

SRI International

May 5, 2005

Noninterference, Transitivity, and Channel-Control Security Policies¹

John Rushby
Computer Science Laboratory
SRI International



Abstract

We consider noninterference formulations of security policies [7] in which the “interferes” relation is intransitive. Such policies provide a formal basis for several real security concerns, such as channel control [17, 18], and assured pipelines [4]. We show that the appropriate formulation of noninterference for the intransitive case is that developed by Haigh and Young for “multidomain security” (MDS) [9, 10]. We construct an “unwinding theorem” [8] for intransitive policies and show that it differs significantly from that of Haigh and Young. We argue that their theorem is incorrect. A companion report [22] presents a mechanically-checked formal specification and verification of our unwinding theorem.

We consider the relationship between transitive and intransitive formulations of security. We show that the standard formulations of noninterference and unwinding [7, 8] correspond exactly to our intransitive formulations, specialized to the transitive case. We show that transitive policies are precisely the “multilevel security” (MLS) policies, and that any MLS secure system satisfies the conditions of the unwinding theorem.

We also consider the relationship between noninterference formulations of security and access control formulations, and we identify the “reference monitor assumptions” that play a crucial role in establishing the soundness of access control implementations.

Contents

1	Introduction	1
2	Basic Noninterference	6
2.1	Access Control Interpretations	12
3	Noninterference and Transitivity	17
3.1	Properties of Transitive Policies	19
4	Intransitive Noninterference	25
5	Comparisons among the Formulations	33
5.1	Intransitive vs. Standard Noninterference	33
5.2	Comparison with Haigh and Young's Formulation	35
6	Summary and Conclusions	40
	Bibliography	41
	References	41

Chapter 1

Introduction

The concept of noninterference was introduced by Goguen and Meseguer [7] in order to provide a formal foundation for the specification and analysis of security policies and the mechanisms that enforce them. Apart from the work of Feiertag, Levitt, and Robinson [6]—which can be seen as a precursor to that of Goguen and Meseguer—previous efforts, among which those of Bell and La Padula [3] were the most influential, formulated security in terms of access control. Access control formulations suffer from a number of difficulties. First, because they are described in terms of a mechanism for enforcing security, they provide no guidance in circumstances where those mechanisms prove inadequate. Second, it is easy to construct perverse interpretations of access control policies that satisfy the letter, but not the intent of the policy, to the point of being obviously unsecure [13,14]. The proponents of access control formulations counter that interpretations or implementations must be “faithful representations” of the model, but they provide no formal definition of that term.

In contrast, noninterference formulations are pure statements of policy, with no commitment to a specific mechanism for enforcing them—although techniques have been developed for demonstrating that specific mechanisms enforce given noninterference policies. Secondly, noninterference policies have the form of a logical theory; *any* implementation that is a model for the theory (i.e., validates its axioms) will be secure.

The idea of noninterference is really rather simple: a security domain u is noninterfering with domain v if no action performed by u can influence subsequent outputs seen by v . Noninterference has been quite successful in providing formal underpinnings for military multilevel security policies and for the methods of verifying their implementations [8,20].

There are, however, a number of practical security problems that seem beyond the scope of noninterference formulations. One of these is “channel-control,” first formulated by Rushby [17, 18]. Channel control security policies can be represented by directed graphs, where nodes represent security domains and edges indicate the direct information flows that are allowed. The paradigmatic example of a channel-control problem is a controller for end-to-end encryption, as portrayed in Figure 1.1 [1, 17].

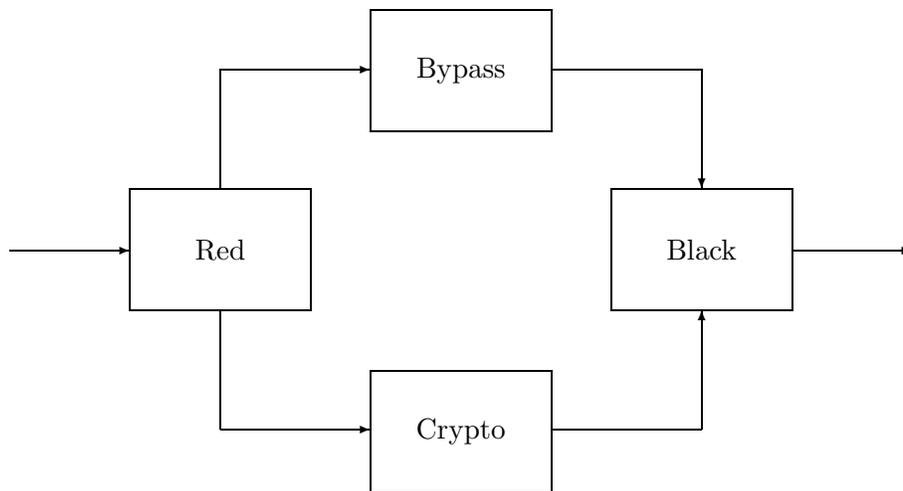


Figure 1.1: End-to-end encryption controller

Plaintext messages arrive at the Red side of the controller; their bodies are sent through the encryption device (Crypto); their headers, which must remain in plaintext so that network switches can interpret them, are sent through the Bypass. Headers and encrypted bodies are reassembled in the Black side and sent out onto the network. The security policy we would like to specify here is the requirement that the *only* channels for information flow from Red to Black must be those through the Crypto and the Bypass.¹ Thus, an important characteristic of many channel control policies is that the edges indicating allowed information flows are not transitive: information is allowed to flow from Red to Black via the Crypto and Bypass, but cannot do so directly.

Another example is shown in Figure 1.2, where transitive and intransitive elements are combined. The edges to the left represent the conventional transitive flow

¹It is a separate problem to specify what those components must do.

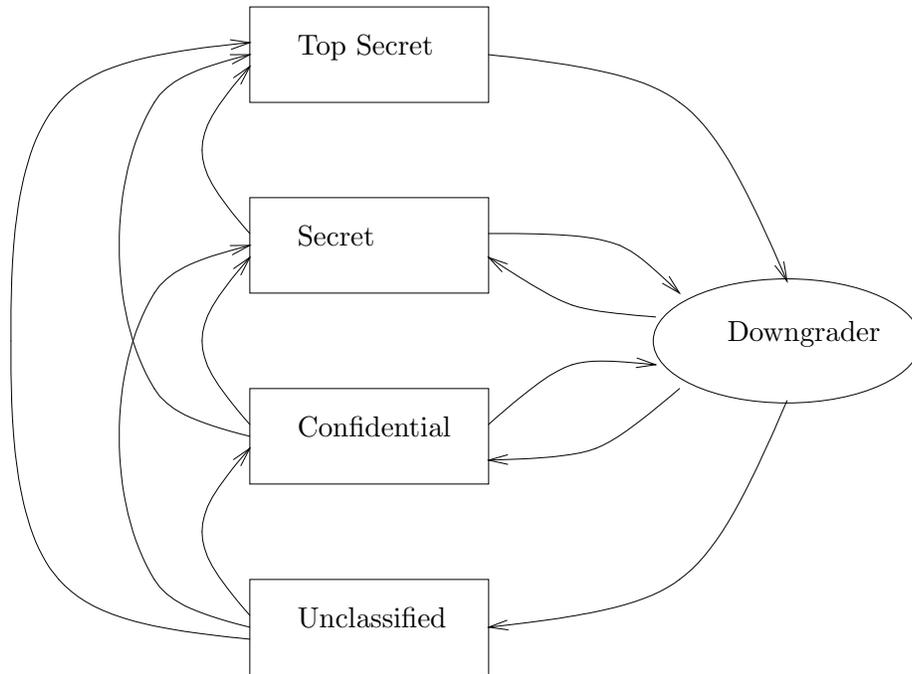


Figure 1.2: Controlled downgrading

relations between the classification levels used in the USA. On the right are edges to and from a special Downgrader domain that are intransitive. The flows represented by these edges are intransitive because, although information can flow, for example, from the Top Secret to the Confidential domain via the Downgrader, it cannot flow directly from Top Secret to Confidential. Thus, information can flow “upward” in security level without restriction, but only flow “downward” through the mediation of the presumably trusted Downgrader domain.

Channel control policies such as those just described seem able to specify a number of security concerns that are beyond the reach of standard security modeling techniques. Boebert and Kain have argued persuasively [4] that a variation on channel-control called “type enforcement” can be used to solve many vexing security problems. A worthwhile challenge, then, is to find an adequate formal foundation for channel-control policies and their ilk.

An early attempt to provide a formal method for verifying, though not specifying, channel-control policies was based on a technique for verifying complete separation [17,19]. The idea was to remove the mechanisms that provided the intended

channels, and then prove that the components of the resulting system were isolated. This approach has recently been shown to be subtly flawed [11], although the method for establishing complete separation has survived fairly intensive scrutiny [12, 24] with only minor emendations.

The success of noninterference formulations in explicating multilevel security policies naturally invites consideration of a noninterference foundation for channel-control. This presents quite a challenge, however. For example, it is clear the Red side of the encryption controller of Figure 1.1 necessarily interferes with the Black; we need to find a way of saying that this interference must only occur through the mediation of the Crypto or the Bypass. Goguen and Meseguer proposed a way of doing this in their original paper on noninterference [7], but the method was incorrect. Goguen and Meseguer recognized this in their second paper on the subject [8] and they introduced several extensions to the basic formulation of noninterference. However, the first really satisfactory treatment of intransitive noninterference policies was given by Haigh and Young [9], with a more polished version the following year [10]. They showed that it was necessary to consider the complete sequence of actions performed subsequent to a given action in order to determine whether that action is allowed to interfere with another domain. For example, an action by the Red domain is allowed to interfere with the Black domain only if there is some intervening action from either the Crypto or the Bypass.

The main purpose of this reportpart of the report is to show that channel-control security policies can be modeled by noninterference policies in which the “interferes” relation is intransitive and in which the definition used for “interference” is that of Haigh and Young. We also show that conventional multilevel policies are a special case of channel-control policies, corresponding to those whose “interferes” relation is transitive. We show that our results collapse to the familiar ones in this special case, thereby providing some additional evidence for their veracity.

An important component of noninterference formulations of security are the “unwinding” theorems [8, 10] that establish conditions on the behavior of individual actions sufficient to ensure security of the system. These unwinding theorems provide the basis of practical methods for formally verifying that an implementation satisfies a noninterference security policy. The main result of this reportpart is the derivation of an unwinding theorem for the channel-control case. We show that this theorem differs significantly from that of Haigh and Young and we argue that their result is incorrect.

The development of noninterference and unwinding for the channel-control case is surprisingly intricate, and in view of the previous history of failed attempts, we present our development rather formally and describe the proofs in detail. An

appendixPart III of this report describes the formal verification of our main theorem using the EHDM formal specification and verification system [23].

This reportpart of the report is organized as follows. In the next chapter we present a development of the standard noninterference formulation of security, and then consider the relationship between noninterference security policies and access control policies. This development is structured to provide a model and a basis for comparison with the generalization given later. Chapter 3 examines the case of intransitive noninterference policies and argues that these have no useful interpretation within the standard formulation of noninterference. The second part of Chapter 3 examines the special properties of transitive policies and shows that they are identical to classical multilevel security. Chapter 4 presents a modified formulation of noninterference that does provide a meaningful interpretation to intransitive policies and derives an unwinding theorem for that interpretation. Chapter 5 compares the transitive and intransitive noninterference formulations, and compares our unwinding theorem with that of Haigh and Young. Chapter 6 presents our conclusions. The appendix presents a formal specification and verification of our Intransitive Unwinding Theorem that has been mechanically checked using the EHDM Verification System [23].

Chapter 2

Basic Noninterference

In this chapter we present the core of Goguen and Meseguer’s formulation of security in terms of noninterference assertions [7], and the unwinding theorem [8] that underlies the associated verification techniques. Our notation differs considerably from that of Goguen and Meseguer and is more similar to that of later authors, such as Haigh and Young [10].

We model a computer system by a conventional finite-state automaton.

Definition 1 A *system* (or *machine*) M is composed of

- a set S of *states*, with an *initial state* $s_0 \in S$,
- a set A of *actions*, and
- a set O of *outputs*,

together with the functions *step* and *output*:

- *step*: $S \times A \rightarrow S$,
- *output*: $S \times A \rightarrow O$.

We generally use the letters $\dots s, t, \dots$ to denote states, letters a, b, \dots from the front of the alphabet to denote actions, and Greek letters α, β, \dots to denote sequences of actions.

Actions can be thought of as “inputs,” or “commands,” or “instructions” to be performed by the machine; $step(s, a)$ denotes the next state of the system when action a is applied in state s , while $output(s, a)$ denotes the result returned by the action.

We derive a function run

- $run: S \times A^* \rightarrow S$,

the natural extension of $step$ to sequences of actions, by the equations

$$\begin{aligned} run(s, \Lambda) &= s, \text{ and} \\ run(s, a \circ \alpha) &= run(step(s, a), \alpha), \end{aligned}$$

where Λ denotes the empty sequence and \circ denotes concatenation.¹

In order to discuss security, we must assume some set of security “domains” and a policy that restricts the allowable flow of information among those domains. The agents or subjects of a particular security domain interact with the system by presenting it with actions, and observing the results obtained. Thus we assume

- a set D of security domains, and
- a function $dom: A \rightarrow D$ that associates a security domain with each action.

We use letters $\dots u, v, w, \dots$ to denote domains.

A *security policy* is specified by a reflexive relation \rightsquigarrow on D . We use $\not\rightsquigarrow$ to denote the complement relation, that is

$$\not\rightsquigarrow = (D \times D) \setminus \rightsquigarrow$$

where \setminus denotes set difference. We speak of \rightsquigarrow and $\not\rightsquigarrow$ as the *interference* and *noninterference* relations, respectively. A policy is said to be *transitive* if its interference relation has that property. \square

We wish to define security in terms of information flow, so the next step is to capture the idea of the “flow of information” formally. The key observation is that information can be said to flow from a domain u to a domain v exactly when actions submitted by domain u cause the behavior of the system perceived by domain v to be different from that perceived when those actions are not present. We therefore define a function that removes, or “purges,” from an action sequence all those actions submitted by domains that are required to be noninterfering with a given domain. The machine is secure if a given domain v is unable to distinguish between the state

¹Observe that we define run using right recursion: that is, we specify $run(s, a \circ \alpha) = run(step(s, a), \alpha)$, rather than the more common left recursive form $run(s, \alpha \circ a) = step(run(s, \alpha), a)$. The choice of right recursion slightly complicates the proof of the basic unwinding theorem (Theorem 1); we employ it here for consistency with the later, more complex development in which its use is essential.

of the machine after it has processed a given action sequence, and the state after processing the same sequence purged of actions required to be noninterfering with v .

Definition 2 For $v \in D$ and α an action sequence in A^* , we define $\text{purge}(\alpha, v)$ to be the subsequence of α formed by deleting all actions associated with domains u such that $u \not\sim v$, that is:

$$\begin{aligned} \text{purge}(\Lambda, v) &= \Lambda \\ \text{purge}(a \circ \alpha, v) &= \begin{cases} a \circ \text{purge}(\alpha, v) & \text{if } \text{dom}(a) \sim v \\ \text{purge}(\alpha, v) & \text{otherwise.} \end{cases} \end{aligned}$$

We identify security with the requirement that

$$\text{output}(\text{run}(s_0, \alpha), a) = \text{output}(\text{run}(s_0, \text{purge}(\alpha, \text{dom}(a))), a).$$

Because we frequently use expressions of the form $\text{output}(\text{run}(s_0, \alpha), a)$, it is convenient to first introduce the functions do and test to abbreviate these forms:

- $\text{do}: A^* \rightarrow S$
- $\text{test}: A^* \times A \rightarrow O$

where

$$\begin{aligned} \text{do}(\alpha) &= \text{run}(s_0, \alpha), \text{ and} \\ \text{test}(\alpha, a) &= \text{output}(\text{do}(\alpha), a). \end{aligned}$$

Then we say a system is *secure* for the policy \sim if

$$\text{test}(\alpha, a) = \text{test}(\text{purge}(\alpha, \text{dom}(a)), a).^2$$

□

The intuition here is that the machine starts off in the initial state s_0 and is presented with a sequence $\alpha \in A^*$ of actions. This causes the machine to produce a series of outputs and to progress through a series of states, eventually reaching the state $\text{do}(\alpha)$. At that point the action a is performed, and the corresponding output $\text{test}(\alpha, a)$ is observed. We can think of presentation of the action a and

²Formulas such as these are to be read as universally quantified over their free variables (here a and α).

observation of its output as an experiment performed by $dom(a)$ in order to learn something about the action sequence α . If $dom(a)$ can distinguish between the action sequences α and $purge(\alpha, dom(a))$ by such experiments, then an action by some domain $u \not\rightsquigarrow dom(a)$ has “interfered” with $dom(a)$ and the system is not secure with respect to policies that specify $u \not\rightsquigarrow dom(a)$.

There are several plausible variations on this notion of security. For example, rather than restricting $dom(a)$ to observe only the individual outputs $test(\alpha, a)$, and $test(purge(\alpha, dom(a)), a)$ in its attempt to distinguish α from $purge(\alpha, dom(a))$, we could allow the whole sequence of outputs produced by actions b in α satisfying $dom(b) \rightsquigarrow dom(a)$ (i.e., the outputs of the actions in α which $dom(a)$ can legitimately observe) to be considered. It is fairly straightforward to prove that such variations are equivalent to the definition used here.

The noninterference definition of security is expressed in terms of sequences of actions and state transitions; in order to obtain straightforward techniques for verifying the security of systems, we would like to derive conditions on individual state transitions. The first step in this development is to partition the states of the system into equivalence classes that all “appear identical” to a given domain. The verification technique will then be to prove that each domain’s view of the system is unaffected by the actions of domains that are required to be noninterfering with it.

Definition 3 A system M is *view-partitioned* if, for each domain $u \in D$, there is an equivalence relation \sim^u on S . These equivalence relations are said to be *output consistent* if

$$s \stackrel{dom(a)}{\sim} t \supset output(s, a) = output(t, a).^3$$

□

Output consistency is required in order to ensure that two states s and t that appear identical to domain u really are indistinguishable in terms of the outputs they produce in response to actions from u .

The definition of security requires that the outputs seen by one domain are unaffected by the actions of other domains that are required to be noninterfering with the first. The next result shows that, for an output consistent system, security is achieved if “views” are similarly unaffected.

Lemma 1 *Let \rightsquigarrow be a policy and M a view-partitioned, output consistent system such that,*

$$do(\alpha) \stackrel{u}{\sim} do(purge(\alpha, u)).$$

³We use \supset to denote implication.

Then M is secure for \rightsquigarrow .

Proof: Setting $u = \text{dom}(a)$ in the statement of the lemma gives

$$\text{do}(\alpha) \stackrel{\text{dom}(a)}{\rightsquigarrow} \text{do}(\text{purge}(\alpha, \text{dom}(a))),$$

and output consistency then provides

$$\text{output}(\text{do}(\alpha), a) = \text{output}(\text{do}(\text{purge}(\alpha, \text{dom}(a))), a).$$

But this is simply

$$\text{test}(\alpha, a) = \text{test}(\text{purge}(\alpha, \text{dom}(a)), a),$$

which is the definition of security for \rightsquigarrow given by Definition 2. \square

Next, we define constraints on individual state transitions.

Definition 4 Let M be a view-partitioned system and \rightsquigarrow a policy. We say that M *locally respects* \rightsquigarrow if

$$\text{dom}(a) \not\rightsquigarrow u \supset s \stackrel{u}{\rightsquigarrow} \text{step}(s, a)$$

and that M is *step consistent* if

$$s \stackrel{u}{\rightsquigarrow} t \supset \text{step}(s, a) \stackrel{u}{\rightsquigarrow} \text{step}(t, a).$$

\square

We now have the local conditions on individual state transitions that are sufficient to guarantee security. This result is a version of the unwinding theorem of Goguen and Meseguer [8].

Theorem 1 (Unwinding Theorem) *Let \rightsquigarrow be a policy and M a view-partitioned system that is*

1. *output consistent,*
2. *step consistent, and*
3. *locally respects \rightsquigarrow .*

Then M is secure for \rightsquigarrow .

Proof: We use proof by induction on the length of α to establish

$$s \stackrel{u}{\sim} t \supset \text{run}(s, \alpha) \stackrel{u}{\sim} \text{run}(t, \text{purge}(\alpha, u)). \quad (2.1)$$

The basis is the case $\alpha = \Lambda$ and is elementary. For the inductive step, we assume the inductive hypothesis for α of length n and consider $a \circ \alpha$. By definition,

$$\text{run}(s, a \circ \alpha) = \text{run}(\text{step}(s, a), \alpha). \quad (2.2)$$

For $\text{run}(t, \text{purge}(a \circ \alpha, u))$, there are two cases to consider.

Case 1: $\text{dom}(a) \rightsquigarrow u$. In this case, the definition of *purge* provides

$$\text{run}(t, \text{purge}(a \circ \alpha, u)) = \text{run}(t, a \circ \text{purge}(\alpha, u)),$$

and the right hand side expands to give

$$\text{run}(t, \text{purge}(a \circ \alpha, u)) = \text{run}(\text{step}(t, a), \text{purge}(\alpha, u)). \quad (2.3)$$

Since $s \stackrel{u}{\sim} t$ and the system is step consistent, it follows that

$$\text{step}(s, a) \stackrel{u}{\sim} \text{step}(t, a)$$

and the inductive hypothesis then gives

$$\text{run}(\text{step}(s, a), \alpha) \stackrel{u}{\sim} \text{run}(\text{step}(t, a), \text{purge}(\alpha, u))$$

which, by virtue of (2.2) and (2.3), completes the inductive step in this case.

Case 2: $\text{dom}(a) \not\rightsquigarrow u$. In this case, the definition of *purge* provides

$$\text{run}(t, \text{purge}(a \circ \alpha, u)) = \text{run}(t, \text{purge}(\alpha, u)) \quad (2.4)$$

and the facts that $\text{dom}(a) \not\rightsquigarrow u$ and that M locally respects \rightsquigarrow ensure

$$s \stackrel{u}{\sim} \text{step}(s, a).$$

Since $s \stackrel{u}{\sim} t$ and $\stackrel{u}{\sim}$ is an equivalence relation, the latter provides

$$\text{step}(s, a) \stackrel{u}{\sim} t$$

and the inductive hypothesis then gives

$$\text{run}(\text{step}(s, a), \alpha) \stackrel{v}{\sim} \text{run}(t, \text{purge}(\alpha, u)),$$

which, by virtue of (2.2) and (2.4), completes the inductive step.

In order to complete the proof, we take $s = t = s_0$ in 2.1 to obtain

$$do(\alpha) \stackrel{u}{\sim} do(purge(\alpha, u))$$

and then, since M is output consistent, invoke Lemma 1 to complete the proof. \square

The unwinding theorem is important because it provides a basis for practical methods for verifying systems that enforce noninterference policies, and also serves to relate noninterference policies to access control mechanisms. We illustrate the latter point by using the unwinding theorem to establish the security of a simple access control mechanism.

2.1 Access Control Interpretations

In order to consider access control mechanisms formally, we need a more elaborate system model. First of all, we need to impose some internal structure on the system state, supposing it to be composed of individual storage locations, or “objects,” each of which has a name and a value. The name of each location is fixed, but its value may change from one state to another. Access control functions determine whether a given security domain may “observe” or “alter” the values in given storage locations. We collect these ideas together and introduce convenient notation in the following definition.

Definition 5 A machine has a *structured state* if there exist

- a set N of *names*,
- a set V of *values*, and a function
- *contents*: $S \times N \rightarrow V$

with the interpretation that $contents(s, n)$ is the value of the object named n in state s . In addition, we require functions

- *observe*: $D \rightarrow \mathcal{P}(N)$, where \mathcal{P} denotes powerset, and
- *alter*: $D \rightarrow \mathcal{P}(N)$

with the interpretation that $observe(u)$ is the set of locations whose values can be observed by domain u , while $alter(u)$ is the set of locations whose values can be changed by u . These functions encode the “access control matrix” that represents the access control policy of the system. An access control policy is enforced when the behavior of the system matches the intended interpretation of the *observe* and *alter* functions. This requires the following three conditions to be satisfied:

Reference Monitor Assumptions

1. First, for $u \in D$ define the relation $\overset{u}{\sim}$ on states by

$$s \overset{u}{\sim} t \text{ iff } (\forall n \in \text{observe}(u): \text{contents}(s, n) = \text{contents}(t, n)).$$

Then, in order for the output of an action a to depend only on the values of objects to which $\text{dom}(a)$ has *observe* access, we require:

$$s \overset{\text{dom}(a)}{\sim} t \supset \text{output}(s, a) = \text{output}(t, a).$$

2. Next, when an action a transforms the system from one state to another, the new values of all changed objects must depend only on the values of objects to which $\text{dom}(a)$ has *observe* access. That is:

$$\begin{aligned} s \overset{\text{dom}(a)}{\sim} t \wedge (\text{contents}(\text{step}(s, a), n) \neq \text{contents}(s, n) & \quad (2.5) \\ \vee \text{contents}(\text{step}(t, a), n) \neq \text{contents}(t, n)) & \\ \supset \text{contents}(\text{step}(s, a), n) = \text{contents}(\text{step}(t, a), n). & \end{aligned}$$

This condition is rather difficult; we discuss it following the complete definition.

3. Finally, if an action a changes the value of object n , then $\text{dom}(a)$ must have *alter* access to n :

$$\text{contents}(\text{step}(s, a), n) \neq \text{contents}(s, n) \supset n \in \text{alter}(\text{dom}(a)).$$

These three conditions are called the ‘‘Reference Monitor Assumptions’’ since they capture the assumptions on the ‘‘reference monitor’’ that performs the access control function in any concrete instantiation of the theory. \square

The second of the Reference Monitor Assumptions is somewhat tricky, so we will now explain it in more detail. The goal is to specify that if action a changes the value of location n , then the only information that may be used in creating the new value should be that provided in variables to which $\text{dom}(a)$ has *observe* access. Thus, if two states s and t have the same values in all the locations to which $\text{dom}(a)$ has *observe* access (i.e., if $s \overset{\text{dom}(a)}{\sim} t$), then it seems we should specify

$$\text{contents}(\text{step}(s, a), n) = \text{contents}(\text{step}(t, a), n) \quad (2.6)$$

for all locations n . The flaw in this specification is that if $\text{dom}(a)$ does *not* have observe access to n , then $s \stackrel{\text{dom}(a)}{\sim} t$ does not prevent $\text{contents}(s, n) \neq \text{contents}(t, n)$. If a does not change the value of location n we will then legitimately have

$$\text{contents}(\text{step}(s, a), n) \neq \text{contents}(\text{step}(t, a), n).$$

The repair to the definition is to require (2.6) to hold only if a does change the value of location n . This is accomplished in (2.5), the second of the Reference Monitor Assumptions specified in Definition 5 above.

This problem of specifying what it means for an operation to “reference” a location has been studied before; Popek and Farber [16], for example, construct the dual notion “NoRef” as follows. First, for $n \in N$, define the equivalence relation $\stackrel{n}{\cong}$ by

$$s \stackrel{n}{\cong} t \stackrel{\text{def}}{=} (\forall m \in N : \text{contents}(s, m) = \text{contents}(t, m) \vee m = n).$$

That is, $s \stackrel{n}{\cong} t$ if the values of all locations, except possibly that of n , are the same in both of states s and t . Then the predicate $\text{NoRef}(a, n)$, which is to be true when action a does not reference location n , is defined by

$$\text{NoRef}(a, n) \stackrel{\text{def}}{=} s \stackrel{n}{\cong} t \supset \text{step}(s, a) \stackrel{n}{\cong} \text{step}(t, a).$$

The motivation for this definition is the idea that if a does not reference the value of location n , then changing the value of that location should have no effect on the values assigned to other locations by action a . It is easy to prove that our notion of reference, as embodied in (2.5), implies the notion embodied in Popek and Farber’s definition. The converse is not true. This is due to a weakness in Popek and Farber’s definition which they discuss in [16, page 742 (footnote 5)]; they suggest a stronger definition whose motivation is identical to that given in our discussion of the formulation of (2.5). Unfortunately, the formal statement of Popek and Farber’s stronger definition contains serious typographical errors and it is impossible to tell what was intended. Nonetheless, we consider the relationship between the description of their definition and ours to be sufficiently close that they provide additional confidence in the correctness of our formulation of the second Reference Monitor Assumption.

Given these definitions, we can now state a theorem that relates noninterference to access control mechanisms.

Theorem 2 *Let M be a system with structured state that satisfies the Reference Monitor Assumptions and the following two conditions.*

1. $u \rightsquigarrow v \supset \text{observe}(u) \subseteq \text{observe}(v)$, and
2. $n \in \text{alter}(u) \wedge n \in \text{observe}(v) \supset u \rightsquigarrow v$.

Then M is secure for \rightsquigarrow .

Proof: We show that the conditions of the theorem satisfy those of the unwinding theorem. We identify the view-partitioning relations $\overset{u}{\sim}$ of the Unwinding Theorem with the corresponding relations defined in the statement of the Reference Monitor Assumptions. Output consistency is then satisfied immediately by the first of the Reference Monitor Assumptions.

To establish step consistency, we must prove

$$s \overset{u}{\sim} t \supset \text{step}(s, a) \overset{u}{\sim} \text{step}(t, a).$$

This can be rewritten as

$$s \overset{u}{\sim} t \supset \text{contents}(\text{step}(s, a), n) = \text{contents}(\text{step}(t, a), n)$$

where $n \in \text{observe}(u)$. There are three cases to consider

Case 1: $\text{contents}(\text{step}(s, a), n) \neq \text{contents}(s, n)$. The third of the Reference Monitor Assumptions gives $n \in \text{alter}(\text{dom}(a))$; since $n \in \text{observe}(u)$, the second of the conditions in the statement of the theorem then gives $\text{dom}(a) \rightsquigarrow u$. The first of the conditions in the statement of the theorem then gives

$$\text{observe}(\text{dom}(a)) \subseteq \text{observe}(u),$$

and $s \overset{u}{\sim} t$ then implies $s \overset{\text{dom}(a)}{\sim} t$. The second of the Reference Monitor Assumptions then provides the conclusion we require.

Case 2: $\text{contents}(\text{step}(t, a), n) \neq \text{contents}(t, n)$. This case is symmetrical with the first.

Case 3: $\text{contents}(\text{step}(t, a), n) = \text{contents}(t, n) \wedge \text{contents}(\text{step}(t, a), n) = \text{contents}(t, n)$. Since $s \overset{u}{\sim} t$ and $n \in \text{observe}(u)$, we have $\text{contents}(s, n) = \text{contents}(t, n)$ and the conclusion follows immediately.

It remains to show that the construction locally respects \rightsquigarrow . That is, we need to show

$$\text{dom}(a) \not\rightsquigarrow u \supset s \overset{u}{\sim} \text{step}(s, a).$$

Taking the contrapositive and expanding the definition of $\overset{u}{\sim}$, this becomes

$$(\exists n \in \text{observe}(u): \text{contents}(s, n) \neq \text{contents}(\text{step}(s, a), n)) \supset \text{dom}(a) \overset{u}{\sim} u.$$

Now if $\text{contents}(s, n) \neq \text{contents}(\text{step}(s, a), n)$, the third condition of the Reference Monitor Assumptions gives $n \in \text{alter}(\text{dom}(a))$. Hence, we have

$$n \in \text{alter}(\text{dom}(a)) \wedge n \in \text{observe}(u)$$

and so the second condition to the theorem requires $\text{dom}(a) \overset{u}{\sim} u$ and the proof is complete. \square

In the following chapter, we will show that transitive noninterference policies satisfy the conditions of Theorem 2 and thereby relate noninterference to the familiar Bell and La Padula [3] formulation of security.

Chapter 3

Noninterference and Transitivity

The only restriction we placed on the relation \rightsquigarrow defining a security policy was that it should be reflexive. However, we will show that, within the formulation presented so far, only relations that are also transitive have a useful interpretation.

In their original paper on the subject, Goguen and Meseguer [7] suggested that intransitive policies could be used to specify channel control policies. For example, the policy of the encryption controller shown in Figure 1.1 could be specified by the four assertions

$$\begin{aligned} \text{Red} &\rightsquigarrow \text{Bypass} \\ \text{Red} &\rightsquigarrow \text{Crypto} \\ \text{Bypass} &\rightsquigarrow \text{Black} \\ \text{Crypto} &\rightsquigarrow \text{Black} \end{aligned}$$

with the understanding that all other combinations, except the reflexive ones, should be noninterfering. In particular, $\text{Red} \not\rightsquigarrow \text{Black}$, even though $\text{Red} \rightsquigarrow \text{Bypass}$ and $\text{Bypass} \rightsquigarrow \text{Black}$, so that the policy \rightsquigarrow is intransitive. This is certainly an intuitively attractive specification of the desired policy; unfortunately, it does not accurately capture the desired properties. The problem is that noninterference is a very strong property: the assertion $\text{Red} \not\rightsquigarrow \text{Black}$ means that there must be *no* way for Black to observe activity by Red. This is not what is required here; Black must certainly be able to observe activity by Red (after all, it is the source of all incoming data), but we want all such observations to be mediated by the Bypass or the Crypto.

If the requirement $\text{Red} \not\rightsquigarrow \text{Black}$ is too strong, it is obvious that the complementary requirement $\text{Red} \rightsquigarrow \text{Black}$ is too weak: it would allow unrestricted communication from Red to Black.

We conclude that noninterference, as formulated so far, cannot specify channel-control policies exemplified by Figure 1.1. The question, then, is what interpretation is to be placed on intransitive policies within the current formulation? In its simplest form, we ask how we are to interpret assertions such as

$$\begin{aligned} A &\not\rightsquigarrow C \\ A &\rightsquigarrow B \\ B &\rightsquigarrow C. \end{aligned}$$

The hope is that this policy describes the “assured pipeline” [4] suggested by

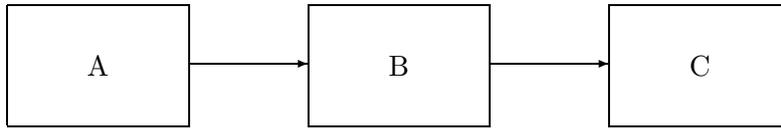


Figure 3.1: Desired interpretation of an intransitive policy

Figure 3.1. But as we have already seen, this hope is not fulfilled: the requirement $A \not\rightsquigarrow C$ precludes *all* interference by domain A on domain C , including that which would use domain B as an intermediary. The only satisfactory interpretation seems to be one in which the intermediate domain B is internally composed of two isolated parts, $B1$ and $B2$ as suggested in Figure 3.2. A can interfere with the $B1$ part of

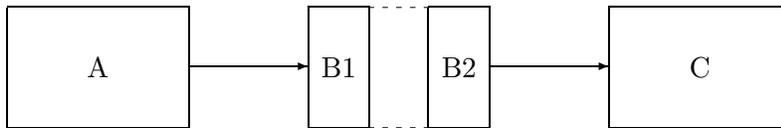


Figure 3.2: Plausible interpretation of an intransitive policy

B (hence $A \rightsquigarrow B$) and the $B2$ part of B can interfere with C (hence $B \rightsquigarrow C$), but the internal dichotomy of B allows $A \not\rightsquigarrow C$. Under this interpretation, however, it is surely more natural to recognize B as two domains and to formulate the policy accordingly:

$$\begin{aligned} A &\not\rightsquigarrow C \\ A &\rightsquigarrow B1 \end{aligned}$$

$$\begin{aligned} B1 &\not\sim B2 \\ B2 &\sim C \end{aligned}$$

But this is (trivially) a transitive policy. We conclude that intransitive policies seem to have no useful interpretation under the present formulation of noninterference.

In the following section, we will develop a formulation of noninterference that does provide a useful interpretation to intransitive policies, and in fact it is an interpretation satisfying the original goal of using noninterference to provide a formal foundation for the specification and verification of channel-control policies. Before we proceed to an examination of intransitive policies, however, we pause to examine the properties of transitive policies.

3.1 Properties of Transitive Policies

To begin, we define the class of *multilevel* security policies that model the systems of clearances and classifications used in the pen-and-paper world.

Definition 6 Let L be a set of *security labels* (comprising “levels,” possibly augmented by “compartments”) with a partial ordering \preceq (usually read as “is dominated by”). The interpretation of $l_1 \preceq l_2$ is that l_2 is more highly classified (in the case of data), or more highly trusted (in the case of individuals), and that information is permitted to flow from l_1 to l_2 , but not vice-versa (unless $l_1 = l_2$).

Let *clearance* : $D \rightarrow L$ be a function that assigns a fixed security label to each domain in D . Then the *multilevel security (MLS) policy* is:

$$u \sim v \text{ iff } \text{clearance}(u) \preceq \text{clearance}(v). \quad (3.1)$$

That is, u may interfere with v if the clearance of v dominates that of u .

An arbitrary security policy given by a relation \sim on D is said to be an *MLS-type policy* if a label set L with a partial ordering \preceq and a function *clearance* : $D \rightarrow L$ can be found such that (3.1) holds. \square

Clearly we have:

Theorem 3 *All MLS-type policies are transitive.*

Proof: This follows directly from the transitivity of the partial order \preceq . \square

The converse is also true. An essentially similar result (using a slightly different construction) was discovered by Dorothy Denning in 1976 [5].

Theorem 4 *All transitive policies are MLS-type policies.*

Proof: Let \rightsquigarrow be a transitive security policy. Define a further relation \leftrightarrow on D by:

$$u \leftrightarrow v \stackrel{\text{def}}{=} u \rightsquigarrow v \wedge v \rightsquigarrow u.$$

The construction ensures that \leftrightarrow is symmetric. Reflexivity and transitivity of \leftrightarrow follow from that of \rightsquigarrow (recall that all policies are reflexive). Thus \leftrightarrow is an equivalence relation. We identify a label set L with the equivalence classes of \leftrightarrow and use $[u]$ to denote the equivalence class of domain u under \leftrightarrow . We define a relation \preceq on L as follows:

$$[u] \preceq [v] \stackrel{\text{def}}{=} \exists \text{ domains } x \in [u] \text{ and } y \in [v] \text{ such that } x \rightsquigarrow y.$$

It is easy to see that \preceq is a partial order on L (i.e., it is reflexive, transitive, and antisymmetric). Finally, we define the function $\text{clearance} : U \rightarrow L$ by

$$\text{clearance}(u) \stackrel{\text{def}}{=} [u].$$

It is then easy to verify that

$$u \rightsquigarrow v \text{ iff } \text{clearance}(u) \preceq \text{clearance}(v),$$

and so it follows that \rightsquigarrow is an MLS-type policy. \square

The access control conditions of Theorem 2 reveal a familiar appearance when they are recast into the notation natural for MLS-type policies. To show this, we must first assign a *classification* label to each storage object by means of a function

- *classification*: $N \rightarrow D$.

Then we have:

Corollary 1 (Bell and La Padula Interpretation) *Let \rightsquigarrow be an MLS-type policy, and M a system with structured state that satisfies the Reference Monitor Assumptions and the following two properties.*

ss-property: $n \in \text{observe}(u) \supset \text{classification}(n) \preceq \text{clearance}(u)$,

***-property:** $n \in \text{alter}(u) \supset \text{clearance}(u) \preceq \text{classification}(n)$.

Then M is secure for \rightsquigarrow .

Proof: Using Theorem 2, we need to prove

$$u \rightsquigarrow v \supset \text{observe}(u) \subseteq \text{observe}(v),$$

and

$$n \in \text{alter}(u) \wedge n \in \text{observe}(v) \supset u \rightsquigarrow v.$$

The first of these can be restated as

$$u \rightsquigarrow v \wedge n \in \text{observe}(u) \supset n \in \text{observe}(v).$$

Using the notation of MLS-type policies and the ss-property, this becomes

$$\begin{aligned} \text{clearance}(u) \preceq \text{clearance}(v) \wedge \text{classification}(n) \preceq \text{clearance}(u) \\ \supset \text{classification}(n) \preceq \text{clearance}(v) \end{aligned}$$

and is satisfied immediately by the transitivity of the partial order \preceq .

Using the notation of MLS-type policies, the second of the conditions in Theorem 2 becomes

$$n \in \text{alter}(u) \wedge n \in \text{observe}(v) \supset \text{clearance}(u) \preceq \text{clearance}(v).$$

Using the ss- and *-properties, the antecedent to this implication becomes

$$\text{clearance}(u) \preceq \text{classification}(n) \wedge \text{classification}(n) \preceq \text{clearance}(v)$$

and the conclusion then follows from the transitivity of the partial order \preceq . \square

The ss- and *-properties named in this result correspond to the “simple-security” and “star” properties of the Bell and La Padula security model [2,3]. The simple-security condition asserts that a subject must only be able to observe objects whose classification is dominated by its own clearance, while the star-property asserts the dual condition that it must only be able to alter objects whose classification dominates its own clearance. Since the corollary establishes that these conditions are adequate to ensure the security of a system that enforces an MLS-type policy, it may seem puzzling that the Bell and La Padula formulation is known to have severe weaknesses [13,14]. In fact, there are two sources for these weaknesses and it may be useful to briefly indicate what they are, and why Corollary 1 is not vulnerable to them.

- One source of weaknesses derives from the lack of a *semantic* characterization of what is meant by “observe” and “alter” in the Bell and La Padula model. It is possible to subvert the model by inverting the intended interpretations of these terms. (So that the simple-security property says the subjects may *alter* only objects of lower classification.) Corollary 1 does not share this weakness because the Reference Monitor Assumptions provide an adequate semantic characterization of the intended interpretation of observe and alter access.
- The other source of weakness concerns the behavior of actions that modify the access control functions. Our notion of a system with structured state is very limited; more realistic models include more implementation detail and also extend the set of access control functions and provide actions for manipulating them. Such actions are called “rules” by Bell and La Padula, who gave a representative set in their Multics interpretation [3]. Two of these rules are known to permit unsecure information flow [15,25]. The reason for this is that the access control “table” and other implementation-level state data of the reference monitor are not treated as objects in the Bell and La Padula model; although the model prevents unsecure information flow through the objects that it explicitly recognizes, it places no constraints on the flow of information through the mechanisms of its own realization.

Corollary 1 does not share this weakness because its system model is very limited and does permit the access control tables to change; thus, it admits no “rules.” In more complex models, that do permit modification to the access control and other internal tables, the “rules” should be individually verified by direct reference to the appropriate unwinding theorem.

The verification of individual “rules” using the unwinding theorem requires identification of the “views” of the machine state held by different security domains. The next result provides some guidance in the identification of such views, by showing that, for a transitive \rightsquigarrow relation, they are “nested” within each other. This is obvious in the Bell and La Padula model (i.e., everything observable by a subject at level l_1 is also observable to a subject of level l_2 where $l_1 \preceq l_2$). What is interesting here is that Theorem 5 shows that this nesting property is inherent, not accidental.

Definition 7 A view-partitioned machine is said to have the *nesting* property if

$$u \rightsquigarrow v \wedge s \overset{v}{\sim} t \supset s \overset{u}{\sim} t.$$

That is, if states s and t appear identical to domain v , then they also appear identical to those domains u that may interfere with v . \square

Theorem 5 *Let \rightsquigarrow be a transitive policy and M a view-partitioned machine which satisfies the conditions of the unwinding theorem. Then there is a nested view-partitioning of M that also satisfies the conditions of the unwinding theorem.*

Proof: Define a new view-partitioning relation $\overset{u}{\simeq}$ on D by

$$s \overset{u}{\simeq} t \stackrel{\text{def}}{=} (\forall v: v \rightsquigarrow u \supset s \overset{v}{\sim} t).$$

That $\overset{u}{\simeq}$ is an equivalence relation follows straightforwardly from the fact that $\overset{u}{\sim}$ is. Output consistency and step consistency of $\overset{u}{\simeq}$ likewise follow from those properties of the $\overset{u}{\sim}$ relation. For $\overset{u}{\simeq}$ to locally respect \rightsquigarrow , we require

$$\forall v \rightsquigarrow u : \text{dom}(a) \not\rightsquigarrow u \supset s \overset{v}{\sim} \text{step}(s, a). \quad (3.2)$$

The transitivity of \rightsquigarrow ensures $\text{dom}(a) \not\rightsquigarrow v$ (since otherwise we could combine $\text{dom}(a) \rightsquigarrow v$ with $v \rightsquigarrow u$ and contradict $\text{dom}(a) \not\rightsquigarrow u$), and (3.2) then follows from the fact that $\overset{v}{\sim}$ locally respects \rightsquigarrow .

For the nesting property, we need to demonstrate $x \rightsquigarrow u \supset s \overset{x}{\sim} t$ given $u \rightsquigarrow v$ and $s \overset{v}{\simeq} t$. Transitivity provides $x \rightsquigarrow v$, and the result then follows from the definition of $s \overset{v}{\simeq} t$. \square

Finally, we prove that unwinding is, in a certain sense, complete: for *any* secure system, we can find a view-partitioning that satisfies the conditions of the unwinding theorem. Note that this result does not depend on the transitivity of \rightsquigarrow , but it does depend on the present interpretation of noninterference which, as we have seen, makes sense only for transitive policies.

Theorem 6 *If M is a secure system, then for each domain $u \in D$ an equivalence relation $\overset{u}{\sim}$ on the set of states can be found that satisfies the conditions of the unwinding theorem.*

Proof: We use the following construction: for $u \in D$ and reachable states s and t , define

$$s \overset{u}{\sim} t \stackrel{\text{def}}{=} (\forall \alpha \in A^*, b \in A : \text{dom}(b) = u \\ \supset \text{output}(\text{run}(s, \alpha), b) = \text{output}(\text{run}(t, \alpha), b)). \quad (3.3)$$

Clearly, $\overset{u}{\sim}$ is an equivalence relation. Output consistency follows by taking $\alpha = \Lambda$ in (3.3). For step consistency, we need

$$s \overset{u}{\sim} t \supset \text{step}(s, a) \overset{u}{\sim} \text{step}(t, a).$$

The conclusion to this implication expands to

$$\text{output}(\text{run}(\text{step}(s, a), \alpha), b) = \text{output}(\text{run}(\text{step}(t, a), \alpha), b)$$

and this is equivalent to

$$\text{output}(\text{run}(s, a \circ \alpha), b) = \text{output}(\text{run}(t, a \circ \alpha), b),$$

which follows directly from the definition of $s \stackrel{u}{\sim} t$.

To show that the construction locally respects \rightsquigarrow , we need to demonstrate

$$\text{dom}(a) \not\rightsquigarrow u \supset s \stackrel{u}{\sim} \text{step}(s, a).$$

The conclusion expands to

$$\text{output}(\text{run}(s, \alpha), b) = \text{output}(\text{run}(\text{step}(s, a), \alpha), b) \quad (3.4)$$

where $\text{dom}(b) = u$. If s is a reachable state, there exists γ such that $s = \text{do}(\gamma)$ and so (3.4) can be written as

$$\text{test}(\gamma \circ \alpha, b) = \text{test}(\gamma \circ a \circ \alpha, b).$$

Since the machine is secure, the definition of noninterference gives

$$\text{test}(\gamma \circ \alpha, b) = \text{test}(\text{purge}(\gamma \circ \alpha, \text{dom}(b)), b)$$

and

$$\text{test}(\gamma \circ a \circ \alpha, b) = \text{test}(\text{purge}(\gamma \circ a \circ \alpha, \text{dom}(b)), b).$$

But, clearly, since $\text{dom}(a) \not\rightsquigarrow u$ and $u = \text{dom}(b)$,

$$\text{purge}(\gamma \circ \alpha, \text{dom}(b)) = \text{purge}(\gamma \circ a \circ \alpha, \text{dom}(b))$$

and the result follows. \square

Chapter 4

Intransitive Noninterference

Goguen and Meseguer recognized the inability of standard noninterference to model channel-control policies and they introduced several extensions to the basic formulation in their second paper on the subject [8]. However, the first really satisfactory treatment of intransitive noninterference policies was given by Haigh and Young [10], with an earlier version the previous year [9].

Both Goguen and Meseguer, and Haigh and Young, recognized that the standard definition of noninterference is too draconian. If $u \not\rightsquigarrow v$, the requirement is that deleting all actions performed by u should produce no change in the behavior of the system as perceived by v . This is too strong if we also have the assertions $u \rightsquigarrow w$ and $w \rightsquigarrow v$. Surely we should only delete those actions of u that are not followed by actions of w : this is the essence of Haigh and Young's reformulation of noninterference. In order to give a formal definition, we need to identify those actions in an action sequence that should not be deleted. This is the purpose of the function *sources*.

Definition 8 We define the function

$$\bullet \text{ sources: } A^* \times D \rightarrow \mathcal{P}(D)$$

by the equations

$$\begin{aligned} \text{sources}(\Lambda, u) &= \{u\} \\ \text{sources}(a \circ \alpha, u)^1 &= \begin{cases} \text{sources}(\alpha, u) \cup \{dom(a)\} & \text{if } \exists v : v \in \text{sources}(\alpha, u) \\ & \wedge dom(a) \rightsquigarrow v \\ \text{sources}(\alpha, u) & \text{otherwise.} \end{cases} \end{aligned}$$

¹This is the definition in which right-recursion is essential.

Our function *sources* corresponds to the function *purgeable* of Haigh and Young [10], although Haigh and Young gave only an informal characterization of their function. In essence $v \in \text{sources}(\alpha, u)$ means either that $v = u$ or that there is a subsequence of α consisting of actions performed by domains w_1, w_2, \dots, w_n such that $w_1 \rightsquigarrow w_2 \rightsquigarrow \dots \rightsquigarrow w_n$, $v = w_1$, and $u = w_n$. In considering whether an action a performed prior to the action sequence α should be allowed to influence u , we ask whether there is any $v \in \text{sources}(\alpha, u)$ such that $\text{dom}(a) \rightsquigarrow v$. Notice that always

$$\text{sources}(\alpha, u) \subseteq \text{sources}(a \circ \alpha, u), \text{ and}$$

$$u \in \text{sources}(\alpha, u).$$

We can now define the function *ipurge* (for *intransitive-purge*):

- $\text{ipurge}: A^* \times D \rightarrow A^*$

by the equations

$$\begin{aligned} \text{ipurge}(\Lambda, u) &= \Lambda \\ \text{ipurge}(a \circ \alpha, u) &= \begin{cases} a \circ \text{ipurge}(\alpha, u) & \text{if } \text{dom}(a) \in \text{sources}(a \circ \alpha, u) \\ \text{ipurge}(\alpha, u) & \text{otherwise.} \end{cases} \end{aligned}$$

Informally, $\text{ipurge}(\alpha, u)$ consists of the subsequence of α with all those actions that should not be able to interfere with u removed. Thus, security is now defined in terms of the *ipurge* function:

A machine is *secure* for the policy \rightsquigarrow if

$$\text{test}(\alpha, a) = \text{test}(\text{ipurge}(\alpha, \text{dom}(a)), a).$$

□

From this point on, our treatment diverges from that of Haigh and Young. We will argue later that their treatment is incorrect. The first step is to establish the revised form of Lemma 1.

Lemma 2 *Let \rightsquigarrow be a policy and M a view-partitioned, output consistent system such that,*

$$\text{do}(\alpha) \stackrel{u}{\sim} \text{do}(\text{ipurge}(\alpha, u)).$$

Then M is secure for \rightsquigarrow .

Proof: The proof is essentially identical to that of Lemma 1.

Setting $u = \text{dom}(a)$ in the statement of the lemma gives

$$\text{do}(\alpha) \stackrel{\text{dom}(a)}{\sim} \text{do}(\text{ipurge}(\alpha, \text{dom}(a))),$$

and output consistency then provides

$$\text{output}(\text{do}(\alpha), a) = \text{output}(\text{do}(\text{ipurge}(\alpha, \text{dom}(a))), a).$$

But this is simply

$$\text{test}(\alpha, a) = \text{test}(\text{ipurge}(\alpha, \text{dom}(a)), a),$$

which is the definition of security for \rightsquigarrow given by Definition 8. \square

Next, we present a series of definitions and lemmas that culminate in the revised form of the unwinding theorem.

Definition 9 Let M be a view-partitioned system and $C \subseteq D$ a set of domains. We define the equivalence relation $\stackrel{C}{\approx}$ on the states of M as follows:

$$s \stackrel{C}{\approx} t \stackrel{\text{def}}{=} (\forall u \in C : s \stackrel{u}{\sim} t).$$

Thus $s \stackrel{C}{\approx} t$ exactly when the states s and t appear identical to all the members of C . \square

Definition 10 Let M be a view-partitioned system and \rightsquigarrow a policy. We say that M is *weakly step consistent* if

$$s \stackrel{u}{\sim} t \wedge s \stackrel{\text{dom}(a)}{\sim} t \supset \text{step}(s, a) \stackrel{u}{\sim} \text{step}(t, a).$$

\square

Lemma 3 Let \rightsquigarrow be a policy and M a view-partitioned system which is weakly step consistent, and locally respects \rightsquigarrow . Then

$$s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t \supset \text{step}(s, a) \stackrel{\text{sources}(\alpha, u)}{\approx} \text{step}(t, a).$$

Proof: Suppose $v \in \text{sources}(\alpha, u)$. We need to show that

$$\text{step}(s, a) \overset{v}{\sim} \text{step}(t, a). \quad (4.1)$$

Note that $v \in \text{sources}(\alpha, u)$ implies $v \in \text{sources}(a \circ \alpha, u)$, and so the hypothesis to the lemma provides

$$s \overset{v}{\sim} t. \quad (4.2)$$

We now consider two cases.

Case 1: $\text{dom}(a) \rightsquigarrow v$. Then, by the definition of the *sources* function, we have $\text{dom}(a) \in \text{sources}(a \circ \alpha, u)$ and the hypothesis to the lemma provides

$$s \overset{\text{dom}(a)}{\sim} t. \quad (4.3)$$

(4.1) then follows from (4.2) and (4.3) by weak step consistency.

Case 2: $\text{dom}(a) \not\rightsquigarrow v$. Then by local respect for \rightsquigarrow ,

$$\begin{aligned} \text{step}(s, a) &\overset{v}{\sim} s, \\ \text{step}(t, a) &\overset{v}{\sim} t \end{aligned}$$

and (4.1) follows from (4.2).

□

Lemma 4 *Let \rightsquigarrow be a policy and M a view-partitioned system that locally respects \rightsquigarrow . Then*

$$\text{dom}(a) \notin \text{sources}(a \circ \alpha, u) \supset s \overset{\text{sources}(\alpha, u)}{\approx} \text{step}(s, a).$$

Proof: We assume the hypothesis and let $v \in \text{sources}(\alpha, u)$. It must be that $\text{dom}(a) \not\rightsquigarrow v$, since otherwise $\text{dom}(a) \in \text{sources}(a \circ \alpha, u)$. Hence, by local respect for \rightsquigarrow ,

$$s \overset{v}{\sim} \text{step}(s, a)$$

and the conclusion follows. □

Lemma 5 *Let \rightsquigarrow be a policy and M a view-partitioned system which is weakly step consistent, and locally respects \rightsquigarrow . Then*

$$s \overset{\text{sources}(\alpha, u)}{\approx} t \supset \text{run}(s, \alpha) \overset{u}{\sim} \text{run}(t, \text{ipurge}(\alpha, u)).$$

Proof: The proof proceeds by induction on the length of α . The basis is the case $\alpha = \Lambda$ and follows straightforwardly by application of definitions. For the inductive step, we assume the result for α of length n , and consider $a \circ \alpha$. We then need to show

$$s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t \supset \text{run}(s, a \circ \alpha) \stackrel{u}{\sim} \text{run}(t, \text{ipurge}(a \circ \alpha, u)).$$

We now consider two cases.

Case 1: $\text{dom}(a) \in \text{sources}(a \circ \alpha, u)$. Then $\text{ipurge}(a \circ \alpha, u) = a \circ \text{ipurge}(\alpha, u)$ and we need to show

$$s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t \supset \text{run}(\text{step}(s, a), \alpha) \stackrel{u}{\sim} \text{run}(\text{step}(t, a), \text{ipurge}(\alpha, u)).$$

Lemma 3 gives

$$s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t \supset \text{step}(s, a) \stackrel{\text{sources}(\alpha, u)}{\approx} \text{step}(t, a)$$

and the result then follows from the inductive hypothesis.

Case 2: $\text{dom}(a) \notin \text{sources}(a \circ \alpha, u)$. Then $\text{ipurge}(a \circ \alpha, u) = \text{ipurge}(\alpha, u)$ and we need to show

$$s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t \supset \text{run}(\text{step}(s, a), \alpha) \stackrel{u}{\sim} \text{run}(t, \text{ipurge}(\alpha, u)).$$

Now Lemma 4 gives

$$\text{dom}(a) \notin \text{sources}(a \circ \alpha, u) \supset s \stackrel{\text{sources}(\alpha, u)}{\approx} \text{step}(s, a)$$

and, since $\text{sources}(\alpha, u) \subseteq \text{sources}(a \circ \alpha, u)$, $s \stackrel{\text{sources}(a \circ \alpha, u)}{\approx} t$ implies

$$s \stackrel{\text{sources}(\alpha, u)}{\approx} t.$$

Because $\stackrel{\text{sources}(\alpha, u)}{\approx}$ is an equivalence relation, it follows that

$$\text{step}(s, a) \stackrel{\text{sources}(\alpha, u)}{\approx} t$$

and the result then follows from the inductive hypothesis.

□

Finally, we can present the unwinding theorem for intransitive noninterference policies.

Theorem 7 (Unwinding Theorem for Intransitive Policies) *Let \rightsquigarrow be a policy and M a view-partitioned system that is*

1. *is output consistent,*
2. *weakly step consistent, and*
3. *locally respects \rightsquigarrow .*

Then M is secure for \rightsquigarrow .

Proof: Taking $s = t = s_0$ in Lemma 5 gives

$$\text{run}(s_0, \alpha) \stackrel{u}{\sim} \text{run}(s_0, \text{ipurge}(\alpha, u)),$$

which can be rewritten in the form

$$\text{do}(\alpha) \stackrel{u}{\sim} \text{do}(\text{ipurge}(\alpha, u)),$$

so that the conclusion follows from Lemma 2. \square

A formal verification of this theorem has been performed using the EHDM specification and verification system and is described in in the Appendix toin Part III of this report. The mechanically checked proof follows the argument of Lemmas 3 to 5 very closely.

In the following chapter, we consider the differences and similarities between this unwinding theorem and both the ordinary unwinding theorem and that of Haigh and Young. Before doing so, however, we note that the access control mechanism described in Definition 5 on page 12 of Chapter 2 works for intransitive noninterference policies as well as for transitive ones.

Theorem 8 *Let M be a system with structured state that satisfies the Reference Monitor Assumptions and the condition*

$$n \in \text{alter}(u) \wedge n \in \text{observe}(v) \supset u \rightsquigarrow v.$$

Then M is secure for \rightsquigarrow .

Proof: The proof is similar to that of Theorem 2. We show that the conditions of the theorem satisfy those of the intransitive unwinding theorem. We identify the view-partitioning relations $\stackrel{u}{\sim}$ of the Intransitive Unwinding Theorem with the corresponding relations defined in the statement of the Reference Monitor Assumptions. Output consistency is then satisfied immediately by the first of the Reference Monitor Assumptions.

To establish weak step consistency, we must prove

$$s \stackrel{u}{\sim} t \wedge s \stackrel{\text{dom}(a)}{\sim} t \supset \text{step}(s, a) \stackrel{u}{\sim} \text{step}(t, a).$$

This can be rewritten as

$$s \stackrel{u}{\sim} t \wedge s \stackrel{\text{dom}(a)}{\sim} t \supset \text{contents}(\text{step}(s, a), n) = \text{contents}(\text{step}(t, a), n)$$

where $n \in \text{observe}(u)$. There are three cases to consider

Case 1: $\text{contents}(\text{step}(s, a), n) \neq \text{contents}(s, n)$. The second of the Reference Monitor Assumptions provides the desired conclusion directly (from the hypothesis $s \stackrel{\text{dom}(a)}{\sim} t$).

Case 2: $\text{contents}(\text{step}(t, a), n) \neq \text{contents}(t, n)$. This case is symmetrical with the first.

Case 3: $\text{contents}(\text{step}(t, a), n) = \text{contents}(t, n) \wedge \text{contents}(\text{step}(t, a), n) = \text{contents}(t, n)$. Since $s \stackrel{u}{\sim} t$, we have $\text{contents}(s, n) = \text{contents}(t, n)$ and the conclusion is immediate.

It remains to show that the construction locally respects \rightsquigarrow . This follows by exactly the same argument as that used in the proof of Theorem 2. \square

It is illuminating to examine the similarity between this access control theorem and the ordinary one (Theorem 2). The only difference between the two theorems is that the ordinary one requires the additional condition

$$u \rightsquigarrow v \supset \text{observe}(u) \subseteq \text{observe}(v).$$

Theorem 8 is able to dispense with this condition because the intransitive unwinding theorem, from which it is derived, requires only *weak* step consistency.

To see how this apparently small difference in formulation allows Theorem 8, but not Theorem 2, to provide an access control interpretation for an intransitive policy, consider the system sketched in Figure 3.1. Theorem 8, allows domain A to have alter access to locations to which domain B has observe access. Similarly, it permits domain B to have alter access to locations to which domain C has observe access. In this way, information can flow from A to B and from B to C . However, A may not have alter access to any locations to which C has observe access; in this way, direct flow of information from A to C is prevented.

The conditions of Theorem 2 also allow domain A to have alter access to locations to which domain B has observe access, but they also require that B have

observe access to every location to which A has observe access. Similarly, considering domains B and C , the conditions of Theorem 2 require that C have observe access to every location to which B has observe access. Transitively, therefore, C has observe access to every location to which A has observe access and so A can have “no secrets” from C . Thus, the additional condition of Theorem 2 forces the transitive completion of the policy, and so allows the direct flow of information from A to C .

Chapter 5

Comparisons among the Formulations

In this chapter we compare our treatment of intransitive noninterference policies with the standard treatment of noninterference and with that of Haigh and Young.

5.1 Intransitive vs. Standard Noninterference

We first compare our treatment of intransitive noninterference policies (Chapter 4) with the standard treatment of noninterference policies (Chapter 2) and the special properties of transitive policies (Chapter 3). We will show that, when restricted to transitive policies, our formulation of noninterference corresponds exactly with the standard treatment. This provides some assurance that our treatment is a natural extension of the standard one. To begin, we establish that the definitions of security coincide in the case of transitive policies.

Lemma 6 *If \rightsquigarrow is transitive, then*

$$v \in \text{sources}(\alpha, u) \supset v \rightsquigarrow u.$$

Proof: The proof is by induction on the length of α . The basis is the case $\alpha = \Lambda$, and reference to Definition 8 shows that

$$\text{sources}(\Lambda, u) = \{u\}$$

and the lemma is satisfied in this case by the reflexivity of \rightsquigarrow .

For the inductive step, Definition 8 gives $v \in \text{sources}(a \circ \alpha, u)$ if either $v \in \text{sources}(\alpha, u)$ or

$$v = \text{dom}(a) \wedge (\exists w \in \text{sources}(\alpha, u) \wedge \text{dom}(a) \rightsquigarrow w).$$

In the first case, the inductive hypothesis provides $v \rightsquigarrow u$ directly; in the second, the inductive hypothesis provides $w \rightsquigarrow u$, we also have $v = \text{dom}(a)$ and $\text{dom}(a) \rightsquigarrow w$, and so transitivity provides $v \rightsquigarrow u$ as required. \square

Lemma 7 *If \rightsquigarrow is transitive, then $\text{ipurge}(\alpha, u) = \text{purge}(\alpha, u)$.*

Proof: Comparison of Definitions 2 and 8 reveals that we only need to demonstrate

$$\text{dom}(a) \rightsquigarrow u \text{ iff } \text{dom}(a) \in \text{sources}(a \circ \alpha, u).$$

The “if” direction was established by the previous lemma. For the “only if” direction, note that $u \in \text{sources}(\alpha, u)$, so that $\text{dom}(a) \in \text{sources}(a \circ \alpha, u)$ follows immediately from Definition 8 and $\text{dom}(a) \rightsquigarrow u$. \square

Theorem 9 *Definitions 2 and 8 of security agree when the relation \rightsquigarrow is transitive.*

Proof: Since the two definitions differ only in their “purge” functions, this result is an immediate consequence of the previous lemma. \square

We now know that the two definitions of security coincide in the case of transitive policies; next, we show that the unwinding theorems do so as well.

Theorem 10 *The Unwinding Theorems 1 and 7 agree when the relation \rightsquigarrow is transitive.*

Proof: The unwinding theorems differ only in that the intransitive version uses weak step consistency where the regular one uses (ordinary) step consistency. Weak step consistency is the condition

$$s \stackrel{u}{\sim} t \wedge s \stackrel{\text{dom}(a)}{\sim} t \supset \text{step}(s, a) \stackrel{u}{\sim} \text{step}(t, a),$$

while ordinary step consistency is the condition

$$s \stackrel{u}{\sim} t \supset \text{step}(s, a) \stackrel{u}{\sim} \text{step}(t, a).$$

Ordinary step consistency obviously implies weak step consistency; thus, we only need to show that weak step consistency implies ordinary step consistency when \rightsquigarrow is transitive. However, it is not necessarily the case that a given view partitioning that satisfies weak step consistency also satisfies ordinary step consistency; thus we must prove that a view partitioning satisfying the intransitive unwinding theorem implies the existence of (another) view partitioning satisfying the ordinary unwinding theorem.

The construction we use is the same as that for the nesting theorem (Theorem 5): we define a new view-partitioning relation $\overset{u}{\simeq}$ on D by

$$s \overset{u}{\simeq} t \stackrel{\text{def}}{=} (\forall v: v \rightsquigarrow u \supset s \overset{v}{\sim} t).$$

The output consistency and local respect for \rightsquigarrow of $\overset{u}{\simeq}$ follow by the same arguments used in Theorem 5, as does the fact that $\overset{u}{\simeq}$ is an equivalence relation. For (ordinary) step consistency, we must show that

$$s \overset{u}{\simeq} t \supset \text{step}(s, a) \overset{u}{\simeq} \text{step}(t, a),$$

or, equivalently,

$$s \overset{u}{\simeq} t \wedge v \rightsquigarrow u \supset \text{step}(s, a) \overset{v}{\sim} \text{step}(t, a).$$

Note that $s \overset{u}{\simeq} t \wedge v \rightsquigarrow u \supset s \overset{v}{\sim} t$. There are now two cases to consider.

Case 1: $\text{dom}(a) \rightsquigarrow u$. In this case, $s \overset{u}{\simeq} t$ implies $s \overset{\text{dom}(a)}{\sim} t$, and since we already have $s \overset{v}{\sim} t$, weak step consistency then supplies $\text{step}(s, a) \overset{v}{\sim} \text{step}(t, a)$ as required.

Case 2: $\text{dom}(a) \not\rightsquigarrow u$. In this case, since we have $v \rightsquigarrow u$, transitivity of \rightsquigarrow requires $\text{dom}(a) \not\rightsquigarrow v$. But then, local respect of \rightsquigarrow by $\overset{v}{\sim}$ requires $\text{step}(s, a) \overset{v}{\sim} s$ and $\text{step}(t, a) \overset{v}{\sim} t$, and so $\text{step}(s, a) \overset{v}{\sim} \text{step}(t, a)$ follows directly from $s \overset{v}{\sim} t$.

□

5.2 Comparison with Haigh and Young's Formulation

The system model used by Haigh and Young [10] differs slightly from that used here. Their *output* function has signature

- *output*: $S \times D \rightarrow O$

whereas we use

- $output: S \times A \rightarrow O$.

Thus, their *output* function allows a domain u to inspect the system state s directly as $output(s, u)$, whereas ours requires the mediation of an action a with $dom(a) = u$ to form $output(s, a)$. Converting our formulation to theirs requires a corresponding change in the definition of the function *test* to signature

- $test: A^* \times D \rightarrow O$

with definition

$$test(\alpha, u) = output(do(\alpha), u).$$

The definition of security becomes

$$test(\alpha, u) = test(ipurge(\alpha, u), u),$$

and that of output consistency changes to

$$s \stackrel{u}{\sim} t \supset output(s, u) = output(t, u).$$

Some small changes are then needed in the proof of Lemma 2 in order to take account of the modified function signatures. No other changes are needed in the development. We have checked this by modifying the formal verification of the AppendixPart III in the manner described above and then re-running all the proofs. The ability to readily check the effect of changed assumptions in this way is one of the great benefits of formal verification: assumptions are recorded with great precision and the “ripple” effect of perturbations can be evaluated mechanically.

Since the slight differences between the system model used here and that used by Haigh and Young have only a trivial impact on the definition of intransitive non-interference, and none at all on our intransitive unwinding theorem, it is reasonable to compare our definitions and theorems with those of Haigh and Young.

Under the proviso that our function *sources* is the same as their informally defined function *purgeable*, our definition for intransitive noninterference is the same as that given by Haigh and Young for “MDS Security.” However, the corresponding unwinding theorems differ and in this section we compare our unwinding theorem for intransitive policies with the “SAT MDS Unwinding Theorem” of Haigh and Young.

In our terminology and notation, the SAT MDS Unwinding Theorem of Haigh and Young is the following.

Proposition 1 (SAT MDS Unwinding Theorem) *Let \rightsquigarrow be a policy and M a view-partitioned system that is*

1. *is output consistent,*
2. *step consistent, and*
3. *MDS-respects \rightsquigarrow .*

Then M is secure for \rightsquigarrow .

That is, Haigh and Young require step consistency where we require *weak* step consistency, and they require a condition we call “MDS-respect” for \rightsquigarrow where we require local respect. The condition MDS-respect is defined as follows by Haigh and Young [10, p. 147, formula (10)].

Definition 11 Let M be a view-partitioned system and \rightsquigarrow a policy. We say that M *MDS-respects* \rightsquigarrow if, for any choice of action a and state s , if a is purgeable with respect to domain u , then

$$s \stackrel{u}{\sim} \text{step}(s, a).$$

□

This definition presents a considerable challenge to interpretation. The function *purgeable* is not defined formally by Haigh and Young, but in its informal definition, and in all previous uses within their paper, it is used in contexts such as “ a is purgeable with respect to u in α .” That is, the purgeability of an action is defined relative to a domain *and* an action sequence. In the definition of MDS-respects, however, there is no reference to an action sequence. Examination of Haigh and Young’s proof of their SAT MDS Unwinding Theorem sheds no light on the interpretation of the crucial notion MDS-respects: the proof is only a sketch and does not employ formal use of definitions.

Any interpretation of MDS-respects that differs from locally respects must be either weaker or stronger than that alternative notion. A stronger notion would require $s \stackrel{u}{\sim} \text{step}(s, a)$ even in some circumstances where $\text{dom}(a) \rightsquigarrow u$. This does not seem very plausible, since the other conditions of the SAT MDS Unwinding Theorem are the same as for the ordinary unwinding theorem, and strengthening one of them must restrict, rather than enlarge, the class of policies admitted. We conclude that MDS-respects must allow $s \not\stackrel{u}{\sim} \text{step}(s, a)$ in some circumstances where $\text{dom}(a) \rightsquigarrow u$. The constraint on the possible values of $\text{step}(s, a)$ in this case must be provided by the other conditions of the theorem, namely output consistency, and step

consistency. However, as these are both the same as in the ordinary noninterference case, it is difficult to see how adequate constraints on the effect of a state transition $step(s, a)$ with $dom(a) \not\rightsquigarrow u$ and $s \not\stackrel{u}{\sim} step(s, a)$ can be achieved by these constraints.

In contrast, our formulation of the unwinding theorem for intransitive policies leaves the locally respects constraint unchanged from the ordinary case, but changes the step consistency constraint to *weak* step consistency. That is, the condition:

$$s \stackrel{u}{\sim} t \supset step(s, a) \stackrel{u}{\sim} step(t, a)$$

of the ordinary case is changed to

$$s \stackrel{u}{\sim} t \wedge s \stackrel{dom(a)}{\sim} t \supset step(s, a) \stackrel{u}{\sim} step(t, a)$$

for the intransitive case.

The second of these conditions is very natural: its intuitive interpretation is that when an action a is performed, those elements of the system state visible to u change in a way that depends only on those same elements, plus those visible to the domain that performed the action.

The ordinary step consistency condition requires that when an action a is performed, those elements of the system state visible to u change in a way that depends on those elements *alone*. This seems more, not less, restrictive than the previous case, until we recall that for transitive policies there is always a view-partitioning that satisfies

$$u \rightsquigarrow v \wedge s \stackrel{v}{\sim} t \supset s \stackrel{u}{\sim} t.$$

In other words, those elements of the state space visible to u include all the elements of the state space visible to domains that may interfere with u .

We should now ask whether a similar explanation can provide a sound interpretation to Haigh and Young's SAT MDS Unwinding Theorem. We believe not, and we use the following example to make our case. Consider a system with four domains U, V, W , and X ; U and V may interfere with W , and W may interfere with X , but U and V must not directly interfere with X . The system state is composed of three internal registers, u, v , and x , all initially zero. Each domain has a single action associated with it: U 's action sets the register u to 1, V 's action sets the register v to 2, W 's action sets the register x to the sum of the contents of u and v , and X 's action outputs the contents of the register x . It should be clear that this system satisfies the stated policy. We need to be able to distinguish it from the insecure variant in which X 's action outputs the sum of the registers u and v directly. In our formulation of intransitive noninterference, U, V and X 's view of the system state is restricted to the registers u, v and x respectively, while W

can view both registers u and v . It is easy to see that our unwinding theorem for intransitive policies is satisfied by this assignment.

Haigh and Young's unwinding theorem is not satisfied, however, since the effect of W 's action on the register x cannot be explained in terms of the objects visible to X . It seems that the set of objects visible to X must be enlarged to include the registers u and v . But how, then, are we to distinguish the system from its insecure variant?

We conclude that all possible interpretations of Haigh and Young's SAT MDS Unwinding Theorem are unsatisfactory. Because there is no precise definition of the crucial requirement that we call "MDS-respects," it is impossible to assign a definitive status to the theorem, and its utility becomes questionable.

We have, we believe, presented adequate evidence that our unwinding theorem for intransitive policies is both true and useful; indeed, we believe it is the strongest theorem possible. We have also presented evidence that Haigh and Young's theorem is essentially different than ours—differing in the crucial step consistency condition, not just the uncertain MDS-respects condition. We therefore believe it unlikely that their theorem, if true, is as generally applicable as ours. Consequently, we consider it likely that their theorem is either false, or true but applicable to a very small class of systems and/or policies.

Chapter 6

Summary and Conclusions

We have examined the issue of transitivity in noninterference security policies. Intransitive noninterference policies would seem, intuitively, to be exactly what is required for the formal specification of channel control and type enforcement policies. We have shown, however, that the standard interpretation of noninterference does not fulfill this expectation. Fortunately, the interpretation of noninterference introduced by Haigh and Young for multidomain security (MDS) does have the properties we require. Our contribution has been the identification of intransitivity of the \rightsquigarrow relation as the key distinction between channel control, type enforcement, and MDS policies on the one hand, and MLS policies on the other.

It can be considered a historical accident that the theory for the transitive case was invented and developed before the intransitive one, and has therefore become regarded as the standard case. We submit that it is now more helpful to regard the intransitive case as the basis for noninterference formulations of security, with the formerly standard treatment regarded as a specialization for the case of transitive policies. The advantage of regarding the development in this light is that one does not have to trouble with the rather difficult and informal argument that the standard treatment makes little sense for intransitive policies; one can simply present the general theory and then show that there is a simpler treatment available in the special case of transitive policies. The attempt to use the standard treatment in the case of intransitive policies simply does not arise with this approach.

Our main technical contributions have been the formulation, rigorous proof, and mechanically-checked formal verification of an unwinding theorem for intransitive policies, a demonstration that the definitions and theorems of the intransitive theory collapse to the standard ones in the case of transitive policies, and an exploration of the properties of transitive policies. Our demonstrations of the equivalence of MLS

and transitive noninterference policies, of the nesting property, and of the result that all MLS secure systems satisfy the conditions of the unwinding theorem, shed some new light on the properties of transitive noninterference security policies.

However, the novel and more interesting case, and the one that prompted this investigation in the first place, is that of intransitive noninterference policies. In future work we hope to explore the practical application of intransitive noninterference formulations to problems of channel control, and to develop effective methods for verifying mechanisms that enforce such policies. We also plan to explore the connection between intransitive noninterference policies and the class of properties, discussed in [21], that can be enforced by kernelization.

Bibliography

- [1] D.H. Barnes. The provision of security for user data on packet switched networks. In *Proc. 1983 IEEE Symposium on Security and Privacy*, pages 121–126, Oakland, CA, April 1983. IEEE Computer Society.
- [2] D.E. Bell and L.J. La Padula. Secure computer systems : Vol. I—mathematical foundations, Vol. II—a mathematical model, Vol III—a refinement of the mathematical model. Technical Report MTR-2547 (three volumes), Mitre Corporation, Bedford, MA, March–December 1973.
- [3] D.E. Bell and L.J. La Padula. Secure computer system: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, Bedford, MA, March 1976.
- [4] W.E. Boebert and R.Y. Kain. A practical alternative to hierarchical integrity policies. In *Proceedings 8th DoD/NBS Computer Security Initiative Conference*, pages 18–27, Gaithersburg, MD, September 1985.
- [5] D.E. Denning. On the derivation of lattice structured information flow policies. Technical Report CSD TR 180, Purdue University, March 1976.
- [6] R.J. Feiertag, K.N. Levitt, and L. Robinson. Proving multilevel security of a system design. In *Proc. 6th ACM Symposium on Operating System Principles*, pages 57–65, November 1977.
- [7] J.A. Goguen and J. Meseguer. Security policies and security models. In *Proc. 1982 Symposium on Security and Privacy*, pages 11–20, Oakland, CA, April 1982. IEEE Computer Society.
- [8] J.A. Goguen and J. Meseguer. Inference control and unwinding. In *Proc. 1984 Symposium on Security and Privacy*, pages 75–86, Oakland, CA, April 1984. IEEE Computer Society.

- [9] J. Haigh and W. Young. Extending the non-interference model of MLS for SAT. In *Proc. 1986 Symposium on Security and Privacy*, pages 232–239, Oakland, CA, April 1986. IEEE Computer Society.
- [10] J. Thomas Haigh and William D. Young. Extending the noninterference version of MLS for SAT. *IEEE Transactions on Software Engineering*, SE-13(2):141–150, February 1987.
- [11] Jeremy Jacob. A note on the use of separability for the detection of covert channels. *Cipher—The Newsletter of the IEEE Technical Committee on Security and Privacy*, pages 25–33, Summer 1989.
- [12] Nancy L. Kelem and Richard J. Feiertag. A separation model for virtual machine monitors. In *Proc. 1991 Symposium on Security and Privacy*, Oakland, CA, May 1991. IEEE Computer Society. To appear.
- [13] John McLean. A comment on the “basic security theorem” of Bell and La Padula. *Information Processing Letters*, 20:67–70, 1985.
- [14] John McLean. Reasoning about security models. In *Proc. 1987 Symposium on Security and Privacy*, pages 123–131, Oakland, CA, April 1987. IEEE Computer Society.
- [15] J.K. Millen and C.M. Cerniglia. Computer security models. Working Paper WP25068, Mitre Corporation, Bedford, MA, September 1983.
- [16] Gerald J. Popek and David R. Farber. A model for verification of data security in operating systems. *Communications of the ACM*, 21(9):737–749, September 1978.
- [17] John Rushby. The design and verification of secure systems. In *Proc. 8th ACM Symposium on Operating System Principles*, pages 12–21, Asilomar, CA, December 1981. (*ACM Operating Systems Review*, Vol. 15, No. 5).
- [18] John Rushby. Verification of secure systems. Technical Report 166, Computing Laboratory, University of Newcastle upon Tyne, Newcastle upon Tyne, UK, August 1981.
- [19] John Rushby. Proof of Separability—a verification technique for a class of security kernels. In *Proc. 5th International Symposium on Programming*, pages 352–367, Turin, Italy, April 1982. Springer-Verlag Lecture Notes in Computer Science, Vol. 137.

- [20] John Rushby. The security model of Enhanced HDM. In *Proceedings 7th DoD/NBS Computer Security Initiative Conference*, pages 120–136, Gaithersburg, MD, September 1984.
- [21] John Rushby. Kernels for safety? In T. Anderson, editor, *Safe and Secure Computing Systems*, chapter 13, pages 210–220. Blackwell Scientific Publications, 1989. (Proceedings of a Symposium held in Glasgow, October 1986).
- [22] John Rushby. Formal verification of the unwinding theorem for intransitive noninterference security policies. Project report, Computer Science Laboratory, SRI International, Menlo Park, CA, March 1991. For Official Use Only.
- [23] John Rushby, Friedrich von Henke, and Sam Owre. An introduction to formal specification and verification using EHDM. Technical Report SRI-CSL-91-2, Computer Science Laboratory, SRI International, Menlo Park, CA, January 1991.
- [24] C.T. Sennett and R. Macdonald. Separability and security models. Technical Report 87020, Royal Signals and Radar Establishment, Malvern, UK, November 1987.
- [25] T. Taylor. Comparison paper between the Bell and La Padula model and the SRI model. In *Proc. 1984 Symposium on Security and Privacy*, pages 195–202, Oakland, CA, April 1984. IEEE Computer Society.

Bibliography

- [1] D.H. Barnes. The provision of security for user data on packet switched networks. In *Proc. 1983 IEEE Symposium on Security and Privacy*, pages 121–126, Oakland, CA, April 1983. IEEE Computer Society.
- [2] D.E. Bell and L.J. La Padula. Secure computer systems : Vol. I—mathematical foundations, Vol. II—a mathematical model, Vol III—a refinement of the mathematical model. Technical Report MTR-2547 (three volumes), Mitre Corporation, Bedford, MA, March–December 1973.
- [3] D.E. Bell and L.J. La Padula. Secure computer system: Unified exposition and Multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, Bedford, MA, March 1976.
- [4] W.E. Boebert and R.Y. Kain. A practical alternative to hierarchical integrity policies. In *Proceedings 8th DoD/NBS Computer Security Initiative Conference*, pages 18–27, Gaithersburg, MD, September 1985.
- [5] D.E. Denning. On the derivation of lattice structured information flow policies. Technical Report CSD TR 180, Purdue University, March 1976.
- [6] R.J. Feiertag, K.N. Levitt, and L. Robinson. Proving multilevel security of a system design. In *Proc. 6th ACM Symposium on Operating System Principles*, pages 57–65, November 1977.
- [7] J.A. Goguen and J. Meseguer. Security policies and security models. In *Proc. 1982 Symposium on Security and Privacy*, pages 11–20, Oakland, CA, April 1982. IEEE Computer Society.
- [8] J.A. Goguen and J. Meseguer. Inference control and unwinding. In *Proc. 1984 Symposium on Security and Privacy*, pages 75–86, Oakland, CA, April 1984. IEEE Computer Society.

- [9] J. Haigh and W. Young. Extending the non-interference model of MLS for SAT. In *Proc. 1986 Symposium on Security and Privacy*, pages 232–239, Oakland, CA, April 1986. IEEE Computer Society.
- [10] J. Thomas Haigh and William D. Young. Extending the noninterference version of MLS for SAT. *IEEE Transactions on Software Engineering*, SE-13(2):141–150, February 1987.
- [11] Jeremy Jacob. A note on the use of separability for the detection of covert channels. *Cipher—The Newsletter of the IEEE Technical Committee on Security and Privacy*, pages 25–33, Summer 1989.
- [12] Nancy L. Kelem and Richard J. Feiertag. A separation model for virtual machine monitors. In *Proc. 1991 Symposium on Security and Privacy*, Oakland, CA, May 1991. IEEE Computer Society. To appear.
- [13] John McLean. A comment on the “basic security theorem” of Bell and La Padula. *Information Processing Letters*, 20:67–70, 1985.
- [14] John McLean. Reasoning about security models. In *Proc. 1987 Symposium on Security and Privacy*, pages 123–131, Oakland, CA, April 1987. IEEE Computer Society.
- [15] J.K. Millen and C.M. Cerniglia. Computer security models. Working Paper WP25068, Mitre Corporation, Bedford, MA, September 1983.
- [16] Gerald J. Popek and David R. Farber. A model for verification of data security in operating systems. *Communications of the ACM*, 21(9):737–749, September 1978.
- [17] John Rushby. The design and verification of secure systems. In *Proc. 8th ACM Symposium on Operating System Principles*, pages 12–21, Asilomar, CA, December 1981. (*ACM Operating Systems Review*, Vol. 15, No. 5).
- [18] John Rushby. Verification of secure systems. Technical Report 166, Computing Laboratory, University of Newcastle upon Tyne, Newcastle upon Tyne, UK, August 1981.
- [19] John Rushby. Proof of Separability—a verification technique for a class of security kernels. In *Proc. 5th International Symposium on Programming*, pages 352–367, Turin, Italy, April 1982. Springer-Verlag Lecture Notes in Computer Science, Vol. 137.

- [20] John Rushby. The security model of Enhanced HDM. In *Proceedings 7th DoD/NBS Computer Security Initiative Conference*, pages 120–136, Gaithersburg, MD, September 1984.
- [21] John Rushby. Kernels for safety? In T. Anderson, editor, *Safe and Secure Computing Systems*, chapter 13, pages 210–220. Blackwell Scientific Publications, 1989. (Proceedings of a Symposium held in Glasgow, October 1986).
- [22] John Rushby. Formal verification of the unwinding theorem for intransitive noninterference security policies. Project report, Computer Science Laboratory, SRI International, Menlo Park, CA, March 1991. For Official Use Only.
- [23] John Rushby, Friedrich von Henke, and Sam Owre. An introduction to formal specification and verification using EHDM. Technical Report SRI-CSL-91-2, Computer Science Laboratory, SRI International, Menlo Park, CA, January 1991.
- [24] C.T. Sennett and R. Macdonald. Separability and security models. Technical Report 87020, Royal Signals and Radar Establishment, Malvern, UK, November 1987.
- [25] T. Taylor. Comparison paper between the Bell and La Padula model and the SRI model. In *Proc. 1984 Symposium on Security and Privacy*, pages 195–202, Oakland, CA, April 1984. IEEE Computer Society.