What, Indeed, Is Intransitive Noninterference? (Extended Abstract)*

Ron van der Meyden

School of Computer Science and Engineering,
University of New South Wales
meyden@cse.unsw.edu.au

Abstract. This paper argues that Haigh and Young's definition of noninterference for intransitive security policies admits information flows that are not in accordance with the intuitions it seeks to formalise. Several alternative definitions are discussed, which are shown to be equivalent to the classical definition of noninterference with respect to transitive policies. Rushby's unwinding conditions for intransitive noninterference are shown to be sound and complete for one of these definitions, TA-security. Access control systems compatible with a policy are shown to be TA-secure, and it is also shown that TA-security implies that the system can be interpreted as an access control system.

1 Introduction

In this paper, we present a new argument against Haigh and Young's [HY87, Rus92] definition of intransitive noninterference, showing that it is too weak for the intuitions it seeks to capture. We present an example that shows that it allows information to flow to an agent, that could not have come from the agents from which it is permitted to acquire information.

This leads us to consider alternative definitions. We show that there is in fact a spectrum of different definitions of noninterference for possibly intransitive policies, including two new notions TA-security and TO-security that we introduce, which are based on intuitions about the transmission of information about actions and observations, respectively. We then study these new definitions from the point of view of proof techniques and an application that have been held in the literature to be of significance for intransitive noninterference. We begin with a discussion of "unwinding conditions," which provide a proof technique for noninterference, but can be taken as a definition of security in their own right. Rushby proved that the classical unwinding conditions of Goguen and Meseguer provide a complete proof technique for noninterference in the transitive case. He proposes a weakening of these conditions for intransitive policies (correcting an earlier proposal by Haigh and Young [HY87]). He establishes soundness of the

* Thanks to the Courant Institute, New York University, for hosting a sabbatical visit during which this research was conducted. Work of the author supported by an Australian Research Council Discovery grant.

J. Biskup and J. Lopez (Eds.): ESORICS 2007, LNCS 4734, pp. 235–250, 2007.
© Springer-Verlag Berlin Heidelberg 2007

weakened unwinding conditions, but not completeness. We give an explanation of this: Rushby's conditions are not complete for the Haigh and Young definition of noninterference. Instead, they are sound and complete for the stronger notion of TA-security. There is a somewhat surprising subtlety in this statement: for completeness, the weakened unwinding conditions must be applied to the appropriate bisimilar system, but the existence of the weak unwindings is not preserved under bisimulation.

We also follow Rushby in considering the behaviour of the definitions on access control systems, the class of applications originally motivating the literature on noninterference. Rushby showed that access control systems satisfying a condition of structural consistency with a policy satisfy Haigh and Young's definition of intransitive noninterference. We argue that Rushby's definition of access control systems can be weakened, and that access control systems consistent with a policy satisfy the stronger notion of TA-security as well as Haigh and Young's definition of security. Moreover, we also show that TA-security implies that there is a way to interpret the system as an access control system in the weakened sense. This shows that TA-security is in some sense equivalent to the existence of an access control implementation of the system.

These results provide strong evidence that TA-security, rather than Haigh and Young's definition, best fits the original objectives for the notion of intransitive noninterference. Nevertheless, the stronger notion of TO-security may well be equally significant for practical purposes. As evidence of this, we prove that access control systems structurally consistent with a policy also satisfy the stronger notion of TO-security, provided we work with an appropriate notion of observation for such systems.

2 Intransitive Noninterference

The notion of *noninterference* was first proposed by Goguen and Meseguer [GM82]. Early work on this area was motivated by multi-level secure systems, and dealt with deterministic systems and partially ordered (hence transitive) information flow policies. A significant body of work has developed since then, with a particular focus on generalization to the case of nondeterministic systems [Sut86, WJ90, McC88, FG01, Rya01] and intransitive policies [Rus92, RG99, Ohe04]. We focus in this paper on intransitive policies in the deterministic case.

Several different types of semantic models have been used in the literature on noninterference. (See [MZ06] for a comparison and a discussion of their relationships.) We work here with the state-observed machine model used by Rushby [Rus92], but similar results would be obtained for other models. This model consists of deterministic machines of the form $\langle S, s_0, A, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$, where S is a set of states, $s_0 \in S$ is the *initial state*, A is a set of actions, $\mathtt{dom} : A \to D$ associates each action to an element of the set D of security domains, $\mathtt{step} : S \times A \to S$ is a deterministic transition function, and $\mathtt{obs} : S \times D \to O$ maps states to an observation in some set O, for each

security domain. We may also refer to security domains more succinctly as "agents". We write $s \cdot \alpha$ for the state reached by performing the sequence of actions $\alpha \in \mathit{Actions}^*$ from state s, defined inductively by $s \cdot \epsilon = s$, and $s \cdot \alpha a = \mathtt{step}(s \cdot \alpha, a)$ for $\alpha \in A^*$ and $a \in A$. Here ϵ denotes the empty sequence.

Noninterference policies, as they are now usually presented, are relations $\rightarrowtail \subseteq D \times D$, with $u \rightarrowtail v$ intuitively meaning that "actions of domain u are permitted to interfere with domain v", or "information is permitted to flow from domain u to domain v". Since, intuitively, a domain should be allowed to interfere with, or have information about, itself, this relation is assumed to be reflexive. In early work on noninterference, it is also assumed to be transitive.

Noninterference is given a formal semantics in the transitive case using a definition based on a "purge" function. Given a policy $\rightarrowtail$, we define the function $\mathtt{purge} : A^* \times D \to A^*$ by taking $\mathtt{purge}(\alpha, u)$ to be the subsequence of all actions a in α with $\mathtt{dom}(a) \rightarrowtail u$. (For clarity, we may use subscripting of agent arguments of functions, writing e.g., $\mathtt{purge}(\alpha, u)$ as $\mathtt{purge}_u(\alpha)$.) The system M is said to be *secure with respect to the transitive policy* $\rightarrowtail$ when for all $\alpha \in A^*$ and domains $u \in D$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \mathtt{purge}_u(\alpha))$. That is, each agent's observations are as if only interfering actions had been performed. An equivalent formulation (which we state more generally for policies that are not necessarily transitive, in anticipation of later discussion) is the following:

Definition 1. *A system M is* P-secure *with respect to a policy $\rightarrowtail$ if for all sequences $\alpha, \alpha' \in A^*$ such that $\mathtt{purge}_u(\alpha) = \mathtt{purge}_u(\alpha')$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \alpha')$.*

This can be understood as saying that agent u's observation depends only on the sequence of interfering actions that have been performed.

Haigh and Young [HY87] generalised the definition of the purge function to intransitive policies as follows. Intuitively, the intransitive purge of a sequence of actions with respect to a domain u is the largest subsequence of actions that could form part of a causal chain of effects (permitted by the policy) ending with an effect on domain u. More formally, the definition makes use of a function $\mathtt{sources} : A^* \times D \Rightarrow \mathcal{P}(D)$ defined inductively by $\mathtt{sources}(\epsilon, u) = \{u\}$ and

$$\mathtt{sources}(a\alpha, u) = \mathtt{sources}(\alpha, u) \cup \{\mathtt{dom}(a) \mid \exists v \in \mathtt{sources}(\alpha, u)(\mathtt{dom}(a) \rightarrowtail v)\}$$

for $a \in A$ and $\alpha \in A^*$. Intuitively, $\mathtt{sources}(\alpha, u)$ is the set of domains v such that there exists a sequence of permitted interferences from v to u within α. The *intransitive purge* function $\mathtt{ipurge} : A^* \times D \to A^*$ is then defined inductively by $\mathtt{ipurge}(\epsilon, u) = \epsilon$ and

$$\mathtt{ipurge}(a\alpha, u) = \begin{cases} a \cdot \mathtt{ipurge}(\alpha, u) & \text{if } \mathtt{dom}(a) \in \mathtt{sources}(a\alpha, u) \\ \mathtt{ipurge}(\alpha, u) & \text{otherwise} \end{cases}$$

for $a \in A$ and $\alpha \in A^*$. An alternative, equivalent formulation that we will find useful is the following: given a set $X \subseteq D$, define $\mathtt{ipurge}_X(\alpha)$ inductively by $\mathtt{ipurge}_X(\epsilon) = \epsilon$ and

$$\mathtt{ipurge}_X(\alpha a) = \begin{cases} \mathtt{ipurge}_{X \cup \{\mathtt{dom}(a)\}}(\alpha) \cdot a & \text{if } \mathtt{dom}(a) \rightarrowtail u \in X \\ \mathtt{ipurge}_X(\alpha) & \text{otherwise} \end{cases}$$

Then $\mathtt{ipurge}_u(\alpha)$ is identical to $\mathtt{ipurge}_{\{u\}}(\alpha)$. The intransitive purge function is then used in place of the purge function in Haigh and Young's definition:

Definition 2. *A system M is IP-secure with respect to a (possibly intransitive) policy $\rightarrowtail$ if for all sequences $\alpha \in A^*$, and $u \in D$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \mathtt{ipurge}_u(\alpha))$.*

Since the function $\mathtt{ipurge}_u$ on A^* is idempotent, this definition, like the definition for the transitive case, can be formulated as: M is IP-secure with respect to a policy $\rightarrowtail$ if for all $u \in D$ and all sequences $\alpha, \alpha' \in A^*$ with $\mathtt{ipurge}_u(\alpha) = \mathtt{ipurge}_u(\alpha')$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \alpha')$. It can be seen that $\mathtt{ipurge}_u(\alpha) = \mathtt{purge}_u(\alpha)$ when $\rightarrowtail$ is transitive, so IP-security is in fact a generalisation of the definition of security for transitive policies.

Roscoe and Goldsmith [RG99] (henceforth, RG) have argued that the Haigh and Young definition is incorrect. However, RG's arguments have not been universally accepted as compelling (see, e.g., [Ohe04]).

Nevertheless, we believe that a case can be made that IP-security is too weak, but on different grounds. Note that the intransitive purge $\mathtt{ipurge}_u(\alpha)$ preserves not just certain actions from the sequence α, but also their *order*. We claim that this allows u to "know" this order in situations where an intuitive reading of the policy would suggest that it ought not to know this order.

The notion of knowledge can be made precise using the the following notion of *view*. The definition uses an absorbtive concatenation function $\circ$, defined over a set X by, for $s \in X^*$ and $x \in X$, by $s \circ x = s$ if x is equal to the final element of s (if any), and $s \circ x = s \cdot x$ (ordinary concatenation) otherwise. Define the view of domain u with respect to a sequence $\alpha \in A^*$ using the function $\mathtt{view}_u : A^* \rightarrow (A \cup O)^*$ (where O is the set of observations in the system) defined by

$$\mathtt{view}_u(\epsilon) = \mathtt{obs}_u(s_0), \text{ and}$$
$$\mathtt{view}_u(\alpha a) = (\mathtt{view}_u(\alpha) \cdot b) \circ \mathtt{obs}_u(s_0 \cdot \alpha),$$

where $b = a$ if $\mathtt{dom}(a) = u$ and $b = \epsilon$ otherwise. That is, $\mathtt{view}_u(\alpha)$ is the sequence of all observations and actions of domain u in the run generated by α, compressed by the elimination of stuttering observations. Intuitively, $\mathtt{view}_u(\alpha)$ is the complete record of information available to agent u in the run generated by the sequence of actions α. The reason we apply the absorbtive concatenation is to capture that the system is asynchronous, with agents not having access to a global clock. Thus, two periods of different length during which a particular observation obtains are not distinguishable to the agent.

We may then say that agent u *knows* a fact ϕ about a sequence α if ϕ is true of all sequences α' such that $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$. Similarly, ϕ is *distributed knowledge* to a group G of agents in a sequence α if ϕ is true of all sequences α' such that $\mathtt{view}_u(\alpha) = \mathtt{view}_u(\alpha')$ for all $u \in G$. These are essentially the definitions of knowledge and distributed knowledge used in the literature on reasoning about knowledge [FHMV95], for an agent with *asynchronous perfect recall* . Intuitively, a fact is *distributed knowledge* to the set of agents G if it could be deduced after combining all the information that these agents have.

We may now present our example illustrating a weakness of IP-security. The essence of the example is that IP-security is consistent with an agent acquiring information that is not distributed knowledge to the agents from which it permitted (by an intransitive policy) to acquire information.

Example 1. Consider the intransitive policy $\rightarrowtail$ given by $H_1 \rightarrowtail D_1$, $H_2 \rightarrowtail D_2$, $D_1 \rightarrowtail L$ and $D_2 \rightarrowtail L$. Intuitively, H_1, H_2 are two High security domains, D_1, D_2 are two downgraders, and L is an aggregator of downgraded information. For this policy, channel control, one of the motivations for intransitive noninterference, would require that any information about L_1 and L_2 available to L must have reached L via the downgraders D_1 and D_2. We may capture this intuition more formally by expecting that if M is a system that is secure with respect to this policy, if a fact about H_1, H_2 is known to L, then it should be distributed knowledge to D_1, D_2. We show that if security is interpreted as $IP - security$, then this expectation can be false.

Define the system M with actions $A = \{h_1, h_2, d_1, d_2, l\}$ with domains H_1, H_2, D_1, D_2, and L, respectively. The set of states of M is the set of all strings in A^*. The transition function is defined by concatenation, i.e. for a state $\alpha \in A^*$ and an action $a \in A$, $\mathtt{step}(\alpha, a) = \alpha a$. The observation functions are defined using the ipurge function associated to the above policy: $\mathtt{obs}_u(\alpha) = [\mathtt{ipurge}(\alpha, u)]$. (Here we put brackets around the sequence of actions when it is interpreted as an observation, to distinguish such occurrences from the actions themselves as they occur in a view.)

It is plain that M is IP-secure. For, if $\mathtt{ipurge}(\alpha, u) = \mathtt{ipurge}(\alpha', u)$ then $\mathtt{obs}_u(s_0 \cdot \alpha) = [\mathtt{ipurge}(\alpha, u)] = [\mathtt{ipurge}(\alpha', u)] = \mathtt{obs}_u(s_0 \cdot \alpha')$.

Consider the sequences of actions $\alpha_1 = h_1 h_2 d_1 d_2$ and $\alpha_2 = h_2 h_1 d_1 d_2$. Note that these differ in the order of the events h_1, h_2. Let ϕ state that there is an occurrence of h_1 before an occurrence of h_2.

Then we have $\mathtt{obs}_L(\alpha_1) = [\mathtt{ipurge}(\alpha_1, L)] = [h_1 h_2 d_1 d_2]$. It follows that in α_1, agent L knows ϕ. We demonstrate that α_2 is a witness showing ϕ is not distributed knowledge to $\{D_1, D_2\}$ in α_1. Plainly α_2 does not satisfy ϕ so we need to show $\mathtt{view}_u(\alpha_1) = \mathtt{view}_u(\alpha_2)$ for $u \in \{D_1, D_2\}$. For this, note

$$
\begin{aligned}
&\mathtt{view}_{D_1}(\alpha_1)\\
&= \mathtt{obs}_{D_1}(\epsilon) \circ \mathtt{obs}_{D_1}(h_1) \circ \mathtt{obs}_{D_1}(h_1 h_2) \circ d_1 \circ \mathtt{obs}_{D_1}(h_1 h_2 d_1) \circ \mathtt{obs}_{D_1}(h_1 h_2 d_1 d_2)\\
&= [\epsilon] \circ [h_1] \circ [h_1] \circ d_1 \circ [h_1 d_1] \circ [h_1 d_1]\\
&= [\epsilon] \circ [\epsilon] \circ [h_1] \circ d_1 \circ [h_1 d_1] \circ [h_1 d_1]\\
&= \mathtt{obs}_{D_1}(\epsilon) \circ \mathtt{obs}_{D_1}(h_2) \circ \mathtt{obs}_{D_1}(h_2 h_1) \circ d_1 \circ \mathtt{obs}_{D_1}(h_2 h_1 d_1) \circ \mathtt{obs}_{D_1}(h_2 h_1 d_1 d_2)\\
&= \mathtt{view}_{D_1}(\alpha_2)
\end{aligned}
$$

The case for $u = D_2$ is symmetric. Thus, L has acquired information that cannot have come from the two sources D_1 and D_2 that are supposed to be, according to the policy, its only sources of information. □

Our example has a rather different character to those of RG. We believe that it more convincingly demonstrates that IP-security allows information flows that contradict the intuitive meaning of the policy, at the level of abstraction at which

the notion of noninterference is intended to operate (rather than the much more detailed level of abstraction to which RG tried to apply it.) We bolster the case for this claim in what follows, by showing that some alternative definitions are better behaved.

3 Alternative Definitions

As a response to Example 1, we consider several alternative definitions of security for intransitive policies.

To begin, let us consider why it was felt to be necessary to modify the definition of P-security for the intransitive case. For this, note that the system of Example 1 is not P-secure. For example, if we take $\alpha = h_1 d_1$ and $\alpha' = d_1$ then $\mathtt{purge}_L(\alpha) = d_1 = \mathtt{purge}_L(\alpha')$ but $\mathtt{obs}_L(s_0 \cdot \alpha) = [\alpha] \neq [\alpha'] = \mathtt{obs}_L(s_0 \cdot \alpha')$. However, this particular instance does not seem like it should be a counterexample to the security of the system. Intuitively, the fact that L observations differ on $s_0 \cdot \alpha$ and $s_0 \cdot \alpha'$ is justifiable, on the grounds that the action d_1 "downgrades" to L the fact that action h_1 has been performed. Thus, whereas IP-security is too weak, P-security seems to be too strong, since it does not permit an agent to forward information that it has acquired.

We are lead to propose two other definitions of security.[1] Both are based on a concrete model of the maximal amount of information that an agent may have after some sequence of actions has been performed, and state that an agent's observation may not give it more than this maximal amount of information. The definitions differ in the modelling of the maximal information, and take the view that an agent increases its information either by performing an action or by receiving information transmitted by another agent.

In the first model of the maximal information, what is transmitted when an agent performs an action is information about the actions performed by other agents. The following definition expresses this in a weaker way than the ipurge function.

Given sets X and A, let the set $\mathcal{T}(X, A)$ be the smallest set containing X and such that if $x, y \in \mathcal{T}(X, A)$ and $z \in A$ then $(x, y, z) \in \mathcal{T}(X, A)$. Intuitively, the elements of $\mathcal{T}(X, A)$ are are binary trees with leaves labelled from X and interior nodes labelled from A.

Given a policy $\rightarrowtail$, define, for each agent $u \in D$, the function $\mathtt{ta}_u : A^* \to \mathcal{T}(\{\epsilon\}, A)$ inductively by $\mathtt{ta}_u(\epsilon) = \epsilon$, and, for $\alpha \in A^*$ and $a \in A$,

1. if $\mathtt{dom}(a) \not\rightarrowtail u$, then $\mathtt{ta}_u(\alpha a) = \mathtt{ta}_u(\alpha)$,
2. if $\mathtt{dom}(a) \rightarrowtail u$, then $\mathtt{ta}_u(\alpha a) = (\mathtt{ta}_u(\alpha), \mathtt{ta}_{\mathtt{dom}(a)}(\alpha), a)$.

Intuitively, $\mathtt{ta}_u(\alpha)$ captures the maximal information that agent u may, consistently with the policy $\rightarrowtail$, have about the past actions of other agents. (The nomenclature is intended to be suggestive of *transmission* of information about

[1] The question of how exactly our definitions relate to RG's definitions is subtle and will be treated elsewhere.

actions.) Initially, an agent has information about what actions have been performed. The recursive clause describes how the maximal information $\mathtt{ta}_u(\alpha)$ permitted to u after the performance of α changes when the next action a is performed. If a may not interfere with u, then there is no change, otherwise, u's maximal permitted information is increased by adding the maximal information permitted to $\mathtt{dom}(a)$ at the time a is performed (represented by $\mathtt{ta}_{\mathtt{dom}(a)}(\alpha)$), as well the fact that a has been performed. Thus, this definition captures the intuition that an agent may only transmit information that it is permitted to have, and then only to agents with which it is permitted to interfere.

Definition 3. *A system M is TA-secure with respect to a policy $\rightarrowtail$ if for all agents u and all $\alpha, \alpha' \in A^*$ such that $\mathtt{ta}_u(\alpha) = \mathtt{ta}_u(\alpha')$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \alpha')$.*

Intuitively, this says that each agent's observations provide the agent with no more than the maximal amount of information that may have been transmitted to it, as expressed by the functions $\mathtt{ta}$.

Example 2. Note that the system of Example 1 is not TA-secure. For,

$$\begin{aligned}\mathtt{ta}_L(h_1h_2d_1d_2) &= (\mathtt{ta}_L(h_1h_2d_1), \mathtt{ta}_{D_2}(h_1h_2d_1), d_2)\\ &= ((\mathtt{ta}_L(h_1h_2), \mathtt{ta}_{D_1}(h_1h_2), d_1), \mathtt{ta}_{D_2}(h_1h_2), d_2)\\ &= ((\mathtt{ta}_L(h_1), \mathtt{ta}_{D_1}(h_1), d_1), (\mathtt{ta}_{D_2}(h_1), \mathtt{ta}_{H_2}(h_1), h_2), d_2)\\ &= ((\epsilon, (\epsilon, \epsilon, h_1), d_1), (\epsilon, \epsilon, h_2), d_2)\end{aligned}$$

and

$$\begin{aligned}\mathtt{ta}_L(h_2h_1d_1d_2) &= (\mathtt{ta}_L(h_2h_1d_1), \mathtt{ta}_{D_2}(h_2h_1d_1), d_2)\\ &= ((\mathtt{ta}_L(h_2h_1), \mathtt{ta}_{D_1}(h_2h_1), d_1), \mathtt{ta}_{D_2}(h_2h_1), d_2)\\ &= ((\mathtt{ta}_L(h_1), (\mathtt{ta}_{D_1}(h_2), \mathtt{ta}_{H_1}(h_2), h_1), d_1), \mathtt{ta}_{D_2}(h_2), d_2)\\ &= ((\epsilon, (\epsilon, \epsilon, h_1), d_1), (\epsilon, \epsilon, h_2), d_2).\end{aligned}$$

So $\mathtt{ta}_L(h_1h_2d_1d_2) = \mathtt{ta}_L(h_2h_1d_1d_2)$, but $\mathtt{obs}_L(h_1h_2d_1d_2) = [h_1h_2d_1d_2] \neq [h_2h_1d_1d_2] = \mathtt{obs}_L(h_2h_1d_1d_2)$. This illustrates that TA-security is in accordance with our intuitions about Example 1. □

The definition of TA-security has one aspect that might plausibly be questioned: it classifies as secure situations in which an agent transmits information to another that it has not actually observed. Whether one considers this to be a violation of security depends on one's attitude to forwarding of unobserved information. IP-security considers this acceptable, as does TA-security. However, it is possible to construct a definition that would consider this as insecure, by changing the definition of the function $\mathtt{ta}$.

Given a policy $\rightarrowtail$, for each domain $u \in D$, define the function $\mathtt{to}_u : A^* \rightarrow \mathcal{T}((A \cup O)^*, A)$ by $\mathtt{to}_u(\epsilon) = \mathtt{obs}_u(s_0)$ and

$$\mathtt{to}_u(\alpha a) = \begin{cases} \mathtt{to}_u(\alpha) & \text{when } \mathtt{dom}(a) \not\rightarrowtail u,\\ (\mathtt{to}_u(\alpha), \mathtt{view}_{\mathtt{dom}(a)}(\alpha), a) & \text{otherwise.}\end{cases}$$

Intuitively, this definition takes the model of the maximal information that an action a may transmit after the sequence α to be the fact that a has occurred, together with the information that $\mathtt{dom}(a)$ *actually* has, as represented by its view $\mathtt{view}_{\mathtt{dom}(a)}(\alpha)$. By contrast, TA-security uses in place of this the maximal information that $\mathtt{dom}(a)$ *may* have. (The nomenclature is intended to be suggestive of *transmission* of information about *observations*.) We may now base the definition of security on the function `to` rather than `ta`.

Definition 4. *The system M is TO-secure with respect to $\rightarrowtail$ if for all domains $u \in D$ and all $\alpha, \alpha' \in A^*$ with $\mathtt{to}_u(\alpha) = \mathtt{to}_u(\alpha')$, we have $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \alpha')$.*

It is possible to give a flatter representation of the information in $\mathtt{to}_u(\alpha)$ that clarifies the relationship of this definition to P-security. Define the *possibly transmitted view* of domain u for a sequence of actions α to be the largest prefix $\mathtt{tview}_u(\alpha)$ of $\mathtt{view}_u(\alpha)$ than ends in an action a with $\mathtt{dom}(a) = u$. Then we have the following result, which intuitively says that u's observations depend only on (1) the parts of the views of other agents which are permitted to pass information to u, that they have actually acted to transmit, and (2) u's knowledge of the ordering of its own actions and the actions of these other agents.

Proposition 1. *M is TO-secure with respect to a policy $\rightarrowtail$ iff for all sequences $\alpha, \alpha' \in A^*$, and domains $u \in D$, if $\mathtt{purge}_u(\alpha) = \mathtt{purge}_u(\alpha')$ and $\mathtt{tview}_v(\alpha) = \mathtt{tview}_v(\alpha')$ for all domains $v \neq u$ such that $v \rightarrowtail u$, then $\mathtt{obs}_u(s_0 \cdot \alpha) = \mathtt{obs}_u(s_0 \cdot \alpha')$.*

The following result describes how these definitions are related. Like IP-security, the notions P-security, TO-security and TA-security are generalizations of the classical notion of noninterference in the transitive case.

Theorem 1. *With respect to a given policy $\rightarrowtail$,*

1. *if M is P-secure then M is TO-secure,*
2. *if M is TO-secure then M is TA-secure,*
3. *if M is TA-secure then M is IP-secure, and*
4. *if $\rightarrowtail$ is transitive then M is P-secure iff M is TO-secure iff M is TA-secure iff M is IP-secure.*

4 Unwinding Relations

In this section we relate our alternative definitions of security for intransitive policies to "unwinding conditions" that have been discussed in the literature as a way to prove noninterference [GM84]. We show that Rushby's proposed unwinding conditions for intransitive noninterference are most closely related to the notion of TA-security (where they provide a sound and complete proof method), although they are also sufficient for TO-security in a special case. We also show the somewhat suprising fact that Rushby's unwinding conditions are not preserved under bisimulation.

We begin by recalling Rushby's results on unwinding for intransitive nonintereference. Suppose we have for each domain u an equivalence relation $\sim_u$ on the states of M. Rushby discusses the following "unwinding" conditions on such equivalence relations.

OC: If $s \sim_u t$ then $\mathtt{obs}_u(s) = \mathtt{obs}_u(t)$. (Output Consistency)
SC: If $s \sim_u t$ then $s \cdot a \sim_u t \cdot a$. (Step Consistency)
LR: If $\mathtt{dom}(a) \not\rightarrowtail u$ then $s \sim_u s \cdot a$. (Left Respect)

If these conditions are satisfied and $\rightarrowtail$ is a transitive policy, then M is P-secure [GM84]. Conversely, consider the particular equivalence relations $\approx_u$ on states, defined by $s \approx_u t$ if for all strings α in A^* we have $\mathtt{obs}_u(s \cdot \alpha) = \mathtt{obs}_u(t \cdot \alpha)$. Rushby uses these equivalence relations to show completeness of the unwinding conditions for transitive noninterference:

Proposition 2. *([Rus92] Theorem 6) Suppose M is P-secure with respect to the transitive policy $\rightarrowtail$. Then the relations $\approx_u$ satisfy OC, SC and LR.*

For intransitive noninterference he introduces the following condition:

WSC: If $s \sim_u t$ and $s \sim_{dom(a)} t$ then $s \cdot a \sim_u t \cdot a$. (Weak Step Consistency)

Define a *weak unwinding* on a system M with respect to a policy $\rightarrowtail$ to be a family of relations $\sim_u$, for $u \in D$, satisfying OC,WSC and LR. It will be convenient to have the following alternate characterization of this notion. Given a system M and a policy $\rightarrowtail$, let $\{\approx_u^{\mathrm{uw}}\}_{u \in D}$ be the smallest family of equivalence relations (under the pointwise containment order) satisfying WSC and LR.

Proposition 3. *There exists a weak unwinding for M with respect to $\rightarrowtail$ iff the relations $\approx_u^{\mathrm{uw}}$ satisfy OC.*

Rushby shows the following:

Proposition 4. *([Rus92], Theorem 7) Suppose that the relations $\{\sim_u\}_{u \in D}$ on a system M satisfy OC,WSC and LR. Then M is IP-secure for $\rightarrowtail$.*

However, he does not establish completeness of these unwinding conditions for IP-security. The following result yields an explanation of this fact.

Theorem 2. *Suppose that there exists a weak unwinding for M with respect to $\rightarrowtail$. Then M is TA-secure with respect to $\rightarrowtail$.*

Since, by Example 2, TA-security is stronger than IP-security, this result implies that the existence of equivalence relations $\sim_u$ satisfying conditions OC, WSC and LR is *not* a necessary condition for IP-security, since if this were the case, then every IP-secure system would be TA-secure.

This raises the question of whether the existence of weak unwindings is equivalent to TA-security instead. We now show that this question can be answered in the positive, provided it is formulated appropriately. The existence of weak unwindings turns out to have a somewhat surprising dependency on the structure of the system.

Given a system $M = \langle S, s_0, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$ with actions A, define the "unfolded" system $\mathtt{uf}(M) = \langle S', s'_0, \mathtt{step}', \mathtt{obs}', \mathtt{dom}\rangle$ with actions A having the same domains as in M, by $S' = A^*$, $s'_0 = \epsilon$, $\mathtt{step}'(\alpha, a) = \alpha a$, and $\mathtt{obs}'_u(\alpha) = \mathtt{obs}_u(s_0 \cdot \alpha)$, where $s_0 \cdot \alpha$ is computed in M. Intuitively, this construction unfolds the graph of M into an infinite tree. Then we have the following.

Theorem 3. *M is TA-secure with respect to $\rightarrowtail$ iff there exists a weak unwinding on $\mathtt{uf}(M)$ with respect to $\rightarrowtail$.*

It is reasonable to give a definition of security on M by reference to $\mathtt{uf}(M)$ since these systems are bisimilar under the obvious notion of bisimulation on the state-observed system model. Bisimilarity of two systems is usually taken to imply their equivalence on all properties of interest. One might therefore expect from Theorem 3 that TA-security implies the existence of a weak unwinding on the system M as well as on $\mathtt{uf}(M)$. It is the case that unwindings on M can be lifted to unwindings on $\mathtt{uf}(M)$.

Proposition 5. *If there exists a weak unwinding for $\rightarrowtail$ on M then there exists a weak unwinding for $\rightarrowtail$ on $\mathtt{uf}(M)$*

However, what we need, given Theorem 3, to deduce the existence of an unwinding on M from TA-security is the converse of this result. The following example shows that the converse does not hold. The reader may obtain some intuition for this example by noting that whereas weak unwinding seems to be sensitive to information about past actions, bisimulation cares only about the future. The essence of the example is that not enough past information is encoded in the states of the system M itself.

Example 3. Consider the system and policy depicted in Figure 1. There are actions a, b, c of domains $\mathtt{A}, \mathtt{B}, \mathtt{C}$ respectively, and s_0 is the initial state. For all domains u other than $\mathtt{D}$, we assume that the observation $\mathtt{obs}_u$ is the same on all states. TA-security therefore depends only on the behaviour of the system with respect to domain $\mathtt{D}$, where there are two possible observations o, o' as indicated. We show that there does not exist a weak unwinding for $\rightarrowtail$ on M, but there does exist one on $\mathtt{uf}(M)$.

For the former, we consider the relation family $\approx_u^{\mathrm{uw}}$ on M. Note that since $\mathtt{B} \not\rightarrowtail \mathtt{D}$ and $s_0 \cdot b = s_1$ we have by LR that $s_0 \approx_{\mathtt{D}}^{\mathrm{uw}} s_1$. Similarly, since $\mathtt{C} \not\rightarrowtail \mathtt{A}$ we have $s_0 \approx_{\mathtt{A}}^{\mathrm{uw}} s_1$. Hence, by WSC, for the action a, we get $s_0 \approx_{\mathtt{D}}^{\mathrm{uw}} s_2$. Since

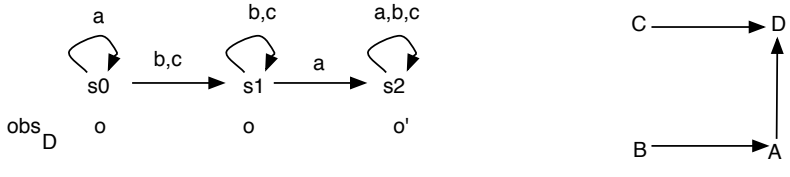

Fig. 1. An example showing TA-security does not imply existence of a weak unwinding

$\mathtt{obs}_{\mathtt{D}}(s_0) = o$ and $\mathtt{obs}_{\mathtt{D}}(s_2) = o'$, we have that $\approx_{\mathtt{D}}^{\mathrm{uw}}$ does not satisfy OC. Since $\approx_u^{\mathrm{uw}}$ is the smallest family satisfying WSC and LR, there can exist no weak unwinding for $\rightarrowtail$ on M.

For the unwinding on $\mathtt{uf}(M)$, consider $\approx_u^{\mathrm{uw}} = \sim_u^{\mathtt{ta}}$. Since this family of equivalence relations satisfies WSC and LR, it suffices to consider the property OC, where we need consider only the domain $\mathtt{D}$, as already noted. Here, the only possible failure of OC is for states α, α' where $\mathtt{ta}_{\mathtt{D}}(\alpha) = \mathtt{ta}_{\mathtt{D}}(\alpha')$, $s_0 \cdot \alpha \in \{s_0, s_1\}$ and $s_0 \cdot \alpha' = s_2$. Now $s_0 \cdot \alpha' = s_2$ implies that α' contains either a b and a later a, or a c and a later a. View $\mathtt{ta}_{\mathtt{D}}(\alpha')$ as a tree with nodes of the form (x, y, e) representing a vertex labelled e with subtrees corresponding to x and y. Then this tree contains a path from a leaf to the root containing either b and later a, or c and later a. The same then applies to the identical tree for $\mathtt{ta}_{\mathtt{D}}(\alpha)$, which implies that α contains either a b and later a or a c and later a. But this means that $s_0 \cdot \alpha = s_2$, a contradiction. Hence the family $\approx_u^{\mathrm{uw}}$ satisfies OC. □

Since $\mathtt{uf}(M)$ and M are bisimilar, this example shows that bisimulation does not preserve existence of a weak unwinding. It is therefore necessary to either abandon the presumption that security properties are preserved under bisimulation, or adopt the stance that existence of a weak unwinding (on the system as presented) is not a sensible notion of security. We prefer the latter, but note that this does not hinder the utility of weak unwinding as a proof technique.

Further evidence of the utility of weak unwinding is the following result, which shows that it can also be used as a proof technique for TO-security. Define the relations $\approx_u^{\mathtt{obs}}$ on states of a system M by $s \approx_u^{\mathtt{obs}} t$ if $\mathtt{obs}_u(s) = \mathtt{obs}_u(t)$. Then we have the following sufficient condition for TO-security:

Proposition 6. *Suppose the relation family $\approx_u^{\mathtt{obs}}$ is a weak unwinding on M with respect to $\rightarrowtail$. Then M is TO-secure with respect to $\rightarrowtail$.*

5 Access Control Systems

As a particular application of the unwinding conditions, Rushby [Rus92] discusses a notion of access control system that he formulates in order to give semantic content to the Bell-La Padula model [BP76] (which has been criticised for lacking semantics). He shows that every access control system satisfying a compatibility condition with respect to a noninterference policy is IP-secure. In this section, we formulate a weaker variant of Rushby's definitions, and show that it implies the stronger notion of TA-security. We also show that our weaker variant implies the even stronger notion of TO-security, provided we work with a specific, but intuitive, definition of observation in access control systems.

Moreover, we also show a converse to the result that access control systems are TA-secure, viz., that every system satisfying TA-security can be interpreted as an access control system. This proves the *equivalence* in some sense of access control and TA-security. We believe that these results, together with the example of Section 3 and the results of the previous section, provide strong evidence that

TA-security, rather than IP-security, is the notion that best realises the original objectives of the notion of intransitive noninterference.

According to Rushby, a *system with structured state* is a machine $\langle S, s_0, A, \mathtt{step}, \mathtt{obs}, \mathtt{dom}\rangle$ together with

1. a set N of *names*,
2. a set V of *values*, and functions
3. $\mathtt{contents} : S \times N \to V$, with $\mathtt{contents}(s, n)$ interpreted as the value of object n in state s,
4. $\mathtt{observe} : D \to \mathcal{P}(N)$, with $\mathtt{observe}(u)$ interpreted as the set of objects that domain u can observe, and
5. $\mathtt{alter} : D \to \mathcal{P}(N)$, with $\mathtt{alter}(u)$ interpreted as the set of objects whose values domain u is permitted to alter.

For a system with structured state, when $u \in D$ and s is a state, write $\mathtt{oc}_u(s)$ for the function mapping $\mathtt{observe}(u)$ to values, defined by $\mathtt{oc}_u(s)(n) = \mathtt{contents}(s, n)$ for $n \in \mathtt{observe}(u)$. Intuitively, $\mathtt{oc}_u(s)$ captures all the content of the state s that is observable to u. Using this, we may define a binary relation $\sim_u^{\mathtt{oc}}$ of *observable content equivalence* on S for each domain $u \in D$, by $s \sim_u^{\mathtt{oc}} t$ if $\mathtt{oc}_u(s) = \mathtt{oc}_u(t)$.

In order to capture the conditions under which the machine operates in accordance with the intuitive interpretations of this extra structure, Rushby defines the following three *Reference Monitor Assumptions*.

RM1. If $s \sim_u^{\mathtt{oc}} t$ then $\mathtt{obs}_u(s) = \mathtt{obs}_u(t)$.
RM2. If $s \sim_{dom(a)}^{\mathtt{oc}} t$ and either $\mathtt{contents}(s \cdot a, n) \neq \mathtt{contents}(s, n)$ or $\mathtt{contents}(t \cdot a, n) \neq \mathtt{contents}(t, n)$ then $\mathtt{contents}(s \cdot a, n) = \mathtt{contents}(t \cdot a, n)$
RM3. If $\mathtt{contents}(s \cdot a, n) \neq \mathtt{contents}(s, n)$ then $n \in \mathtt{alter}(dom(a))$.

The first of these says that an agent's observation depends only on the values of the objects observable to the agent. The third says that if an action can change the value of an object, then the agent of that action is in fact permitted to alter that object. The condition RM2 is more subtle. The following provides a possibly more perspicuous formulation of this condition:

Proposition 7. *RM2 is equivalent to the following: For all states s, either*

1. *for all $t \sim_{dom(a)}^{\mathtt{oc}} s$, we have $\mathtt{contents}(t \cdot a, n) = \mathtt{contents}(t, n)$, or*
2. *for all $t \sim_{dom(a)}^{\mathtt{oc}} s$, we have $\mathtt{contents}(s \cdot a, n) = \mathtt{contents}(t \cdot a, n)$*

That is, with the choice depending only on information observable to $dom(a)$, the effect of the action is either to make no change to n or to assign a new value to n that depends only on information observable to $dom(a)$.

In addition to the reference monitor assumptions, Rushby considers the condition:

AOI. If $\mathtt{alter}(u) \cap \mathtt{observe}(v) \neq \emptyset$ then $u \rightarrowtail v$.

Intuitively, this says that the ability to write to a value that an agent can observe counts as a way to interfere with that agent. Rushby shows the following:

Proposition 8. *([Rus92], Theorems 2,8) Suppose M is a system with structured state that satisfies RM1-RM3 and AOI. Then the family of relations $\sim_u^{\mathtt{oc}}$ on M is a weak unwinding with respect to $\rightarrowtail$. Hence M is IP-secure for $\rightarrowtail$.*

By the results of the previous section, Rushby's result in fact yields the stronger conclusion that access control systems consistent with a policy are TA-secure. We can further strengthen this result by weakening the precondition.

Note that the condition RM2 says that the next value of n produced on performing an action a depends only on the values of names observable to $\mathtt{dom}(a)$. If n is not observable to $\mathtt{dom}(a)$, this may be too strong. Consider, for example, the situation where n represents a block of memory, and the action a writes to a single location within this block. Here the successor value depends on the value written (which will typically depend on the values of names observable to $\mathtt{dom}(a)$), but also on the previous value of n. Similarly, if the name n is an object in an object-oriented system, and the effect of the action is to call a method of this object, then the successor value will depend of the input parameters of the call (which will depend on values of names observable to $\mathtt{dom}(a)$), but also on the value of n. Thus, the condition RM2 can plausibly be weakened to the following.

> [RM2′] For all actions a, states s, t and names $n \in \mathtt{alter}(dom(a))$, if $s \sim_{\mathtt{dom}(a)}^{\mathtt{oc}} t$ and $\mathtt{contents}(s, n) = \mathtt{contents}(t, n)$ we have $\mathtt{contents}(s{\cdot}a, n) = \mathtt{contents}(t \cdot a, n)$.

That is, for $n \in \mathtt{alter}(dom(a))$, the value $\mathtt{contents}(s \cdot a, n)$ is a function of both $\mathtt{contents}(s, n)$ and $\mathtt{oc}_{\mathtt{dom}(a)}(s)$. Using Proposition 7 it can be seen that RM2 implies RM2′. The converse does not hold.

We now weaken Rushby's notion of access control system by replacing RM2 by RM2′. We define a system with structured states to be a *weak access control system* if it satisfies conditions RM1,RM2′, and RM3.

We also introduce a related notion on systems without structured states, that expresses that the system behaves as if it were an access control system. Say that a system M with states S *admits a weak access control implementation consistent with* $\rightarrowtail$ if there exists a set of names N, a set of values V and functions $\mathtt{observe} : D \times S \rightarrow \mathcal{P}(N)$, $\mathtt{alter} : D \times S \rightarrow \mathcal{P}(N)$ and $\mathtt{contents} : N \times S \rightarrow V$, with respect to which M is a weak access control system satisfying the condition AOI.

The following shows that weak access control systems compatible with a policy satisfy Rushby's unwinding conditions for intransitive noninterference:

Proposition 9. *Suppose M is a weak access control system consistent with $\rightarrowtail$. Then the family of relations $\sim_u^{\mathtt{oc}}$ is a weak unwinding on M with respect to $\rightarrowtail$.*

We may also show a converse to this result, which leads to the conclusion that unwinding and weak access control systems are essentially equivalent.

Proposition 10. *Suppose that there exists a weak unwinding on M with respect to $\rightarrowtail$. Then M admits a weak access control interpretation consistent with $\rightarrowtail$.*

Combining these results with those of the previous section, we see that there is a close correspondence between TA-security, weak access control interpretations, and weak unwindings.

Corollary 1. *The following are equivalent*

1. *M is TA-secure with respect to $\rightarrowtail$,*
2. *$\mathtt{uf}(M)$ admits a weak access control interpretation consistent with $\rightarrowtail$,*
3. *there exists a weak unwinding on $\mathtt{uf}(M)$ with respect to $\rightarrowtail$.*

From Theorem 2 and Proposition 9, we also obtain the following.

Corollary 2. *If M is a weak access control system consistent with $\rightarrowtail$ then M is TA-secure for $\rightarrowtail$.*

This conclusion is a more general result than Proposition 8, in which we have both weakened the antecedent and strengthened the consequent. The following example shows that we cannot further strengthen the conclusion to TO-security.

Example 4. Consider the system for the policy $A \rightarrowtail B \rightarrowtail C$ with structured states for the set of names n_{AB}, n_{BC}, taking boolean values. Intuitively, these variables represent channels between the agents, so that $n_{AB} \in \mathtt{alter}(A) \cap \mathtt{observe}(B)$ and $n_{BC} \in \mathtt{alter}(B) \cap \mathtt{observe}(C)$. Plainly this is consistent with AOI. We represent states as tuples $s = (n_{AB}, n_{BC})$ with the obvious interpretation for `contents`. The initial state of the system is $(0, 0)$. Domain A has actions a with semantics $n_{AB} := 1$ and B has action b with semantics $n_{BC} := n_{AB}$. The observation functions are defined on the state $s = (n_{AB}, n_{BC})$ by $\mathtt{obs}_A(s) = \mathtt{obs}_B(s) = \bot$ and $\mathtt{obs}_C(s) = n_{BC}$. It can be verified that this system satisfies RM1,RM2′, RM3. However, it does not satisfy TO-security. To see this, consider the sequences $\alpha = b$ and $\alpha' = ab$. Here we have $\mathtt{purge}_C(\alpha) = b = \mathtt{purge}_C(\alpha')$, and $\mathtt{tview}_B(\alpha) = \bot b = \mathtt{tview}_B(\alpha')$ but $\mathtt{obs}_C(s_0 \cdot \alpha) = 0 \neq 1 = \mathtt{obs}_C(s_0 \cdot \alpha')$. □

Notice that in this example, not all of the names observable to a domain have their contents visible in the observation of the domain. Say that a system with structured states is *fully observable* if in all states s we have $\mathtt{obs}_u(s) = \mathtt{oc}_u(s)$. Note that this means that the relations $\sim_u^{\mathtt{oc}}$ and $\approx_u^{\mathtt{obs}}$ coincide. We now obtain the following from Propositions 6 and 9. This shows that, modulo the reasonable assumption of full observability, we can derive a result similar to Corollary 2, but with the yet stronger conclusion of TO-security.

Corollary 3. *If M is a fully observable weak access control system consistent with $\rightarrowtail$ then M is TO-secure with respect to $\rightarrowtail$.*

A similar result does not hold with P-security in place of TO-security.

6 Conclusion

Our results have left open a number of technical questions. We have shown that weak unwindings provide a complete proof technique for TA-security, but

have not provided a complete technique for TO-security. The reason for this is that there is inherently no tractable set of conditions on the states of the system that characterizes TO-security. We will treat this topic in a followup paper [Mey07] which deals with the complexity of the notions of security discussed in this paper. Another area requiring investigation is the generalization of our definitions to nondeterministic systems and systems that are not input-enabled, as has been studied for IP-security by von Oheimb [Ohe04]. More generally, one could consider extensions to the richer semantic framework of process algebra.

Both the fact, as argued by RG, that the notion of (intransitive) noninterference on its own falls short of expressing the correctness properties of downgraders that they sought to capture, and the fact, as we have shown, that there are several plausible notions of noninterference for intransitive policies, suggests that the notion of noninterference policy expressed by a relation $\rightarrowtail$ on domains lacks expressiveness that will be required in applications. We believe further work on richer formats for the expression of causality and information flow policies is warranted. The approach we have followed in this paper, of comparing an agent's actual information to an intuitive concrete operational model of the maximal information that an agent is permitted to have and transmit, could well be useful in this enterprise.

The specific case of downgrading policies has received some recent attention. Chong and Myers [CM04] have proposed a flexible language that attaches downgrading conditions to data items. Mantel and Sands [MS04] have proposed to introduce a programming annotation for downgrading, enabling the programmer to explicitly mark regions of code that are permitted to violate a transitive policy. They apply a definition based on IP-security. Bossi et al [BPR04] develop a theory of downgrading grounded in bisimulation-based notions of unwinding. Sabelfeld and Sands [SS05] lay out some general principles and direction for research in this area. It would be of interest to reconsider these contributions in the light of our results in this paper.

References

[BP76] Bell, D.E., La Padula, L.J.: Secure computer system: unified exposition and multics interpretation. Technical Report ESD-TR-75-306, Mitre Corporation, Bedford, MA (March 1976)

[BPR04] Bossi, A., Piazza, C., Rossi, S.: Modelling downgrading in information flow security. In: Proc. IEEE Computer Security Foundations Workshop, pp. 187–201. IEEE Computer Society Press, Los Alamitos (2004)

[CM04] Chong, S., Myers, A.C.: Security policies for downgrading. In: 11th ACM Conf. on Computer and Communications Security (CCS), ACM Press, New York (October 2004)

[FG01] Focardi, R., Gorrieri, R.: Classification of security properties (Part I: information flow). In: Focardi, R., Gorrieri, R. (eds.) Foundations of Security Analysis and Design. LNCS, vol. 2171, pp. 331–396. Springer, Heidelberg (2001)

[FHMV95] Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning About Knowledge. MIT Press, Cambridge (1995)

[GM82] Goguen, J.A., Meseguer, J.: Security policies and security models. In: Proc. IEEE Symp. on Security and Privacy, Oakland, pp. 11–20. IEEE Computer Society Press, Los Alamitos (1982)

[GM84] Goguen, J.A., Meseguer, J.: Unwinding and inference control. In: IEEE Symp. on Security and Privacy, IEEE Computer Society Press, Los Alamitos (1984)

[HY87] Haigh, J.T., Young, W.D.: Extending the noninterference version of MLS for SAT. IEEE Trans. on Software Engineering SE-13(2), 141–150 (1987)

[MB94] Moses, Y., Bloom, B.: Knowledge, timed precedence and clocks (preliminary report). In: Proc. ACM Symp. on Principles of Distributed Computing, pp. 294–303. ACM Press, New York (1994)

[McC88] McCullough, D.: Noninterference and the composability of security properties. In: Proc. IEEE Symp. on Security and Privacy, pp. 177–186. IEEE Computer Society Press, Los Alamitos (1988)

[Mey07] van der Meyden, R.: The complexity of notions of intransitive noninterference (unpublished manuscript, 2007)

[MS04] Mantel, H., Sands, D.: Controlled declassification based on intransitive noninterference. In: Chin, W.-N. (ed.) APLAS 2004. LNCS, vol. 3302, pp. 129–145. Springer, Heidelberg (2004)

[MZ06] van der Meyden, R., Zhang, C.: A comparison of semantic models for noninterference. In: Proc. Workshop on Formal Aspects of Security and Trust, Hamilton, Ontario, Canada, August 2006. LNCS, Springer, Heidelberg (to appear), extended version at `http://www.cse.unsw.edu.au/~meyden/research/publications.html`

[Ohe04] von Oheimb, D.: Information flow control revisited: Noninfluence = Noninterference + Nonleakage. In: Samarati, P., Ryan, P.Y.A., Gollmann, D., Molva, R. (eds.) ESORICS 2004. LNCS, vol. 3193, pp. 225–243. Springer, Heidelberg (2004)

[RG99] Roscoe, A.W., Goldsmith, M.H.: What is intransitive noninterference? In: IEEE Computer Security Foundations Workshop, pp. 228–238. IEEE Computer Society Press, Los Alamitos (1999)

[Rus92] Rushby, J.: Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International (December 1992)

[Rya01] Ryan, P.Y.: Mathematical models of computer security. In: Focardi, R., Gorrieri, R. (eds.) Foundations of Security Analysis and Design. LNCS, vol. 2171, pp. 1–62. Springer, Heidelberg (2001)

[SS05] Sabelfeld, A., Sands, D.: Dimensions and principles of declassification. In: Proceedings of the 18th IEEE Computer Security Foundations Workshop, pp. 255–269. IEEE Computer Society Press, Los Alamitos (2005)

[Sut86] Sutherland, D.: A model of information. In: Proc. 9th National Computer Security Conf., pp. 175–183 (1986)

[WJ90] Wittbold, J.T., Johnson, D.M.: Information flow in nondeterministic systems. In: IEEE Symposium on Security and Privacy, pp. 144–161. IEEE Computer Society Press, Los Alamitos (1990)