c(a, \rho)$  for each file  $a \in F$  to be the sequence of classes  $c_1 c_2 \dots c_m$  to which  $a$  is assigned during the execution sequence  $\rho$ . Letting  $\lambda$  denote the empty sequence, we have:

$$c(a, \lambda) = (\underline{a})_0$$

$$c(a, r_1 \dots r_k) = \begin{cases} (c(a, r_1 \dots r_{k-1}) (\underline{a})_k) & \text{if } r_k = p \text{ and } a \in R_p \\ c(a, r_1 \dots r_{k-1}) & \text{otherwise} \end{cases}$$

where  $(\underline{a})_i$  denotes the class of  $a$  in state  $s_i$ . A condition for determinacy is formulated in terms of the final value of a class sequence  $c(a, \rho) = c_1 c_2 \dots c_m$ . Letting  $v(a, \rho)$  denote  $c_m$ , a protection system is defined to be determinate if  $v(a, \rho) = v(a, \rho')$  for all files  $a \in F$  and all execution sequences  $\rho$  and  $\rho'$ . An example of a system with two independent processes that is not determinate is illustrated in Figure 2.5-2. Here  $v(b, \rho) \neq v(b, \rho')$  for file  $b$  and sequences  $\rho$  and  $\rho'$ .

We have intentionally defined determinacy in terms of final classes rather than class sequences (as is done analogously in Coffman and Denning). Determinacy with respect to final classes is important for two reasons. First, the final class of a file  $a$  written by processes  $p_1, \dots, p_n$  should be uniquely determined by the information  $p_1, \dots, p_n$  write into  $a$  and not by the order in which  $p_1, \dots, p_n$  execute. Second, an observer should not be able to obtain unauthorized information by testing the final class of file  $a$  (e.g., by attempting a read or write operation on  $a$ ). Determinacy with respect to class sequences is not important if no observer is allowed to examine the class sequence of file  $a$ .

Theorem 2.5-1. A process system  $\Pi = \langle P, \prec \rangle$  is determinate if for all  $p$  and  $p' \in P$  one of the following conditions is true:

- a)  $p \prec p'$  or  $p' \prec p$ , or
- b)  $R_p \cap D_{p'} = D_p \cap R_{p'} = \emptyset$ .

$$\Pi = \langle \{p_1, p_2\}, \langle \rangle \rangle$$

$$F = \{a, b, c\}; \quad (\underline{a})_0 = (\underline{b})_0 = L, \quad (\underline{c})_0 = H$$

$$D_{p_1} = a, R_{p_1} = b \text{ (i.e., } p_1 \text{ inputs from } a \text{ and outputs to } b)$$

$$D_{p_2} = c, R_{p_2} = b \text{ (i.e., } p_2 \text{ inputs from } c \text{ and outputs to } a)$$

$$\rho = [p_1 \ p_1] \ [p_2 \ p_2]$$

$$\rho' = [p_2 \ p_2] \ [p_1 \ p_1]$$

Tables showing the classes of the files a, b, and c and of the working stores  $w_{p_1}$  and  $w_{p_2}$  after each operation is performed for  $\rho$  and  $\rho'$ :

	$\overset{\rho}{[p_1 \ p_1] \ [p_2 \ p_2]}$				
a	L	L	L	L	H
b	L	L	L	L	L
c	H	H	H	H	H
$w_{p_1}$	L	L	L	L	L
$w_{p_2}$	L	L	L	H	H

	$\overset{\rho'}{[p_2 \ p_2] \ [p_1 \ p_1]}$				
a	L	L	H	H	H
b	L	L	L	L	H
c	H	H	H	H	H
$w_{p_1}$	L	L	L	H	H
$w_{p_2}$	L	H	H	H	H

Class sequences for b:

$$c(b, \rho) = L \ L \Rightarrow \quad v(b, \rho) = L$$

$$c(b, \rho') = L \ H \Rightarrow \quad v(b, \rho') = H$$

Figure 2.5-2. Example of a Non-Deterministic System.

proof. The proof proceeds by considering a sequence of subsystems  $\Pi_1, \dots, \Pi_n$  in which  $\Pi_1 = \langle \{p_1\}, \emptyset \rangle$  for an initial process  $p_1$  of  $P$  (i.e.,  $p_1$  has no predecessors),  $\Pi_n = \Pi$ , and  $\Pi_{k+1} = \langle P_{k+1}, \prec_{k+1} \rangle$  where  $P_{k+1}$  is obtained from  $P_k$  by adding any process  $p \in P - P_k$  all of whose predecessors (if any) are in  $P_k$  and  $\prec_{k+1}$  is the restriction of  $\prec$  on the subset  $P_{k+1}$ .  $\Pi_1$  is determinate under the assumption that individual processes are determinate. Assume that  $\Pi_k$  is determinate, and let  $p$  be the process added to  $P_k$  to obtain  $P_{k+1}$ . Let  $a \in F$  have initial class  $A$  and suppose  $\Pi_k$  leaves file  $a$  with final class  $\underline{a}$ . We wish to show that, in  $\Pi_{k+1}$ , the final class of  $a$  is  $\underline{a}$  if  $a \notin R_p$  and  $\underline{a} \oplus \underline{W}_p$  if  $a \in R_p$ , where  $\underline{W}_p$  is  $p$ 's working store. We must also show that  $\underline{a}$  is not affected by the presence of  $p$  and that  $\underline{W}_p$  is unique. Conditions (a) and (b) both imply that  $p$  cannot alter domain classes of any  $p' \in P_k$  prior to the initiation of  $p'$ ; therefore,  $\underline{W}_{p'}$  is unique for all  $p' \in P_k$ . Hence,  $\underline{a}$  cannot be altered by the presence of  $p$ . Since the classes of  $D_p$  are produced by  $\Pi_k$  and since  $p$  is determinate, the class  $\underline{W}_p$  when  $p$  performs its last write operation on  $a$  is unique.

Now, for  $a \notin R_p$  it is clear that, in  $\Pi_{k+1}$ , the final class of  $a$  is  $\underline{a}$ . For  $a \in R_p$  the final class of  $a$ , in  $\Pi_{k+1}$ , is

$$\begin{aligned} & A \oplus \left[ \bigoplus \{ \underline{W}_{p'} \mid p' \in P_{k+1} \text{ and } a \in R_{p'} \} \right] \\ &= A \oplus \left[ \bigoplus \{ \underline{W}_{p'} \mid p' \in P_k \text{ and } a \in R_{p'} \} \right] \oplus \underline{W}_p \\ &= \underline{a} \oplus \underline{W}_p. \end{aligned}$$

Therefore, the subsystem  $\Pi_{k+1}$  produces unique final classes for every file and is determinate. Repeating the argument for  $k = 1, 2, \dots, n-1$  establishes the determinacy of  $\Pi$ .

If one wants to guarantee class sequence determinacy, it is sufficient to add the condition  $R_p \cap R_{p'} = \emptyset$  to condition (a) of the theorem statement. In this case, the determinacy proof resembles that given in Coffman and Denning [CD73: Theorem 2.1, p. 39].

### 2.5.2 Single-Class, Static and Dynamic Binding

Let  $F^S$  denote those files statically bound to security classes, and let  $F^d$  denote those files dynamically bound to security classes. Also, let  $P^S$  denote those processes whose working stores are bound statically, and let  $P^d$  denote those processes whose working stores are bound dynamically. We first observe that any file in  $F^S$  can be ignored from determinacy considerations, since its class is invariant and therefore completely independent of the behavior of all processes. Second, the interaction between two processes  $p$  and  $p'$  both in  $P^S$  can be ignored since  $W_p$  and  $W_{p'}$  are predetermined by clearance( $p$ ) and clearance( $p'$ ) respectively (see Section 2.4.1), and therefore unique for all execution sequences. Hence it is not necessary in this case to account for the conditions  $R_p \cap D_{p'} \neq \emptyset$  or  $D_p \cap R_{p'} \neq \emptyset$ . However, the interaction between a process  $p \in P^S$  and  $p' \in P^d$  cannot be ignored since  $W_{p'}$  is determined by  $D_{p'}$ , and therefore is affected by the behavior of  $p$  if  $R_p \cap D_{p'} \cap F^d \neq \emptyset$  (although the condition  $D_p \cap R_{p'} \neq \emptyset$  cannot affect  $W_{p'}$ ). Finally, the interaction between two process  $p$  and  $p'$  both in  $P^d$  cannot be ignored since  $R_p \cap D_{p'} \cap F^d \neq \emptyset$  affects  $W_{p'}$  and  $D_p \cap R_{p'} \cap F^d \neq \emptyset$  affects  $W_p$  as before. These observations lead to the following theorem.



Theorem 2.5-2. A process system  $\Pi = \langle P, \prec \rangle$  is determinate if for all  $p$  and  $p' \in P$  one of the following conditions is true:

- a)  $p, p' \in P^S$ , or
- b)  $p \prec p'$  or  $p' \prec p$ , or
- c)  $R_p \cap D_{p'} \cap F^d = \emptyset$  for  $p \in P^S$  and  $p' \in P^d$ , or
- d)  $R_p \cap D_{p'} \cap F^d = D_p \cap R_{p'} \cap F^d = \emptyset$  for  $p, p' \in P^d$ .

### 2.5.3 Multi-Class, Dynamic Binding

This type of system would appear to be more difficult to analyze since processes may share objects in their working stores (i.e.,  $W_p \cap W_{p'} \neq \emptyset$  is possible) and since a transfer operation replaces (rather than extends) the contents of the receiving object in the case of a scalar variable, for example. We deal with the first problem by partitioning the set of objects in the working store  $W$  into those that are shared, denoted  $W^{sh}$ , and those that are private, denoted  $W^{pr}$ , and defining each process  $p$  to be a computation from a domain  $D_p \subseteq F \cup W^{sh}$  to a range  $R_p \subseteq F \cup W^{sh}$  using its private working store,  $W_p \cap W^{pr}$ , for intermediate results. We deal with the second by taking into account the interaction between two processes  $p$  and  $p'$  when  $R_p \cap R_{p'} \cap W^{rep} \neq \emptyset$ , where  $W^{rep} \subseteq W$  is the set of replacable objects (since the final class of an object  $a \in W^{rep}$  may depend on the order in which the processes write into  $a$ ). The determinacy requirements now follow:

Theorem 2.5-3. A process system  $\Pi = \langle P, \prec \rangle$  is determinate if for all  $p$  and  $p' \in P$  one of the following conditions is true:

- a)  $p \prec p'$  or  $p' \prec p$ , or
- b)  $R_p \cap D_{p'} = D_p \cap R_{p'} = R_p \cap R_{p'} \cap W^{rep} = \emptyset$ .

Note that the above requirements are identical to those of the single-class, dynamic binding case except that processes must also be mutually non-interfering with respect to replacable objects.

## PROGRAM CERTIFICATION UNDER STATIC BINDING

## Chapter 3

In Section 2.4.4 we specified the security requirements of a multi-class system with static binding of security classes. We examined a run-time mechanism proposed by Fenton [Fe73,Fe74a] that guarantees the secure execution of a program in this environment, and we stated that security can also be guaranteed with a compile-time mechanism that certifies the secure execution of a program before it executes. The purpose of this chapter is to present such a mechanism.

Program certification mechanisms have several advantages over run-time verification mechanisms in protection systems. One important advantage is that the execution of a program is guaranteed to be secure before it executes; hence a program cannot leak information by purposely causing security violations (see Section 2.2). Another is that it does not impair the execution speed of a program, since all security checks are performed prior to program execution (although some run-time support may still be necessary or desirable for reliability reasons). Of course this would be of dubious value if the certification mechanism added substantially to the cost of compilation! The mechanism proposed here, however, does not and indeed adds suprisingly little overhead to the compilation process. A third advantage is that the certification process itself can be specified in terms of higher level language

structures, rather than low level hardware instructions. In this form it is more understandable, and correctness is more easily established.

The principles of program certification were first formalized by Floyd [F167] and Naur [Na66] in an effort to develop a methodology for proving program correctness. The basic idea is that the programmer specifies "invariant assertions" about the state of the variables at various steps of the program. Program certification then is the process of proving that these assertions will hold for every possible execution of the program (i.e., for every set of input values). This can be done manually or (in principle at least) automatically by a theorem-proving program. Since it is a slow and tedious manual task, a considerable amount of research effort has been directed toward the construction of efficient automatic certification mechanisms. However, progress in this area has been slower than expected. A complete survey of techniques for proving program correctness has been given by Elspas et al. [EL72].

This chapter demonstrates the construction of an automatic compile-time mechanism that certifies the secure execution of a program under static binding of objects to security classes. In this case, only one form of assertion must be proved at each step in the program, namely, that the information  $\alpha$  stored in a variable  $a$  satisfies the property  $\underline{\alpha} \rightarrow a$ .

This is not the first time compiler certification mechanisms to verify security properties have been proposed. Conway et al. [CM72] investigated compile-time procedures for controlling a program's data independent accesses to objects in a data base (data dependent accesses

are controlled by run-time procedures). Moore [Mo74] investigated a compile-time mechanism that constructs an "information flow graph" showing the flow of information between variables of a program and the conditions under which such flow occurs. Our results complement Moore's. In our case, the flow relation on the security classes is given, and the objective is to verify that the program satisfies these flow requirements. In his case, the objective is to construct a representation of possible flows among the variables of the program. Other related work in the general area of protection and programming languages can be found in [GS73,MC74,Mr73].

The approach we take is to show how to incorporate a certification mechanism into the analysis phase of a syntax-driven compiler. We do this by examining various syntactic constructs typical of high level languages (most of our constructs are taken from PASCAL [Wr71], Algol-60 [Na63], and PL/I [IB71]). For each syntactic type, we specify one or more actions to be performed (along with other semantic actions such as type checking and code generation) when a string of that type is recognized; these actions are referred to as the certification semantics of the language.

To control the complexity of the presentation, the mechanism will be introduced stepwise in six stages:

- 1) Expressions, assignment, and I/O of elementary data structures,
- 2) Conditional statements,
- 3) Gotos,
- 4) Interrupts,
- 5) Complex data structures, and
- 6) Procedure calls.

The objective of stage (1) is to certify all explicit flows among simple data objects. Implicit flows, such as those arising from conditional statements and gotos, are deferred to the later stages (see Section 2.4.4 for a discussion of implicit v. explicit flow). At stage (1) it is also necessary to assume that programs are uninterruptible — i.e., their execution cannot be interrupted for exceptional conditions such as end-file or overflow. The reason is that the occurrence of an interrupt may cause the subsequent execution of statements to be conditioned on the data that caused the interrupt and thereby result in an implicit flow.

At stage (2) we consider the certification of implicit flows that arise during the execution of three common conditional control statements: if-then-else, while-do, and repeat-until. As will be seen, the technique used actually applies to any conditional statement, such as the PASCAL case statement [Wr71] or the Algol for statement [Na63]. At this stage, we are therefore able to certify simple "structured programs".

At stage (3) we consider the rather unattractive complications begotten by the allowance of a goto statement. Since the goto statement permits unrestricted transfer of control, locating all statements conditioned on a particular expression becomes considerably more difficult. Indeed, it is necessary to perform a complete control flow analysis of a program before its security can be certified. Since the use of unrestricted gotos is of questionable value anyway, we do not consider all of the intricacies of the requisite control flow analysis. The purpose of this analysis is to illustrate the general principle

involved in certifying general control flow structures. The fact of its excessive complexity further justifies the practicality of the simpler, structured approach taken at stage (2).

At stage (4) we suggest two viable approaches for guaranteeing security in an environment in which interrupts can occur. The first approach — program structuring — requires a program to conform to a rather restricted structure: it must assign values to security classes in a monotonically increasing order (i.e., it must first compute all its lowest security results, then its next lowest, and so on). The second approach — interrupt inhibition — permits a program to define its own "interrupt handlers" with an on-statement and inhibits all interrupts for which no such action is defined. This latter approach allows greater flexibility than the former, though it does require modest run-time support.

At stage (5) we consider three complex data structures: arrays, records, and lists. Whereas the elements of a record can belong to different security classes, the elements of arrays and lists must belong to the same class, since in the latter cases there is no way of determining at compile-time which element is being referenced.

At stage (6) we consider procedure calls and observe that it is not generally possible for the compiler to complete the certification of a program which calls external procedures. However, the compiler can generate information which when subsequently used by a linker (i.e., a linkage editor or linking loader), will verify that the collection of subroutines bound together is secure. In the special case where a call is made to an external procedure that derives all of its results from

the input parameters and data in the lowest class L, the compiler can complete the verification of the procedure call.

### 3.1 Expressions, Assignment, and I/O of Elementary Data Structures

Consider a language which allows only the elementary data types known as "constants" - i.e., integers, reals, characters, and Booleans - and which allows only the elementary data objects known as "scalar variables" and "files" of the elementary data types. All constants are assumed to belong to the lowest security class L, and all variables and files are assumed to be assigned to a security class by an explicit declaration, such as

a: integer of class A.

(Observe that we are in effect treating security classes as data types.) Arithmetic and Boolean expressions are formed from these data structures in the usual way.

We shall employ the notation <exp> to denote the security class of an expression; for example, the security class of the expression "a \* b + c" is <exp> = a \* b + c. In general, we denote by <x> the security class associated with a string of syntactic type <x>.

There are only three statement types to consider at this point: assignment, input, and output. An assignment statement has the structure

<var> := <exp> .

Since this statement represents a flow from <exp> to <var>, security requires that

<exp> + <var> .

For example, the statement "c := a + b" has the structure



$\langle \text{exp} \rangle = "a + b"$  and  $\langle \text{var} \rangle = "c"$ , with corresponding security requirement  $\underline{a} \oplus \underline{b} \rightarrow \underline{c}$ . An input statement has the structure

input  $\langle \text{inlist} \rangle$  from  $\langle \text{file} \rangle$

where  $\langle \text{inlist} \rangle$  is a list of input variables  $\langle \text{var} \rangle_1, \dots, \langle \text{var} \rangle_n$ . Since this statement represents a flow from  $\langle \text{file} \rangle$  to  $\langle \text{var} \rangle_i$  ( $1 \leq i \leq n$ ), security requires that

$\langle \text{file} \rangle \rightarrow \langle \text{var} \rangle_i$  ( $1 \leq i \leq n$ ).

Letting  $\langle \text{inlist} \rangle = \langle \text{var} \rangle_1 \oplus \dots \oplus \langle \text{var} \rangle_n$ , (by Proposition 2.3-5) this requirement is equivalent to

$\langle \text{file} \rangle \rightarrow \langle \text{inlist} \rangle$ .

For example, the statement "input  $a, b$  from  $c$ " has the structure  $\langle \text{file} \rangle = "c"$  and  $\langle \text{inlist} \rangle = "a, b"$ , with corresponding security requirement  $\underline{c} \rightarrow \underline{a} \oplus \underline{b}$ . An output statement has the structure

output  $\langle \text{outlist} \rangle$  to  $\langle \text{file} \rangle$

where  $\langle \text{outlist} \rangle$  is a list of expressions  $\langle \text{exp} \rangle_1, \dots, \langle \text{exp} \rangle_n$ . Since this statement represents a flow from  $\langle \text{exp} \rangle_i$  to  $\langle \text{file} \rangle$  ( $1 \leq i \leq n$ ), security requires that

$\langle \text{exp} \rangle_i \rightarrow \langle \text{file} \rangle$  ( $1 \leq i \leq n$ ).

Letting  $\langle \text{outlist} \rangle = \langle \text{exp} \rangle_1 \oplus \dots \oplus \langle \text{exp} \rangle_n$ , (by Proposition 2.3-2) this becomes

$\langle \text{outlist} \rangle \rightarrow \langle \text{file} \rangle$ .

For example, the statement "output  $a, b * c$  to  $d$ " has the structure  $\langle \text{outlist} \rangle = "a, b * c"$  and  $\langle \text{file} \rangle = "d"$ , with corresponding security requirement  $\underline{a} \oplus \underline{b} \oplus \underline{c} \rightarrow \underline{d}$ .

Syntax and certification semantics for the above statement types is summarized in Figure 3.1-1. If a rule contains two occurrences of a

### Certification Semantics

Syntax

$\langle \text{addop} \rangle ::= + \mid - \mid \vee$   
 $\langle \text{mulop} \rangle ::= * \mid / \mid \wedge$   
 $\langle \text{relop} \rangle ::= = \mid \neq \mid < \mid \leq \mid > \mid \geq$

$\langle \text{factor} \rangle ::= \langle \text{var} \rangle$   
 $\langle \text{factor} \rangle ::= \langle \text{const} \rangle$   
 $\langle \text{factor} \rangle ::= ( \langle \text{exp} \rangle )$   
 $\langle \text{factor} \rangle ::= \neg \langle \text{factor} \rangle$

$\langle \text{term} \rangle ::= \langle \text{factor} \rangle$   
 $\langle \text{term} \rangle ::= \langle \text{term} \rangle \langle \text{mulop} \rangle \langle \text{factor} \rangle$

$\langle \text{aexp} \rangle ::= \langle \text{term} \rangle \mid \langle \text{addop} \rangle \langle \text{term} \rangle$   
 $\langle \text{aexp} \rangle ::= \langle \text{aexp} \rangle \langle \text{addop} \rangle \langle \text{term} \rangle$

$\langle \text{exp} \rangle ::= \langle \text{aexp} \rangle$   
 $\langle \text{exp} \rangle ::= \langle \text{aexp} \rangle \langle \text{relop} \rangle \langle \text{aexp} \rangle$

$\langle \text{inlist} \rangle ::= \langle \text{var} \rangle$   
 $\langle \text{inlist} \rangle ::= \langle \text{inlist} \rangle , \langle \text{var} \rangle$

$\langle \text{outlist} \rangle ::= \langle \text{exp} \rangle$   
 $\langle \text{outlist} \rangle ::= \langle \text{outlist} \rangle , \langle \text{exp} \rangle$

$\langle \text{assign stmt} \rangle ::= \langle \text{var} \rangle := \langle \text{exp} \rangle$   
 $\langle \text{input stmt} \rangle ::= \text{input } \langle \text{inlist} \rangle \text{ from } \langle \text{file} \rangle$   
 $\langle \text{output stmt} \rangle ::= \text{output } \langle \text{outlist} \rangle \text{ to } \langle \text{file} \rangle$

$\langle \text{factor} \rangle := \langle \text{var} \rangle$   
 $\langle \text{factor} \rangle := \text{L}$   
 $\langle \text{factor} \rangle := \langle \text{exp} \rangle$   
 $\langle \text{factor} \rangle := \langle \text{factor} \rangle$

$\langle \text{term} \rangle := \langle \text{factor} \rangle$   
 $\langle \text{term} \rangle := \langle \text{term} \rangle \langle \text{mulop} \rangle \langle \text{factor} \rangle$

$\langle \text{aexp} \rangle := \langle \text{term} \rangle$   
 $\langle \text{aexp} \rangle := \langle \text{aexp} \rangle \langle \text{addop} \rangle \langle \text{term} \rangle$

$\langle \text{exp} \rangle := \langle \text{aexp} \rangle$   
 $\langle \text{exp} \rangle := \langle \text{aexp} \rangle \langle \text{relop} \rangle \langle \text{aexp} \rangle$

$\langle \text{inlist} \rangle := \langle \text{var} \rangle$   
 $\langle \text{inlist} \rangle := \langle \text{inlist} \rangle \langle \text{var} \rangle$

$\langle \text{outlist} \rangle := \langle \text{exp} \rangle$   
 $\langle \text{outlist} \rangle := \langle \text{outlist} \rangle \langle \text{exp} \rangle$

if  $\langle \text{exp} \rangle \neq \langle \text{var} \rangle$  then SECURITY ERROR

if  $\langle \text{file} \rangle \neq \langle \text{inlist} \rangle$  then SECURITY ERROR  
 if  $\langle \text{outlist} \rangle \neq \langle \text{file} \rangle$  then SECURITY ERROR

Figure 3.1-1. Syntax and Certification Semantics for Assignment and I/O.

syntactic type  $\langle x \rangle$ , we denote the second occurrence by  $\langle x \rangle'$  so that the semantics are not ambiguous. The interpretation of the certification semantics for a rule

$$\langle x \rangle ::= \omega$$

where  $\omega$  is a string of terminals and non-terminals is straightforward. When a string of form  $\omega$  is recognized—i.e., a "syntax tree" with root node  $\langle x \rangle$  is constructed, the action specified by the semantics is performed. If this action includes an assignment of the form  $\langle x \rangle := c$ , where  $c$  is a security class, then the class  $c$  (derived from nodes of subtrees of  $\langle x \rangle$ ) is associated with the root node  $\langle x \rangle$ . If the action includes a security test and that test fails, the security error is reported and the program is not certified. The reader unfamiliar with such specifications may wish to consult Gries [Gi71: Section 12.2] for similar semantic routines that transform infix expressions to quadruples. The implementation of the parser and certification routines is not important (though clearly, some parsing techniques are more suitable than others).

To illustrate the certification process, we consider a bottom-up parse of the assignment statement " $d := a * b + c$ ". As the expression on the right side is recognized, the classes of the operands  $a$ ,  $b$ , and  $c$  are combined using  $\oplus$ . This proceeds as follows: when " $a * b$ " is reduced to a  $\langle \text{term} \rangle$ ,  $a \oplus b$  is assigned to  $\langle \text{term} \rangle$ ; then when " $\langle \text{term} \rangle + c$ " is reduced to an  $\langle \text{exp} \rangle$ ,  $\langle \text{term} \rangle \oplus c$  is assigned to  $\langle \text{exp} \rangle$ . The security of the assignment statement is easily verified when " $\langle \text{var} \rangle := \langle \text{exp} \rangle$ " is reduced to an  $\langle \text{assign stmt} \rangle$  by testing  $\langle \text{exp} \rangle \rightarrow \langle \text{var} \rangle$ . The details of the process are illustrated by Figure 3.1-2 which shows the syntax tree for

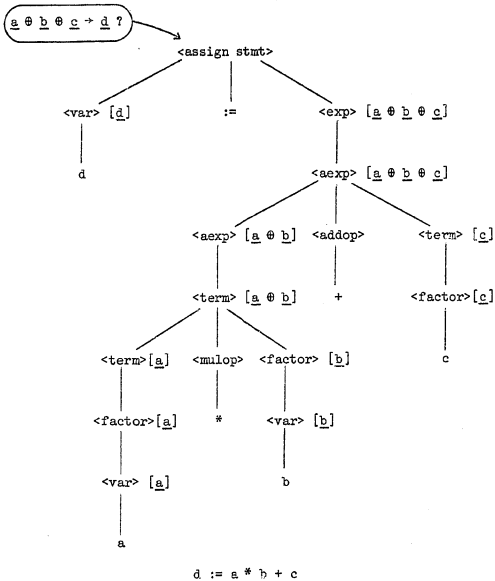


Figure 3.1-2. Certification Tree of an Assignment Statement.

the expression; the security classes (in brackets [ ]) have been associated with each node of the tree, and the nodes at which flows must be verified are indicated. Such a tree is defined to be a certification tree.

The certification of input and output statements is similar. In the case of an input statement, the variables constituting the input list are combined using @ as the reduction to <inlist> is performed. Figure 3.1-3 displays the certification tree for the statement "input a, b from c". In the case of an output statement, the expressions constituting the output list are combined using @ as the reduction to <outlist> is performed.

### 3.2 Conditional Statements

We now extend the syntax and the corresponding certification semantics described in the previous section to include three conditional statements:

if <exp> then <stmt><sub>1</sub> [else <stmt><sub>2</sub>]

while <exp> do <stmt><sub>1</sub>

repeat <stmt><sub>1</sub> until <exp>

where <stmt> is an elementary statement (assignment or I/O), conditional statement, or compound statement (begin-end statement). These statements all have the underlying abstract structure (see Section 2.4.4):

<exp>: <stmt><sub>1</sub> [, <stmt><sub>2</sub>].

Since this structure specifies an implicit flow from <exp> to all variables and files assigned values in <stmt><sub>1</sub> [and <stmt><sub>2</sub>], security requires (in addition to the independent security of the substatements) that all implicit flows from <exp> be secure. Letting V(<stmt>) denote

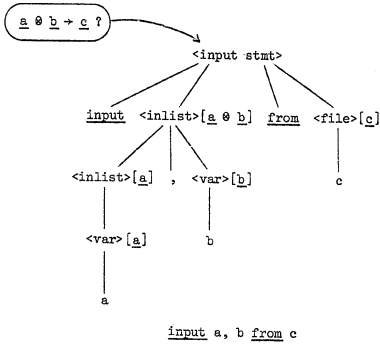


Figure 3.1-3. Certification Tree of an Input Statement.

the set of variables and files assigned values in a string of syntactic type  $\langle \text{stmt} \rangle$  (see Section 2.4.4), this means that for a conditional statement  $\langle \text{stmt} \rangle$  with structure

$$\langle \text{stmt} \rangle = \langle \text{exp} \rangle : \langle \text{stmt} \rangle_1 [, \langle \text{stmt} \rangle_2]$$

the following property must hold:

$$\langle \text{exp} \rangle \rightarrow \underline{b} \quad \forall b \in V(\langle \text{stmt} \rangle_1) [\cup V(\langle \text{stmt} \rangle_2)]$$

or equivalently:

$$\langle \text{exp} \rangle \rightarrow \underline{b} \quad \forall b \in V(\langle \text{stmt} \rangle).$$

For if-then-else and while-do statements, this property could be verified during the analysis of  $\langle \text{stmt} \rangle_1$  [and  $\langle \text{stmt} \rangle_2$ ] since  $\langle \text{exp} \rangle$  is recognized before these substatements are parsed. However, this technique is not suitable for the repeat-until statement since  $\langle \text{exp} \rangle$  is recognized after  $\langle \text{stmt} \rangle_1$  is analyzed. Therefore, it is necessary to verify the above property after the complete conditional statement is recognized.

To see how this can be done, observe that (by Proposition 2.3-5) the above property is equivalent to

$$\langle \text{exp} \rangle \rightarrow \otimes \{ \underline{b} \mid b \in V(\langle \text{stmt} \rangle) \}.$$

This suggests a set of actions that forms  $\otimes \{ \underline{b} \mid b \in V(\langle \text{stmt} \rangle) \}$  for each statement (non-conditional as well as conditional) as it is recognized. These actions, shown in Figure 3.2-1, supplement the actions specified in Figure 3.1-1. In Figure 3.2-1 we have used  $\langle \text{stmt} \rangle$  to denote the class  $\otimes \{ \underline{b} \mid b \in V(\langle \text{stmt} \rangle) \}$  since no other class was previously associated with  $\langle \text{stmt} \rangle$ .

The alert reader may recognize a syntactic ambiguity in the grammar for an  $\langle \text{if stmt} \rangle$ ; this does not affect our analysis and can be dealt with by the usual techniques of resolving ambiguities in programming language syntax [Gi71].

Syntax	Certification Semantics
$\langle \text{assign stmt} \rangle ::= \langle \text{var} \rangle := \langle \text{exp} \rangle$	$\langle \text{assign stmt} \rangle := \langle \text{var} \rangle$
$\langle \text{input stmt} \rangle ::= \langle \text{input} \rangle \langle \text{inlist} \rangle \text{ from } \langle \text{file} \rangle$	$\langle \text{input stmt} \rangle := \langle \text{inlist} \rangle$
$\langle \text{output stmt} \rangle ::= \text{output } \langle \text{outlist} \rangle \text{ to } \langle \text{file} \rangle$	$\langle \text{output stmt} \rangle := \langle \text{file} \rangle$
$\langle \text{if stmt} \rangle ::= \text{if } \langle \text{exp} \rangle \text{ then } \langle \text{stmt} \rangle$	$\langle \text{if stmt} \rangle := \langle \text{stmt} \rangle$
$\langle \text{if stmt} \rangle ::= \text{if } \langle \text{exp} \rangle \text{ then } \langle \text{stmt} \rangle \text{ else } \langle \text{stmt} \rangle$	if $\langle \text{exp} \rangle \neq \langle \text{if stmt} \rangle$ then SECURITY ERROR if $\langle \text{exp} \rangle \neq \langle \text{if stmt} \rangle$ then SECURITY ERROR
$\langle \text{while stmt} \rangle ::= \text{while } \langle \text{exp} \rangle \text{ do } \langle \text{stmt} \rangle$	$\langle \text{while stmt} \rangle := \langle \text{stmt} \rangle$ if $\langle \text{exp} \rangle \neq \langle \text{while stmt} \rangle$ then SECURITY ERROR
$\langle \text{repeat stmt} \rangle ::= \text{repeat } \langle \text{stmt} \rangle \text{ until } \langle \text{exp} \rangle$	$\langle \text{repeat stmt} \rangle := \langle \text{stmt} \rangle$ if $\langle \text{exp} \rangle \neq \langle \text{repeat stmt} \rangle$ then SECURITY ERROR
$\langle \text{stmt list} \rangle ::= \langle \text{stmt} \rangle$	$\langle \text{stmt list} \rangle := \langle \text{stmt} \rangle$
$\langle \text{stmt list} \rangle ::= \langle \text{stmt list} \rangle' ; \langle \text{stmt} \rangle$	$\langle \text{stmt list} \rangle := \langle \text{stmt list} \rangle' \emptyset \langle \text{stmt} \rangle$
$\langle \text{compound stmt} \rangle ::= \text{begin } \langle \text{stmt list} \rangle \text{ end}$	$\langle \text{compound stmt} \rangle := \langle \text{stmt list} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{assign stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{assign stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{input stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{input stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{output stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{output stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{if stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{if stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{while stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{while stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{repeat stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{repeat stmt} \rangle$
$\langle \text{stmt} \rangle ::= \langle \text{compound stmt} \rangle$	$\langle \text{stmt} \rangle := \langle \text{compound stmt} \rangle$

Figure 3.2-1. Syntax and Certification Semantics for Conditional Statements.



The certification process is illustrated in Figure 3.2-2 for the statements

```

a)  while a > 0 do
      begin
        b := b + 1;
        output b to c;
        a := a - 1
      end

b)  if a = 0
      then while b < c do b := b + 1
      else d := 1.

```

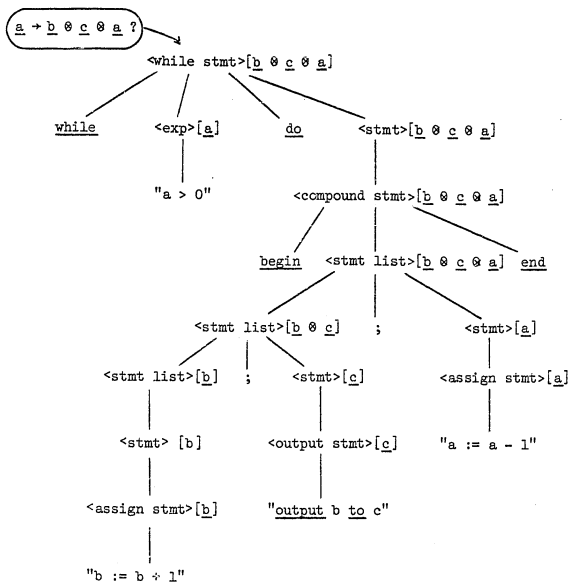
It should be clear that this certification process is equally applicable to all statements with an underlying abstract structure

$\langle \text{exp} \rangle: \langle \text{stmt} \rangle_1, \dots, \langle \text{stmt} \rangle_n \ (n \geq 1).$

This includes case statements, for statements, etc. (see also Section 2.4.4).

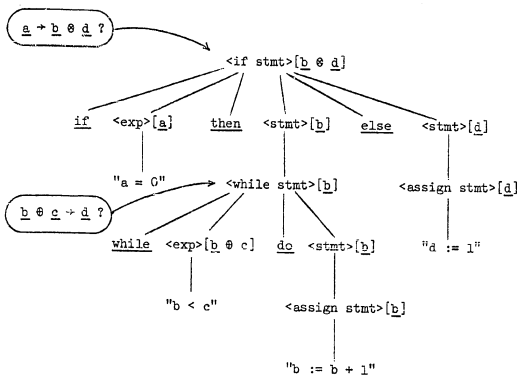
### 3.3 Gotos

If the language is extended to include a goto statement, the task of certifying implicit flows becomes considerably more complex. This is because goto statements give rise to arbitrary control flow structures that are not isolated by syntactic elements, making it more difficult to determine the set of statements immediately conditioned on an expression. To illustrate, consider the program in Figure 3.3-1. Because of the extensive use of gotos, it is practically impossible to tell which statements are conditioned on the test "b = 0" at the statement labelled "4". In general, with gotos, the set of statements conditioned on an expression may be scattered throughout an entire program. (It is interesting to note that these considerations forced Fenton to introduce a "call" instruction into his Minsky machine, so that the



a) while a > 0 do begin b := b + 1; output b to c; a := a - 1 end

Figure 3.2-2. Certification Trees of Conditional Statements.



b) if a = 0 then while b < c do b := b + 1 else d := 1

Figure 3.2-2, cont.

1: <u>input</u> a, b <u>from</u> f1;	}	$\phi_1; \xi_1 = "a = 0"$
c := 0;		
<u>if</u> a = 0 <u>then goto</u> 4;	}	$\phi_2; \xi_2 = \text{null}$
2: a := a + 1;		
3: <u>output</u> a <u>to</u> f2;	}	$\phi_3; \xi_3 = "a < b"$
<u>if</u> a < b <u>then goto</u> 2 <u>else goto</u> 6;		
4: <u>if</u> b = 0 <u>then goto</u> 3;	}	$\phi_4; \xi_4 = "b = 0"$
5: b := a;		
c := 1;	}	$\phi_5; \xi_5 = \text{null}$
6: <u>output</u> a, b <u>to</u> f2		
	}	$\phi_6; \xi_6 = \text{null}$

Program

Partitioning into Basic  
Blocks

Figure 3.3-1. Example of a Program with Gotos.

scope of a condition could be limited - see Section 2.4.4) For this reason, it is not possible to certify implicit flows until the entire program is analyzed and a control flow analysis performed. Only then can the underlying abstract conditional structures

$$\xi: \phi_1, \dots, \phi_m$$

be determined, and implicit flows from the expression  $\xi$  to the objects (i.e., variables and files) assigned values in the program substructures  $\phi_1, \dots, \phi_m$  be verified.

We now outline a general procedure for performing this control flow analysis and verifying the security of all implicit flows. This analysis is performed not because we believe the mechanism should support gotos, but rather to illustrate the general characteristics of all conditional structures and to justify the approach taken in the preceding section.

The first step is to partition the program into a set of substructures  $\phi_1, \dots, \phi_n$  each of which is a basic block. A basic block is a sequence of statements with exactly one entry point (the first statement in the sequence) and exactly one exit point (the last statement in the sequence). It is important to note that each basic block  $\phi_i$  contains at most one conditional statement (e.g., if-then-else statement), which must be the last statement of  $\phi_i$ . The conditional expression in this statement (if it exists) will be denoted  $\xi_i$ . Figure 3.3-1 shows the partitioning for the sample program. A thorough treatment of basic blocks and the control flow analysis to be described can be found in [A170, LM69, Ta74].

Let  $V(\phi)$  denote the set of variables assigned values in basic block  $\phi$ . Associate with each block  $\phi$  a class  $\underline{\phi}$ ,

$$\underline{\phi} = \begin{cases} \otimes \{ \underline{b} \mid b \in V(\phi) \} & \text{if } V(\phi) \neq \emptyset \\ H & \text{otherwise} \end{cases}$$


where  $H$  is the highest security class and therefore an identity element over  $\otimes$ . Note that the process of partitioning a program into basic blocks and of associating classes with the basic blocks can be performed during the analysis phase of compilation, using techniques similar to those described in the preceding section.

The next step is to construct a control flow graph of the program; that is, a directed graph whose nodes are the basic blocks and whose edges represent the control flow paths. There is an edge from  $\phi_j$  to  $\phi_k$  if control can flow directly from  $\phi_j$  to  $\phi_k$ . Figure 3.3-2 shows the control flow graph of the sample program of Figure 3.3-1. Here we have labelled each node  $\phi$  with its class  $\underline{\phi}$  (in brackets [ ]). For those nodes with more than one outgoing edge, we have also shown the conditional expression that determines which edge will be selected.

The final step is to use the control flow graph to determine the set of basic blocks (nodes) conditioned on each expression  $\xi_i$ . Let  $\Psi_i$  denote the set of basic blocks conditioned on  $\xi_i$ . All implicit flow from  $\xi_i$  is then certified for security by verifying that

$$\xi_i \rightarrow \otimes \{ \underline{\phi} \mid \phi \in \Psi_i \}.$$

For basic block  $\phi_i$  associated with expression  $\xi_i$ , the immediate forward dominator  $\text{IFD}(\phi_i)$ , is the closest block to  $\phi_i$  among the set of all blocks lying on every path from  $\phi_i$  to the exit block of the graph.



$$\text{IFD}(\phi_1) = \phi_6$$

$$\text{IFD}(\phi_3) = \phi_6$$

$$\text{IFD}(\phi_4) = \phi_6$$

$$\Psi_1 = \{\phi_2, \phi_3, \phi_4, \phi_5\}$$

$$\Psi_3 = \{\phi_2, \phi_3\}$$

$$\Psi_4 = \{\phi_2, \phi_3, \phi_5\}$$

Figure 3.3-2. Control Flow Graph of Program of Figure 3.3-1.

It follows that

$$\Psi_i = \{\phi \mid \phi \text{ lies on a path from } \phi_i \text{ to } \text{IFD}(\phi_i)\}.$$

This follows from the observation that all control flow paths originating with  $\phi_i$  join at  $\text{IFD}(\phi_i)$ ; hence, the only blocks directly conditioned on  $\xi_i$  are those that lie on some path from  $\phi_i$  to  $\text{IFD}(\phi_i)$ . (See Figure 3.3-2.) Of course, a block  $\phi$  executed after  $\text{IFD}(\phi_i)$  can be indirectly conditioned on  $\xi_i$  if an assignment is made to some variable  $b$  in  $\phi \in \Psi_i$ , and  $\phi$  is directly conditioned on the value of  $b$ . Because the flow relation is transitive, we do not have to worry about such indirect conditioning: the flows from  $\xi_i$  to  $b$  and from  $b$  to all variables assigned values in  $\phi$  will be verified.

It is now easy to justify the mechanism of the preceding section for certifying the conditional statements

if <exp> then <stmt><sub>1</sub> [else <stmt><sub>2</sub>],

while <exp> do <stmt><sub>1</sub>, and

repeat <stmt><sub>1</sub> until <exp>

since in all cases, the immediate forward dominator (of  $\phi_{\langle \text{exp} \rangle}$ ) is precisely the next statement, and therefore  $\Psi_{\langle \text{exp} \rangle}$  is precisely the conditional statement itself. Hence, it is possible to determine the underlying abstract conditional structures of these statements without performing a global flow analysis.

### 3.4 Interrupts

Additional mechanisms are needed to guarantee security in the presence of interrupts. This is because an interrupt may cause the execution of subsequent statements to be conditioned on the data that caused the interrupt, giving rise dynamically to implicit flows that



go undetected by our present mechanism. Consider, for example, the execution of the following program in an environment where an overflow interrupt can result in program termination:

```

sum := 0;
i := 0;
repeat
  begin
    sum := sum + x;
    i := i + 1;
    output i to b
  end
until false

```

Assume  $\underline{i} = \underline{b} = L$  (the lowest security class) and  $\underline{x} = \underline{\text{sum}} = H$  (the highest security class). Our present mechanism would certify this program, since the iteration appears to be conditioned on the Boolean constant false which is associated with  $L$ , and  $L \rightarrow \underline{i} \oplus \underline{\text{sum}} \oplus \underline{b}$ . Should the variable `sum` overflow, however, the iteration would be terminated by the overflow interrupt, whereupon the iteration will have been conditioned on `sum` and an invalid flow from `sum` to `b` will have occurred (since  $\underline{\text{sum}} = H \neq \underline{b} = L$ ). Indeed, in this case file `b` contains information about `x`: the value of `x` can be approximated by computing  $\text{MAXVALUE}/\text{LASTb}$ , where  $\text{MAXVALUE}$  is a constant representing the largest value a variable could assume before overflowing, and  $\text{LASTb}$  is the last value of `i` written into file `b`. However, this violates security since  $\underline{x} \neq \underline{b}$ . Our present certification mechanism would have detected this had the programmer made explicit his intentions:

```

repeat
  begin
    sum := sum + x;
    i := i + 1;
    output i to b
  end
until sum overflows

```

There appear to be four general approaches to solving this problem. The first two are too drastic to be viable; the third reduces the problem to the single-class case of Section 2.4.1; and the fourth, which inhibits all interrupts except those for which the programmer has specified explicit actions, is the most attractive in the context of the mechanism under consideration in this chapter. Each will be considered below.

The first possible solution is a run-time mechanism that receives control when a program interrupts. It would abort the program, erase its entire memory (including all files it had written), and "sound an alarm". This solution is unacceptable for at least two reasons: first, there is no good way of determining at run-time whether or not an interrupt caused a security violation; and second, a program should be able to process, without "sounding an alarm", any interrupt that does not cause a violation.

The second possible solution is based on the principle of restructuring a program so that the certification mechanism can verify the correctness of interrupts. The idea is to have the compiler insert conditional expressions, corresponding to conditions that might cause interrupts, at the appropriate points in the program. These conditions are then used to verify the security of all interrupts. For example, the preceding program could be modified to:

```

sum := 0;
i := 0;
repeat
  begin
    if sum + x < MAXVALUE then
      begin
        sum := sum + x;
        if i + 1 < MAXVALUE then
          begin
            i := i + 1;
            output i to b
          end
        end
      else halt
    end
  end
until false

```

The compiler would now detect the security violation in this program since the dependency of  $i$  and  $b$  on  $x$  is explicit. This solution is unattractive because it seriously complicates the compilation process, possibly to the point of rendering the entire certification mechanism too cumbersome to be useful.

The third possible solution is based on the principle of imposing so many restrictions on program structure that interrupts cannot cause the problem discussed above. One method of doing this is to require that the program be specified as a collection of modules, in one-to-one correspondence with the security classes. If  $\phi_A$  is the program module corresponding to class A, the following constraints must be observed:

- a)  $\phi_A$  may assign values only to objects in class A;
- b) If  $x$  is an input parameter to  $\phi_A$ , then  $x \rightarrow A$ ;
- c) If  $y$  is an output parameter from  $\phi_A$ , then  $A \rightarrow y$ ;
- d) If  $\phi_A$  transfers control to  $\phi_B$ , then  $A \rightarrow B$ .

Under these restrictions, an interrupt occurring in  $\phi_A$  can be conditioned only on those objects  $z$  for which  $z \rightarrow A$ , so that the problem is

eliminated. It should be noted that these restrictions make the problem isomorphic to the single-class case considered in Section 2.4.1, where files are used to transmit parameters among processes. In this case there is no need for compiler certification of the modules, so we shall not consider it further.

The fourth solution is based on the simple observation that no security violation can occur in an environment that inhibits interrupts, since the only way a statement can in this case be conditioned on the state of a variable is by explicitly testing its value. However, a completely interrupt-free environment is impractical, as a program should be allowed to process certain interrupts such as end-file or overflow, should it desire to do so. We therefore propose to 1) extend the language to include a construct for specifying interrupt conditions and procedures for handling them ("interrupt handlers"), 2) modify the certification process to certify programs in the extended language, and 3) include a run-time mechanism that invokes (as a procedure) the interrupt handler of a specified interrupt but inhibits an unspecified interrupt.

The language construct we add is similar to the PL/I ON statement [IB71], and has the following structure:

on <condition> <ident> do <stmt> ,

where <condition> names an interrupt condition (e.g., overflow, zero-divide, or end-file), <ident> is the variable or file name to which it applies, and <stmt> is the statement to be called when the interrupt occurs (after execution of the statement, control returns to the place where the interrupt occurred). Examples of on statements are:

on enafile b do moredata := false

on overflow x do begin x := MAXVALUE; flag := true end

The scope of an on statement is not important; i.e., it could be a single statement, a block, or the entire program.

It should be clear from our preceding analysis of conditional structures (Sections 3.2 and 3.3) that the on statement is just another example of a conditional statement; hence the security requirements are simply

<var> → <stmt>,

which leads to the certification semantics

<on stmt> := <stmt>  
if <var> ≠ <on stmt> then SECURITY ERROR.

To illustrate the mechanism, consider again the program that divulged the value of variable x by causing the variable sum to overflow. In the present context, failure to specify an on statement for the condition "sum overflows" will not enable this interrupt, whereupon the program will go on executing forever (or until the program exhausts its ration of processor time). In this case, the last value of i written into file b contains no information whatsoever about x. On the other hand, if the overflow condition is specified explicitly:

```

on overflow sum do flag := false;
flag := true;
sum := 0;
i := 0;
repeat
  begin
    sum := sum + x;
    i := i + 1;
    output i to b
  end
until ¬ flag

```

the security violation will be detected when the compiler verifies the conditions  $\text{sum} \rightarrow \text{flag}$  and  $\text{flag} \rightarrow \text{sum} \oplus i \oplus b$ .

This last solution is the most attractive, since it permits a program to process selected interrupts without requiring substantial overhead or imposing undue restrictions on program structure.

### 3.5 Complex Data Structures

#### 3.5.1 Arrays

Let  $a$  be an array having  $n$  dimensions (subscripts). The elements of  $a$  must all belong to the same security class, since there is no way of determining at compile-time which element (or even which row or column) is referenced by an expression like  $a[i_1, \dots, i_n]$ . Therefore, an array is bound to a single class.

The certification semantics for array references, shown in Figure 3.5-1, are based on the principle that an array reference

$\langle \text{array} \rangle [ \langle \text{sublist} \rangle ]$

is secure only if

$\langle \text{sublist} \rangle \rightarrow \langle \text{array} \rangle$  ,

where

$\langle \text{sublist} \rangle = \langle \text{exp} \rangle_1 \oplus \dots \oplus \langle \text{exp} \rangle_n$

and  $\langle \text{exp} \rangle_1, \dots, \langle \text{exp} \rangle_n$  are the subscripts constituting  $\langle \text{sublist} \rangle$ .

The semantics for the rule " $\langle \text{array ref} \rangle ::= \langle \text{array} \rangle [ \langle \text{sublist} \rangle ]$ " include the generation of code to check that  $\langle \text{sublist} \rangle$  specifies subscripts within the range of  $\langle \text{array} \rangle$ . Without this check, it would be possible for a statement like " $a[i] := b$ " to cause an invalid from variable  $b$  to a variable  $c$  if  $i$  is out of range (i.e.,  $a[i]$  references

<p>Syntax</p> <p><u>&lt;sublist&gt; ::= &lt;exp&gt;</u></p> <p><u>&lt;sublist&gt; ::= &lt;sublist&gt; , &lt;exp&gt;</u></p> <p><u>&lt;array ref&gt; ::= &lt;array&gt; [ &lt;sublist&gt; ]</u></p> <p><u>&lt;var&gt; ::= &lt;simple var&gt;</u></p> <p><u>&lt;var&gt; ::= &lt;array ref&gt;</u></p>	<p>Certification Semantics</p> <p><u>&lt;sublist&gt; := &lt;exp&gt;</u></p> <p><u>&lt;sublist&gt; := &lt;sublist&gt; , # &lt;exp&gt;</u></p> <p>if &lt;sublist&gt; ≠ &lt;array&gt; then SECURITY ERROR</p> <p><u>&lt;array ref&gt; := &lt;array&gt;</u></p> <p>Generate code to check range of &lt;sublist&gt;</p> <p><u>&lt;var&gt; := &lt;simple var&gt;</u></p> <p><u>&lt;var&gt; := &lt;array ref&gt;</u></p>
--	---

Figure 3.5-1. Syntax and Certification Semantics for Array References.

variable  $c$  rather than an element of  $a$ ). In case the programmer has specified actions for handling invalid subscripts, the methods of Section 3.4 can be used to verify the security of the corresponding interrupt. In case the programmer has not specified such actions, the compiler can generate code that causes the invalid reference to refer to the first element of the array, for example, so that no invalid flow can occur.

Note that we have also defined (Figure 3.5-1) the syntactic type  $\langle \text{var} \rangle$  to be a simple variable or array reference now. Whereas  $\langle \text{array ref} \rangle$  is determined when the structure is recognized,  $\langle \text{simple var} \rangle$  is obtained directly from the symbol table.

It should be noted that, in principle, there is a slight difference between array references appearing in an expression and those appearing on the left side of an assignment statement. In the former case, it is necessary only that

$$\langle \text{array ref} \rangle := \langle \text{array} \rangle \text{ @ } \langle \text{sublist} \rangle$$

but not that

$$\langle \text{sublist} \rangle \rightarrow \langle \text{array} \rangle$$

(which is a stronger requirement). We did not consider this distinction important enough to present the more complicated semantics required to identify whether the array reference is the left part of an assignment or not. However, the semantics of Figure 3.5-1 can be extended to this case if need be.

### 3.5.2 Records

Define a record  $r$  to be a structure consisting of  $n$  elements or fields,  $x_1, \dots, x_n$ , the element  $x_i$  being identified by the compound name



$r.x_i$ . Unlike an array, the elements of a record can belong to different security classes, since they are always referenced by unique names. In fact, record elements can be treated exactly as other variables of syntactic type  $\langle \text{simple var} \rangle$ . However, doing so requires some care in the specification of the certification semantics for records (see Figure 3.5-2).

The statement "input  $r$  from  $a$ " where  $a$  is a file evidently requires  $\underline{a} \rightarrow \underline{r.x_i}$  for  $1 \leq i \leq n$ , or equivalently that

$$\underline{a} \rightarrow \underline{r.x_1} \oplus \dots \oplus \underline{r.x_n} \equiv \bigotimes r.$$

Similarly, the statement "output  $r$  to  $b$ " requires that

$$\bigoplus r \equiv \underline{r.x_1} \oplus \dots \oplus \underline{r.x_n} \rightarrow \underline{b}.$$

An assignment " $r := s$ ", which is valid syntactically only if  $r$  and  $s$  have identical structure, is secure only if

$$\underline{s.x_i} \rightarrow \underline{r.x_i} \quad (1 \leq i \leq n).$$

Note that the apparently simpler check  $\bigoplus \underline{s} \rightarrow \bigotimes \underline{r}$  does not work in this case since, for example, it will reject assignments " $r := s$ " in which  $\underline{r.x_i} = \underline{s.x_i}$  for all  $i$  but  $\underline{s.x_i} \neq \underline{s.x_j}$  for some  $i \neq j$ .

### 3.5.3 List Structures

Define a list structure  $\ell$  to be a set of one or more records of identical structure. A record in  $\ell$  is referenced by a pointer variable  $p$ , associated with  $\ell$  in the declaration. We denote by  $\ell p$  the record of  $\ell$  pointed to by  $p$ , and by  $\ell.xp$  the element  $x$  in the record of  $\ell$  pointed to by  $p$ .

The security requirements are a combination of those for arrays and records. As for arrays, a pointer  $p$  can select an arbitrary member from the aggregate; therefore, it is necessary that  $\ell.xp$  be a single class,

Syntax

$\langle \text{rec input stmt} \rangle ::= \text{input } \langle \text{record} \rangle \text{ from } \langle \text{file} \rangle$

$\langle \text{rec output stmt} \rangle ::= \text{output } \langle \text{record} \rangle \text{ to } \langle \text{file} \rangle$

$\langle \text{rec assign stmt} \rangle ::= \langle \text{record} \rangle := \langle \text{record} \rangle$

Certification Semantics

$\langle \text{rec input stmt} \rangle := \emptyset \langle \text{record} \rangle$

if  $\langle \text{file} \rangle \neq \langle \text{rec input stmt} \rangle$  then SECURITY ERROR

$\langle \text{rec output stmt} \rangle := \langle \text{file} \rangle$

if  $\emptyset \langle \text{record} \rangle \neq \langle \text{rec output stmt} \rangle$  then SECURITY ERROR

if  $\langle \text{record} \rangle$  and  $\langle \text{record} \rangle$  have corresponding elements  $x_1, \dots, x_n$ ,

then  $\left\{ \begin{array}{l} \text{if } \exists i (1 \leq i \leq n): \langle \text{record} \rangle.x_i \neq \langle \text{record} \rangle.x_i \\ \text{then SECURITY ERROR} \\ \langle \text{rec assign stmt} \rangle := \emptyset \langle \text{record} \rangle \end{array} \right.$

else TYPE ERROR

Figure 3.5-2. Syntax and Certification Semantics for Records.

irrespective of the value of  $p$ , for any given field  $x$ . Also, the security of the reference  $\underline{l}.x \uparrow p$  requires that  $\underline{p} \rightarrow \underline{l}.x$ , and the security of  $\underline{l} \uparrow p$  requires  $\underline{p} \rightarrow \otimes \underline{l}$ . Since  $\underline{p} \rightarrow \otimes \underline{l}$  implies  $\underline{p} \rightarrow \underline{l}.x$  for each field  $x$ , verifying that  $\underline{p} \rightarrow \otimes \underline{l}$  is sufficient for all references to elements of list  $\underline{l}$ . In addition to these security requirements for record references, the security requirements for flows to and from the records of  $\underline{l}$  (see Figure 3.5-2) apply here also. Because of these similarities we have not shown the certification semantics for list structures in a separate figure. As with arrays, it is necessary to generate code that checks whether the pointer  $p$  always designates an element of the list  $\underline{l}$ , after each assignment to  $p$ .

### 3.6 Procedure Calls

Let  $f$  be a procedure with formal input parameters  $x_1, \dots, x_m$  and formal output parameters  $y_1, \dots, y_n$ . Consider an elementary statement of the form

call  $f(a_1, \dots, a_m; b_1, \dots, b_n)$ ,

where  $a_1, \dots, a_m$  are the actual input parameters of  $f$ , and  $b_1, \dots, b_n$  are the actual output parameters of  $f$ . It is easy to see that the security requirements of this statement are simply:

- a)  $f$  must be secure,
- b)  $\underline{a_i} \rightarrow \underline{x_i}$  ( $1 \leq i \leq m$ ), and
- c)  $\underline{y_j} \rightarrow \underline{b_j}$  ( $1 \leq j \leq n$ ).

As usual, if the call statement is conditioned on an expression  $\xi$ , it is necessary to verify the implicit flow from  $\xi$  to the output parameters:

- d)  $\underline{\xi} \rightarrow \underline{b_1} \otimes \dots \otimes \underline{b_n}$ .

For the moment we assume  $f$  does not write into any objects that are global or have a "lifetime" exceeding the activation of  $f$  (e.g., Algol own variables [Na63]), so that there are no other implicit flows from an expression  $\xi$  as a result of the call to procedure  $f$ . Later we show how to remove this restriction. Now, condition (d) can be verified by the same technique used to verify any implicit flow, namely by computing  $\underline{b}_1 \otimes \dots \otimes \underline{b}_n$  as the output parameters are recognized and associating this with the representative statement type,  $\langle \text{proc call stmt} \rangle$  (see Figure 3.6-1). However, conditions (a) - (c) cannot be verified until after procedure  $f$  is compiled and its security established. Hence, it is necessary to save the classes of the parameters (see Figure 3.6-1) and check these conditions at the time  $f$  is linked to the calling program. To handle the case where  $f$  is compiled independently, it is therefore necessary to extend the certification mechanism to include a mechanism in the linker for certifying explicit flows not "bound" at compile-time. This implies that the certification of a program is not really complete until all external procedure calls are verified by the linker.

There is a serious limitation to the above approach when the procedure  $f(x_1, \dots, x_m; y_1, \dots, y_n)$  is supposed to handle arbitrary classes of information (as is typical of library routines). Then  $x_1, \dots, x_m$  must be assigned to the highest security class  $H$  so that  $\underline{a} \rightarrow \underline{x}_i$  ( $1 \leq i \leq m$ ) is guaranteed to hold for all calls. But this implies that  $y_1, \dots, y_n$  must be assigned to the highest security class  $H$ , since their values are derived from  $x_1, \dots, x_m$ . This in turn requires that the caller of  $f$  must also assign  $b_1, \dots, b_n$  to  $H$  (even if  $a_1, \dots, a_m$  are in

<p>Syntax</p> <p>&lt;input parms&gt; ::= &lt;exp&gt;   &lt;input parms&gt; , &lt;exp&gt;</p> <p>&lt;output parms&gt; ::= &lt;var&gt;</p> <p>&lt;output parms&gt; ::= &lt;output parms&gt; ' , &lt;var&gt;</p> <p>&lt;proc call stmt&gt; ::= &lt;proc name&gt; ( &lt;input parms&gt; ; &lt;output parms&gt; )</p>	<p>Certification Semantics</p> <p>save <u>exp</u> for later verification</p> <p>&lt;output parms&gt; ::= &lt;var&gt; save &lt;var&gt; for later verification</p> <p>&lt;output parms&gt; ::= &lt;output parms&gt; ' @ &lt;var&gt; save &lt;var&gt; for later verification</p> <p>&lt;proc call stmt&gt; ::= &lt;output parms&gt;</p>
--	--

Figure 3.6-1. Syntax and Certification Semantics for Procedure Calls.

<p>Syntax</p> <p>&lt;input parms&gt; ::= &lt;exp&gt;</p> <p>&lt;input parms&gt; ::= &lt;input parms&gt; ' , &lt;exp&gt;</p> <p>&lt;output parms&gt; ::= &lt;var&gt;</p> <p>&lt;output parms&gt; ::= &lt;output parms&gt; ' , &lt;var&gt;</p> <p>&lt;proc call stmt&gt; ::= &lt;proc name&gt; ( &lt;input parms&gt; ; &lt;output parms&gt; )</p> <p>&lt;fun call&gt; ::= &lt;proc name&gt; ( &lt;input parms&gt; )</p> <p>&lt;factor&gt; ::= &lt;fun call&gt;</p>	<p>Certification Semantics</p> <p>&lt;input parms&gt; ::= &lt;exp&gt;</p> <p>&lt;input parms&gt; ::= &lt;input parms&gt; ' @ &lt;exp&gt;</p> <p>&lt;output parms&gt; ::= &lt;var&gt;</p> <p>&lt;output parms&gt; ::= &lt;output parms&gt; ' @ &lt;var&gt;</p> <p>if &lt;input parms&gt; ≠ &lt;output parms&gt; then SECURITY ERROR</p> <p>&lt;proc call stmt&gt; ::= &lt;output parms&gt;</p> <p>&lt;fun call&gt; ::= &lt;input parms&gt;</p> <p>&lt;factor&gt; ::= &lt;fun call&gt;</p>
--	--

Figure 3.6-2. Syntax and Certification Semantics for Restricted Procedure Calls.

the lowest class  $L_l$ ) so that  $y_j \rightarrow b_j$  ( $1 \leq j \leq n$ ) also holds. By now the limitation should be obvious — it is impossible to obtain low security results from unrestricted procedures designed to handle data in arbitrary security classes.

One solution to this problem defines different versions of the procedure  $f$  for different security classes; the appropriate version is then selected at the time  $f$  is linked to the calling program. This is analogous to the GENERIC procedure in PL/I [IB71] which operates on different data types (i.e., integer, real, etc.).

Another solution restricts the procedure  $f$  so that the output parameters are derived entirely from the input parameters and data belonging to the lowest security class  $L$  (e.g., constants). If  $f$  is known to satisfy this property, then security can be established completely at compile-time by verifying

$$a_1 \oplus \dots \oplus a_m \rightarrow b_1 \oplus \dots \oplus b_n$$

(see Figure 3.6-2). An interesting special case of such procedures is the "function" type procedure (SQRT, SIN, LOG, etc.). Here a procedure  $f$  is called during expression evaluation (by a call  $f(a_1, \dots, a_m)$ ) and returns with a single result, derived entirely from the input parameters  $a_1, \dots, a_m$  and constants. Since there are no output parameters, there are no side effects and no implicit flows. Figure 3.6-2 also shows syntax and certification semantics for such function calls.

Consider now the case where the procedure  $f$  writes into objects  $z_1, \dots, z_k$  that are global or have a "lifetime" exceeding the activation of  $f$ . For example,  $f$  could be a "monitor" procedure for controlling access to a shared message buffer [Ho74]. To guarantee the secure

execution of  $f$ , it is now necessary to also verify implicit flows to  $z_1, \dots, z_k$  that result when a call to  $f$  is conditioned on one or more expressions  $\xi_1, \dots, \xi_r$ . This can be done at the time  $f$  is linked to the calling procedure as follows. Associate with the procedure  $f$  an additional formal parameter  $\underline{z}$ , where  $\underline{z} = \underline{z}_1 \oplus \dots \oplus \underline{z}_k$ , and with the call to  $f$  an additional actual parameter  $\underline{\xi}$ , where  $\underline{\xi} = \xi_1 \oplus \dots \oplus \xi_r$ . The linking mechanism then verifies that  $\underline{\xi} \rightarrow \underline{z}$ .

### 3.7 Applications

The ability to certify programs and procedures for security using the methods of this chapter resolves several open questions about information security. It implies that it is possible to prevent leakage on many of the so-called "covert channels". It implies that the selective confinement problem has a solution. It implies that it is possible to design processes that can process confidential files, but cannot output information giving correlations between these files. It implies that data base query programs having security clearances can be properly designed. What is most important, all of the above can be handled with a simple, automatic certification procedure. These points are developed in the following subsections.

#### 3.7.1 Covert Channels

The mechanism we have proposed will prevent leakage on any "stored information channel". By this we mean that if a process  $q$  is not permitted access to information in some class  $A$ , then  $q$  will be unable to obtain information in  $A$  by examining the contents of some storage object. This implies that it is possible to prevent leakage on many

"covert channels". For example, a process  $p$  would not be able to leak a confidential value  $x$  by opening  $x$  files  $a_1, \dots, a_x$  under the observation of a process  $q$  not authorized to see  $x$  (i.e.,  $x \not\rightarrow q$ ). To see why this is so, let  $b_i$  denote the "state variable" associated with file  $a_i$ ; that is,  $b_i$  indicates whether or not  $a_i$  has been opened for reading or writing, etc.. We assume that  $a_i = b_i$ . If the certification mechanism certifies  $p$ , it must be true that  $x \rightarrow b_i$  ( $1 \leq i \leq x$ ), since the value of each  $b_i$  depends on the value of  $x$ . But then  $q$  will not be able to access these variables as a means for determining the value of  $x$  unless  $b_i \rightarrow q$  and, therefore,  $x \rightarrow q$ .

Of course this does not prevent  $p$  from leaking the value of  $x$  to  $q$  if  $p$  can alter the performance of the system in a way that is observable to  $q$  but not to a "security monitor". This is still an open problem.

### 3.7.2 Confinement

Lampson defines the confinement problem to be that of constraining a "service process" so that it cannot leak confidential information about a "customer process" [La73]. He outlines a solution to the problem that constrains the service process from retaining any information, confidential or not, after it ceases to operate on behalf of the customer process. This is also true of the solutions proposed in [An74, Jo73]. For this reason, the confined service process is referred to as "memoryless".

A more flexible solution to the problem would be to permit a confined service process to save non-confidential information in global variables and files. For example, a confined income tax computing



service could then save address and billing information on its use by customers, though not information on its customers' incomes. Schroeder has suggested that a solution to this problem may be possible with a mechanism for certifying the process to be confined [Sc72]. We now show how this can be done with the mechanism proposed in this chapter. (Fenton also shows how this problem can be solved with his run-time mechanism [Fe74a]).

Let  $q(XN, XC; YN, YC)$  be a service process, where  $XN$  and  $XC$  are sets of input parameters for receiving non-confidential and confidential data respectively and  $YN$  and  $YC$  are sets of output parameters for returning non-confidential and confidential results respectively ( $XN, XC, YN, YC$  may contain single data elements, complex data structures, or files). Also let the information class of non-confidential data be  $A$  and that of confidential data be  $B$  (i.e.,  $x \in XN \Rightarrow \underline{x} = A$  and  $x \in XC \Rightarrow \underline{x} = B$ ), where  $A \rightarrow B$  but not  $B \rightarrow A$ . Then  $q$  is guaranteed to not leak confidential information about a customer  $p$  if it is secure and does not write into any global variables or files in a class  $C$  such that  $B \rightarrow C$ . (We will assume its local memory is cleared and deleted when it terminates.) The methods of this chapter show that a compiler can certify  $q$  not only for security but also for confinement. Note that  $q$  may invoke other procedures, as long as they are also certified for confinement, and that the linkage to such procedures can be guaranteed secure.

The mechanism is also applicable to a more general case in which there may be more than two class of confidential information and the problem is to constrain a service process from leaking data belonging

to some of these classes. Here we let  $q(X_1, \dots, X_n; Y_1, \dots, Y_n)$  be a service process, where  $X_1, \dots, X_n$  are sets of input parameters for receiving data in classes  $A_1, \dots, A_n$  respectively and  $Y_1, \dots, Y_n$  are sets of output parameters for returning results in classes  $A_1, \dots, A_n$  respectively. Suppose that  $q$  is not permitted to leak data belonging to some set  $Z$  of classes ( $Z$  does not necessarily have to be a subset of  $\{A_1, \dots, A_n\}$ ). Then  $q$  may be certified for confinement if it is established that  $q$  is secure and does not write into any global variables or files in a class  $B$  such that  $\otimes Z \rightarrow B$ .

A problem where this more general confinement mechanism is useful is that of constraining a service process from leaking any information that links information in two classes  $A_1$  and  $A_2$  together (i.e., information derived from data in both  $A_1$  and  $A_2$ ) while permitting it to retain information in classes  $A_1$  and  $A_2$  separately. This can be achieved by letting  $Z = \{A_1 \otimes A_2\}$ , so that the confined process cannot write into any class  $B$  such that  $A_1 \otimes A_2 \rightarrow B$ . For example,  $A_1$  may contain a list of names and  $A_2$  a list of salaries. A confined service process, such as the above, would be able to process data in both  $A_1$  and  $A_2$ , but it would not be allowed to output information that associated names with salaries. The utility of this, in solving the usually vexing problems of preventing outputs resulting from cross-correlations of files in data banks, is obvious.

### 3.7.3 Data Bases and Data Banks

Suppose a system (or network of systems) has a large data base containing different classes of information about individuals. For

example, one class may contain employment records, a second health records, a third credit reports, a fourth criminal records, and so forth forth. We assume that all access to the data base is restricted to a set of procedures  $P$ , so that if a process  $q$  requires information from the data base, it must issue a query that in turn invokes a procedure  $p \in P$ . Now, suppose that associated with  $q$  is a clearance level,  $q$ , meaning that  $q$  is not permitted access to information in class  $A$  unless  $A \rightarrow q$  (see Section 2.4.1). To guarantee the security of the system, it is therefore necessary to prove that no response  $r$  returned to  $q$  is derived from information in classes  $A_1, \dots, A_n$  unless  $A_1 \oplus \dots \oplus A_n \rightarrow q$ . Clearly, this property is guaranteed to hold as long as each  $p \in P$  is certified secure. The significance of this result must not be overlooked — it says that we can guarantee the security of the system.

We do not claim our mechanism is sufficient to handle all of the security requirements of a data base system. For example, it does not appear to be applicable to certain "data dependent" access requirements where access to a datum (e.g., the SALARY field of an employee's record) may depend on the value of the datum (e.g., SALARY < 10000). See [CH75, CM72, Hf71, Ts74] for a general discussion of the security requirements of data base systems. Future research is needed to determine the applicability of our mechanism to these systems.

## PROGRAM CERTIFICATION UNDER DYNAMIC BINDING

## Chapter 4

In Section 2.4.5 we examined the security requirements of a multi-class system with dynamic binding of objects to security classes. We observed that this type of system requires both a run-time mechanism that dynamically updates security classes and a compile-time mechanism that inserts special code into a program (so that all implicit flows are accounted for) and certifies the secure execution of the transformed program. The purpose of this chapter is to describe these mechanisms.

As in the previous chapter, we restrict attention to highly structured programs (even for which, as will be seen, certification is not easy). Define a program (or statement)  $\phi$  recursively by:

1.  $\phi$  is an assignment statement " $b := f(a_1, \dots, a_n)$ ", where  $f(a_1, \dots, a_n)$  denotes an expression with operands  $a_1, \dots, a_n$ .
2.  $\phi$  is a compound statement (or concatenation) "begin  $\phi_1; \dots; \phi_m$  end", where  $\phi_1, \dots, \phi_m$  are programs.
3.  $\phi$  is a selection statement "if  $e$  then  $\phi_1$  [else  $\phi_2$ ]", where  $e$  is a Boolean variable and  $\phi_1$  and  $\phi_2$  are programs.
4.  $\phi$  is an iteration statement "while  $e$  do  $\phi_1$ ", where  $e$  is a Boolean variable and  $\phi_1$  is a program.

No generality is lost by restricting the " $e$ " in selection and iteration statements to a Boolean variable, since "if  $e$  then  $\phi_1$  else  $\phi_2$ "

(where  $\xi$  is an expression) is equivalent to "begin  $e := \xi$ ; if  $e$  then  $\phi_1$  else  $\phi_2$  end" and "while  $\xi$  do  $\phi_1$ " is equivalent to "begin  $e := \xi$ ; while  $e$  do begin  $\phi_1$ ;  $e := \xi$  end end". Likewise, no generality is lost by excluding a case statement, for statement, or repeat-until statement: a case statement can be transformed into a sequence of nested if-then-else statements; for statements and repeat-until statements can be transformed into while-do statements. (For example, the statement "repeat  $\phi_1$  until  $e$ " is equivalent to "begin  $\phi_1$ ; while  $e$  do  $\phi_1$  end".)

A given statement  $\phi$  of a program is said to be at nesting level  $k$  ( $k \geq 0$ ) if it is textually enclosed by  $k$  selection or iteration statements. Note that if a program  $\phi$  is at nesting level  $k$ , then execution of  $\phi$  is conditioned on the values of exactly  $k$  Boolean variables, and all implicit flows to variables assigned values in  $\phi$  occur from these  $k$  Boolean variables.

The hardware support is described in Section 4.1. It consists of tags for dynamically binding objects to security classes and a stack for storing the classes of Boolean variables. Each register of the machine contains a tag field that holds the class of the information stored in it. Immediately prior to execution of a statement  $\phi$  at nesting level  $k$ , the stack contains the classes of the  $k$  Boolean variables on which the execution of  $\phi$  is conditioned. A special hardware mechanism updates the tag field of a register receiving a result, using the tag field of the input operands (explicit flows) and the classes on the stack (implicit flows) to derive the new tag setting.

The software support is described in Section 4.2. It consists of a compile-time procedure (called "Algorithm T1") that transforms a given

program into a secure sequence of instructions for the hardware. This is done by inserting update instructions into the program. These instructions cause the tag field of a register receiving an implicit flow from a Boolean variable to be updated even if no explicit assignment is made to the register. Algorithm T1 is easily incorporated into the analysis phase of a compiler and does not substantially increase the compilation time of a program. However, it may substantially increase its execution time (because of the update instructions).

The run-time inefficiencies wrought by Algorithm T1 lead us in Section 4.3 to investigate a more sophisticated transformation procedure ("Algorithm T2") that analyzes the flow of information through a program and generates update instructions only where security cannot otherwise be guaranteed. This works as follows. Beginning with assumptions about the initial classes of the input parameters of the program, Algorithm T2 simulates information flow as it will occur during program execution until it is able to determine upper and lower bounds on the class of each variable after each statement is executed. It then inserts an update instruction into any statement which can cause an implicit flow from a Boolean variable  $e$  to a variable  $b$  and the bounds on  $e$  and  $b$  are such that  $e \rightarrow b$  is not otherwise guaranteed to hold. In this case, run-time support is necessary to verify the correctness of the initial classes of the input variables when the program is entered. Algorithm T2 has the advantage of adding a smaller factor to the time and space requirements of a program (indeed, it may not be necessary to generate any update instructions at all) but the disadvantage of being considerably more complex.

In Section 4.4 we consider possible extensions to the mechanism for supporting files, arrays, procedures, interrupts, static binding, and a proposed verify statement. Files and arrays are easily handled with minor modifications to the hardware and software mechanisms. Procedures are also easily handled provided they do not reference global variables. Interrupts are more difficult, however, since the possibility of an interrupt must be considered during the transformation of a program. Support for objects that are statically bound is easily added, so that the mechanism is applicable to programs which reference both statically bound and dynamically bound variables. The proposed verify statement is a language extension allowing greater flexibility in the specification of the initial classes of the input parameters (for Algorithm T2). Although useful, this feature greatly complicates the implementation of T2. Further investigation is needed to determine its feasibility.

#### 4.1 Hardware Mechanism

Under dynamic binding we need an efficient mechanism for obtaining and updating the security class of an object. This objective is achieved with a tagged architecture [Fu73] that includes tags, a stack, and a class updating mechanism.

With each register  $u$  there is associated a tag field whose contents identify the security class bound to  $u$  - i.e., the value of the function tag( $u$ ) described in Section 2.1. For an object  $a$  stored in location loc( $a$ ), we shall employ the usual notation "a" to denote the tag field of loc( $a$ ) and simply " $a$ " to denote its value field. Tags are associated with both arithmetic and storage registers.

A hardware stack HS is used for saving and restoring the classes of Boolean variables. As noted earlier, a statement  $\phi$  at nesting depth  $k$  is conditioned on the values of  $k$  Boolean variables  $e_1, \dots, e_k$  (unless  $k = 0$ ), and information flows from these variables to any variable  $b$  assigned a value in  $\phi$ . Rather than storing the classes  $e_1, \dots, e_k$  separately in HS, it is convenient to initialize HS with  $e_0 = L$  (the lowest class) and place  $e_0 \oplus \dots \oplus e_i$  at stack position  $i$  ( $0 \leq i \leq k$ ), so that the top of the stack is the least upper bound of  $\{e_1, \dots, e_k\}$ . This is done as follows. Denote by HS the class on the top of HS, and let push( $e$ ) be an operation which computes HS  $\oplus e$  and adds this to the stack. (To an observer, this has the effect of setting HS := HS  $\oplus e$ .) After  $e_1, \dots, e_k$  have been pushed onto HS, HS =  $e_1 \oplus \dots \oplus e_k$ . It is then possible to update the class of the variable  $b$  to reflect the implicit flows from  $e_1, \dots, e_k$  to  $b$  by performing the single operation  $b := b \oplus \text{HS}$ . The operation pop removes and discards the value on top of HS. (To an observer, this has the effect of restoring HS to the value it had prior to the corresponding push operation.)

The push and pop operations are to be performed when (and only when) a selection or iteration statement is executed as follows:

<u>Statement</u>	<u>Execution</u>
<u>if</u> $e$ <u>then</u> $\phi_1$ [ <u>else</u> $\phi_2$ ]	<u>push</u> ( $e$ ); <u>if</u> $e$ <u>then</u> $\phi_1$ [ <u>else</u> $\phi_2$ ]; <u>pop</u>
<u>while</u> $e$ <u>do</u> $\phi_1$	<u>push</u> ( $e$ ); <u>while</u> $e$ <u>do begin</u> $\phi_1$ ; <u>pop</u> ; <u>push</u> ( $e$ ) <u>end pop</u>

For the present we leave unspecified how the push and pop operations are associated with these statements; in the next section we shall show how this can be done with a software mechanism.



Assume  $\underline{HS} = A$  immediately prior to the execution of a program  $\phi$ . It is easily verified from the specifications above and the definition of a program that the following propositions hold:

Proposition 4.1-1.  $\underline{HS} = A$  immediately after execution of  $\phi$  (since selection and iteration statements both restore  $\underline{HS}$  on termination).

Proposition 4.1-2.  $A \rightarrow \underline{HS}$  during execution of  $\phi$  (since the operation  $\text{push}(e)$  sets  $\underline{HS} := \underline{HS} \oplus e$ ).

Proposition 4.1-3. If  $\phi = \text{"begin } \phi_1; \dots; \phi_m \text{ end"}$ , then  $\underline{HS} = A$  immediately before and after execution of each  $\phi_i$  ( $1 \leq i \leq m$ ).

Proposition 4.1-4. If  $\phi = \text{"if } e \text{ then } \phi_1 \text{ [else } \phi_2 \text{]}"$ , then  $\underline{HS} = A \oplus e$  immediately before and after execution of  $\phi_1$  [or  $\phi_2$ ]. Furthermore, since  $e \rightarrow A \oplus e$ ,  $e \rightarrow \underline{HS}$  during execution of  $\phi_1$  [or  $\phi_2$ ].

Proposition 4.1-5. If  $\phi = \text{"while } e \text{ do } \phi_1 \text{"}$ , then  $\underline{HS} = A \oplus e^i$  immediately before and after execution of  $\phi_1$  on the  $i^{\text{th}}$  iteration, where  $e^i$  is the class of  $e$  on the  $i^{\text{th}}$  iteration. Furthermore, since  $e^i \rightarrow A \oplus e^i$ ,  $e^i \rightarrow \underline{HS}$  during execution of  $\phi_1$ . Since  $A \oplus e^i$  is popped from  $\underline{HS}$  at the end of the  $i^{\text{th}}$  iteration, the size of  $\underline{HS}$  is always  $k + 2$  (where  $k$  is the nesting depth of  $\phi$ ) immediately before and after execution of  $\phi_1$ .

Proposition 4.1-6.  $A = L$  if  $\phi$  is at nesting level 0.

The class updating mechanism automatically updates the class of an object  $b$  when an assignment statement  $\phi = \text{"b := f(a_1, \dots, a_n)"}$  is executed, using the classes of the operands and the top of the stack. Specifically, it sets  $\underline{b} := \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{HS}$ .

As an illustration of how this can be done, suppose the statement "a := b \* c + d" is translated into the following code for a single accumulator machine:

```
LOAD b
MULT c
ADD d
STO a
```

Letting ACC denote the accumulator, the hardware would be designed to perform the following operations for each instruction:

<u>Instruction</u>	<u>Operation on Data</u>	<u>Operation on Tags</u>
LOAD b	ACC := b	ACC := $b \oplus \underline{HS}$
MULT c	ACC := ACC * c	ACC := $\underline{ACC} \oplus \underline{c} \oplus \underline{HS}$
ADD d	ACC := ACC + d	ACC := $\underline{ACC} \oplus \underline{d} \oplus \underline{HS}$
STO a	a := ACC	a := $\underline{ACC} \oplus \underline{HS}$

It is especially important to note that the class  $\underline{HS}$  is used at every step to update tags. This insures that implicit flows are properly accounted for when new values are produced.

Given an assignment statement  $\phi = "b := f(a_1, \dots, a_n)"$ , the update operation " $\underline{b} := \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{HS}$ " clearly has the effect of setting  $\underline{b}$  to be at least as great as the classes of all variables that explicitly flow into b. That it also sets  $\underline{b}$  to be at least as great as the classes of all variables that implicitly flow into b has been suggested earlier and will be established rigorously below by Theorem 4.1-1, which states that  $\underline{HS}$  is the least upper bound of the classes of all values on which  $\phi$  is conditioned. To prove the theorem, we need the result of the following lemma which states that for an iteration "while e do  $\phi_1$ ", e can only increase from one iteration to the next.

Lemma 4.1-1. Consider the execution of an iteration statement  $\phi =$  "while  $e$  do  $\phi_1$ ". If  $\underline{e}^i$  is the class of  $e$  immediately prior to the  $i^{\text{th}}$  iteration of  $\phi_1$  ( $i \geq 1$ ), then  $\underline{e}^{i-1} \rightarrow \underline{e}^i$  ( $i > 1$ ).

proof. By Proposition 4.1-5,  $\underline{e}^{i-1} \rightarrow \underline{HS}$  during the  $(i-1)^{\text{st}}$  iteration of  $\phi_1$ . We then have

$$\underline{e}^i = \begin{cases} \underline{e}^{i-1} & \text{if } e \text{ is not assigned a value on the } (i-1)^{\text{st}} \text{ iteration,} \\ \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus A & \text{if the last assignment to } e \text{ on the} \\ & (i-1)^{\text{st}} \text{ iteration associated with } e \text{ some value} \\ & f(a_1, \dots, a_n) \text{ when } \underline{HS} = A. \end{cases}$$

But  $\underline{e}^{i-1} \rightarrow A$ , so that  $\underline{e}^{i-1} \rightarrow \underline{e}^i$ .

Theorem 4.1-1. Consider the execution of a program, and let  $\phi$  be the next statement to be executed. Let  $E(\phi)$  be the set of values on which the execution of  $\phi$  is conditioned. Then  $\underline{HS} = \underline{E(\phi)}$  immediately before and after execution of  $\phi$ , where

$$\underline{E(\phi)} = \begin{cases} \{\underline{e} \mid e \in E(\phi)\} & \text{if } E(\phi) \neq \emptyset \\ L & \text{otherwise} \end{cases}$$

proof. If  $\phi$  is a member of a sequence "begin  $\phi_1; \dots; \phi_m$  end" (where  $m = 1$  implies that begin and end need not be present),  $E(\phi) = E(\phi_1)$ ; and by Proposition 4.1-3,  $\underline{HS} = \underline{E(\phi)}$  prior to executing  $\phi$  if and only if  $\underline{HS} = \underline{E(\phi_1)}$  prior to executing  $\phi_1$ . Therefore we may assume without loss of generality that  $\phi$  is the first member of a sequence.

Now, if the nesting level  $k$  of  $\phi$  is 0, the initial conditions imply  $E(\phi) = \emptyset$  and  $\underline{HS} = L$  (by Proposition 4.1-6), so that the theorem holds. As an induction hypothesis, assume the theorem holds for  $\phi$  at nesting levels  $k < K$  and consider  $\phi$  at level  $K$ .

There are two possibilities. First, suppose there is a statement  $\phi' = \text{"if } e \text{ then } \phi_1 \text{ [else } \phi_2\text{]}$ , where  $\phi$  is the first member of sequence  $\phi_1$  [or  $\phi_2$ ]. Clearly  $E(\phi) = E(\phi') \cup \{e\}$ . Since  $\phi'$  is at nesting level  $K-1$ , the induction hypothesis implies  $\underline{HS} = \underline{E(\phi')}$  prior to executing  $\phi'$ . Since  $\underline{HS} \oplus \underline{e}$  is pushed onto the stack  $\underline{HS}$  before executing  $\phi_1$  [or  $\phi_2$ ], we have

$$\underline{E(\phi)} = \underline{E(\phi')} \cup \{e\} = \underline{E(\phi')} \oplus \underline{e} = \underline{HS} \oplus \underline{e}$$

as required. Second, suppose there is a statement  $\phi' = \text{"while } e \text{ do } \phi_1\text{"}$ , where  $\phi$  is the first member of sequence  $\phi_1$ . At the  $r^{\text{th}}$  iteration,

$$E(\phi) = E(\phi') \cup \{e^1, \dots, e^i, \dots, e^r\}$$

(since execution of each iteration is conditioned on the preceding), where  $e^i$  is the value of  $e$  prior to the  $i^{\text{th}}$  iteration. Whence,

$$\underline{E(\phi)} = \underline{E(\phi')} \oplus \underline{e}^1 \oplus \dots \oplus \underline{e}^r.$$

By Lemma 4.1-1,  $\underline{e}^{i-1} \rightarrow \underline{e}^i$  ( $1 < i \leq r$ ) so that  $\underline{e}^1 \oplus \dots \oplus \underline{e}^r = \underline{e}^r$ .

By Proposition 4.1-5,  $\underline{HS} \oplus \underline{e}^r$  will be on the stack just prior to the  $r^{\text{th}}$  iteration; and by induction  $\underline{HS} = \underline{E(\phi')}$ , so that

$$\underline{E(\phi)} = \underline{E(\phi')} \oplus \underline{e}^r = \underline{HS} \oplus \underline{e}^r$$

as required.

Figure 4.1-1 illustrates the operation of the hardware mechanism by tracing the execution of the insecure program cited in Section 2.4.5. The two traces shown illustrate the claim made earlier, that  $\underline{b} = \underline{a}$  on termination of the program, while  $\underline{b} = \underline{L}$  even though  $\underline{a} = \underline{H}$  (hence,  $\underline{a} \neq \underline{b}$  as required). This program demonstrates that the hardware mechanism is insufficient to guarantee security, and that a software mechanism is required.

Source Statement	Execution Trace	Stack (HS at right)	<u>a</u> , <u>a</u>	<u>b</u> , <u>b</u>	<u>c</u> , <u>c</u>	<u>e</u> , <u>e</u>
Initial conditions for the case <u>a</u> = 0:						
<u>begin</u>		L	H,0	L,0	L,0	L,false
<u>e</u> := ( <u>a</u> = 0);	<u>e</u> := ( <u>a</u> = 0)					
<u>if</u> <u>e</u> <u>then</u> <u>c</u> := 1;	<u>push</u> ( <u>e</u> )	L, H				H,true
	<u>c</u> := 1					H,1
<u>e</u> := ( <u>c</u> = 0);	<u>pop</u>	L				H,false
<u>if</u> <u>e</u> <u>then</u> <u>b</u> := 1	<u>e</u> := ( <u>c</u> = 0)	L,H				
	<u>push</u> ( <u>e</u> )					
	<u>pop</u>					
<u>end</u>						
Final conditions ( <u>b</u> = <u>a</u> = 0):						
		L	H,0	L,0	H,1	H,false
Initial conditions for the case <u>a</u> = 1:						
<u>begin</u>		L	H,1	L,0	L,0	L,false
<u>e</u> := ( <u>a</u> = 0);	<u>e</u> := ( <u>a</u> = 0)					
<u>if</u> <u>e</u> <u>then</u> <u>c</u> := 1;	<u>push</u> ( <u>e</u> )	L, H				H,false
	<u>pop</u>	L				
<u>e</u> := ( <u>c</u> = 0);	<u>e</u> := ( <u>c</u> = 0)	L,L				L,true
<u>if</u> <u>e</u> <u>then</u> <u>b</u> := 1	<u>push</u> ( <u>e</u> )					
	<u>b</u> := 1					
	<u>pop</u>	L				
<u>end</u>						
Final conditions ( <u>b</u> = <u>a</u> = 1):						
		L	H,1	L,1	L,0	L,true

Figure 4.1-1. Execution Traces of an Insecure Program.

## 4.2 Software Mechanism

The hardware mechanism outlined in the previous section is not sufficient to guarantee security since an implicit flow from a Boolean variable  $e$  to a variable  $b$  can occur in the absence of any explicit flow to  $b$  ( $b$  is updated only when an explicit assignment is made, but  $e$  flows implicitly to  $b$  even when the statements assigning values to  $b$  are not executed). In general, this may occur during execution of a selection statement, "if  $e$  then  $\phi_1$  [else  $\phi_2$ ]", if there is an assignment statement " $b := f(a_1, \dots, a_n)$ " in  $\phi_1$  [or  $\phi_2$ ] and that statement is not executed. This may also occur during execution of an iteration statement, "while  $e$  do  $\phi_1$ ", if there is an assignment statement " $b := f(a_1, \dots, a_n)$ " in  $\phi_1$ , since no statement in  $\phi_1$  is executed after the final evaluation of  $e$  (which necessarily leaves  $e = \text{false}$ ).

This problem can be solved with a compile-time mechanism that inserts into the compiled code an operation "update  $b$ ", having the effect

$$\underline{b} := \underline{b} \oplus \underline{HS},$$

at all locations where an assignment to  $b$  is conditioned on one or more Boolean variables  $e_1, \dots, e_k$  (recall from Theorem 4.1-1 that  $\underline{HS} = \underline{e_1} \oplus \dots \oplus \underline{e_k}$  at that time). Then all implicit flow to  $b$  is guaranteed secure irrespective of whether or not the assignment statement is executed. The mechanism also inserts into the compiled code all necessary push and pop instructions.

Rather than show the details of the compilation process, we have expressed the mechanism as a high level recursive procedure  $Tl(\phi, V)$  which transforms an arbitrary program  $\phi$  into an equivalent secure one

by inserting all necessary push, pop, and update instructions (see Figure 4.2-1). The parameter  $V$  returns the set of variables that appear on the left side of assignment statements in  $\phi$  (i.e.,  $V$  is equivalent to the set  $V(\phi)$  introduced in Section 2.4.4 and referenced in Chapter 3); it is used to generate update instructions. (We have written update  $X$  to denote a sequence of update instructions, one for each member of the set  $X$ .) The language used to express  $T1$  resembles PASCAL [Wr71] and includes the types "program" and "variable"; their use should be clear in Figure 4.2-1. The initial call to  $T1$  is  $T1(\phi, V)$  where  $\phi$  is an entire program.  $T1$  returns the transformed (secure) program in  $\phi$  and the set of variables that appear on the left side of assignment statements in  $V$ . Update instructions are generated as needed for elements of  $V$  (the proof that this is done correctly is given below).

Figure 4.2-2 (a) shows the results of the transformation when  $T1$  is applied to the sample program of Figure 4.1-1. The reader may easily verify that the update instructions cause  $b$  to be identical to  $a$  on termination of the program, whereupon the "leak" has been plugged. Figure 4.2-2 (b) shows an equivalent program, where the same leak can occur using iteration statements; in this case also, the transformed program is secure.

Theorem 4.2-1. Let  $\phi$  be a given program. After execution of a  $T1(\phi, V)$  the following will be true:

- 1)  $V$  will be the set of all variables appearing on the left side of assignment statements in  $\phi$ ; and

```

Tl: procedure ( $\phi$ : program, V: powerset variable);

  V1, V2: powerset variable;
  i: integer;

  do case  $\phi$  of

     $\phi$  = "b := f(a1, ..., an)":      {Assignment}
      V := b;

     $\phi$  = "begin  $\phi_1$ ; ...;  $\phi_m$  end":    {Concatenation}
      begin
        V :=  $\emptyset$ ;
        for i := 1 to m do begin
          Tl( $\phi_i$ , V1);
          V := V  $\cup$  V1 end
        end;

     $\phi$  = "if e then  $\phi_1$  [else  $\phi_2$ ]": {Selection}
      begin
        Tl( $\phi_1$ , V1);
        if  $\phi_2 \neq \emptyset$  then Tl( $\phi_2$ , V2) else V2 :=  $\emptyset$ ;
        V := V1  $\cup$  V2;
         $\phi$  := "begin push(e); if e then  $\phi_1$  [else  $\phi_2$ ];
              update V - V1  $\cap$  V2; pop end";
      end;

     $\phi$  = "while e do  $\phi_1$ ":            {Iteration}
      begin
        Tl( $\phi_1$ , V);
         $\phi$  := "begin push(e); while e do begin  $\phi_1$ ; pop; push(e) end
              update V; pop end";
      end

    endcase

  end Tl

```

Figure 4.2-1. Transformation Algorithm Tl.



<u>Insecure Source Programs</u>	<u>Transformed Secure Program</u>
a) <u>begin</u> b := 0; c := 0; e := (a = 0); <u>if e then c := 1;</u> e := (c = 0); <u>if e then b := 1</u> <u>end</u>	<u>begin</u> b := 0; c := 0; e := (a = 0); <u>begin push(e); if e then c := 1;</u> <u>update c; pop end;</u> e := (c = 0); <u>begin push(e); if e then b := 1;</u> <u>update b; pop end</u> <u>end</u>
b) <u>begin</u> b := 0; c := 0; e := (a = 0); <u>while e do begin</u> c := 1; e := <u>false end</u> e := (c = 0); <u>while e do begin</u> b := 1; e := <u>false end</u> <u>end</u>	<u>begin</u> b := 0; c := 0; e := (a = 0); <u>begin push(e);</u> <u>while e do begin begin c := 1;</u> e := <u>false end; pop; push(e) end;</u> <u>update c, e; pop end;</u> e := (c = 0); <u>begin push(e);</u> <u>while e do begin begin b := 1;</u> e := <u>false end; pop; push(e) end;</u> <u>update b, e; pop end</u> <u>end</u>

Figure 4.2-2. Transformation of Insecure Programs.

- 2)  $\phi$  will be secure. Specifically, all necessary push, pop, and update instructions will have been inserted into  $\phi$  so that after any execution of  $\phi$ ,  $\underline{HS} \rightarrow \underline{b}$  will hold for all variables  $b \in V$ . Therefore, if  $\phi$  is conditioned on Boolean variables  $e_1, \dots, e_k$ ,  $\underline{e_i} \rightarrow \underline{b}$  ( $1 \leq i \leq k$ ) is guaranteed to hold for all  $b \in V$  (since  $\underline{HS} = \underline{e_1} \oplus \dots \oplus \underline{e_k}$  on entry to  $\phi$  by Theorem 4.1-1). Moreover, if  $\phi$  is a selection or iteration statement with Boolean variable  $e$ ,  $\underline{HS} \oplus \underline{e} \rightarrow \underline{b}$  will hold for all  $b \in V$ .

proof. Define the "structure index"  $d$  of a program to be the number of concatenation, selection, and iteration structure operations used to construct it. We proceed by induction on the value  $d$  of  $\phi$ .

BASIS ( $d=0$ ): In this case  $\phi$  is an assignment " $b := f(a_1, \dots, a_n)$ ".

On termination, T1 leaves  $\phi$  unchanged and  $V = \{b\}$ . Since the hardware updating mechanism guarantees that  $\underline{HS} \rightarrow \underline{b}$  after the assignment is performed, T1 correctly processes  $\phi$ .

INDUCTION: Let  $\phi$  have structure index  $d = D$  and suppose the theorem holds for  $d < D$ . There are three cases.

- i)  $\phi = \text{"begin } \phi_1; \dots; \phi_m \text{ end"}$ . In this case, T1 may be assumed to correctly process  $\phi_i$  ( $1 \leq i \leq m$ ), since  $\phi_i$  has structure index  $d_i < D$ . By Proposition 4.1-3, if  $\underline{HS} = A$  prior to execution of  $\phi$ , then  $\underline{HS} = A$  prior to execution of each  $\phi_i$ . Therefore, for  $b \in V_i$  (where  $V_i$  is just the set of variables assigned values in  $\phi_i$ ),  $\underline{HS} \rightarrow \underline{b}$  will hold after execution of  $\phi_i$ . Hence, for  $b \in V_1 \cup \dots \cup V_m$ ,  $\underline{HS} \rightarrow \underline{b}$  will hold after execution of  $\phi$ . Since T1 returns  $V = V_1 \cup \dots \cup V_m$  and no

push, pop, or update instructions need be inserted into  $\phi$ ,  
T1 correctly processes  $\phi$ .

- ii)  $\phi = \text{"if } e \text{ then } \phi_1 \text{ [else } \phi_2 \text{]"}.$  In this case, T1 may be assumed to correctly process  $\phi_1$  [and  $\phi_2$ ] since  $\phi_1$  and  $\phi_2$  have  $d_1 < D$  and  $d_2 < D$ . Therefore  $V = V_1 \cup V_2$ . The transformed program is "begin push(e); if e then  $\phi_1$  [else  $\phi_2$ ]; update  $V - V_1 \cap V_2$ ; pop end". Therefore, if  $\phi_1$  is executed, then for  $b \in V_1$ ,  $\underline{HS} \oplus e \rightarrow b$  will hold after execution of  $\phi_1$ ; if  $\phi_2$  is executed, then for  $b \in V_2$ ,  $\underline{HS} \oplus b \rightarrow b$  will hold after execution of  $V_2$ . Hence, the update instruction guarantees that for  $b \in V$ ,  $\underline{HS} \rightarrow \underline{HS} \oplus e \rightarrow b$  will hold after execution of  $\phi$ , irrespective of whether or not  $\phi_1$  [or  $\phi_2$ ] is executed.
- iii)  $\phi = \text{"while } e \text{ do } \phi_1 \text{"}.$  In this case, T1 may be assumed to correctly process  $\phi_1$  since  $d_1 < D$ . Therefore  $V = V_1$ . The transformed program is "begin push(e); while e do begin  $\phi_1$ ; pop; push(e) end; update V; pop end". In this case it is trivially true that for  $b \in V$ ,  $\underline{HS} \rightarrow \underline{HS} \oplus e \rightarrow b$  will hold after execution of  $\phi$ , since the class of each  $b \in V$  is updated by the update instruction.

The reader may have wondered if insertion of push and pop instructions is necessary. If the hardware is designed to automatically push the stack on execution of any conditional branch, there will be no need for the compiler to insert push instructions into the compiled code. To see this, consider the translation of selection and iteration statements. Assume that Boolean false is represented internally by 0 and

true by any non-zero value. Let B be an unconditional branch instruction and BZ a branch on zero instruction. The translation of selection and iteration statements into code for a machine with a single accumulator (ACC) is:

- a) begin push(e); if e then  $\phi_1$  else  $\phi_2$ ; update  $V - V_1 \cap V_2$ ; pop  
end

```

LOAD e
PUSH
BZ L1
code for  $\phi_1$ 
B L2
L1 code for  $\phi_2$ 
L2 code for update
POP

```

- b) begin push(e); while e do begin  $\phi_1$ ; pop; push(e) end; update V;  
pop end

```

L1 LOAD e
PUSH
BZ L2
code for  $\phi_1$ 
POP
B L1
L2 code for update
POP

```

The PUSH instruction pushes HS @ ACC onto the stack. Since no other statement types generate conditional branch instructions, it is clear from the above that the PUSH operation can be performed automatically when the BZ instruction is executed. Note, however, that the compiler must still generate code for the POP operation, since there is no instruction analogous to the BZ instruction that is always paired with POP.

Hereafter, we will assume without loss of generality that the operation push(e) is performed automatically and adopt the simpler forms for transformed iteration and selection statements:

```

begin if e then  $\phi_1$  [else  $\phi_2$ ]; update  $V - V_1 \cap V_2$ ; pop end
begin while e do begin  $\phi_1$ ; pop end; update V; pop end

```

The mechanism outlined above may substantially increase the time and space requirements of a program. For example, the program

```

if e1
  then if e2 then a := 0 else b := 0
  else c := 0

```

would be transformed into

```

begin
  if e1
    then begin if e2 then a := 0 else b := 0; update a, b; pop end
    else c := 0
  update a, b, c;
  pop
end

```

Execution of the transformed program results in as many as five update operations being performed, when only one assignment operation is performed. Furthermore, none of the update operations need be performed at all if e1 = e2 = L (since they will have no effect). The software mechanism introduced in the next section provides a means for eliminating some (or all) of the update operations when they are not required.

#### 4.3 Software Mechanism for Selective Updating

The mechanism of this section improves on the one of the previous section in two ways: it generates fewer update instructions, and it permits the programmer to declare initial classes for input parameters. Again, our approach is to describe an algorithm that transforms a given program into a secure one.

The syntax of PASCAL is convenient for declaring initial classes for input parameters. For example,

$f$ : procedure ( $x$ : integer of class A,  $y$ : integer of class L;  $z$ : integer)

specifies initial class of A and L for formal input parameters  $x$  and  $y$  respectively. A run-time check must verify the classes of the actual parameters. Thus, if  $f$  is invoked by an operation "call  $f(a, b, c)$ ", the run-time mechanism would check that  $\underline{a} = A$  and  $\underline{b} = B$  before calling  $f$ .

The transformation algorithm T2 accepts a program  $\phi$  as input, transforms  $\phi$  into a secure program, and then returns the set  $V$  of variables to which statements of  $\phi$  assign values. It also returns tables, denoted  $\ell$  and  $h$ , specifying the lowest and highest possible classes for each variable. In particular, if  $\ell[a]$  and  $h[a]$  are the lowest and highest classes of variable  $a$  before T2 processes  $\phi$ , then  $\ell[a]$  and  $h[a]$  will be updated by T2 so that they represent the lowest and highest possible classes of  $a$  after an execution of  $\phi$ . The initial call on T2 for a program  $\phi$  is of the form  $T2(\phi, V, \ell, h)$ , where  $\ell$  and  $h$  are initialized as follows: if  $a$  is an input variable of declared class  $A$ , then  $\ell[a] = h[a] = A$ ; otherwise  $\ell[a] = L$  and  $h[a] = H$ . Algorithm T2 is given in Figure 4.3-1.

For each call, T2 simulates information flow during possible executions of the program  $\phi$  until bounds on the classes of all variables are uniquely determined. This is done by simulating the effect of  $\phi$  on the machine's stack and the classes of all variables. A global variable SS ("software stack") is used by T2 to represent the hardware stack HS, but with one difference: SS contains pairs of classes corresponding to the lowest and highest classes that could be pushed onto HS during an

```

T2: procedure ( $\phi$ : program; V: powerset variable;
              $\ell$ , h: array [variable] of class);

V1, V2: powerset variable;
 $\ell_1$ ,  $\ell_2$ , h1, h2: array [variable] of class;
i: integer;
change: boolean;

UPDATE: procedure (X, Y: powerset variable);

Y := {b  $\in$  X | SS.h  $\neq$   $\ell$ [b]};
for all b $\in$ Y do begin  $\ell$ [b] :=  $\ell$ [b]  $\oplus$  SS. $\ell$ ; h[b] := h[b]  $\oplus$  SS.h end
end UPDATE;

do case  $\phi$  of           {Assignment, Concatenation, Selection, Iteration}

 $\phi$  = "b := f(a1, ..., an)":           {Assignment}

    begin
         $\ell$ [b] :=  $\ell$ [a1]  $\oplus$  ...  $\oplus$   $\ell$ [an]  $\oplus$  SS. $\ell$ ;
        h[b] := h[a1]  $\oplus$  ...  $\oplus$  h[an]  $\oplus$  SS.h;
        V := b
    end;

 $\phi$  = "begin  $\phi_1$ ; ...;  $\phi_m$  end":           {Concatenation}

    begin
        V :=  $\emptyset$ ;
        for i := 1 to m do
            begin
                T2( $\phi_i$ , V1,  $\ell$ , h);
                V := V  $\cup$  V1
            end
        end;

```

Figure 4.3-1. Transformation Algorithm T2.

$\phi = \text{"if } e \text{ then } \phi_1 \text{ [else } \phi_2\text{]"} \text{ or } \quad \{\text{Selection}\}$   
"begin if e then  $\phi_1$  [else  $\phi_2$ ]; update U; pop end":

begin  
 PUSH(e);  
 $l_1, l_2 := l; h_1, h_2 := h;$   
 $T2(\phi_1, V_1, l_1, h_1);$   
if  $\phi_2 \neq \emptyset$  then  $T2(\phi_2, V_2, l_2, h_2)$  else  $V_2 := \emptyset;$   
 $V := V_1 \cup V_2;$   
for all  $b \in V$  do  
begin  
 $l[b] := l_1[b] \otimes l_2[b];$   
 $h[b] := h_1[b] \otimes h_2[b]$   
end;  
 UPDATE( $V - V_1 \cap V_2, U$ );  
 POP;  
 $\phi = \text{"begin if e then } \phi_1 \text{ [else } \phi_2\text{]; update U; pop end"}$   
end;

$\phi = \text{"while e do } \phi_1 \text{" or } \quad \{\text{Iteration}\}$   
"begin while e do begin  $\phi_1$ ; pop end; update U; pop end":

begin  
 $l_1 := l; h_1 := h;$   
repeat  
begin  
 PUSH(e);  $V := \emptyset;$   
 $T2(\phi_1, V, l_1, h_1);$   
for all  $b \in V$  do  
begin  
if  $h[b] \neq h[b] \otimes h_1[b]$  then  
begin  $h[b] := h[b] \otimes h_1[b];$  change := true end;  
if  $l[b] \neq l[b] \otimes l_1[b]$  then  
begin  $l[b] := l[b] \otimes l_1[b];$  change := true end  
end;  
 POP; PUSH(e)  
end  
until  $\neg$  change;  
 UPDATE( $V, U$ );  
 POP;  
 $\phi := \text{"begin while e do begin } \phi_1\text{; pop end; update U; pop end"}$   
end

endcase

end T2

Figure 4.3-1, cont.



execution of  $\phi$ . Denoting by  $\underline{SS.l}$  and  $\underline{SS.h}$  the lowest and highest pair on top of SS, the function  $\text{PUSH}(e)$  pushes  $(\underline{SS.l} \oplus l[e], \underline{SS.h} \oplus h[e])$  onto SS. Corresponding,  $\text{POP}$  removes the top pair from SS and discards it. The main complication of the simulation process occurs for iteration statements,  $\phi = \text{"while } e \text{ do } \phi_1"$ . For such a statement, the value of  $e$  (and thus  $\underline{HS}$ ) may increase from one iteration to the next. Therefore, a program  $\phi$  already transformed by T2 may require reprocessing to guarantee all variables in need of updating have been detected. Such reprocessing will terminate as soon as the  $l$  and  $h$  tables have the same values before and after processing an iteration statement — that is, as soon as the values of  $l$  and  $h$  "stabilize".

The details of the operation of T2 are explained fully by the following theorem which establishes its correctness.

Theorem 4.3-1. Let  $\phi$ ,  $l$ ,  $h$ , and SS be given such that prior to any execution of  $\phi$

$l[a] \rightarrow \underline{a} \rightarrow h[a]$  for each variable  $a$  of  $\phi$ , and

$\underline{SS.l} \rightarrow \underline{HS} \rightarrow \underline{SS.h}$ .

After execution of a call  $T2(\phi, V, l, h)$  the following will be true:

- 1)  $V$  will be the set of all variables appearing on the left side of assignment statements in  $\phi$ ;
- 2)  $\phi$  will be secure; that is, all necessary push, pop, and update operations will have been inserted into  $\phi$  so that after any execution of  $\phi$ ,  $\underline{HS} \rightarrow \underline{b}$  will hold for all  $b \in V$ . Therefore, if  $\phi$  is conditioned on Boolean variables  $e_1, \dots, e_k$ ,  $\underline{e_i} \rightarrow \underline{b}$  ( $1 \leq i \leq k$ ) is guaranteed to hold for all  $b \in V$  (since  $\underline{HS} = \underline{e_1} \oplus \dots \oplus \underline{e_k}$  on entry to  $\phi$  by Theorem 4.1-1). Moreover,

if  $\phi$  is a selection or iteration statement with Boolean variable  $e$ ,  $\underline{HS} \oplus \underline{e} \rightarrow \underline{b}$  will hold for all  $b \in V$ .

- 3)  $\ell$  and  $h$  will be updated by T2 so that after any execution of  $\phi$

$$\ell[a] \rightarrow \underline{a} \rightarrow h[a] \text{ for all variables } a \text{ of } \phi.$$

proof. The proof is by induction on the "structure index"  $d$  of  $\phi$ , as in Theorem 4.2-1.

BASIS ( $d = 0$ ): In this case  $\phi = "b := f(a_1, \dots, a_n)"$ , and  $V = \{b\}$ .

The hardware guarantees that  $\underline{b} = \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{HS}$  after an execution of  $\phi$ , so that  $\underline{HS} \rightarrow \underline{b}$  is guaranteed to hold. Since  $\underline{SS}.\ell \rightarrow \underline{HS} \rightarrow \underline{SS}.h$  and  $\ell[a_i] \rightarrow \underline{a_i} \rightarrow h[a_i]$  ( $1 \leq i \leq n$ ) prior to an execution of  $\phi$ , the lattice properties (see Section 2.3) imply that  $\ell[b] = \ell[a_1] \oplus \dots \oplus \ell[a_n] \oplus \underline{SS}.\ell \oplus \underline{b} \rightarrow h[a_1] \oplus \dots \oplus h[a_n] \oplus \underline{SS}.h = h[b]$  after an execution of  $\phi$ , establishing the correctness of the  $\ell$  and  $h$  tables returned by T2.

INDUCTION: Let  $\phi$  have structure index  $d = D$ , and assume T2 satisfies (1) - (3) for any algorithm whose level is  $d < D$ . There are three possibilities.

- i)  $\phi = "\underline{begin} \phi_1; \dots; \phi_m \underline{end}"$ . In this case each  $\phi_i$  has  $d_i < D$ , so we may assume that T2 transforms it correctly. Therefore, if  $\ell[a] \rightarrow \underline{a} \rightarrow h[a]$  prior to an execution of  $\phi_i$  ( $1 \leq i \leq m$ ), the call  $T2(\phi_i, V_1, \ell, h)$  properly updates the  $\ell$  and  $h$  tables so that  $\ell[a] \rightarrow \underline{a} \rightarrow h[a]$  after an execution of  $\phi_i$ . Hence, the  $\ell$  and  $h$  tables returned by T2 for  $\phi$  are correct. Moreover,

$V = V_1 \cup \dots \cup V_m$  when T2 completes processing  $\phi$ . Finally, the induction assumption guarantees that  $\underline{HS} \rightarrow \underline{b}$  for  $b \in V_1$  ( $1 \leq i \leq m$ ), whereupon  $\underline{HS} \rightarrow \underline{b}$  for all  $b \in V$  (since HS is not modified by begin-end statements).

ii)  $\phi = \text{"if } e \text{ then } \phi_1 \text{ [else } \phi_2\text{]"} \text{ or}$

$\phi = \text{"begin if } e \text{ then } \phi_1 \text{ [else } \phi_2\text{]; update U; pop end"}.$

Since  $\phi_1$  and  $\phi_2$  have  $d_1 < D$  and  $d_2 < D$ , the induction assumption implies that  $\ell_1$  and  $h_1$  contain correct class bounds after T2 processes  $\phi_1$ , and  $\ell_2$  and  $h_2$  contain correct class bounds after T2 processes  $\phi_2$ . Since  $\ell_1$  (or  $\ell_2$ ) and  $\ell$ , and  $h_1$  (or  $h_2$ ) and  $h$  can differ only for  $b \in V_1 \cup V_2$ ,

$$\ell[b] = \ell_1[b] \oplus \ell_2[b] \rightarrow \underline{b} \rightarrow h_1[b] \oplus h_2[b] = h[b]$$

must hold if either  $\phi_1$  or  $\phi_2$  is executed, establishing the correctness of the  $\ell$  and  $h$  tables returned by T2. Now, T2 correctly simulates HS, so that  $\underline{SS.h} \oplus h[e]$  is a correct upper bound on  $\underline{HS} \oplus \underline{e}$  when UPDATE is called. Since UPDATE generates an update instruction for any  $b \in V - V_1 \cap V_2$  for which  $\underline{SS.h} \oplus h[e] \neq \ell[b]$ , this guarantees that

$$\underline{HS} \rightarrow \underline{HS} \oplus \underline{e} \rightarrow \underline{SS.h} \oplus h[e] \rightarrow \ell[b] \rightarrow \underline{b}$$

will hold after an execution of  $\phi$ . (Note that, if  $b \in V_1 \cap V_2$  the induction assumption implies that  $\underline{HS} \rightarrow \underline{HS} \oplus \underline{e} \rightarrow \underline{b}$  must hold if either  $\phi_1$  or  $\phi_2$  is executed; therefore update instructions will not be generated in this case, as for algorithm T1). The call on UPDATE also corrects the tables  $\ell$  and  $h$  to reflect changes that can occur when an update instruction is executed, establishing the correctness of the  $\ell$  and  $h$  tables returned by T2.

iii)  $\Phi = \text{"while } e \text{ do } \Phi_1 \text{"}$  or

$\Phi = \text{"begin while } e \text{ do begin } \Phi_1; \text{ pop end; update } U; \text{ pop end"}$ .

Since  $\Phi_1$  had  $d_1 < D$ , the induction assumption implies that the first call  $T2(\Phi_1, V_1, l_1, h_1)$  correctly processes  $\Phi_1$ . Therefore, if  $\Phi_1$  is executed (or iterated) at least once, the first iteration of  $\Phi_1$  is guaranteed secure, and for all  $b \in V = V_1$ ,  $l_1[b] \rightarrow \underline{b} \rightarrow h_1[b]$  will hold after the execution of  $\Phi_1$ . Hence, if  $\Phi_1$  is iterated exactly 0 or 1 times,

$$l[b] \otimes l_1[b] \rightarrow \underline{b} \rightarrow h[b] \otimes h_1[b]$$

will hold after the execution of  $\Phi$ , establishing the correctness of the  $l$  and  $h$  tables after  $T2$  has processed  $\Phi_1$  once.

Now, to determine bounds on  $\underline{b}$  after  $m \geq 0$  iterations of  $\Phi_1$ , it is necessary to repeat this procedure until  $l[b]$  and  $h[b]$  stabilize; that is, until further calls to  $T2$  yield bounds for  $l_1$  and  $h_1$  that fall between  $l$  and  $h$ . Clearly this process must terminate after a finite number of steps, since  $l[b]$  and  $h[b]$  cannot change beyond  $L$  and  $H$  in a finite lattice. After  $l$  and  $h$  have been stabilized, the top of the stack  $SS$  contains the upper bound  $SS.h \otimes h[e]$ , where  $h[e]$  is the (stabilized, maximal) class of  $e$ . The call on UPDATE then guarantees that update instructions are generated for all  $b \in V$  for which  $SS.h \otimes h[e] \neq l[b]$ , so that

$$HS \rightarrow HS \otimes e \rightarrow SS.h \otimes h[e] \rightarrow l[b] \rightarrow \underline{b}$$

will hold after an execution of  $\Phi$ . The call on UPDATE also corrects the  $l$  and  $h$  tables to account for bound changes caused by the update instructions.

Figure 4.3-2 illustrates the result of applying T1 to a sample program, and Figure 4.3-3 illustrates T2 applied to the same program. In Figure 4.3-3 we have traced the simulation process showing the changes to the outermost  $\ell$  and  $h$  tables and the stack SS. We assume the classes A, B, and C are linearly related by  $A \rightarrow B \rightarrow C$ . Note that the body of the while statement requires two applications of T2 for stabilization. T2 produces three less update instructions than T1 for this program.

The proof of Theorem 4.3-1 suggests that the execution time of T2 can be considerable in comparison with that of T1. The following theorem gives a bound on the most time consuming part of T2.

Theorem 4.3-2. Let  $\phi$  be an iteration "while  $e$  do  $\phi_1$ ". Then  $T2(\phi, V, \ell, h)$  will complete after at most  $(\delta+1)|V|$  simulations of  $\phi_1$ , where  $\delta$  is the maximum path length (i.e., depth) in the precedence graph of the security class lattice, and  $|V|$  is the number of elements in  $V$ .

proof. As noted in the proof of Theorem 4.3-1, each simulation of  $\phi_1$  can alter  $\ell[b]$  or  $h[b]$  for  $b \in V$  only. Let  $\ell^m[b]$  and  $h^m[b]$  denote the value of  $\ell[b]$  and  $h[b]$  after the  $m^{\text{th}}$  iteration. Since  $\ell^m[b] = \ell^{m-1}[b] \oplus \ell_1[b]$  and  $h^m[b] = h^{m-1}[b] \oplus h_1[b]$ , we have

$$\begin{aligned} \ell &\rightarrow \ell^1[b] \rightarrow \ell^{m-1}[b] \rightarrow \dots \rightarrow \ell^1[b] \rightarrow \\ &h^1[b] \rightarrow \dots \rightarrow h^{m-1}[b] \rightarrow h^m[b] \rightarrow h. \end{aligned}$$

From this it is clear that for each  $b \in V$ ,  $\ell[b]$  and  $h[b]$  together can assume at most  $\delta$  distinct classes, and  $\ell[b] = h[b]$  is possible for at most one class. Hence  $m \leq (\delta+1)|V|$ .

Source Program

```
begin
  e:=a>0;
  while e do
    begin c:=c+1; t:=b-c; e:=b>0 end;
  end;
```

```
e:=a>0;
e1:=b<0;
if e
  then begin if e1 then c:=0;
            b:=1 end
        else c:=1
        end;
```

end

Transformed Program

```
begin
  e:= b>0;
  begin
    while e do
      begin
        begin c:=c+1; b:=b-c; e:=b>0 end;
      pop
      end;
    update b,c,e; pop
  end;
  e :=a>0;
  e1:=b<0;
  begin
    if e
      then begin begin if e1 then c:=0; update c;pop end;
              b:=1 end
            else c:=1;
            update b; pop
          end
    end
  end
```

Figure 4.3-2. Example of a Program Transformed by T1.

Source Program	Simulation	SS (SS at right)	l, h values at outer level					Transformed Program
initial conditions:			a	b	c	e	el	
<u>begin</u>		(L,L)	A,A	B,B	C,C	L,H	L,H	<u>begin</u>
e:=b>0;	e:=b>0							e:=b>0;
while e do	PUSH(e)	(L,L),(B,B)				B,B		<u>begin</u>
<u>begin</u>								while e do
c:=c+1;	c:=c+1							<u>begin</u>
b:=b-c;	b:=b-c							begin c:=c+1; b:=b-c;
e:=b>0	e:=b>0							e:=b>0 end;
<u>end;</u>	POP	(L,L),(B,C)		B,C				<u>end;</u>
	PUSH(e)							pop
c:=c+1	c:=c+1							<u>end;</u>
b:=b-c	b:=b-c							update b,e; pop
e:=b>0	e:=b>0							<u>end;</u>
UPDATE	UPDATE							
POP	POP	(L,L)						
e:=a>0;	e:=a>0					A,A		e:=a>0;
el:=b<0;	el:=b<0	(L,L),(A,A)					B,C	el:=b<0;
if e	PUSH(e)							<u>begin</u>
then <u>begin</u>								if e
if el then	PUSH(el)	(L,L),(A,A),(B,C)						then <u>begin</u> <u>begin</u>
c:=0;	c:=0							if el then c:=0; pop
UPDATE	UPDATE							end; b:=1 end
POP	POP	(L,L),(A,A)						else c:=1;
b:=1 end	b:=1							pop end
else c:=1	c:=1							<u>end</u>
UPDATE	UPDATE							
POP	POP	(L,L)						<u>end</u>
<u>end</u>								<u>end</u>

Figure 4.3-3. Example of a Program Transformed by T2.

Therefore, we shall make no claims to the effect that T2 is efficient! At this stage in our research, we are primarily interested in the feasibility of performing such a transformation, and so we have sought a representation relatively easy to understand and prove correct. Making an implementation of T2 practical is a matter for future study. An investigation of the techniques used by compilers to analyze the flow of data in a program for the purpose of code optimization (see, for example, [Al70,Gi71]) may suggest some approaches.

It is interesting that despite the complexity of T2, it cannot be guaranteed to generate the minimal number of update instructions; that is, it may insert update instructions into a program that have no effect when executed, independent of the values of the input parameters. This could result, for example, by simulating the body of a while statement more times than it could ever execute, and thereby arriving at  $\ell$  and  $h$  bounds on a variable that are either too high or too low. Another way in which this could result is illustrated by the following program:

```
begin
   $b := a;$ 
   $e := (a = 0);$ 
  if  $e$  then  $b := 1$ 
end;
```

where  $\ell[a] = L$ ,  $h[a] = H$ , and  $\underline{SS}.h = \underline{SS}.\ell = L$  on entry to T2. After T2 processes the first two assignment statements,  $\ell[b] = \ell[e] = L$  and  $h[b] = h[e] = H$ . Therefore, when T2 processes the if statement, it pushes  $(L,H)$  onto  $\underline{SS}$ . Since  $\underline{SS}.h \neq \ell[b]$ , it will generate an instruction "update  $b$ " to be executed on completion of the if statement. However, this is unnecessary since  $\underline{b}$  will be the same as  $\underline{e}$ , and therefore HS, for any execution of the program. This problem may be solved with a more complicated mechanism that keeps additional information on the



classes of the variables, but we leave this for future study. Here again, the techniques used by optimizing compilers to analyze data flow may prove applicable to this problem.

#### 4.4 Extensions

##### 4.4.1 Files and I/O

Files and I/O operations are easily handled with minor extensions to the basic hardware and software mechanisms. Within the tagged architecture, hardware tags can be associated with files (e.g., as part of a file descriptor [Fu73]) just as with scalar variables. For a file  $a$ , we let  $\underline{a}$  denote the tag associated with  $a$ .

For an input statement:

input  $b_1, \dots, b_n$  from  $a$ ,

the class updating mechanism sets

$\underline{b_i} := \underline{a} \oplus \underline{HS} \quad (1 \leq i \leq n).$

For an output statement

output  $a_1, \dots, a_n$  to  $b$ ,

the class updating mechanism sets

$\underline{b} := \underline{b} \oplus \underline{a_1} \oplus \dots \oplus \underline{a_n} \oplus \underline{HS}.$

Note that in the case of output the prior class of  $b$  derives the new class of  $b$ , because the contents of  $b$  are extended (rather than replaced) to include  $a_1, \dots, a_n$ . The requisite modifications to the certification algorithms T1 and T2 are straightforward. For example, update instructions must be generated for files as for scalar variables.

#### 4.4.2 Arrays

Although it is possible under dynamic binding to permit the elements of an array to be bound to different classes, this approach can be rejected for at least two reasons. First, the usual semantic definition of an array in programming languages requires all elements to be of the same data type — e.g., in PASCAL

a: array [index] of type.

The same semantic reasons motivating this can be used to motivate uniformity of security classes within an array. Second, the implementation of T2 would become hopelessly complex, for the  $l$  and  $h$  tables would have to contain entries for each array element. We therefore require arrays to be bound to a single class and assume that the tag of the array is a component of the array descriptor. As usual, the value of the tag for array  $a$  is denoted  $\underline{a}$ .

Associate with array  $a$  the class variable  $\underline{ra}$ , which is set to contain the class of an array reference  $a[i_1, \dots, i_n]$ ; corresponding to any such reference,

$\underline{ra} := \underline{a} @ i_1 @ \dots @ i_n$ .

The variable  $\underline{ra}$  might be stored in the array descriptor and updated automatically by hardware when an array element is addressed [Fu73]. Whenever  $a[i_1, \dots, i_n]$  appears on the left side of an assignment statement, the hardware performs the operation

$\underline{a} := \underline{a} @ \underline{ra} @ HS$ .

As in the case of files,  $\underline{a}$  is updated from a prior value of  $\underline{a}$ , since the array is being modified rather than replaced.

Extensions to T2 are simple in this case. Array references must be treated as outlined above.

#### 4.4.3 Procedure Calls

Let  $f(x_1, \dots, x_m; y_1, \dots, y_n)$  be a procedure with formal input parameters  $x_1, \dots, x_m$  and formal output parameters  $y_1, \dots, y_n$ . Consider a statement

call  $f(a_1, \dots, a_m; b_1, \dots, b_n)$ ,

where  $a_1, \dots, a_m$  are the actual input parameters and  $b_1, \dots, b_n$  are the actual output parameters. On entry to  $f$ , the run-time mechanism sets

$x_i := a_i \oplus \underline{HS}$  ( $1 \leq i \leq m$ )

and verifies that  $a_i = A_i$  if  $x_i$  was declared to have initial class  $A_i$  in procedure  $f$ . On exit, the run-time mechanism sets

$b_j := y_j \oplus \underline{HS}$  ( $1 \leq j \leq n$ ).

As usual, if the call statement is conditioned on the value of one or more Boolean variables, it may be necessary to update some of the  $b_j$  by HS on completion of the conditional statement(s). As in Section 3.6, for the moment we assume that procedure  $f$  does not write into objects that are global or have a "lifetime" exceeding the activation of  $f$ .

Although the modifications to T1 to allow procedure calls are straightforward, those for T2 are more complex if T2 does not have bounds for  $y_1, \dots, y_n$  when the call statement is processed (since it is then unable to determine the upper bounds for  $b_1, \dots, b_n$ ). The difficulty is resolved in the restricted case where the output parameters are derived entirely from the input parameters and data belonging to the lowest class  $L$ . In this case, the bounds on  $b_1, \dots, b_n$  are determined by

$l[b_j] := \underline{SS}.l$ , and

$h[b_j] := h[a_1] \oplus \dots \oplus h[a_m] \oplus \underline{SS}.h \quad (1 \leq j \leq n)$ .

Consider now the case where  $f$  writes into objects  $z_1, \dots, z_r$  that are global or have a "lifetime" exceeding the activation of  $f$ . There are two reasons why our present mechanism is inadequate to deal with this case. First, if  $f$  writes into such objects and the call to  $f$  is conditioned on Boolean variables  $e_1, \dots, e_k$ , then  $z_1, \dots, z_r$  must be updated by  $\underline{e_1} \oplus \dots \oplus \underline{e_k}$  even if the call is not executed. To do so requires a considerably more complex mechanism than we have outlined here. Second, if  $f$  writes into global objects, the  $l$  and  $h$  tables computed by T2 for a program  $\phi$  cannot be guaranteed to give lower and upper bounds on the classes of all variables of  $\phi$ , whence the execution of  $\phi$  cannot be guaranteed secure. Clearly further research is needed to determine whether or not it is feasible to support global variables that are bound dynamically.

#### 4.4.4 Interrupts

Following the approach taken in Section 3.4, we consider the alternative whereby all interrupts are inhibited except those specified by an on statement

$\phi = \underline{\text{on}} \ \zeta \ \underline{\text{do}} \ \phi_1$ ,

where  $\zeta$  names an interrupt condition (e.g., overflow or end-file),  $a$  is the variable to which it applies, and  $\phi_1$  is the statement to be invoked when the interrupt occurs.

Since execution of  $\phi_1$  is conditioned on (some property of) variable  $a$ , the operation push( $a$ ) must be performed when  $\phi_1$  is invoked (i.e.,

when variable  $a$  causes an interrupt of type  $\zeta$ ) and pop must be performed when  $\phi_1$  completes and returns. This guarantees that all objects assigned values in  $\phi_1$  are automatically updated by  $a$ .

Just as with any other conditional structure, it is necessary to insert update instructions into a program to account for implicit flows caused by the nonoccurrence of an interrupt. As an example, consider the insecure program used earlier, modified to use on statements; it copies  $a \in \{0,1\}$  to  $b$  while preserving  $b = L$  even when  $a = H$ :

```

on overflow x do c := 0;
on overflow y do b := 0;
b := 1; c := 1;
x := 0; y := 0;
x := a + MAXVALUE;
y := c + MAXVALUE

```

To eliminate the problem, the program must be modified as follows:

```

on overflow x do c := 0;
on overflow y do b := 0;
b := 1; c := 1;
x := 0; y := 0;
x := a + MAXVALUE; push(x); update c; pop;
y := c + MAXVALUE; push(y); update b; pop

```

The example suggests what must be done to deal with the general case. For  $\phi = \text{"on } \zeta \text{ a do } \phi_1\text{"}$ , the certification mechanism must process  $\phi_1$  and determine  $V$ ; it must then identify all points in the program at which " $\zeta \text{ a}$ " can occur and insert there the sequence:

push(a); update V; pop.

(Of course T2 can be used to reduce the update set  $V$ , but the principle is the same.) This obviously poses a serious complication for either T1 or T2: these algorithms must now determine if any declared interrupt can occur during the execution of each statement. At present, our research has yielded little insight into this problem. A solution to

the dynamic binding case will evidently depend on a viable solution to the interrupt problem.

#### 4.4.5 Static Binding

It is possible to extend the basic hardware and software mechanism to guarantee the secure execution of a program that references some statically bound objects. Let  $N^S$  and  $N^d$  denote, respectively, the statically and dynamically bound objects. The hardware extensions include an extension of the tag field of a register (or descriptor) for distinguishing objects in  $N^S$  from those in  $N^d$ . It also includes a modification to the class updating mechanism so that for an object  $a \in N^S$ ,  $a$  is not changed when information flows into  $a$ . It could also verify the validity of the flow; however, this is unnecessary if we assume that some certification program verifies that only valid flows into objects of  $N^S$  can occur.

Since algorithm T1 performs no information flow analysis, it is not easily extended to handle the certification of flows into objects of  $N^S$ . On the other hand, only two modifications to T2 are needed. (See Figure 4.4-1.) They assume that for an object  $a \in N^S$ , an initial (fixed) class is specified on the declaration for  $a$ , and that this class is assigned to  $l[a]$  and  $h[a]$  (which remain constant). The first modification is to the code for processing assignment statements " $b := f(a_1, \dots, a_n)$ ". For  $b \in N^S$  it is necessary to verify the security of the explicit flow from the expression to  $b$ . This is done by testing that the upper bound on  $a_1 @ \dots @ a_n @ HS$  flows into  $b$ . The second modification alters the UPDATE procedure responsible for guaranteeing the security of all implicit flows that arise from selection and iteration statements. Here

## a) Revised Code for Processing Assignment Statements:

```

 $\phi = "b := f(a_1, \dots, a_n)";$ 

begin
  if  $b \in N^d$ 
    then
      begin
         $l[b] := l[a_1] \oplus \dots \oplus l[a_n] \oplus SS.l;$ 
         $h[b] := h[a_1] \oplus \dots \oplus h[a_n] \oplus SS.h$ 
      end
      else if  $h[a_1] \oplus \dots \oplus h[a_n] \oplus SS.h \neq l[b]$  then
        SECURITY ERROR;
      end
       $v := b$ 
    end
  end

```

## b) Revised Code for UPDATE Procedure:

```

UPDATE: procedure (X, Y: powerset variable);

 $Y := \{b \in X \cap N^d \mid SS.h \neq l[b]\};$ 
for all  $b \in Y$  do
  begin  $l[b] := l[b] \oplus SS.l; h[b] := h[b] \oplus SS.h$  end;
if  $\{b \in X \cap N^d \mid SS.h \neq l[b]\} \neq \emptyset$  then SECURITY ERROR
end UPDATE

```

Figure 4.4-1. Modifications to T2 to Support Static Binding.

it is necessary to verify that the upper bound on HS flows into b for all  $b \in N^S$  that are assigned a value in the selection or iteration statement under consideration.

#### 4.4.6 Verify Statement

Consider the following program:

```
max: integer procedure (x, y, z: integer);
```

```
    if x > y then max := x else max := y;
```

```
    if z > max then max := z
```

```
end max.
```

Since the initial values for x, y, and z are not specified, T2 will insert the instruction "update max" after the second if statement so that  $z \rightarrow \text{max}$  is guaranteed to hold when the program terminates ( $x \rightarrow \text{max}$  and  $y \rightarrow \text{max}$  will hold after any execution sequence). Clearly this is not necessary if  $x = y = z$  irrespective of the value of x. However, our present mechanism has no provision for specifying such a constraint unless the value of x is known.

Generalizing, we see that the need for update instructions may be reduced, even when the precise initial classes of the input parameters are unknown, if additional information is available about the flow relations that hold on the classes of these variables. We propose that such additional information be declared by means of a new statement

```
verify  $\xi_1, \dots, \xi_n$ ,
```

where each  $\xi_i$  ( $1 \leq i \leq n$ ) is a "class expression" specifying constraints on the input classes. Figure 4.4-2 shows the syntax of this statement. The operands of a class expression may be either "absolute" or "relative"; absolute operands reference specific security classes (e.g., L)



```

<class op> ::= @ | @
<flow op> ::= → | =
<class factor> ::= <security class> | class ( <ident> )
<class term> ::= <class factor> | <class term> <class op> <class factor>
<class exp> ::= <class term> | <class exp> <flow op> <class term>
<verify list> ::= <class exp> | <verify list> , <class exp>
<verify stmt> ::= verify <verify list>

```

Figure 4.4-2. Syntax of Verify Statement.

and relative operands reference the class of objects (e.g., class(a)).

Examples of verify statements are:

verify class(x) = class(y) = class(z),

verify class(a) = L, class(b) → B,

verify class(a) ⊕ class(b) → class(c) = A.

The first of the above examples corresponds to a relation on classes which, if true, obviates the need for the update instruction in the earlier program.

The class expressions of a verify statement are evaluated on procedure entry, and execution of the procedure is conditioned on all constraints being satisfied. (It might also be possible to use a verify statement to control entry to a block of statements, but we have not explored this possibility.)

Although the verify statement is appealing, further research is needed to determine its viability. Since precise classes are not known for the input variables, it is necessary to find an alternative method for determining upper and lower bounds on the classes of all variables. One possible approach would be to develop an "algebra" of class expressions. Here the entries in the l and h tables would be class expressions, and T2 would manipulate these entries according to the rules of the algebra. For example, consider the following program:

```
sum: procedure (a,b: integer);
     verify class(a) → class(b);
     sum := a + b
end.
```

To compute l[sum], T2 would form the expression

$$l[a] \oplus l[b] \oplus \underline{SS.l} = \underline{class}(a) \oplus \underline{class}(b) \oplus \underline{SS.l}$$

If  $SS.l = L$  initially, this would then be reduced to

class(b)

(since  $L \rightarrow \text{class}(a) \rightarrow \text{class}(b)$ ). The computation of  $H[\text{sum}]$  would be the same. We leave for future research the task of determining the feasibility of this approach (or some other).

#### 4.5 Applications

The applications discussed in Section 3.7 for static binding, such as the confinement problem, are also suitable for the mechanism proposed in this chapter. In fact, more flexible solutions to these problems are in principle possible since the security classes of all objects do not have to be bound at compile-time. For instance, the compiler could certify a program for confinement with respect to some of its input parameters, without knowing their precise security classes (which could vary from one call to the next). However, the increased flexibility may be more than offset by the increased cost for the tagged architecture and the more complex certification procedure. Further research is required to resolve these matters.

## CONCLUSIONS

## Chapter 5

5.1 Summary

In Chapter 2 we developed a model of information flow consisting of three components: an information flow structure specifying security classes and valid flow paths between these classes, states defining possible system configurations, and transition operators specifying state changes. We defined a secure system to be one that did not permit communication along invalid flow paths and showed that security can be obtained by keeping the system in a safe state. An investigation of the mathematical properties of the model led to the conclusion that under certain assumptions, justified by the semantics of the problem, the flow structure forms a lattice. This result is significant in two respects: 1) it permits concise formulations of the security requirements of different systems and 2) it admits efficient implementation mechanisms. It serves as a foundation for the results reported in this thesis.

We examined the security requirements of five increasingly complex systems: 1) single-class, static binding, 2) single-class, dynamic binding, 3) single-class, static and dynamic binding, 4) multi-class, static binding, and 5) multi-class, dynamic binding. The results of this analysis showed that our model provides a unifying theory of secure information flow in that all known systems that restrict information

flow fall into these classes. Of particular interest are the two multi-class cases, previously investigated only by Fenton in the context of a Minsky machine [Fe73,Fe74a,Fe74b]. Fenton proposed a hardware mechanism for performing run-time checks during the execution of a program, proved that his mechanism guarantees security under static binding, but conceded that it did not guarantee security under dynamic binding. In Chapters 3 and 4, we made significant extensions to, and improvements on, these results.

In Chapter 3 we introduced a compile-time mechanism for certifying the secure execution of a program under static binding. We formulated the mechanism in terms of semantic actions to be performed during the analysis phase of compilation. The significant contribution of this work is twofold: 1) It permits the security of a program to be established before it executes. This eliminates the need for costly hardware support and/or inefficient software run-time mechanisms. It also eliminates the possibility of programs purposely causing security violations as a means of leaking information. 2) It demonstrates the feasibility and practicality of constructing automatic program certification mechanisms for proving security properties.

In Chapter 4 we investigated the more complex problem of guaranteeing the secure execution of a program under dynamic binding. We outlined a hardware mechanism, based on tagged architecture, for dynamically binding objects to security classes. We presented two possible compile-time mechanisms for transforming an arbitrary program into an equivalent secure one. Both mechanisms operate by inserting instructions into the original program. The first mechanism is considerably

less complex than the second, but the second produces better code. The significant contribution of this work is that it proves the feasibility of guaranteeing security under dynamic binding; whether or not it is practical, however, is an open question we leave for future study.

## 5.2 Areas of Future Research

1. Application. We have not offered a complete set of guidelines for applying our model and mechanisms to the design of a system. For example, we have not suggested a method for constructing a suitable set of security classes when the usual government and military classification system is not applicable. We have not specified the design of an "operating system kernel" that guarantees the correct implementation of the mechanism. We have said little about the applicability of the mechanism to data base systems.
2. Dynamic Binding. Although we demonstrated the feasibility of certifying the security of a program under dynamic binding, we left many issues unresolved. For example, the question of whether or not it is possible to construct a transformation procedure producing optimal code (in the sense of inserting the minimal number of update instructions) is open. We did not provide an adequate mechanism for dealing with interrupts. We left unresolved the problem of handling global variables. We barely hinted at a method for dealing with the verify statement.
3. Program Development and Debugging. The program certification procedures analyze the flow of information in a program for the purpose of verifying security. The procedures may also be applicable

to a completely different problem: program development and debugging. For example, it would be possible to determine the variables affected by a change to some other variable. It may be possible to determine the best modularization of a program in the sense of minimizing the number of variables common to several modules, while keeping the code expansion due to modularization as small as possible [Ba73].

# LIST OF REFERENCES



## LIST OF REFERENCES

- [Ab69] Abbott, J. C., Sets, Lattices, and Boolean Algebras, Allyn and Bacon, 1969.
- [Al70] Allen, F. E., "Control Flow Analysis," Proc. of a Symposium on Compiler Optimization, SIGPLAN Notices, 5, 7, July 1970.
- [An74] Andrews, G. R., "COPS - A Protection Mechanism for Computer Systems," Ph.D. Dissertation, Univ. of Wash., July 1974.
- [Ba73] Bayer, R., "On Program Volume and Program Modularization," CSD TR 105, Purdue Univ., Sept. 1973.
- [BB74] Bell, D. E. and Burke, E. L., "A Software Validation Technique for Certification: The Methodology," MTR-2932, 1, The MITRE Corp., Bedford, Mass., Nov. 1974.
- [BL73a] Bell, D. E. and LaPadula, L. J., "Secure Computer Systems: Mathematical Foundations", MTR-2547, ESD-TR-73-278-I, AD-770-768, 1, The MITRE Corp., Bedford, Mass., Mar. 1973.
- [BL73b] Bell, D. E. and LaPadula, L. J., "Secure Computer Systems: Mathematical Foundations and Model," M74-244, The MITRE Corp., Bedford, Mass., May 1973.
- [Bi67] Birkhoff, G. Lattice Theory, Amer. Math. Soc. Col. Pub., XXV, 3rd ed., 1967.
- [CH75] Cash, J., Haseman, W. D., and Winston, A. B., "Security for the GPLAN System," Purdue Univ., Jan. 1975.
- [CD73] Coffman, E. G. and Denning, P. J., Operating System Principles, Prentice Hall, 1973.
- [CM72] Conway, R. W., Maxwell, W. L., and Morgan, H. L., "Implementation of Security Structures in Information Systems," Comm. ACM, 15, 4, Apr. 1972, 211-220.
- [DD74] Denning, D. E., Denning, P. J., and Graham, G. S., "Selectively Confined Subsystems," Proc. International Workshop on Protection in Operating Systems, IRIA, Aug. 1974.

- [De71] Denning, P. J., "Third Generation Computer Systems," Computing Surveys, 3, 4, Dec. 1971.
- [De74] Denning, P. J., "Structuring Operating Systems for Reliability," CSD TR 119, Purdue Univ., 1974.
- [DV66] Dennis, J. B. and VanHorn, E. C., "Programming Semantics for Multiprogrammed Computations," Comm. ACM, 9, 3, Mar. 1966.
- [EL72] Elspas, B., Levit, K., Waldinger, R. and Waksman, A., "An Assessment of Techniques for Proving Program Correctness," Computing Surveys, 4, 2, June 1972, 97-147.
- [En72] England, D. M., "Architectural Features of System 250", Infotech State of the Art Report 14, Operating Systems, 1972.
- [EL67] Evans, D. C. and LeClerc, J. Y., "Address Mapping and the Control of Access in an Interactive Computer," AFIPS Conf. Proc., 30, SJCC, 1967, 23-30.
- [Fa68] Fabry, R. S., "Preliminary Description of A Supervisor for a Machine Oriented Around Capabilities," ICR Quarterly Report 18, Univ. of Chicago, Aug. 1968.
- [Fa71] Fabry, R. S., "List Structured Addressing," Ph.D. Thesis, Univ. of Chicago, Mar. 1971.
- [Fa73] Fabry, R. S., "Dynamic Verification of Operating System Decisions," Comm. ACM, 16, 11, Nov. 1973, 659-668.
- [Fa74] Fabry, R. S., "The Case for Capability Based Computers," Comm. ACM, 17, 6, June 1974.
- [Fe73] Fenton, J. S., "Information Protection Systems," Ph.D. Dissertation, Univ. of Cambridge, 1973.
- [Fe74a] Fenton, J. S., "Memoryless Subsystems," Computer Journal, 17, 2, Feb. 1974, 143-147.
- [Fe74b] Fenton, J. S., "An Abstract Computer Model Demonstrating Directional Information Flow," Univ. of Cambridge, 1974.
- [Fu73] Feustel, E. A., "On the Advantages of Tagged Architecture," IEEE Trans. on Computers, C-22, 7, July 1973, 646-656.
- [Fl67] Floyd, R. W., "Assigning Meanings to Programs," Math. Aspects of Computer Science, 19, Amer. Math. Soc., (ed. J. T. Schwartz), 1967, 19-32.
- [Ga72] Gaines, R. S., "An Operating System Based on the Concept of a Supervisory Computer," Comm. ACM, 15, 3, Mar. 1972, 150-156.

- [G373] George, J. E. and Sagar, G. R., "Variables - Binding and Protection," SIGPLAN Notices, 8, 12, Dec. 1973.
- [GD72] Graham, G. S. and Denning, P. J., "Protection - Principles and Practice," AFIPS Conf. Proc., 40, SJCC, 1972, 417-429.
- [Gr68] Graham, R. M., "Protection in an Information Processing Utility," Comm. ACM, 11, 5, May 1968, 365-369.
- [Gi71] Gries, D., Compiler Construction for Digital Computers, Wiley, 1971.
- [Ha70] Hansen, P. B., "The Nucleus of a Multiprogramming System," Comm. ACM, 13, 4, Apr. 1970, 238-250.
- [HR75] Harrison, M. A., Ruzzo, W. L., and Ullman, J. D., "On Protection in Operating Systems," Univ. of Ca., Berkeley, Nov. 1974.
- [Ho74] Hoare, C. A. R., "Monitors: An Operating System Structuring Concept," Comm. ACM, 17, 10, Oct. 1974.
- [Hf69] Hoffman, L. J., "Computers and Privacy: A Survey," Computing Surveys, 1, 2, June 1969, 85-104.
- [Hf71] Hoffman, L. J., "The Formulary Model for Flexible Privacy and Access Control," AFIPS Conf. Proc., 39, FJCC, 1971, 587-601.
- [IB68] IBM, "System/360 Principles of Operation," IBM Report No. GC28-6821, Sept. 1968.
- [IB71] IBM, "System/360 PL/I (F) Language Reference Manual," IBM Report No. GC28-8201-3, 1971.
- [Il68] Iliffe, J. K., Basic Machine Principles, MacDonald Elsevier, 1968.
- [Jo73] Jones, A. K., "Protection in Programmed Systems," Ph.D. Thesis, Carnegie-Mellon Univ., June 1973.
- [La69] Lampson, B. W., "Dynamic Protection Structures," AFIPS Conf. Proc., 35, FJCC, 1969, 27-38.
- [La71] Lampson, B. W., "Protection," Proc. Fifth Princeton Symposium on Information Sciences and Systems, Princeton Univ., Mar. 1971, 437-443.
- [La73] Lampson, B. W., "A Note on the Confinement Problem," Comm. ACM, 16, 10, Oct. 1973, 613-615.
- [Li73] Lindsay, B., "Suggestions for an Extensible Capability-Based Machine Architecture," International Workshop on Computer Architecture, Grenoble, June 1973.

- [LM69] Lowry, E. S. and Medlock, C. W., "Object Code Optimization," Comm. ACM, 12, 1, Jan. 1969, 13-22.
- [Mi67] Minsky, M. L., Computation; Finite and Infinite Machines, Prentice-Hall, 1967.
- [Mo74] Moore, C. G. III, "Potential Capabilities in ALGOL-like Programs," TR 74-211, Dept. of Computer Science, Cornell Univ., Sept. 1974.
- [MC74] Moore, C. G. III and Conway, R., "Program Predictability and Data Security," TR 74-212, Dept. of Computer Science, Cornell Univ., Sept. 1974.
- [Mr73] Morris, J. H., "Protection in Programming Languages," Comm. ACM, 16, 1, Jan. 1973, 15-21.
- [Na63] Naur, P. et.al., "Revised Report of the Algorithmic Language Algol-60," Comm. ACM, 6, 1, 1963.
- [Na66] Naur, P., "Proof of Algorithms by General Snapshots," BIT, 6, 4, 1966, 310-316.
- [Ne72] Needham, R. M., "Protection Systems and Protection Implementations," AFIPS Conf. Proc., 41, FJCC, 1972, 572-578.
- [NF74] Neumann, P. G. and Fabry, R. S., "On the Design of Provably Secure Operating Systems," Proc. of the International Workshop on Protection in Operating Systems, IRIA, Aug. 1974.
- [Or72] Organick, E. I., The MULTICS System: An Examination of Its Structure, MIT Press, 1972.
- [Po73] Popek, G. J., "Access Control Models," Ph.D. Thesis, Harvard Univ., 1973.
- [Po74] Popek, G. J., "Protection Structures," IEEE Computer, June 1974.
- [PK74] Popek, G. J. and Kline, C. S., "Verifiable Secure Operating System Software," AFIPS Conf. Proc., 43, NCC, 1974.
- [Ro74] Rotenberg, L. J., "Making Computers Keep Secrets," Ph.D. Thesis, MIT, MAC TR-115, Feb. 1974.
- [Sc72] Schroeder, M. D., "Cocperation of Mutually Suspicious Subsystems in a Computer Utility," Ph. D. Thesis, MIT, TR-104, Sept. 1972.
- [SS72] Schroeder, M. D. and Saltzer, J. H., "A Hardware Architecture for Implementing Protection Rings," Comm. ACM, 15, 3, Mar. 1972.

- [Sc71] Scott, D., "The Lattice of Flow Diagrams," Semantics of Algorithmic Languages, Springer Lecture Notes in Math., 188, (ed. E. Engeler), 1971, 311-366.
- [St73] Stone, H. S., Discrete Mathematical Structures and Their Applications, Science Research Associates, 1973.
- [Su74] Sturgis, H. E., "A Postmortum for a Time Sharing System," Ph.D. Thesis, Univ. of Ca., Berkeley, Jan. 1974.
- [Ta74] Tarjan, R., "Finding Dominators in Directed Graphs," SIAM J. Computing, 3, 1, Mar. 1974.
- [Ts73] Tsichritzis, D., "Protection Implementation," Proc. of the 7th Annual Princeton Conf. on Information Sciences and Systems, Princeton Univ., Mar. 1973.
- [Ts74] Tsichritzis, D., "A Note on Protection in Data Base Systems," Proc. of the International Workshop on Protection in Operating Systems, IRIA, Aug. 1974.
- [Wa74] Walter, K. G. et. al., "Modeling the Security Interface," Jennings Computing Center, Case Western Reserve Univ., Report No. 1158, Aug. 1974.
- [We69] Weissman, C., "Security Controls in the ADEPT-50 Time-Sharing System," AFIPS Conf. Proc., 35, FJCC, 1969, 119-133.
- [Wi72] Wilkes, M. V., Time Sharing Computer Systems, Amer. Elsevier, 2nd ed., 1972.
- [Wr71] Wirth, N., "The Programming Language Pascal," Acta Informatica, 1, 1971, 35-63.
- [Wu74] Wulf et. al., "HYDRA: The Kernel of a Multiprocessor System," Comm. ACM, 17, 6, June 1974.

VITA

## VITA

Dorothy Elizabeth Robling Denning was born in Grand Rapids, Michigan on August 12, 1945. She graduated from Ottawa Hills High School there in 1963. From 1963-1969 she studied mathematics at the University of Michigan, receiving the A.B. degree with distinction in 1967 and the A.M. degree in 1969. While at Michigan, she was employed by the Radio Astronomy Observatory as an Assistant Research Mathematician, by the Computing Center as a Consultant, and by the Computer and Communication Sciences Department as a Teaching Fellow.

From 1969-1972 she was a Systems Programmer at the University of Rochester Computing Center. While at Rochester, she taught a course in compiler design for the Graduate School of Management, and she developed and taught courses in programming languages and compiler design for the Department of Electrical Engineering. She was also an active participant at SHARE meetings, heading the Command Languages Project in 1972.

In 1972 she left Rochester to pursue a Ph.D. degree in Computer Science at Purdue University. From 1972-1974 she was a Graduate Instructor, teaching a course in programming languages. In 1974 she was awarded a Fellowship from IBM. Mrs. Denning has published papers in both Radio Astronomy and Computer Science. She is a member of ACM, Alpha Lambda Delta, and Phi Kappa Phi.