

What, Indeed, is Intransitive Noninterference? (Preliminary Report)*

Ron van der Meyden

School of Computer Science and Engineering,
University of New South Wales
`meijden@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 also 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

The term ‘noninterference’ is used in the computer security literature to refer to formal definitions of information flow or causality between security domains. The classical theory of noninterference [GM82] dealt with transitive policies, which are closely related to partially ordered security levels. This theory is unable to deal with certain systems requiring channel control and downgrading [Rus81]. To overcome these limitations, Haigh and Young [HY87] proposed a variant of the classical definition for channel control applications, which was further promulgated by Rushby [Rus92] for intransitive policies more generally. Roscoe and Goldsmith [RG99] have argued that the variant definition is in fact inadequate for dealing with downgrading, and have proposed an alternate definition. Their arguments do not seem to have been universally accepted, however, and the Haigh and Young definition remains in use and continues to be the subject of research [HALL⁺05,MS04,Ohe04].

In this paper, we present a new argument against this definition, showing that it may be 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

* Version of February 6, 2007. Thanks to the Courant Institute, New York University, for hosting a sabbatical visit during which this work was conducted.

policies, all having some intuitive plausibility, and all equivalent to the classical definition in the case of transitive policies. The weakest of these definitions is the classical purge-based definition, the strongest is Haigh and Young’s definition. We define two notions of intermediate strength. These definitions state that the agent should not have more information than the maximal amount it is permitted to have, but differ in the modelling of the maximal amount of information that may be transmitted by an action. One of these notions, TO-security, states that an agent may not transmit information that it has not directly observed. The other notion, TA-security, takes a more liberal view, in which an agent’s actions may transmit information that it is permitted to have, even if it has not actually observed this information.

We include 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 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. Further, 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.

The structure of the paper is as follows. Section 2 recalls the Haigh and Young definition, and presents our new argument against the definition. In Section 3 we introduce two new definitions of noninterference for intransitive noninterference, and show how they are related. We discuss unwinding conditions in Section 4. The results on access control systems are presented in Section 5. We close in Section 6 with some remarks concerning related research and issues requiring investigation.

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.¹ This model consists of deterministic machines of the form $\langle S, s_0, A, \mathbf{step}, \mathbf{obs}, \mathbf{dom} \rangle$, where S is a set of states, $s_0 \in S$ is the *initial state*, A is a set of actions, $\mathbf{dom} : A \rightarrow D$ associates each action to an element of the set D of security domains, $\mathbf{step} : S \times A \rightarrow S$ is a deterministic transition function, and $\mathbf{obs} : S \times D \rightarrow O$ 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 $s \cdot a\alpha = \mathbf{step}(s \cdot \alpha, a)$ for $\alpha \in A^*$ and $a \in A$. Here ϵ denotes the empty sequence.

Noninterference policies, as they are now usually presented², are relations $\rightsquigarrow \subseteq D \times D$, with $u \rightsquigarrow v$ 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 \rightsquigarrow , we define the function $\mathbf{purge} : A^* \times D \rightarrow A^*$ such that $\mathbf{purge}(\alpha, u)$ is the subsequence of all actions a in α such that $\mathbf{dom}(a) \rightsquigarrow u$. (For clarity, we may use subscripting of agent arguments of functions, writing e.g., $\mathbf{purge}(\alpha, u)$ as $\mathbf{purge}_u(\alpha)$.) The system M is then said to be *secure with respect to the policy* \rightsquigarrow when for all $\alpha \in A^*$ and domains $u \in D$, we have $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \mathbf{purge}_u(\alpha))$. That is, each agent’s observations are as if only interfering actions had been performed. An equivalent formulation is the following: for all sequences $\alpha, \alpha' \in A^*$ such that $\mathbf{purge}_u(\alpha) = \mathbf{purge}_u(\alpha')$, we have $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \alpha')$. This can be understood as saying that agent u ’s observation depends only on the sequence of interfering actions that have been performed.

¹ In a subsequent version of this paper we will also treat Rushby’s action-observed model, and show that the corresponding definitions in that model are related to those in the state-observed model by means of a natural mapping from action observed systems to state-observed systems.

² Goguen and Meseguer used a slightly richer notion.

Haigh and Young [HY87] generalised the definition of the purge function to intransitive policies as follows. Intuitively, given a sequence of actions $a_1 \dots a_n \in A^*$, the intransitive purge of this sequence with respect to a domain u is the largest subsequence $b_1 \dots b_k$ such that for each $i < k$, either $b_i \rightsquigarrow u$ or for some $j \leq k$ we have $\text{dom}(b_i) \rightsquigarrow \text{dom}(b_j)$ and $i < j$. More formally, the definition makes use of a function $\text{sources} : A^* \times D \Rightarrow \mathcal{P}(D)$ defined inductively by $\text{sources}(\epsilon, u) = \{u\}$ and

$$\text{sources}(a \cdot \alpha, u) = \text{sources}(a \cdot \alpha, u) \cup \{\text{dom}(a) \mid \exists v \in \text{sources}(v, \alpha)(\text{dom}(a) \rightsquigarrow v)\}$$

for $a \in A$ and $\alpha \in A^*$. Intuitively, $\text{sources}(\alpha, u)$ 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 $\text{ipurge} : A^* \times D \rightarrow A^*$ is then defined inductively by $\text{ipurge}(\epsilon, u) = \epsilon$ and

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

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 $\text{ipurge}_X(\alpha)$ inductively by $\text{ipurge}_X(\epsilon) = \epsilon$ and

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

Then $\text{ipurge}_u(\alpha)$ is identical to $\text{ipurge}_{\{u\}}(\alpha)$.

Using the intransitive purge function, a system M is defined to be secure with respect to a possibly intransitive policy \rightsquigarrow if for all sequences $\alpha \in A^*$, and $u \in D$, we have $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \text{ipurge}_u(\alpha))$. We will refer to this notion as *IP-security*. Since the function 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 possibly intransitive policy \rightsquigarrow if for all $u \in D$ and all sequences $\alpha, \alpha' \in A^*$ with $\text{ipurge}_u(\alpha) = \text{ipurge}_u(\alpha')$, we have $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$. It can be seen that $\text{ipurge}_u(\alpha) = \text{purge}_u(\alpha)$ when \rightsquigarrow 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. They present a number of concrete examples to make this case, including a policy $H \rightsquigarrow D \rightsquigarrow L$, intended to represent that D is a “downgrader” process that decides which of the High (H) secrets may safely be revealed to a Low process L . The essence of their argument is that the definition has the effect that a downgrader D ’s action permits *all* information about preceding H actions to become known to L , even if the intent of the downgrading action was to release only some specific information about the preceding H actions.

That is, this definition does not enable the intent of specific downgrading actions to be specified. But it might be countered in that intransitive noninterference was not intended to make such fine grained distinctions, but only to

express some coarse architectural constraints on information flow. In RG’s examples, the information released to L is information that was known to D at the time of release. One might sensibly hold the view that therefore RG’s examples are not counterexamples to the intuitive reading of the policy. Indeed, 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 $\text{ipurge}_u(\alpha)$ 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 an agent’s knowledge in a system can be given formal meaning using 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 $\alpha \in A^*$ by $\text{view}_u(\epsilon) = \text{obs}_u(s_0)$, and $\text{view}_u(\alpha \cdot a) = (\text{view}_u(\alpha) \cdot b) \circ \text{obs}_u(s_0 \cdot \alpha)$, where $b = a$ if $\text{dom}(a) = u$ and $b = \epsilon$ otherwise. That is, $\text{view}_u(\alpha)$ 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, $\text{view}_u(\alpha)$ 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.

Using the notion of view, we may define for each agent u an equivalence relation \equiv_u on sequences of actions by $\alpha \equiv_u \alpha'$ if $\text{view}_u(\alpha) = \text{view}_u(\alpha')$. We may then say that in the run generated by a sequence of actions α , agent u *knows* a fact ϕ about α if ϕ is true of all sequences α' such that $\alpha \equiv_u \alpha'$. This is essentially the definition of knowledge for an agent with *asynchronous perfect recall* used in the literature on reasoning about knowledge [FHMV95]. Another notion used in this literature is the notion of the *distributed knowledge*³ of a group of agents. Intuitively, a fact is distributed knowledge to a set of agents G if it could be deduced after combining all the information that these agents have. More formally, we may define the relations \equiv_G on sequences of actions by $\alpha \equiv_G \alpha'$ if $\alpha \equiv_u \alpha'$ for all $u \in G$. A fact ϕ is then distributed knowledge to group G in a sequence α if it holds of all sequences α' such that $\alpha \equiv_G \alpha'$.

³ Distributed knowledge has been argued to be too strong a notion to apply in asynchronous systems, such as those we consider in this paper, and a weaker notion called *inherent knowledge* has been proposed [MB94]. However, the notion of distributed knowledge will suffice for our purposes since our arguments would also hold if distributed knowledge were to be replaced by a weaker notion. Indeed, we will argue in a longer version of the paper that in the present setting, it is appropriate to consider a novel notion of group knowledge (which relies on the group knowing the order of the actions of agents within the group) which is yet stronger than distributed knowledge.

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 \succrightarrow given by $H_1 \succrightarrow D_1$, $H_2 \succrightarrow D_2$, $D_1 \succrightarrow L$ and $D_2 \succrightarrow L$. 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.

Define the system M with actions $A = \{h_1, h_2, d_1, d_2, l\}$ of domains H_1, H_2, D_1, D_2, 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 $a \in A$, $\text{step}(\alpha, a) = \alpha a$. The observation functions are defined using the ipurge function associated to the above policy: $\text{obs}_u(\alpha) = [\text{ipurge}(\alpha, u)]$. (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 $\text{ipurge}(\alpha, u) = \text{ipurge}(\alpha', u)$ then $O_u(s_0 \cdot \alpha) = [\text{ipurge}(\alpha, u)] = [\text{ipurge}(\alpha', u)] = O_u(s_0 \cdot \alpha')$. We claim that the system nevertheless has an undesirable information flow.

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 . Then we have $\text{obs}_L(\alpha_1) = [\text{ipurge}(\alpha_1, L)] = [\alpha_1]$, so in α_1 agent L knows the entire run, and in particular, the ordering of the events h_1, h_2 . However, $\text{obs}_{D_1}(\alpha_1) = [\text{ipurge}(\alpha_1, D_2)] = [h_1 d_1] = \text{obs}_{D_1}(\alpha_2)$, so the agent's final observation does not enable it to distinguish α_1 from α_2 . Moreover, this holds even when we consider views, since

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

i.e., $\alpha_1 \equiv_{D_1} \alpha_2$. By symmetry, we also have $\alpha_1 \equiv_{D_2} \alpha_2$, hence $\alpha_1 \equiv_{\{D_1, D_2\}} \alpha_2$. This means that D_1 and D_2 do not have distributed knowledge of the ordering of the events h_1, h_2 , even with respect to the asynchronous perfect recall interpretation of knowledge, in which they reason based on everything that they learn in the run.

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. \square

This 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.)

3 Alternative Definitions

Roscoe and Goldsmith [RG99], follow their criticism of Rushby’s definition with a proposed alternative definition. They work in the framework of the process algebra CSP, which lacks a distinction between observations and actions, so it is not entirely clear how best to translate this proposal into the framework of state machines. Ignoring the role of observations for the moment, RG’s Theorem 1 suggests the following translation:

Definition: M is P-secure with respect to a policy \succrightarrow if for all sequences $\alpha, \alpha' \in A^*$ and domains $u \in D$ if $\text{purge}_u(\alpha) = \text{purge}_u(\alpha')$, then $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$.

Intuitively, this says that u ’s observations depend only on the prior actions in domains that may interfere with u . Observe that this is precisely the variant of the classical definition of security discussed above, so it is obvious that this definition, like IP-security, is a generalisation of the classical definition for transitive policies.

Note that the system of Example 1 is not P-secure. For example, if we take $\alpha = h_1 h_2 d_1$ and $\alpha' = h_2 h_1 d_1$ then $\text{purge}_L(\alpha) = d_1 = \text{purge}_L(\alpha')$ but $\text{obs}_L(\alpha) = \alpha \neq \alpha' = \text{obs}_L(\alpha')$.

However (and, presumably, this is why it was felt necessary to vary the definition in the intransitive case), P-security may be too strong, since it does not permit an agent to transmit to others information that it has acquired from others. The only way to enable this consistently with the definition of P-security is to define a distinct action for each piece of observed information that is to be transmitted. While this is precisely what RG advocate, based on their aim to express detailed reasoning about downgrading, this entails a very high modelling overhead, and it is not possible to define such actions in the context of input-enabled systems, where they cannot be correlated with the information they are meant to transmit.

We therefore propose two other definitions of security.⁴ 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 agents 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, we take the view that what is transmitted when an agent performs an action is information about the actions

⁴ The question of how exactly our definitions relate to RG’s definition requires a treatment of mappings between state machine models and CSP, for which there are a number of plausible candidates. In a longer version of the paper we show that one of these yields that RG’s definition is in fact P-security, but on another RG’s definition is not equivalent to any of ours: the difference relates to the treatment of an interfering agent’s observations after its last action.

performed by other agents. The following definition does so in a weaker way than the `ipurge` function.

Let the set $HT(A^*)$ of *hierarchical triples over A^** be the smallest set containing A^* and such that if $x, y, z \in HT(A^*)$ then $(x, y, z) \in HT(A^*)$. Given a policy \rightsquigarrow , define, for each agent $u \in D$, the function $\mathbf{ta}_u : A^* \rightarrow HT(A^*)$ inductively by $\mathbf{ta}_u(\epsilon) = \epsilon$, and, for $\alpha \in A^*$ and $a \in A$,

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

Intuitively, $\mathbf{ta}_u(\alpha)$ captures the maximal information that agent u may, consistently with the policy \rightsquigarrow , have about the past actions of other agents. In particular, the definition says that nothing new is transmitted to u by an action that may not interfere with u , and an action a that may interfere with u may add to u 's information $\mathbf{ta}_u(\alpha)$ everything that is potentially known to $\text{dom}(a)$, represented by $\mathbf{ta}_{\text{dom}(a)}(\alpha)$, as well as the fact that the action a has been performed.

We now define a system M to be TA-secure with respect to a policy \rightsquigarrow if for all agents u and all $\alpha, \alpha' \in A^*$ such that $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$, we have $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$. 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 \mathbf{ta} .

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

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

and

$$\begin{aligned} \mathbf{ta}_L(h_2 h_1 d_1 d_2) &= (\mathbf{ta}_L(h_2 h_1 d_1), \mathbf{ta}_{D_2}(h_2 h_1 d_1), d_2) \\ &= ((\mathbf{ta}_L(h_2 h_1), \mathbf{ta}_{D_1}(h_2 h_1), d_1), \mathbf{ta}_{D_2}(h_2 h_1), d_2) \\ &= ((\mathbf{ta}_L(h_1), (\mathbf{ta}_{D_1}(h_1), \mathbf{ta}_{H_1}(h_2), h_1), d_1), \mathbf{ta}_{D_2}(h_2), d_2) \\ &= ((\epsilon, (\epsilon, \epsilon, h_1), d_1), (\epsilon, \epsilon, h_2), d_2) \end{aligned}$$

So $\mathbf{ta}_L(h_1 h_2 d_1 d_2) = \mathbf{ta}_L(h_2 h_1 d_1 d_2)$, but $O_L(h_1 h_2 d_1 d_2) = [h_1 h_2 d_1 d_2] \neq [h_2 h_1 d_1 d_2] = O_L(h_2 h_1 d_1 d_2)$. \square

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.

Example 3. Consider a downgrading system with policy $H \rightsquigarrow D \rightsquigarrow L$, actions $A = \{h, d, l\}$ of domains H, D, L , respectively, states A^* with $s_0 = \epsilon$ and transitions $\text{step}(\alpha, a) = \alpha a$. Let the observation functions be given by $\text{obs}_L(\alpha) = \mathbf{ta}_L(\alpha)$ and $\text{obs}_u(\alpha) = \perp$ for all $\alpha \in A^*$ and $u \in \{H, D\}$. Let

$\alpha, \alpha' \in \{h\}^*$. Then we have $\mathbf{ta}_D(\alpha) = \mathbf{ta}_D(\alpha')$ iff $\alpha = \alpha'$. Thus, from $\mathbf{obs}_L(\alpha d) = (\mathbf{ta}_L(\alpha), \mathbf{ta}_D(\alpha), d)$, it is possible for L to deduce the exact sequence of H actions performed in α . On the other hand, since always $\mathbf{obs}_D(\alpha) = \perp$, agent D itself has no information about this sequence. Thus, in this system D can be viewed as causing information to be transmitted to L that it does not itself have. \square

Whether one considers this example to illustrate 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.

Given a policy \succrightarrow , for each domain u , define the function $\mathbf{to}_u : A^* \rightarrow HT(A^*)$ by $\mathbf{to}_u(\epsilon) = \mathbf{obs}_u(s_0)$ and $\mathbf{to}_u(\alpha a) = \mathbf{to}_u(\alpha)$ when $\mathbf{dom}(a) \not\rightarrow u$ and $\mathbf{to}_u(\alpha a) = (\mathbf{to}_u(\alpha), \mathbf{view}_{\mathbf{dom}(a)}(\alpha), a)$ otherwise. We then define M to be TO-secure with respect to \succrightarrow if for all domains u and all $\alpha, \alpha' \in A^*$ with $\mathbf{to}_u(\alpha) = \mathbf{to}_u(\alpha')$, we have $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \alpha')$. 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 $\mathbf{dom}(a)$ *actually* has, as represented by its view $\mathbf{view}_{\mathbf{dom}(a)}(\alpha)$. By contrast, TA-security uses in place of this the maximal information that $\mathbf{dom}(a)$ *may* have.

It is possible to give a flatter representation of the information in $\mathbf{to}_u(\alpha)$ 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 $\mathbf{tview}_u(\alpha)$ of $\mathbf{view}_u(\alpha)$ than ends in an action a with $\mathbf{dom}(a) = u$. Then we have the following:

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

We note the following property of these definitions:

Proposition 2. *Let $\alpha, \alpha' \in A^*$ and $u \in D$.*

1. *If M is TO-secure and $\mathbf{to}_u(\alpha) = \mathbf{to}_u(\alpha')$ then $\mathbf{view}_u(\alpha) = \mathbf{view}_u(\alpha')$.*
2. *If M is TA-secure and $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$ then $\mathbf{view}_u(\alpha) = \mathbf{view}_u(\alpha')$.*

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.

1. *If M is P-secure with respect to \succrightarrow then M is TO-secure with respect to \succrightarrow .*
2. *If M is TO-secure with respect to \succrightarrow then M is TA-secure with respect to \succrightarrow .*
3. *If M is TA-secure with respect to \succrightarrow then M is IP-secure with respect to \succrightarrow .*

4. If \succrightarrow is transitive then M is P -secure with respect to \succrightarrow iff M is TO -secure with respect to \succrightarrow iff M is TA -secure with respect to \succrightarrow iff M is IP -secure with respect to \succrightarrow .

Proof. Part (1) is immediate from Proposition 1. Part (4) follows from parts (1)-(3), using the fact that P -security and IP -security are equivalent with respect to transitive policies (Rushby [Rus92] Theorem 9).

We now prove part (2). We claim that for all $u \in D$ and $\alpha, \alpha' \in A^*$, if $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$ then $\mathbf{to}_u(\alpha) = \mathbf{to}_u(\alpha')$. The result is then immediate from the definition of TO -security. The proof of the claim is by induction on $|\alpha| + |\alpha'|$. The base case of $\alpha = \alpha' = \epsilon$ is trivial. Supposing that the claim holds for strings of shorter combined length, suppose $\mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\alpha')$, where $\alpha, \alpha' \in A^*$ and $a \in A$. We consider two cases:

1. Case 1: $\text{dom}(a) \not\succrightarrow u$. Then $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\alpha')$, so by the inductive hypothesis we have $\mathbf{to}_u(\alpha) = \mathbf{to}_u(\alpha')$. Since $\mathbf{to}_u(\alpha a) = \mathbf{to}_u(\alpha)$ in this case it is immediate that $\mathbf{to}_u(\alpha a) = \mathbf{to}_u(\alpha')$.
2. Case 2: $\text{dom}(a) \succrightarrow u$. We may assume without loss of generality that the last action in α' also interferes with u , else we may apply the previous case. So let $\alpha' = \beta b$. Then $\mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\beta b)$, from which we get that $a = b$ and $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\beta)$ and $\mathbf{ta}_{\text{dom}(a)}(\alpha) = \mathbf{ta}_{\text{dom}(a)}(\beta)$. By induction we have that $\mathbf{to}_u(\alpha) = \mathbf{to}_u(\beta)$ and $\mathbf{to}_{\text{dom}(a)}(\alpha) = \mathbf{to}_{\text{dom}(a)}(\beta)$. From the latter, we obtain by TO -security and Proposition 2 that $\mathbf{view}_{\text{dom}(a)}(\alpha) = \mathbf{view}_{\text{dom}(a)}(\beta)$. It is now immediate from the definition of \mathbf{ta} that $\mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\beta b)$.

To show part (3), we claim that if $u \in X \subseteq D$ then $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\text{ipurge}_X(\alpha))$. In particular, $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\text{ipurge}_u(\alpha))$. Note that this straightforwardly implies that if $\text{ipurge}_u(\alpha) = \text{ipurge}_u(\alpha')$ then $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$, and so if M is TA -secure with respect to \succrightarrow then M is IP -secure with respect to \succrightarrow . To prove the claim we proceed by induction. The case of $\alpha = \epsilon$ is trivial. Assuming the claim holds for α , consider αa where $a \in A$. We consider two cases: depending on whether $\text{dom}(a) \succrightarrow u$.

If $\text{dom}(a) \not\succrightarrow u$, then $\mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\alpha)$. We consider two further subcases. Write $v \succrightarrow X$ if there exists $w \in X$ such that $v \succrightarrow w$.

1. Case 1: $\text{dom}(a) \not\succrightarrow X$. Then $\text{ipurge}_X(\alpha a) = \text{ipurge}_X(\alpha)$. Hence

$$\begin{aligned} \mathbf{ta}_u(\text{ipurge}_X(\alpha a)) &= \mathbf{ta}_u(\text{ipurge}_X(\alpha)) \\ &= \mathbf{ta}_u(\alpha) && \text{(by induction)} \\ &= \mathbf{ta}_u(\alpha a). \end{aligned}$$

2. Case 2: $\text{dom}(a) \succrightarrow X$. Then $\text{ipurge}_X(\alpha a) = \text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha) \cdot a$. Hence, from the induction hypothesis and the definitions,

$$\begin{aligned} \mathbf{ta}_u(\text{ipurge}_X(\alpha a)) &= \mathbf{ta}_u(\text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha) \cdot a) \\ &= \mathbf{ta}_u(\text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha)) \\ &= \mathbf{ta}_u(\alpha) && \text{(by induction)} \\ &= \mathbf{ta}_u(\alpha a). \end{aligned}$$

If $\text{dom}(a) \mapsto u$, then $\text{ipurge}_X(\alpha a) = \text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha) \cdot a$. Thus

$$\begin{aligned} \text{ta}_u(\text{ipurge}_u(\alpha a)) &= \text{ta}_u(\text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha) \cdot a) \\ &= (\text{ta}_u(\text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha)), \text{ta}_{\text{dom}(a)}(\text{ipurge}_{X \cup \{\text{dom}(a)\}}(\alpha)), a) \\ &= (\text{ta}_u(\alpha), \text{ta}_{\text{dom}(a)}(\alpha), a) \quad (\text{by induction}) \\ &= \text{ta}_u(\alpha a). \end{aligned}$$

This completes the proof of the claim. \square

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 surprising fact that Rushby’s unwinding conditions are not preserved under bisimulation.

We begin by recalling Rushby’s results on unwinding for intransitive noninterference. 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.

$$\begin{aligned} \text{OC:} & \text{ If } s \sim_u t \text{ then } \text{obs}_u(s) = \text{obs}_u(t). && \text{(Output Consistency)} \\ \text{SC:} & \text{ If } s \sim_u t \text{ then } s \cdot a \sim_u t \cdot a. && \text{(Step Consistency)} \\ \text{LR:} & \text{ If not } \text{dom}(a) \mapsto u \text{ then } s \sim_u s \cdot a. && \text{(Left Respect)} \end{aligned}$$

If these conditions are satisfied then M is secure with respect to a transitive policy [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 $\text{obs}_u(s \cdot \alpha) = \text{obs}_u(t \cdot \alpha)$. Rushby uses these equivalence relations to show completeness of the unwinding conditions for transitive noninterference:

Proposition 3. ([Rus92] Theorem 6) *Suppose M is P -secure with respect to the transitive policy \mapsto . Then the relations \approx_u satisfy OC, SC and LR.*

For intransitive noninterference he introduces the following condition:

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

Define a *weak unwinding* on a system M with respect to a policy \mapsto 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 \mapsto , let $\{\approx_u^{\text{uw}}\}_{u \in D}$ be the smallest family of equivalence relations (under the pointwise containment order) satisfying WSC and LR.

Proposition 4. *There exists a weak unwinding for M with respect to \mapsto iff the relations \approx_u^{uw} satisfy OC.*

Proof. The implication from right to left is trivial. For the implication from right to left, suppose that $\{\sim_u\}_{u \in D}$ is a weak unwinding for M with respect to \mapsto . It is immediate from the definition of \approx_u^{uw} and fact that the property of being an equivalence relation, WSC and LR are defined by Horn formulas that $\approx_u^{\text{uw}} \subseteq \sim_u$. The fact that \sim_u satisfies OC now implies that \approx_u^{uw} satisfies OC. \square

Rushby shows the following:

Proposition 5. (*[Rus92], Theorem 7*) *Suppose that the relations \sim_u on a system M satisfy OC, WSC and LR. Then M is IP-secure for \mapsto .*

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 \mapsto . Then M is TA-secure with respect to \mapsto .*

Proof. We show that for $u \in D$ and $\alpha, \alpha' \in A^*$, if $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$ then $s_0 \cdot \alpha \sim_u s_0 \cdot \alpha'$. By OC, it also follows that $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \alpha')$, which is what we need for TA-security. We proceed by induction on $|\alpha| + |\alpha'|$. The base case of $\alpha = \alpha' = \epsilon$ is trivial. Supposing that the result holds for sequences of shorter combined length, consider sequences αa and α' , where $a \in A$ and $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$.

If $\mathbf{dom}(a) \not\mapsto u$, then $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha a) = \mathbf{ta}_u(\alpha')$. Hence, by induction, $s_0 \cdot \alpha \sim_u s_0 \cdot \alpha'$. Also, by LR, we have $s_0 \cdot \alpha a \sim_u s_0 \cdot \alpha$. Thus $s_0 \cdot \alpha a \sim_u s_0 \cdot \alpha'$ by transitivity of \sim_u .

If $\mathbf{dom}(a) \mapsto u$, then $\mathbf{ta}_u(\alpha a) = (\mathbf{ta}_u(\alpha), \mathbf{ta}_{\mathbf{dom}(a)}(\alpha), a)$, which implies that the action a also occurs in α' as the last action in a domain interfering with u . If there are any subsequent noninterfering actions, we may switch the role of αa and α' and apply the previous case. Hence, we may assume $\alpha' = \beta a$ for some sequence of actions β , so $\mathbf{ta}_u(\alpha') = (\mathbf{ta}_u(\beta), \mathbf{ta}_{\mathbf{dom}(a)}(\beta), a)$. It follows from the equality $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\alpha')$ that $\mathbf{ta}_u(\alpha) = \mathbf{ta}_u(\beta)$ and $\mathbf{ta}_{\mathbf{dom}(a)}(\alpha) = \mathbf{ta}_{\mathbf{dom}(a)}(\beta)$. By the inductive hypothesis, we obtain $s_0 \cdot \alpha \sim_u s_0 \cdot \beta$ and $s_0 \cdot \alpha \sim_{\mathbf{dom}(a)} s_0 \cdot \beta$. It follows from this by WSC that $s_0 \cdot \alpha a \sim_u s_0 \cdot \beta a$. \square

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 $M = \langle S, s_0, \text{step}, \text{obs}, \text{dom} \rangle$ with actions A , define the “unfolded” system $\text{uf}(M) = \langle S', s'_0, \text{step}', \text{obs}', \text{dom} \rangle$ with actions A having the same domains as in M , by $S' = A^*$, $s'_0 = \epsilon$, $\text{step}'(\alpha, a) = \alpha a$, and $\text{obs}'_u(\alpha) = \text{obs}_u(s_0 \cdot \alpha)$, 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 \rightarrow iff there exists a weak unwinding on $\text{uf}(M)$ with respect to \rightarrow .*

Proof. Define the equivalence relations \sim_u^{ta} on $\text{uf}(M)$ by $\alpha \sim_u^{\text{ta}} \alpha'$ iff $\text{ta}_u(\alpha) = \text{ta}_u(\alpha')$. We first note that for the relations \approx_u^{uw} on $\text{uf}(M)$, we have $\alpha \approx_u^{\text{uw}} \alpha'$ iff $\alpha \sim_u^{\text{ta}} \alpha'$. For, the relations \sim_u^{ta} are equivalence relations and satisfy WSC and LR by definition of ta_u . Thus, we have $\approx_u^{\text{uw}} \subseteq \sim_u^{\text{ta}}$ by definition of \approx_u^{uw} . Conversely, we show that $\text{ta}_u(\alpha) = \text{ta}_u(\alpha')$ implies $\alpha \approx_u^{\text{uw}} \alpha'$, by induction on $|\alpha| + |\alpha'|$. The case of $\alpha = \alpha' = \epsilon$ is clear. Suppose $\text{ta}_u(\alpha a) = \text{ta}_u(\alpha')$. If $\text{dom}(a) \not\rightarrow u$, we have $\text{ta}_u(\alpha) = \text{ta}_u(\alpha a) = \text{ta}_u(\alpha')$, so by the inductive hypothesis, we have $\alpha \approx_u^{\text{uw}} \alpha'$. Since \approx_u^{uw} satisfies LR, we obtain $\alpha a \approx_u^{\text{uw}} \alpha'$ and it follows that $\alpha a \approx_u^{\text{uw}} \alpha'$. For the case where $\text{ta}_u(\alpha a) = (\text{ta}_u(\alpha), \text{ta}_{\text{dom}(a)}(\alpha), a)$, we may assume without loss of generality that the final action in α' may interfere with u , and derive that $\alpha' = \beta a$ where $\text{ta}_u(\alpha) = \text{ta}_u(\beta)$ and $\text{ta}_{\text{dom}(a)}(\alpha) = \text{ta}_{\text{dom}(a)}(\beta)$. It follows by the inductive hypothesis that $\alpha \approx_u^{\text{uw}} \beta$ and $\alpha \approx_{\text{dom}(a)}^{\text{uw}} \beta$, hence by the fact that \approx_u^{uw} satisfies WSC that $\alpha a \approx_u^{\text{uw}} \beta a$, as required.

Suppose M is TA-secure with respect to \rightarrow . Then WSC and LR for the relations \sim_u^{ta} are immediate from the definition of the functions ta_u , and OC follows directly from TA-security. Hence the family \sim_u^{ta} is a weak unwinding on $\text{uf}(M)$. Conversely, suppose that there exists a weak unwinding \sim_u on $\text{uf}(M)$ with respect to \rightarrow . Then $\sim_u^{\text{ta}} = \approx_u^{\text{uw}} \subseteq \sim_u$ satisfies OC, hence M is TA-secure with respect to \rightarrow . \square

It is reasonable to give a definition of security on M by reference to $\text{uf}(M)$ 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 $\text{uf}(M)$. It is the case that unwindings on M can be lifted to unwindings on $\text{uf}(M)$.

Proposition 6. *If there exists a weak unwinding for \rightarrow on M then there exists a weak unwinding for \rightarrow on $\text{uf}(M)$*

Proof. Suppose the relations \sim_u are a weak unwinding for \rightarrow on M . Define \sim'_u on $\text{uf}(M)$ by $\alpha \sim'_u \alpha'$ if $s_0 \cdot \alpha \sim_u s_0 \cdot \alpha'$ in M . It is easily checked that the relations \sim'_u are a weak unwinding on $\text{uf}(M)$. \square

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.




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

Example 4. Consider the system and policy depicted in Figure 1. There are actions a, b, c of domains A, B, C respectively, and s_0 is the initial state. For all domains u other than D , we assume that the observation obs_u is the same on all states. TA-security therefore depends only on the behaviour of the system with respect to domain D , where there are two possible observations o, o' as indicated. We show that there does not exist a weak unwinding for \succrightarrow on M , but there does exist one on $\text{uf}(M)$.

For the former, we consider the relation family \approx_u^{uw} on M . Note that since $B \not\rightarrow D$ and $s_0 \cdot b = s_1$ we have $s_0 \approx_D^{\text{uw}} s_1$. Similarly, since $C \not\rightarrow A$ we have $s_0 \approx_A^{\text{uw}} s_1$. Hence, by WSC, for the action a , we get $s_0 \approx_D^{\text{uw}} s_2$. Since $\text{obs}_D(s_0) = o$ and $\text{obs}_D(s_2) = o'$, we have that \approx_D^{uw} does not satisfy OC. Since \approx_u^{uw} is the smallest family satisfying WSC and LR, there can exist no weak unwinding for \succrightarrow on M .

For the unwinding on $\text{uf}(M)$, consider $\approx_u^{\text{uw}} = \sim_u^{\text{ta}}$. Since this family of equivalence relations satisfies WSC and LR, it suffices to consider the property OC, where we need consider only the domain D , as already noted. Here, the only possible failure of OC is for states α, α' where $\text{ta}_D(\alpha) = \text{ta}_D(\alpha')$, $s_0 \cdot \alpha \in \{s_0, s_1\}$ 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 $\text{ta}_D(\alpha')$ 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 $\text{ta}_D(\alpha)$, 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^{uw} satisfies OC. \square

Since $\text{uf}(M)$ and M are bisimilar, this example shows that bisimulation does not preserve 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^{obs} on states of a system M by $s \approx_u^{\text{obs}} t$ if $\text{obs}_u(s) = \text{obs}_u(t)$. Then we have the following sufficient condition for TO-security:

Proposition 7. *Suppose the relation family \approx_u^{obs} is a weak unwinding on M with respect to \rightarrow . Then M is TO-secure with respect to \rightarrow .*

Proof. We claim that $\text{to}_u(\alpha) = \text{to}_u(\alpha')$ implies $s_0 \cdot \alpha \approx_u^{\text{obs}} s_0 \cdot \alpha'$. The result then follows immediately using OC. The proof of the claim is by induction on the combined length of the sequences α, α' . In case both are ϵ , the claim plainly holds. Suppose that it holds for sequences of shorter combined length, and consider the sequences αa and α' , where $a \in A$. We consider two cases, depending on whether $\text{dom}(a) \rightarrow u$.

1. If $\text{dom}(a) \not\rightarrow u$, then it follows from the definitions that $\text{to}_u(\alpha a) = \text{to}_u(\alpha a) = \text{to}_u(\alpha')$, so by the inductive hypothesis, we have $s_0 \cdot \alpha \approx_u^{\text{obs}} s_0 \cdot \alpha'$. By LR we moreover have $s_0 \cdot \alpha a \approx_u^{\text{obs}} s_0 \cdot \alpha$, and we conclude $s_0 \cdot \alpha a \approx_u^{\text{obs}} s_0 \cdot \alpha'$.
2. If $\text{dom}(a) \rightarrow u$, then we may assume without loss of generality that $\alpha' = \beta b$ where $\text{dom}(b) \rightarrow u$, else we may apply the previous case. It then follows from $\text{to}_u(\alpha a) = \text{to}_u(\beta b)$ that $a = b$ and $\text{to}_u(\alpha) = \text{to}_u(\beta)$ and $\text{to}_{\text{dom}(a)}(\alpha) = \text{to}_{\text{dom}(a)}(\beta)$. By the inductive hypothesis, we obtain $s_0 \cdot \alpha \approx_u^{\text{obs}} s_0 \cdot \beta$ and $s_0 \cdot \alpha \approx_{\text{dom}(a)}^{\text{obs}} s_0 \cdot \beta$. By WSC and $a = b$ we obtain $s_0 \cdot \alpha a \approx_{\text{dom}(a)}^{\text{obs}} s_0 \cdot \beta b$, as required.

□

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 intentions of the notion of intransitive noninterference.

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

1. a set N of *names*,

2. a set V of *values*, and functions
3. $\mathbf{contents} : S \times N \rightarrow V$, with $\mathbf{contents}(s, n)$ interpreted as the value of object n in state s ,
4. $\mathbf{observe} : D \rightarrow \mathcal{P}(N)$, with $\mathbf{observe}(u)$ interpreted as the set of objects that domain u can observe, and
5. $\mathbf{alter} : D \rightarrow \mathcal{P}(N)$, with $\mathbf{alter}(u)$ 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 $\mathbf{oc}_u(s)$ for the function mapping $\mathbf{observe}(u)$ to values, defined by $\mathbf{oc}_u(s)(n) = \mathbf{contents}(s, n)$ for $n \in \mathbf{observe}(u)$. Intuitively, $\mathbf{oc}_u(s)$ captures all the content of the state s that is observable to u . Using this, we may define a binary relation $\sim_u^{\mathbf{oc}}$ of *observable content equivalence* on S for each domain $u \in D$, by $s \sim_u^{\mathbf{oc}} t$ if $\mathbf{oc}_u(s) = \mathbf{oc}_u(t)$.

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*.⁵

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

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

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

That is, with the choice depending only on information observable to $\mathbf{dom}(a)$, 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 $\mathbf{dom}(a)$.

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

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

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

⁵ Strictly, Rushby works with action-observed systems. We express here a version for state-observed systems.

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 $\text{dom}(a)$. If n is not observable to $\text{dom}(a)$, 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 $\text{dom}(a)$), 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 $\text{dom}(a)$), but also on the value of n . Thus, the condition RM2 can plausibly be weakened to the following.⁶

[RM2'] For all actions a , states s, t and names $n \in \text{alter}(\text{dom}(a))$, if $s \sim_{\text{dom}(a)}^{\text{oc}} t$ and $\text{contents}(s, n) = \text{contents}(t, n)$ we have $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$.

That is, for $n \in \text{alter}(\text{dom}(a))$, the value $\text{contents}(s \cdot a, n)$ is a function of both $\text{contents}(s, n)$ and $\text{oc}_{\text{dom}(a)}(s)$. Using Proposition 8 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* \rightarrow if there exists a set of names N , a set of values V and functions $\text{observe} : D \times S \rightarrow \mathcal{P}(N)$, $\text{alter} : D \times S \rightarrow \mathcal{P}(N)$ and $\text{contents} : N \times S \rightarrow V$, 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:

⁶ A weakened condition resembling RM2' has also been used in a slightly different context by Greve, Wilding and vanFleet [GWV03]. A similar weakening has also been proposed by von Oheimb [Ohe04], who also shows that the definition of access control system consistent with a policy can be weakened while still implying the existence of unwinding conditions implying IP-security. He adds the conditions $\text{dom}(a) \rightarrow u$, $n \in \text{observe}(u)$ and $s \sim_u^{\text{oc}} t$ to the preconditions of RM2. We would argue that there are some difficulties giving intuitive content to von Oheimb’s proposal. First, whether or not a system is an access control system should be independent of the policy that may be applied to it. The fact $\text{dom}(a) \rightarrow u$ in the precondition is in any case derivable from $n \in \text{observe}(u)$ in the context of AOI and RM3, so can be omitted without loss of generality. Second, since u does not occur in the consequent, the quantification over u has a rather obscure meaning.

Proposition 10. *Suppose M is a weak access control system consistent with \mapsto . Then the family of relations \sim_u^{oc} is a weak unwinding on M with respect to \mapsto .*

Proof. OC is direct from RM1 and LR follows easily from RM3 and AOI. For WSC, suppose that $s \sim_{\text{dom}(a)}^{\text{oc}} t$ and $s \sim_u^{\text{oc}} t$, i.e., $\text{oc}_u(s) = \text{oc}_u(t)$. We need to show $s \cdot a \sim_u^{\text{oc}} t \cdot a$, which amounts to showing $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$ for all $n \in \text{observe}(u)$. We consider the two possibilities $n \in \text{alter}(\text{dom}(a))$ and $n \notin \text{alter}(\text{dom}(a))$:

1. If $n \in \text{alter}(\text{dom}(a))$ then we have $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$ by RM2'.
2. If $n \notin \text{alter}(\text{dom}(a))$, then using RM3 we have $\text{contents}(s \cdot a, n) = \text{contents}(s, n) = \text{contents}(t, n) = \text{contents}(t \cdot a, n)$.

Thus, we have $s \cdot a \sim_u^{\text{oc}} t \cdot a$. □

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 11. *Suppose that there exists a weak unwinding on M with respect to \mapsto . Then M admits a weak access control interpretation consistent with \mapsto .*

Proof. Write $[s]_u$ for the equivalence class of s under \sim_u . We define the access control interpretation on M as follows:

1. $N = D$
2. $\text{observe}(u) = \{u\}$
3. $\text{alter}(u) = \{v \in D \mid u \mapsto v\}$
4. $\text{contents}(s, u) = [s]_u$.

RM1 is immediate from the fact that \sim satisfies output consistency.

For RM2', we proceed by contradiction. Note that $s \sim_u t$ iff $[s]_u = [t]_u$ iff $s \sim_u^{\text{oc}} t$. If RM2' fails then there exists an action a , a name $n \in \text{alter}(\text{dom}(a))$ and states s, t such that $s \sim_{\text{dom}(a)}^{\text{oc}} t$ and $\text{contents}(s, n) = \text{contents}(t, n)$ and $\text{contents}(s \cdot a, n) \neq \text{contents}(t \cdot a, n)$. Thus, $n = u$ for some $u \in D$ such that $\text{dom}(a) \mapsto u$, and we have $s \sim_{\text{dom}(a)}^{\text{oc}} t$ and $s \sim_u t$ but not $s \cdot a \sim_u t \cdot a$. This contradicts WSC.

For RM3, the only value of n to be checked is $n = u$, and $u \in \text{alter}(\text{dom}(a))$ iff $\text{dom}(a) \mapsto u$. Also $\text{contents}(\text{step}(s, a), u) = \text{contents}(s, u)$ iff $\text{step}(s, a) \approx_u s$. Hence, the contrapositive of RM3 states that if not $\text{dom}(a) \mapsto u$ then $\text{step}(s, a) \approx_u s$, which is just the condition LR.

Plainly, if $n \in \text{alter}(u)$ and $n \in \text{observe}(v)$ then $n = v$ and $u \mapsto v$, so AOI also holds. □

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 \rightsquigarrow ,
2. $\mathbf{uf}(M)$ admits a weak access control interpretation consistent with \rightsquigarrow ,
3. there exists a weak unwinding on $\mathbf{uf}(M)$ with respect to \rightsquigarrow .

Proof. The equivalence of (1) and (3) is from Theorem 3, and that between (2) and (3) follows from Propositions 10 and 11. \square

From Theorem 2 and Proposition 10, we also obtain the following.

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

This conclusion is a more general result than Proposition 9, 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 5. Consider the system for the policy $A \rightsquigarrow B \rightsquigarrow C$ 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 $n_{AB} \in \mathbf{alter}(A) \cap \mathbf{observe}(B)$ and $n_{BC} \in \mathbf{alter}(B) \cap \mathbf{observe}(C)$. 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 $\mathbf{obs}_A(s) = \mathbf{obs}_B(s) = \perp$ and $\mathbf{obs}_C(s) = n_{BC}$. 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 $\mathbf{purge}_C(\alpha) = b = \mathbf{purge}_C(\alpha')$, and $\mathbf{tview}_B(\alpha) = \perp b = \mathbf{tview}_B(\alpha')$ but $\mathbf{obs}_C(s_0 \cdot \alpha) = 0 \neq 1 = \mathbf{obs}_C(s_0 \cdot \alpha')$. \square

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 M is a system with structured states is *fully observable* if in all states s we have $\mathbf{obs}_u(s) = \mathbf{oc}_u(s)$. Note that this means that the relations $\sim_u^{\mathbf{oc}}$ and $\approx_u^{\mathbf{obs}}$ coincide. We now obtain the following from Propositions 7 and 10. 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 \rightsquigarrow then M is TO-secure with respect to \rightsquigarrow .*

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

Example 6. Note that if in Example 5 we change the definition of $\mathbf{obs}_B(s)$ to n_{AB} , then the system continues to satisfy RM1, RM2', RM3 and AOI, and not P-security. The modified system has $\mathbf{obs}_u(s) = \mathbf{oc}_u(s)$ for all states s . (So it also satisfies TO-security.) \square

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 does not facilitate the expression of correctness of downgraders, 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 \rightarrow 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 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. Sabelfeld and Sands [SS05] lay out some general principles and direction for research in this area. It would be of interest to reconsider some of these recent contributions in the light of our results in this paper.

References

- [BP76] 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, M.A., March 1976.
- [CM04] S. Chong and A. C. Myers. Security policies for downgrading. In *11th ACM Conf. on Computer and Communications Security (CCS)*, Oct 2004.
- [FG01] R. Focardi and R. Gorrieri. Classification of security properties (Part I: information flow). In R. Focardi and R. Gorrieri, editors, *Foundations of Security Analysis and Design, FOSAD 2000, Bertinoro, Italy, September 2000*, volume 2171 of *LNCS*, pages 331–396. Springer, 2001.
- [FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning About Knowledge*. MIT-Press, 1995.
- [GM82] J.A. Goguen and J. Meseguer. Security policies and security models. In *Proc. IEEE Symp. on Security and Privacy*, pages 11–20, Oakland, 1982.
- [GM84] J.A. Goguen and J. Meseguer. Unwinding and inference control. In *IEEE Symp. on Security and Privacy*, 1984.

- [GWW03] D. Greve, M. Wilding, and M. Vanfleet. A separation kernel formal security policy. In *Proc. Fourth International Workshop on the ACL2 Theorem Prover and Its Applications*, 2003. Paper and associated proof scripts at <http://www.cs.utexas.edu/users/moore/ac12/workshop-2003/>.
- [HALL⁺05] N.B. Hadj-Alouane, S. Lafrance, Feng Lin, J. Mullins, and M.M. Yeddes. On the verification of intransitive noninterference in multilevel security. *IEEE Trans. on Systems, Man and Cybernetics, Part B*, 35(5):948–958, Oct. 2005.
- [HY87] J.T. Haigh and W.D. Young. Extending the noninterference version of MLS for SAT. *IEEE Trans. on Software Engineering*, SE-13(2):141–150, Feb 1987.
- [MB94] Y. Moses and B. Bloom. Knowledge, timed precedence and clocks (preliminary report). In *Proc. ACM Symp. on Principles of Distributed Computing*, pages 294–303, 1994.
- [McC88] D. McCullough. Noninterference and the composability of security properties. In *Proc. IEEE Symp. on Security and Privacy*, pages 177–186, 1988.
- [Mey07] R. van der Meyden. The complexity of notions of intransitive noninterference. unpublished manuscript, 2007.
- [MS04] H. Mantel and D. Sands. Controlled declassification based on intransitive noninterference. In *Proc. Asian Symp. on Programming Languages and Systems*, volume 3302 of *LNCS*, pages 129–145. Springer-Verlag, November 2004.
- [MZ06] R. van der Meyden and C. Zhang. A comparison of semantic models for noninterference. In *Proc. Workshop on Formal Aspects of Security and Trust, Hamilton, Ontario, Canada*, LNCS. Springer, August 2006. to appear, extended version at <http://www.cse.unsw.edu.au/~meyden/research/publications.html>.
- [Ohe04] D. von Oheimb. Information flow control revisited: Noninfluence = Noninterference + Nonleakage. In *Computer Security – ESORICS 2004*, volume 3193 of *LNCS*, pages 225–243, 2004.
- [Pin95] S. Pinsky. Absorbing covers and intransitive non-interference. In *Proc. IEEE Symp. on Security and Privacy*, pages 102–113, 1995.
- [Pos46] E.L. Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52, 1946.
- [RG99] A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *IEEE Computer Security Foundations Workshop*, pages 228–238, 1999.
- [Rus81] J. Rushby. Design and verification of secure systems. In *Proc. 8th Symposium on Operating Systems Principles*, pages 12–21, Asilomar CA, Dec 1981. (ACM Operating Systems Review, Vol 15, No. 1).
- [Rus92] J. Rushby. Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International, Dec 1992.
- [Rya01] P.Y. Ryan. Mathematical models of computer security. In R. Focardi and R. Gorrieri, editors, *Foundations of Security Analysis and Design, FOSAD 2000, Bertinoro, Italy, September 2000*, volume 2171 of *LNCS*, pages 1–62. Springer, 2001.
- [SS05] Andrei Sabelfeld and David Sands. Dimensions and principles of declassification. In *Proceedings of the 18th IEEE Computer Security Foundations Workshop*, pages 255–269. IEEE Computer Society Press, 2005.
- [Sut86] D. Sutherland. A model of information. In *Proc. 9th National Computer Security Conf.*, pages 175–183, 1986.

- [WJ90] J. T. Wittbold and D. M. Johnson. Information flow in nondeterministic systems. In *IEEE Symposium on Security and Privacy*, pages 144–161, 1990.