

Architectural Refinement and Notions of Intransitive Noninterference*

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

May 23, 2012

Abstract

This paper deals with architectural designs that specify components of a system and the permitted flows of information between them. In the process of systems development, one might refine such a design by viewing a component as being composed of subcomponents, and specifying permitted flows of information between these subcomponents and others in the design. The paper studies the soundness of such refinements with respect to a spectrum of different semantics for information flow policies, including Goguen and Meseguer's purge-based definition, Haigh and Young's intransitive purge-based definition, and some more recent notions TA-security, TO-security and ITO-security defined by van der Meyden. It is shown that all these definitions support the soundness of architectural refinement, for both a state- and an action-observed model of systems. A notion of systems refinement in which the information content of observations is reduced is also studied. It is also shown that refinement preserves weak access control structure, an implementation mechanism that ensures TA-security.

1 Introduction

Architectural design is a high level of systems specification, concerned with identifying the components of a system and the patterns of their interaction. In this paper, we consider the relationship between information flow security policies and the architecture development process. We use ideas from the literature

*Author version of a paper to appear in Formal Aspects of Computing. Copyright BCS. The original publication is available at springerlink.com. Version of May 23, 2012. Work supported by Australian Research Council Discovery grants DP0451529 and DP0987769. An extended abstract of this work appeared in International Symposium on Engineering Secure Software and Systems, February 04-06, 2009 Leuven, Belgium, Springer LNCS No 5429, pp. 60-74. The present version adds proofs and the content of Sections 6 and 7.

on information flow security to give semantics to architectures, and study how such semantics support a systems development process that refines high level architectural designs to more detailed architectures.

The type of information flow security policies that form the basis of this work place constraints on the permitted flows of information, or causal effects, between system components, and are referred to in the literature as *non-interference* policies. These policies can be represented as a binary relation on the set of components. For classical multi-level security policies, this relation is transitive. It has been proposed that extensions to multi-level security, such as downgraders, require that the policy be intransitive [HY87, Rus92]. An architectural interpretation of intransitive noninterference policies is gaining increased prominence through such efforts as the MILS (Multiple Independent Levels of Security and Safety) approach to high-assurance systems design [AFHOT06, VBC⁺05], which envisages the utilization of recent advances in the efficiency of separation kernels to increase the degree of componentization of systems, enabling secure systems to be built from a mix of small, trusted and more complex, untrusted components [RR83], with global security properties assured from the separation property and a verification effort focussed on the trusted components.

During the process of system design, one may *refine* an architectural diagram by specifying internal structure for some of the systems components, breaking them down into sub-components, and specifying the permitted interferences of these sub-components with each other and with other components in the design. This leads to the following question: if one now builds a system according to the refined architecture, is it guaranteed to be compliant to the original architectural diagram? This property needs to hold in order for the process of architectural refinement to be sound for designs of information flow secure systems.

Our contribution in this paper is to answer this question for a range of meanings of the notion of a system being compliant with an architectural design. For the meaning of architectures, we consider several different semantics for security with respect to intransitive noninterference policies that have previously been proposed in the literature. The first of these (which we call here P-security) is Goguen and Meseguer’s original semantics [GM82] for transitive policies, which has generally been felt to be inappropriate for intransitive policies, although such application is not without its adherents [RG99]. An alternate semantics for intransitive noninterference policies (called here IP-security) was proposed by Haigh and Young [HY87], and propounded by Rushby [Rus92]. We have argued recently that this definition is flawed: it considers some systems to be secure that have flows of information that are contrary to the intuitions for intransitive policies [Mey08]. In response, we have defined [Mey08] several alternate semantics for noninterference policies — TO-security, ITO-security and TA-security — that are better behaved, and all equivalent to the classical definition in case the policy is transitive. Moreover, TA-security can be shown [Mey08] to correspond in a precise sense to Rushby’s “unwinding” proof technique for intransitive non-interference. We remark that under certain circumstances, the definitions given by Roscoe and Goldsmith [RG99] correspond either to P-security or to ITO-

security— for details, we refer the reader to [Mey07], which is concerned with a detailed comparison of the above definitions of security in a number of different semantic frameworks.

The answer to our question turns out to be that a common notion of architectural refinement is sound with respect to *all* these semantics for intransitive noninterference policies. Indeed, we show this for two types of systems models, one in which observations are made at a state, and the other in which observations take the form of outputs returned on the invocation of an action.

In order to deal with the action-observed model, it proves to be convenient to first consider also a notion of refinement, on systems rather than on architectural designs, in which the amount of information in observations is reduced. This notion is of some independent interest since it is intuitive that security of a system should be preserved when we reduce the amount of information in the observations that agents make in the system. In fact, we show that while the intuition holds for P-security, IP-security and TA-security, not all of the semantics of intransitive policies are preserved by this notion of systems refinement. The ones that do not, TO-security and ITO-security, state both upper and lower bounds on permitted information. Our results in this regard help to clarify the content of these semantics. Nevertheless, we do identify a useful sufficient condition under which all semantics are preserved under reduction of observation content, and this condition aids in the derivation of the results on architectural refinement with respect to the action-observed model.

In the next step of systems development beyond architectural design, one might apply specific mechanisms to enforce the constraints imposed by the architecture. One implementation mechanism that is known to support the satisfaction of information flow policies is the use of a type of access control structure, in which domains are restricted with respect to the objects that they may read and write. A simple statically checkable condition on these restrictions entails IP-security [Rus92] as well as the stronger notion of TA-security [Mey08]. We show that a notion of refinement, closely related to architectural refinement, can also be defined for access control structure, and that refinements at the architectural level can be satisfied by refinements at the level of the access control implementation. Moreover, we define a more expressive notion of access control structure than has previously been considered in the context of intransitive policies, in which read and write constraints may apply differentially to individual actions performed within a domain, rather than uniformly across the domain. We show how such a policy at the action level can implement a policy at the domain level.

Taken together, these results provide the justification for a systems development process that proceeds from high level architectural design, via refinements, though to implementations using access control mechanisms. We give an example that illustrates the process.

The structure of the paper is as follows. In Section 2 we give a formal model for architectures and define architectural refinement. Section 3 defines our state-observed system model, noninterference policies, and the spectrum of different semantics for these policies. In Section 4 we show that architectural re-

refinement preserves all the policy semantics for the state-observed model. Next, in Section 5, we consider access control as a mechanism for the implementation of architectures, and show that access control structure is also preserved under architectural refinement. The example of a refinement from architecture through to access control implementation is presented in this section. Section 6 defines a notion of systems refinement on state-observed systems, in which the amount of information in observations is reduced, and identifies conditions under which this preserves the architectural semantics we consider. Section 7 deals with architectural refinement in action-observed systems. Related literature is discussed in section 8, and we make some concluding remarks in Section 9.

2 Architectural Refinement

We begin by formalising the notion of architecture, and defining a relation of refinement between architectures. In this section, we leave the question of the formal semantics of architectures open. In the following sections, we will show that architectural refinement is sound with respect to several different semantic interpretations of architectures.

The notion of architecture that we consider expresses only the highest levels of a system design. Intuitively, an architecture specifies that the system is comprised of a set of components that generate and hold information, and constrains the permitted flows of information between these components. Using terminology from the security literature, we refer to these components as domains.

Define an *architecture* to be a pair $\mathcal{A} = (D, \succrightarrow)$, where D is a set of domains and \succrightarrow is a reflexive binary relation over D . We call the relation \succrightarrow the *information flow policy* of the architecture. Several related intuitions may be associated with this relation. One is that information is permitted to flow from a domain u to a domain v only if $u \succrightarrow v$. Another is that actions in domain u may have directly observable effects in domain v only if $u \succrightarrow v$. The relation is assumed to be reflexive because information flow from a domain to itself can never be prevented. The security literature has frequently assumed information flow policies to be transitive. The following example illustrates a case where this assumption is not desirable.

Example 1: Consider the architecture $\mathcal{A} = (\{H, D, L\}, \succrightarrow)$ where \succrightarrow is the downgrader policy depicted in the upper part of Figure 1 (we omit reflexive edges). Here H represents the high security domain, L the low security domain and D the downgrader. Information may flow from L to H , but any flow in the other direction needs to be mediated by the downgrader (so we would not want transitivity of \succrightarrow). Information flow from L to D is permitted, to allow for requests by L to D for H information, e.g., freedom-of-information requests. \square

The process of systems design may take a component in a high level design, and specify that it is to be implemented as the composition of a set of lower level

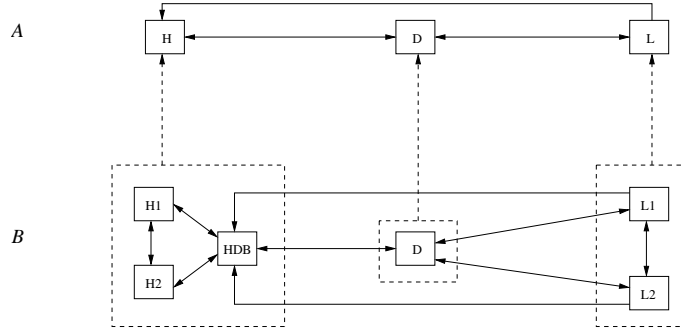


Figure 1: Refinement of a downgrader architecture

components. In this design step, the permitted flows of information between the lower level components, and others in the design, should also be specified. One way to understand this process is to view the design step as establishing a relationship of *refinement* between low level and high level architectures.

Example 2: The architecture $\mathcal{B} = (\{H_1, H_2, HDB, D, L_1, L_2\}, \succrightarrow)$ (where the policy \succrightarrow is depicted in the lower part of Figure 1) represents a refinement of \mathcal{A} . We refine H into three components, two high level users H_1 and H_2 and a database HDB . Similarly, we refine L into two low level users L_1 and L_2 . We assume that these may transmit information to each other, the high database, and D , but that all information flow from H_1 and H_2 to L_1 and L_2 is mediated by HDB and D . \square

To formalise the notion of architectural refinement, let $\mathcal{A}_1 = (D_1, \succrightarrow_1)$ and $\mathcal{A}_2 = (D_2, \succrightarrow_2)$ be architectures. A *refinement mapping* from \mathcal{A}_1 to \mathcal{A}_2 is a function $r : D_1 \rightarrow D_2$ such that

1. r is *onto* D_2 , and
2. for all $u, v \in D_1$, if $u \succrightarrow_1 v$ then $r(u) \succrightarrow_2 r(v)$.

In this case, we write $\mathcal{A}_1 \leq_r \mathcal{A}_2$, and say that \mathcal{A}_1 is a *refinement* of \mathcal{A}_2 .

The requirement that r be surjective captures that intuition that if a design calls for the existence of a component, then a more detailed design should include an implementation of that component. Intuitively, each domain u in the high level design \mathcal{A}_2 is implemented by the set of domains in the preimage $r^{-1}(\{u\})$ in the detailed design. (For brevity, we henceforth write $r^{-1}(u)$ for $r^{-1}(\{u\})$.) The second condition expresses that the lower level policy should not permit information flow between two subdomains that was prohibited between their superdomains. That is, if in a high level design, information is not permitted to flow directly from component U to a component V , it should be incorrect to implement U and V using lower level components u and v , respectively, with information permitted to flow from u directly to v . (Note that the refinement mapping in Figure 1 satisfies this condition.)

We note that the definition of refinement is transitive in the following sense: if $\mathcal{A}_1 \leq_r \mathcal{A}_2$ and $\mathcal{A}_2 \leq_s \mathcal{A}_3$, then $\mathcal{A}_1 \leq_{s \circ r} \mathcal{A}_3$. This is a key property usually considered to be desirable in theories of refinement, since it permits the development of an architectural design to proceed in a number of stages, each of which refines the previous, by guaranteeing that the final design is a refinement of the original design.

Our main result in this paper is that the process of replacing a high level architecture by a lower level refinement is a sound design step, in the sense that any concrete system that implements the refined architecture also implements the higher level architecture. In order to make this claim formally precise, we first need to define what counts as a concrete implementation of an architecture. As a step towards this, we first consider a range of possible semantic interpretations of information flow policies. We use these in the semantics of architectures in Section 4.

3 Semantics of Information Flow Policies

To give semantics to architectures, we recall in this section several classical semantics for information flow policies [GM82, HY87, Rus92], and several new definitions proposed by van der Meyden [Mey08]. These definitions can be given for both state- and action-observed machines. We consider here the state-observed versions. The action-observed variants are discussed in a later section. The content of this section is largely definitional, and drawn from [Mey08].

The *state-observed* machine model [Rus92] for these definitions 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 write $s \cdot \alpha$ for the state reached by performing the sequence of actions $\alpha \in A^*$ 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.

Transitive information flow policies have been given a formal semantics using a definition based on a “purge” function [GM82]. Given a set $E \subseteq D$ of domains and a sequence $\alpha \in A^*$, we write $\alpha \upharpoonright E$ for the subsequence of all actions a in α with $\mathbf{dom}(a) \in E$. Given a policy \rightsquigarrow , define the function $\mathbf{purge} : A^* \times D \rightarrow A^*$ by

$$\mathbf{purge}(\alpha, u) = \alpha \upharpoonright \{v \in D \mid v \rightsquigarrow u\}.$$

For clarity, we may use subscripting of domain arguments of functions, writing e.g., $\mathbf{purge}(\alpha, u)$ as $\mathbf{purge}_u(\alpha)$.

Definition 1 *A system M is P-secure with respect to a policy \rightsquigarrow if 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 domain u 's observation depends only on the sequence of interfering actions that have been performed. By idempotence of \mathbf{purge}_u , this definition is equivalent to the classical formulation, according to which the system M is secure when for all $\alpha \in A^*$ and domains $u \in D$, we have $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \mathbf{purge}_u(\alpha))$. This formulation can be understood as saying that each domain's observations are as if only interfering actions had been performed. We note that we will apply P-security to intransitive policies as well as the transitive policies for which it was originally intended.

While P-security is well accepted for transitive policies, it has been felt to be inappropriate for intransitive policies, since it does not permit L to *ever* learn about H actions if the policy is $H \rightsquigarrow D \rightsquigarrow L$, even if D is intended to be a trusted downgrader of H information. To address this deficiency, Haigh and Young [HY87] proposed to generalise the definition of the purge function to intransitive policies. (We follow the presentation of [Rus92].) 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 $\mathbf{sources} : A^* \times D \rightarrow \mathcal{P}(D)$ defined inductively by $\mathbf{sources}(\epsilon, u) = \{u\}$ and

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

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

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

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

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

Then $\mathbf{ipurge}_u(\alpha)$ is identical to $\mathbf{ipurge}_{\{u\}}(\alpha)$. Haigh and Young's definition can then be formulated as the following variant of P-security in which we use the \mathbf{ipurge} function in place of the \mathbf{purge} function.

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

It can be seen that $\mathbf{ipurge}_u(\alpha) = \mathbf{purge}_u(\alpha)$ when \rightsquigarrow is transitive, so IP-security is in fact a generalisation of the definition of security for transitive policies.

These definitions are critiqued in [Mey08], where it is shown that IP-security sometimes allows quite unintuitive flows of information. In response, several alternative definitions are proposed. Each is based on a concrete model of the maximal amount of information that a domain may have after some sequence of actions has been performed, and states that a domain'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 a domain increases its information either by performing an action or by receiving information transmitted by another domain.

In the first model of the maximal information, what is transmitted when an domain performs an action is information about the actions performed by other domains. 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 \mathcal{T} containing X and such that if $x, y \in \mathcal{T}$ and $z \in A$ then $(x, y, z) \in \mathcal{T}$. Intuitively, the elements of $\mathcal{T}(X, A)$ are binary trees with leaves labelled from X and interior nodes labelled from A .

Given a policy \rightsquigarrow , define, for each domain $u \in D$, the function $\mathbf{ta}_u : A^* \rightarrow \mathcal{T}(\{\epsilon\}, A)$ 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 domain u may, consistently with the policy \rightsquigarrow , have about the past actions of other domains. (The nomenclature is intended to be suggestive of *transmission* of information about *actions*.) Initially, a domain has no information about what actions have been performed. The recursive clause describes how the maximal information $\mathbf{ta}_u(\alpha)$ 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 $\text{dom}(a)$ at the time a is performed (represented by $\mathbf{ta}_{\text{dom}(a)}(\alpha)$), as well the fact that a has been performed. Thus, this definition captures the intuition that a domain may only transmit information that it is permitted to have, and then only to domains with which it is permitted to interfere.

Definition 3 *A system M is TA-secure with respect to a policy \rightsquigarrow if for all domains 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 domain's observations provide the domain with no more than the maximal amount of information that may have been transmitted to it, as expressed by the functions \mathbf{ta} .

The notion of TA-security can be shown to be a better fit to the intended applications and theory of IP-security. On the other hand, it may still be too weak for some applications. For example, it considers to be secure a system where a downgrader transmits to L an email attachment that it received from H ,

without opening the attachment first (so that it does not know what information it is transmitting!) The second of van der Meyden’s definitions is intended to address this potential deficiency.

The definition uses the following notion of *view*. The definition uses an absorptive concatenation function \circ , defined over a set X by, for $s \in X^*$ and $x \in X$, by $s \circ x = s$ if $s \neq \epsilon$ and x is equal to the final element of s , and $s \circ x = sx$ (ordinary concatenation) otherwise. Represent the view of domain u with respect to a sequence $\alpha \in A^*$ using the function $\text{view}_u : A^* \rightarrow O(A \cup O)^*$ (where O is the set of observations in the system), defined inductively by

$$\begin{aligned} \text{view}_u(\epsilon) &= \text{obs}_u(s_0), \text{ and} \\ \text{view}_u(\alpha a) &= \begin{cases} \text{view}_u(\alpha) a \text{obs}_u(s_0 \cdot \alpha) & \text{if } \text{dom}(a) = u \\ \text{view}_u(\alpha) \circ \text{obs}_u(s_0 \cdot \alpha) & \text{otherwise} \end{cases} \end{aligned}$$

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 domain u in the run generated by the sequence of actions α . The absorptive concatenation is intended to capture that the system is asynchronous, with domains not having access to a global clock. Thus, two periods of different length during which a particular observation obtains are not distinguishable to the domain.

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

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

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 $\text{dom}(a)$ *actually* has, as represented by its view $\text{view}_{\text{dom}(a)}(\alpha)$. By contrast, TA-security uses in place of this the maximal information that $\text{dom}(a)$ *may* have. (The nomenclature ‘ to ’ is intended to be suggestive of *transmission* of information about *observations*.)

Another variant of these definitions is mentioned in [Mey08] because of its relationship to a definition of Roscoe and Goldsmith [RG99]. Given a policy \rightsquigarrow , for each domain $u \in D$, define the function $\text{ito}_u : A^* \rightarrow \mathcal{T}(O(A \cup O)^*, A)$ by $\text{ito}_u(\epsilon) = \text{obs}_u(s_0)$ and

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

This definition is just like that of to , with the difference that the information that may be transmitted by an action a to domains u such that $\text{dom}(a) \rightsquigarrow u$ but $\text{dom}(a) \neq u$, includes the observation $\text{obs}_{\text{dom}(a)}(s_0 \cdot \alpha a)$ produced by the action a . Intuitively, the definition of security based on this notion will allow that the




Figure 2: A system for the downgrader policy

action a transmits not just the information observable to $\text{dom}(a)$ at the time that it is invoked, but also the new information that it computes and makes observable in $\text{dom}(a)$. This information is not included in the value $\text{ito}_{\text{dom}(a)}(\alpha)$ itself, since the definition of security will state that the new observation may depend only on this value. The nomenclature in this case is intended to be suggestive of *immediate* transmission of information about observations.

We may now base the definition of security on either the function to or ito rather than ta .

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

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

We remark that under certain circumstances, the definitions given by Roscoe and Goldsmith [RG99] correspond either to P-security or to ITO-security— see [Mey07] for details. The following result shows how these definitions are related:

Theorem 1 ([Mey08]) *With respect to a given policy \rightsquigarrow , P-security implies TO-security implies ITO-security implies TA-security implies IP-security.*

Examples showing that all these notions are distinct are presented in [Mey08]. We give just one example here to illustrate how the definitions work.

Example 3: Figure 2 illustrates a system for the downgrader architecture \mathcal{A} of Figure 1. There are two actions h, d , with $\text{dom}(h) = H$ and $\text{dom}(d) = D$. Transitions on taking an action are indicated by edge labels, and the observation made at a state is indicated below the state for each domain. For example, at the initial (leftmost) state s_0 we have $\text{obs}_L(s_0) = 0$, and at the (rightmost) state t reached from s_0 by the sequence of actions hd we have $\text{obs}_D(t) = 1$.

This system is not TO-secure (hence also not P-secure). To see this, consider

the sequences $\alpha = d$ and $\beta = hd$. We have

$$\begin{aligned} \mathbf{to}_L(d) &= (\mathbf{to}_L(\epsilon), \mathbf{view}_D(\epsilon), d) \\ &= (\epsilon, 0, d) \\ &= (\mathbf{to}_L(\epsilon), \mathbf{view}_D(h), d) \\ &= \mathbf{to}_L(hd). \end{aligned}$$

but $\mathbf{obs}_L(s_0 \cdot d) = 0$ and $\mathbf{obs}_L(s_0 \cdot hd) = 1$, so this is a violation of TO-security. On the other hand, these sequences do not yield a violation of ITO-security, because

$$\begin{aligned} \mathbf{ito}_L(d) &= (\mathbf{ito}_L(\epsilon), \mathbf{view}_D(d), d) \\ &= (\epsilon, 0d0, d). \end{aligned}$$

and

$$\begin{aligned} \mathbf{ito}_L(hd) &= (\mathbf{ito}_L(h), \mathbf{view}_D(hd), d) \\ &= (\mathbf{ito}_L(\epsilon), \mathbf{view}_D(hd), d) \\ &= (\epsilon, 0d1, d). \end{aligned}$$

so $\mathbf{ito}_L(d) \neq \mathbf{ito}_L(hd)$. It can be shown that the system is in fact ITO-secure (and therefore also TA-secure and IP-secure). Intuitively, L learns that the H action has been performed, but this (according to ITO-security) is considered secure because D also acquires this information (at the same time), and D is in fact responsible for having transmitted the information to L . TO-security involves the more stringent requirement that D *know* (by seeing in its view) that h has been performed *before* D performs the action d that transmits the information to L . On the other hand TA-security is more liberal: the system would remain TA-secure even if we were to replace the final D observation of 1 by the observation 0, so that D never learns whether H has performed h . Intuitively, TA-security requires just that D be causally involved in any transmission of information about H activity to L . \square

Since definitions of security like those introduced above quantify over the infinite set of sequences of actions, it is desirable to have a proof technique for security that avoids having to examine an infinite number of possibilities. One approach that is prominent in the literature is the use of “unwinding” relations [GM84]. Rushby [Rus92] discusses the following unwinding relations for intransitive noninterference

Definition 5 A weak unwinding for a system M with respect to a policy \succrightarrow is an indexed collection of equivalence relations $\{\sim_u\}_{u \in D}$ on the states of M satisfying the following conditions:

$$OC: \text{ If } s \sim_u t \text{ then } \mathbf{obs}_u(s) = \mathbf{obs}_u(t). \quad (\text{Output Consistency})$$

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

$$LR: \text{ If } \text{dom}(a) \not\prec_u \text{ then } s \sim_u s \cdot a. \quad (\text{Left Respect})$$

Rushby shows that this notion provides a sound proof technique for IP-security. The following result of van der Meyden [Mey08] shows that it is in fact also a sound proof technique for the stronger notion TA-security.

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

The converse implication does not quite hold, but it turns out that unwinding is in fact a complete proof technique for TA-security if we allow consideration of a bisimilar variant of M . 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. It is easy to see that $\text{uf}(M)$ is bisimilar to M in an obvious sense. The following result of van der Meyden [Mey08] provides a sense in which unwindings are a complete proof technique for TA-security.

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

We use this characterization in the results that follow.

4 Soundness of Architectural Refinement

We are now in a position to assign a formal meaning to architectures, and to state precisely the main result of the paper.

Let X -security be a notion of security for information flow policies, like those discussed in the previous section. We say that a system M is X -compliant with an architecture $\mathcal{A} = (D, \rightsquigarrow)$, where X is a definition of security, if

1. D is the set of domains of M , and
2. M is X -secure with respect to \rightsquigarrow .

Intuitively, M is X -compliant with an architecture if it has an implementation for each of the components required by the architecture, and the information flows between these implementation components is consistent with the architecture’s policy (with consistency defined by X -security.)

Consider now the definition of architectural refinement introduced in Section 2. In a system design process we may have started with a high level architecture \mathcal{A} , refined this to a lower level architecture \mathcal{B} , and then constructed a system that implements \mathcal{B} . In what sense have we then implemented the architecture \mathcal{A} that formed the highest level specification of the system? For this, we need to be able to view the system from the perspective of the set of domains of the higher level architecture \mathcal{A} rather than those of \mathcal{B} . This is the intent of the following definition.

Let $M = \langle S, s_0, A, \text{step}, \text{obs}^1, \text{dom}^1 \rangle$ be a system with set of domains D_1 , and suppose $r : D_1 \rightarrow D_2$ is surjective. Then we may construct a system $r(M) = \langle S, s_0, A, \text{step}, \text{obs}^2, \text{dom}^2 \rangle$ as follows:

1. the actions A , set of states S , initial state s_0 , and transition function step of $r(M)$ are exactly as in M ,
2. the set of domains of $r(M)$ is D_2
3. $\text{dom}^2(a) = r(\text{dom}^1(a))$ for each $a \in A$,
4. for $u \in D_2$ and $s \in S$, the observation $\text{obs}_u^2(s)$ is the function $f : r^{-1}(u) \rightarrow O$, given by $f(v) = \text{obs}_v^1(s)$ for $v \in r^{-1}(u)$.

Intuitively, each domain $u \in D_2$ is viewed by the refinement mapping as being broken down into the collection of subdomains $r^{-1}(u)$. An action of a subdomain in M is interpreted in $r(M)$ as belonging to its superdomain. The observation of a superdomain in $r(M)$ is taken to be the collection of all observations of its subdomains.

We can now give a formal meaning to soundness of architectural refinement. Say that a notion of compliance X is *preserved under architectural refinement* if whenever $r : \mathcal{A}_1 \rightarrow \mathcal{A}_2$ is a refinement mapping, if M is a system that is X -compliant with \mathcal{A}_1 , then $r(M)$ is X -compliant with \mathcal{A}_2 . That, is, if the goal of design was to construct a system that complies with architecture \mathcal{A}_2 , we may do so by building a system M that complies with architecture \mathcal{A}_1 and viewing this system from the perspective of \mathcal{A}_2 through the mapping r .

Our main result is the following. (We give the proof later.)

Theorem 4 *The following notions are all preserved under architectural refinement: P-compliance, TO-compliance, ITO-compliance, TA-compliance and IP-compliance.*

One might also view a system to be secure if it has a weak unwinding, although there are some difficulties with this in view of the fact that weak unwinding is not preserved under bisimulation [Mey08]. Nevertheless, the following result is useful for the proof of Theorem 4, and is of independent interest in that it shows that a sound proof technique for security is preserved by architectural refinement.

Theorem 5 *Let $(D_1, \succ_1) \leq_r (D_2, \succ_2)$ and suppose that M is a system with set of domains D_1 . If there exists a weak unwinding on M then there exists a weak unwinding on $r(M)$.*

Proof: Let $\{\sim_u\}_{u \in D}$ be a weak unwinding on M . Define the family of binary relations $\{\approx_u\}_{u \in D}$ on the states of $r(M)$ by $s \approx_u t$ if $s \sim_v t$ for all $v \in r^{-1}(u)$. We show that this is a weak unwinding on $r(M)$, i.e., that it satisfies OC, WSC and LR.

OC: If $s \approx_u t$ then $s \sim_v t$ for all $v \in r^{-1}(u)$. Since \sim satisfies OC, this means that $\text{obs}_v^M(s) = \text{obs}_v^M(t)$ for all $v \in r^{-1}(u)$. This in turn implies that $\text{obs}_u^{r(M)}(s) = \text{obs}_u^{r(M)}(t)$, as required for OC.

WSC: Suppose $s \approx_u t$ and $s \approx_{\text{dom}^{r(M)}(a)} t$. Note that $\text{dom}^M(a) \in r^{-1}(\text{dom}^{r(M)}(a))$, so the latter implies $s \sim_{\text{dom}^M(a)} t$. The former gives that for all $v \in r^{-1}(u)$, we

have $s \sim_v t$. Since \sim satisfies WSC, we conclude that $s \cdot a \sim_v t \cdot a$ for all $v \in r^{-1}(u)$. This yields that $s \cdot a \approx_u t \cdot a$.

LR: Suppose $\text{dom}^{r(M)}(a) \not\rightarrow_2 u$. Then also $\text{dom}^M(a) \not\rightarrow_1 v$, for all $v \in r^{-1}(u)$. By LR for \sim , we obtain that $s \cdot a \sim_v s$ for all $v \in r^{-1}(u)$. Thus, $s \cdot a \approx_u s$. \square

To prove Theorem 4, we first prove a number of technical lemmas. The following lemma captures a uniform pattern of explanation for the proof of Theorem 4. Note that the definitions of security are expressed in terms of mappings X taking as arguments a system M , a domain u of this system, a sequence of actions α in this system, and an information flow policy \rightarrow on the set of domains of the system. (In our notation above, the parameters M and \rightarrow have generally been left implicit). The system is classified as secure according to the notion X with respect to policy \rightarrow if for all systems M , domains u and sequences of actions α, α' , if $X(M, u, \alpha, \rightarrow) = X(M, u, \alpha', \rightarrow)$ then $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$.

Lemma 1 *Let $(D_1, \rightarrow_1) \leq_r (D_2, \rightarrow_2)$ and suppose that M is a system with set of domains D_1 . Suppose M is X -compliant with (D_1, \rightarrow_1) and that for all domains $v \in D_1$ and all sequences of actions α, α' , if*

$$X(r(M), r(v), \alpha, \rightarrow_2) = X(r(M), r(v), \alpha', \rightarrow_2)$$

then

$$X(M, v, \alpha, \rightarrow_1) = X(M, v, \alpha', \rightarrow_1).$$

Then $r(M)$ is X -compliant with (D_2, \rightarrow_2) .

Proof: Suppose M is X -compliant with (D_1, \rightarrow_1) . To show that $r(M)$ is X -compliant with (D_2, \rightarrow_2) , suppose that α, α' are sequences of actions of $r(M)$ (hence also of M), and u a domain of $r(M)$, such that

$$X(r(M), u, \alpha, \rightarrow_2) = X(r(M), u, \alpha', \rightarrow_2).$$

We need to show that $\text{obs}_u^{r(M)}(s_0 \cdot \alpha) = \text{obs}_u^{r(M)}(s_0 \cdot \alpha')$. By definition of $r(M)$, this amounts to showing that for all $v \in r^{-1}(u)$, we have $\text{obs}_v^M(s_0 \cdot \alpha) = \text{obs}_v^M(s_0 \cdot \alpha')$. For this, note that we have, by the assumed property, that

$$X(M, v, \alpha, \rightarrow_1) = X(M, v, \alpha', \rightarrow_1)$$

for all $v \in r^{-1}(u)$. Since M is X -compliant with (D_1, \rightarrow_1) , this implies that $\text{obs}_v^M(s_0 \cdot \alpha) = \text{obs}_v^M(s_0 \cdot \alpha')$, as required. \square

In practice, we will often establish the condition of the lemma by constructing for all domains $v \in D_1$ a mapping F_v such that

$$X(M, v, \alpha, \rightarrow_1) = F_v(X(r(M), r(v), \alpha, \rightarrow_2))$$

for all sequences of actions α of M_1 .

The following lemma is useful when showing preservation of TO-security and ITO-security.

Lemma 2 *Let $r : D_1 \rightarrow D_2$ and suppose that M is a system with set of domains D_1 . Then for all sequences of actions α, α' and domains v of M , if $\mathbf{view}_{r(v)}^{r(M)}(\alpha) = \mathbf{view}_{r(v)}^{r(M)}(\alpha')$ then $\mathbf{view}_v^M(\alpha) = \mathbf{view}_v^M(\alpha')$.*

Proof: Note that views with respect to $r(v)$ in $r(M)$ are sequences consisting of actions a such that $\mathbf{dom}^{r(M)}(a) = r(v)$, and observations which are mappings f from $r^{-1}(r(v))$ to observations in M . Define the function F_v mapping such sequences with respect to $r(v)$ in $r(M)$ to views with respect to v in M by the following induction: $F_v(\epsilon) = \epsilon$,

$$F_v(\sigma a) = \begin{cases} F_v(\sigma) a & \text{if } \mathbf{dom}^M(a) = v \\ F_v(\sigma) & \text{otherwise} \end{cases}$$

when a is an action, and $F_v(\sigma f) = F_v(\sigma) \circ f(v)$ when f is an observation. We claim that for all sequences of actions α , we have $F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha)) = \mathbf{view}_v^M(\alpha)$. The result then follows.

The proof is by induction on the length of α . In case $\alpha = \epsilon$, we have $F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha)) = F_v(\mathbf{obs}_{r(v)}^{r(M)}(s_0)) = \mathbf{obs}_v^M(s_0) = \mathbf{view}_v^M(\alpha)$. For the inductive case of a sequence αa where a is an action and the result holds for α , we consider two cases, depending on whether $\mathbf{dom}^M(a) = v$.

Suppose first that $\mathbf{dom}^M(a) = v$. Then $\mathbf{dom}^{r(M)}(a) = r(v)$. Hence

$$\begin{aligned} F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha a)) &= F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha) a \mathbf{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha a)) \\ &= F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha) a \mathbf{obs}_v^M(s_0 \cdot \alpha a)) \\ &= \mathbf{view}_v^M(\alpha) a \mathbf{obs}_v^M(s_0 \cdot \alpha a) \\ &= \mathbf{view}_v^M(\alpha a) \end{aligned}$$

Next, suppose that $\mathbf{dom}^M(a) \neq v$. Here we consider two subcases, depending on whether $\mathbf{dom}^{r(M)}(a) = r(v)$. If $\mathbf{dom}^{r(M)}(a) = r(v)$ then we have

$$\begin{aligned} F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha a)) &= F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha) a \mathbf{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha a)) \\ &= F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha) a \mathbf{obs}_v^M(s_0 \cdot \alpha a)) \\ &= \mathbf{view}_v^M(\alpha) \circ \mathbf{obs}_v^M(s_0 \cdot \alpha a) \\ &= \mathbf{view}_v^M(\alpha a) \end{aligned}$$

If $\mathbf{dom}^M(a) \neq v$ and $\mathbf{dom}^{r(M)}(a) \neq r(v)$ then we consider two further subcases, depending on whether $\mathbf{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha a) = \mathbf{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha)$. If this condition holds, then it follows that $\mathbf{obs}_v^M(s_0 \cdot \alpha a) = \mathbf{obs}_v^M(s_0 \cdot \alpha)$, hence

$$\begin{aligned} F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha a)) &= F_v(\mathbf{view}_{r(v)}^{r(M)}(\alpha)) \\ &= \mathbf{view}_v^M(\alpha) \\ &= \mathbf{view}_v^M(\alpha a). \end{aligned}$$

On the other hand, if $\text{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha a) \neq \text{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha)$, then we have

$$\begin{aligned}
F_v(\text{view}_{r(v)}^{r(M)}(\alpha a)) &= F_v(\text{view}_{r(v)}^{r(M)}(\alpha) \text{obs}_{r(v)}^{r(M)}(s_0 \cdot \alpha a)) \\
&= F_v(\text{view}_{r(v)}^{r(M)}(\alpha)) \circ \text{obs}_v^M(s_0 \cdot \alpha a) \\
&= \text{view}_v^M(\alpha) \circ \text{obs}_v^M(s_0 \cdot \alpha a) \\
&= \text{view}_v^M(\alpha a).
\end{aligned}$$

Hence, in all cases we have the desired equation. \square

We can now give the proof of Theorem 4.

Proof: Let $(D_1, \succrightarrow_1) \leq_r (D_2, \succrightarrow_2)$ and suppose that M is a system with set of domains D_1 . We show in each case that the conditions of Lemma 1 hold. If σ is a sequence and S is a set, write $\sigma \upharpoonright S$ for the subsequence of σ consisting of all elements in S .

For v a domain of a system M , and \succrightarrow a policy, let $I(M, \succrightarrow, v)$ be the set of actions a of M with $\text{dom}(a) \succrightarrow v$. Note that $I(M, \succrightarrow_1, v) \subseteq I(r(M) \succrightarrow_2, r(v))$. For, if $a \in I(M, \succrightarrow_1, v)$ then $\text{dom}_1(a) \succrightarrow_1 v$, so by the definition of refinement, $\text{dom}_2(a) = r(\text{dom}_1(a)) \succrightarrow_2 r(v)$, hence $a \in I(r(M), \succrightarrow_2, r(v))$.

P-security: For $v \in D_1$, define $F_v(\alpha) = \alpha \upharpoonright I(M, \succrightarrow_1, v)$. To see that this function has the required property, note that

$$\begin{aligned}
F_v(\text{purge}_{r(v)}^{r(M)}(\alpha)) &= \text{purge}_{r(v)}^{r(M)}(\alpha) \upharpoonright I(M, \succrightarrow_1, v) \\
&= (\alpha \upharpoonright I(r(M), \succrightarrow_2, r(v))) \upharpoonright I(M, \succrightarrow_1, v) \\
&= \alpha \upharpoonright I(M, \succrightarrow_1, v) \\
&= \text{purge}_v^M(\alpha)
\end{aligned}$$

where we use the fact that $I(M, \succrightarrow_1, v) \subseteq I(r(M) \succrightarrow_2, r(v))$ in the third step.

TO-security: For TO-security we show that for all sequences of actions α, α' , and $v \in D_1$, if $\text{to}_{r(v)}^{r(M)}(\alpha) = \text{to}_{r(v)}^{r(M)}(\alpha')$ then $\text{to}_v^M(\alpha) = \text{to}_v^M(\alpha')$. The proof is by induction on the combined length of α and α' . The base case of $\alpha = \alpha' = \epsilon$ is trivial.

Consider sequences αa and α' , where a is an action, a domain $v \in D_2$, and suppose that $\text{to}_{r(v)}^{r(M)}(\alpha a) = \text{to}_{r(v)}^{r(M)}(\alpha')$. We consider two cases, depending on whether $\text{dom}^{r(M)}(a) \succrightarrow_2 r(v)$.

If $\text{dom}^{r(M)}(a) \not\succrightarrow_2 r(v)$, then $\text{to}_{r(v)}^{r(M)}(\alpha) = \text{to}_{r(v)}^{r(M)}(\alpha a) = \text{to}_{r(v)}^{r(M)}(\alpha')$. Hence, by induction, we have $\text{to}_v^M(\alpha) = \text{to}_v^M(\alpha')$. Since $\text{dom}^{r(M)}(a) \not\succrightarrow_2 r(v)$, also $\text{dom}^M(a) \not\succrightarrow_1 v$. Hence $\text{to}_v^M(\alpha a) = \text{to}_v^M(\alpha)$. It follows that $\text{to}_v^M(\alpha a) = \text{to}_v^M(\alpha')$.

Alternately, if $\text{dom}^{r(M)}(a) \succrightarrow_2 r(v)$, then it follows from $\text{to}_{r(v)}^{r(M)}(\alpha a) = \text{to}_{r(v)}^{r(M)}(\alpha')$ that α' is not ϵ . Let $\alpha' = \beta b$, where b is an action. If $\text{dom}^{r(M)}(b) \not\succrightarrow_2 r(v)$, then we may apply the previous case with the roles of αa and α' swapped. We may assume, therefore, that $\text{dom}^{r(M)}(b) \succrightarrow_2 r(v)$. It then follows that $a = b$ and $\text{to}_{r(v)}^{r(M)}(\alpha) = \text{to}_{r(v)}^{r(M)}(\beta)$ and $\text{view}_{\text{dom}^{r(M)}(a)}^{r(M)}(\alpha) = \text{view}_{\text{dom}^{r(M)}(a)}^{r(M)}(\beta)$. From the former, we have by induction that $\text{to}_v^M(\alpha) = \text{to}_v^M(\beta)$. From the latter, we

have by Lemma 2 that $\mathbf{view}_{\mathbf{dom}^M(a)}^{r(M)}(\alpha) = \mathbf{view}_{\mathbf{dom}^M(a)}^{r(M)}(\beta)$. In both the case that $\mathbf{dom}^M(a) \mapsto_1 v$ and the case that $\mathbf{dom}^M(a) \not\mapsto_1 v$, it follows from these that $\mathbf{to}_v^M(\alpha a) = \mathbf{to}_v^M(\beta a)$.

ITO-security: In the case of ITO-security, we show that if M is ITO-secure, then $\mathbf{ito}_{r(v)}^{r(M)}(\alpha) = \mathbf{ito}_{r(v)}^{r(M)}(\alpha')$ implies $\mathbf{ito}_v^M(\alpha) = \mathbf{ito}_v^M(\alpha')$. The proof is by induction on the combined length of α and α' . The base case of $\alpha = \alpha' = \epsilon$ is trivial. Consider sequences αa and α' , where a is an action, a domain $v \in D_2$, and suppose that $\mathbf{ito}_{r(v)}^{r(M)}(\alpha a) = \mathbf{ito}_{r(v)}^{r(M)}(\alpha')$. The case of $\mathbf{dom}^{r(M)}(a) \mapsto_2 r(v)$ proceeds exactly along the lines of the argument for TO-security, and is left to the reader.

In case $\mathbf{dom}^{r(M)}(a) \mapsto_2 r(v)$, it follows from $\mathbf{ito}_{r(v)}^{r(M)}(\alpha a) = \mathbf{ito}_{r(v)}^{r(M)}(\alpha')$ that α' is not ϵ . Let $\alpha' = \beta b$, where b is an action. If $\mathbf{dom}^{r(M)}(b) \not\mapsto_2 r(v)$, then we may apply the previous case with the roles of αa and α' swapped. We may assume, therefore, that $\mathbf{dom}^{r(M)}(b) \mapsto_2 r(v)$. It then follows that $a = b$ and $\mathbf{ito}_{r(v)}^{r(M)}(\alpha) = \mathbf{ito}_{r(v)}^{r(M)}(\beta)$ and either

1. $\mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\alpha a) = \mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\beta a)$ (in case $\mathbf{dom}^{r(M)}(a) \neq r(v)$) or
2. $\mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\alpha) = \mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\beta)$. (in case $\mathbf{dom}^{r(M)}(a) = r(v)$).

By induction, we have that $\mathbf{ito}_v^M(\alpha) = \mathbf{ito}_v^M(\beta)$. Note that in the case that $\mathbf{dom}^{r(M)}(a) = r(v)$, we also have by the induction hypothesis that $\mathbf{ito}_{\mathbf{dom}^M(a)}^M(\alpha) = \mathbf{ito}_{\mathbf{dom}^M(a)}^M(\beta)$. Hence

$$\begin{aligned} \mathbf{ito}_{\mathbf{dom}^M(a)}^M(\alpha a) &= (\mathbf{ito}_{\mathbf{dom}^M(a)}^M(\alpha), \mathbf{view}_{\mathbf{dom}^M(a)}^M(\alpha), a) \\ &= (\mathbf{ito}_{\mathbf{dom}^M(a)}^M(\beta), \mathbf{view}_{\mathbf{dom}^M(a)}^M(\beta), a) \\ &= \mathbf{ito}_{\mathbf{dom}^M(a)}^M(\beta a). \end{aligned}$$

By ITO-security of M , it then follows that $\mathbf{obs}_{\mathbf{dom}^M(a)}(s_0 \cdot \alpha a) = \mathbf{obs}_{\mathbf{dom}^M(a)}(s_0 \cdot \beta a)$. Together with $\mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\alpha) = \mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\beta)$, this gives the conclusion that $\mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\alpha a) = \mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\beta a)$. Thus, we in fact have $\mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\alpha a) = \mathbf{view}_{\mathbf{dom}^{r(M)}(a)}^{r(M)}(\beta a)$ irrespective of whether $\mathbf{dom}^{r(M)}(a) = r(v)$.

Thus, in both the case that $\mathbf{dom}^M(a) \mapsto_1 v$ and the case that $\mathbf{dom}^M(a) \not\mapsto_1 v$, it follows from these that $\mathbf{ito}_v^M(\alpha a) = \mathbf{ito}_v^M(\beta a)$.

TA-security: By Theorem thm:ta-unwind-equiv, a system N is TA-secure if there exists a weak unwinding on the unfolded system $\mathbf{uf}(N)$. It is easy to verify that $\mathbf{uf}(r(M)) = r(\mathbf{uf}(M))$. Hence if M is TA-secure then there exists a weak unwinding on $\mathbf{uf}(M)$. By Theorem 5, there exists a weak unwinding $r(\mathbf{uf}(M))$, hence on $\mathbf{uf}(r(M))$. It follows that $r(M)$ is TA-secure.

IP-security: We claim that if $X \subseteq D_1$ and $Y \subseteq D_2$ are sets of domains such that $r(X) \subseteq Y$, then $\mathbf{ipurge}_X^M(\alpha) = \mathbf{ipurge}_X^M(\mathbf{ipurge}_Y^{r(M)}(\alpha))$. Taking $Y =$

$\{r(v)\}$ and $X = \{v\}$, we then obtain $\text{ipurge}_v^M(\alpha) = \text{ipurge}_v^M(\text{ipurge}_{r(v)}^{r(M)}(\alpha))$, so we may take $F_v(\alpha) = \text{ipurge}_v^M(\alpha)$.

To prove the claim, we proceed by induction on the length of α . If S is set of domains and u a domain, we write $u \mapsto S$ if there exists $v \in S$ such that $u \mapsto v$. The base case is trivial. Consider a sequence of actions αa , where the claim holds for α , and suppose $r(X) \subseteq Y$. We consider a number of cases, depending on whether $\text{dom}^{r(M)}(a) \mapsto_2 Y$.

Suppose first that $\text{dom}^{r(M)}(a) \not\mapsto_2 Y$. Note that if we were to have $\text{dom}^M(a) \mapsto_1 v \in X$, then by the definition of refinement we would have $\text{dom}^{r(M)}(a) = r(\text{dom}^M(a)) \mapsto_2 r(v) \in r(X)$, hence $\text{dom}^{r(M)}(a) \mapsto_2 Y$. Thus, we must in fact have $\text{dom}^M(a) \not\mapsto_1 X$. Hence

$$\begin{aligned} & \text{ipurge}_X^M(\text{ipurge}_Y^{r(M)}(\alpha a)) \\ &= \text{ipurge}_X^M(\text{ipurge}_Y^{r(M)}(\alpha)) \quad \text{since } \text{dom}^{r(M)}(a) \not\mapsto_2 Y \\ &= \text{ipurge}_X^M(\alpha) \quad \text{by induction} \\ &= \text{ipurge}_X^M(\alpha a) \end{aligned}$$

This proves the claim for the sequence αa .

Next, suppose that $\text{dom}^{r(M)}(a) \mapsto_2 Y$. We further break this case into two subcases, depending on whether $\text{dom}^M(a) \mapsto_1 X$. Suppose first that $\text{dom}^M(a) \not\mapsto_1 X$. Note that $r(X) \subseteq Y$ implies $r(X) \subseteq Y \cup \{\text{dom}^M(a)\}$. Hence

$$\begin{aligned} & \text{ipurge}_X^M(\text{ipurge}_Y^{r(M)}(\alpha a)) \\ &= \text{ipurge}_X^M(\text{ipurge}_{Y \cup \{\text{dom}^{r(M)}(a)\}}^{r(M)}(\alpha) a) \quad \text{since } \text{dom}^{r(M)}(a) \mapsto_2 Y \\ &= \text{ipurge}_X^M(\text{ipurge}_{Y \cup \{\text{dom}^{r(M)}(a)\}}^{r(M)}(\alpha)) \quad \text{since } \text{dom}^M(a) \not\mapsto_1 X \\ &= \text{ipurge}_X^M(\alpha) \quad \text{by induction} \\ &= \text{ipurge}_X^M(\alpha a) \quad \text{since } \text{dom}^M(a) \not\mapsto_1 X. \end{aligned}$$

Alternately, suppose that $\text{dom}^M(a) \mapsto_1 X$. Note that $r(X \cup \{\text{dom}^M(a)\}) \subseteq Y \cup \{r(\text{dom}^M(a))\} = Y \cup \{\text{dom}^{r(M)}(a)\}$. Hence

$$\begin{aligned} & \text{ipurge}_X^M(\text{ipurge}_Y^{r(M)}(\alpha a)) \\ &= \text{ipurge}_X^M(\text{ipurge}_{Y \cup \{\text{dom}^{r(M)}(a)\}}^{r(M)}(\alpha) a) \quad \text{since } \text{dom}^{r(M)}(a) \mapsto_2 Y \\ &= \text{ipurge}_{X \cup \{\text{dom}^M(a)\}}^M(\text{ipurge}_{Y \cup \{\text{dom}^{r(M)}(a)\}}^{r(M)}(\alpha)) a \quad \text{since } \text{dom}^M(a) \mapsto_1 X \\ &= \text{ipurge}_{X \cup \{\text{dom}^M(a)\}}^M(\alpha) a \quad \text{by induction} \\ &= \text{ipurge}_X^M(\alpha a) \quad \text{since } \text{dom}^M(a) \mapsto_1 X. \end{aligned}$$

This completes the proof of the claim. \square

We remark that it can be shown independently of other assumptions that if X is one of **purge**, **ipurge**, **ta**, or **to** then (with all parameters made explicit), if

$$X(r(M), r(v), \alpha, \mapsto_2) = X(r(M), r(v), \alpha', \mapsto_2)$$

then

$$X(M, v, \alpha, \mapsto_1) = X(M, v, \alpha', \mapsto_1).$$




Figure 3: ito not preserved under refinement

The following example shows this is not true for $X = \text{ito}$, but the proof of Theorem 4 handles this case by showing that it does hold under the further assumption that M is ITO-secure.

Example 4: Consider system in Figure 3, where there are domains H, I, J with actions h of domain H and i of domain I . Observations in the domains are given under the states in the order H, I, J . Intuitively, action i is used to test whether there has been an occurrence of h . If so, this fact is made observable to both I and J simultaneously.

Let \rightarrow_1 be the reflexive closure on $\{H, I, J\}$ of the fact $I \rightarrow_1 J$, and let r be given by $r(H) = H$ and $r(I) = r(J) = K$. Let \rightarrow_2 be the smallest reflexive relation on $\{H, K\}$. Then r is a refinement mapping. Take $\alpha = hi$ and $\beta = i$. Then

$$\begin{aligned}
\text{ito}_K^{r(M)}(\alpha) &= (\text{ito}_K^{r(M)}(h), \text{view}_K^{r(M)}(h), i) \\
&= (\text{ito}_K^{r(M)}(\epsilon), (0^I 0^J), i) \\
&= (\epsilon, (0^I 0^J), i) \\
&= (\text{ito}_K^{r(M)}(\epsilon), \text{view}_K^{r(M)}(\epsilon), i) \\
&= \text{ito}_K^{r(M)}(\beta)
\end{aligned}$$

where $(0^I 0^J)$ denotes the function mapping I to 0 and J to 0. However,

$$\begin{aligned}
\text{ito}_J^M(\alpha) &= (\text{ito}_J^{r(M)}(h), \text{view}_I^{r(M)}(h), i) \\
&= (\text{ito}_J^M(\epsilon), 0i1, i)
\end{aligned}$$

and

$$\begin{aligned}
\text{ito}_J^M(\beta) &= (\text{ito}_J^M(\epsilon), \text{view}_I^{r(M)}(i), i) \\
&= (\text{ito}_J^M(\epsilon), 0i0, i)
\end{aligned}$$

so although $r(J) = K$, we do not have that $\text{ito}_J^M(\alpha) = \text{ito}_J^M(\beta)$. \square

5 Access Control

The interpretations of policies discussed above are somewhat abstract. We now consider a further, more concrete interpretation, that is closer to the level of *mechanisms* for control of information flow. One of the key mechanisms for implementation of secure systems is access control. Rushby [Rus92] defined the notion of access control system and showed that access control systems are IP-secure with respect to a policy \mapsto if they are consistent with the policy in an appropriate sense. We present here a variant of Rushby’s definitions due to van der Meyden, that strengthens Rushby’s result (see [Mey08] for an explanation of how this variant improves on Rushby’s.)

Define an *access control structure* for a machine $\langle S, s_0, A, \text{step}, \text{obs}, \text{dom} \rangle$ with domains D to be a tuple $\text{AC} = (N, V, \text{contents}, \text{alter}, \text{observe})$, where

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

We call the pair (M, AC) a system with structured state. For a system with structured state, define for each domain $u \in D$, an equivalence relation on the states S of M , of *observable content equivalence*, by $s \sim_u^{\text{oc}} t$ if $\text{contents}(s, n) = \text{contents}(t, n)$ for all $n \in \text{observe}(u)$. That is, two states are related for u if they are identical with respect to the values of objects that u may observe.

The following conditions are a variant of Rushby’s “Reference Monitor” conditions.

WAC1. If $s \sim_u^{\text{oc}} t$ then $\text{obs}_u(s) = \text{obs}_u(t)$.

WAC2. For all actions a , states s, t and objects $n \in \text{alter}(\text{dom}(a))$, if $s \sim_{\text{dom}(a)}^{\text{oc}} t$ and $\text{contents}(s, n) = \text{contents}(t, n)$ then $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$.

WAC3. If $\text{contents}(s \cdot a, n) \neq \text{contents}(s, n)$ then $n \in \text{alter}(\text{dom}(a))$.

Intuitively, WAC1-WAC3 capture the conditions under which the machine operates in accordance with the intuitive interpretations of the structure AC . WAC1 and WAC3 are identical to Rushby’s RM1 and RM3, respectively. WAC1 says that a domain’s observation depends only on the values of the objects observable to it. WAC2 (a modification of Rushby’s condition RM2) says that if an action in domain u is permitted to change the value of an object n , then the

new value of n depends only on its old value and the values of objects of domain u . WAC3 says that if an action can change the value of an object, then the domain of that action is in fact permitted to alter that object. We call the pair (M, \mathbf{AC}) a *weak access control system* if it satisfies WAC1-WAC3. We say that M *admits a weak access control interpretation* if there exists an access control structure \mathbf{AC} such that (M, \mathbf{AC}) is a weak access control system.

Plainly, if there is an object that domain u may alter and domain v may observe, then information flow from domain u to v cannot be prevented. We say that a policy \succrightarrow is *consistent* with a weak access control system if the following condition is satisfied:

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

Access control interpretations are closely related to TA-security in a way that is similar to the connection between weak unwinding and TA-security already noted above in Theorem 3. The following result is shown in [Mey08].

Theorem 6 *The following are equivalent*

1. M is TA-secure with respect to \succrightarrow ,
2. $\mathbf{uf}(M)$ admits a weak access control interpretation consistent with \succrightarrow ,

Moreover, if M admits a weak access control interpretation consistent with \succrightarrow , then so does $\mathbf{uf}(M)$ (hence M is TA-secure).

This result shows that TA-security captures, in a precise sense, the notion of information flow security that may be enforced by access control mechanisms.

We now consider the interaction of refinement and access control structure. Suppose that M is a system for the set of domains D_1 , with access control structure $\mathbf{AC}_1 = (N_1, V_1, \mathbf{contents}_1, \mathbf{observe}_1, \mathbf{alter}_1)$ with respect to which M is a weak access control system. Let $r : (D_1, \succrightarrow_1) \rightarrow (D_2, \succrightarrow_2)$ be a refinement mapping. Define $r(\mathbf{AC}_1) = (N_2, V_2, \mathbf{contents}_2, \mathbf{observe}_2, \mathbf{alter}_2)$ to be the access control structure given by

1. $N_2 = N_1$ and $V_2 = V_1$ and $\mathbf{contents}_2 = \mathbf{contents}_1$,
2. $\mathbf{observe}_2(u) = \bigcup_{v \in r^{-1}(u)} \mathbf{observe}_1(v)$,
3. $\mathbf{alter}_2(u) = \bigcup_{v \in r^{-1}(u)} \mathbf{alter}_1(v)$.

Intuitively, \mathbf{AC}_2 has the same set of objects, with the same contents at each state of M as in \mathbf{AC}_1 , and each domain of D_2 observes and alters the objects in all of its subdomains in D_1 . The following result shows that weak access control structure is preserved under architectural refinement.

Theorem 7 *Let $r : (D_1, \succrightarrow_1) \rightarrow (D_2, \succrightarrow_2)$ be a refinement mapping.*

1. *If (M, \mathbf{AC}_1) is a weak access control system then $(r(M), r(\mathbf{AC}_1))$ is a weak access control system.*

2. If AC_1 is consistent with \succrightarrow_1 then $r(\text{AC}_1)$ is consistent with \succrightarrow_2 .

Proof: Write \sim^1 and \sim^2 for the relations of observable content equivalence on the states of M with respect to AC_1 and $\text{AC}_2 = r(\text{AC}_1)$, respectively. Let O^2 be the observation function on $r(M)$. Since $\text{contents}_1 = \text{contents}_2$, we write just contents in either case. For part (1), we prove WAC1-WAC3.

WAC1: We need to show that $s \sim_u^2 t$ implies $O_u^2(s) = O_u^2(t)$, for all $u \in D_2$. By definition, $s \sim_u^2 t$ implies that $s \sim_v^1 t$ for all $v \in r^{-1}(u)$. By WAC1 for M and AC_1 , we have $O_v^1(s) = O_v^1(t)$ for all $v \in r^{-1}(u)$. By construction of O^2 we obtain that $O_u^2(s) = O_u^2(t)$.

WAC3: Suppose that $n \notin \text{alter}_2(\text{dom}^2(a))$. Then for all $v \in r^{-1}(u)$, we have $n \notin \text{alter}_1(v)$. In particular, since $r(\text{dom}^1(a)) = \text{dom}^2(a)$, we have $n \notin \text{alter}_1(\text{dom}^1(a))$. By WAC3 for AC_1 , we obtain that $\text{contents}(s \cdot a, n) = \text{contents}(s, n)$, as required.

WAC2: Let a be an action of M with $\text{dom}^1(a) = v$ and $\text{dom}^2(a) = u = r(v)$. Suppose that $n \in \text{alter}_2(u)$ and $s \sim_u^2 t$ and $\text{contents}(s, n) = \text{contents}(t, n)$. We need to show that $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$. There are two possibilities: $n \in \text{alter}_1(v)$ or $n \notin \text{alter}_1(v)$. In case $n \notin \text{alter}_1(v)$, then we have by what we proved above for WAC3 that $\text{contents}(s \cdot a, n) = \text{contents}(s, n)$ and $\text{contents}(t, n) = \text{contents}(t \cdot a, n)$, from which it is immediate that $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$, as required. Alternately, suppose that $n \in \text{alter}_1(v)$. By definition of observe^2 and contents^2 and the fact that $v \in r^{-1}(u)$, it follows that $s \sim_v^1 t$. Thus, by WAC2 for AC_1 , we again obtain that $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$.

For part (2) suppose AC_1 is consistent with \succrightarrow_1 . Let $\text{alter}_2(u_1) \cap \text{observe}_2(u_2) \neq \emptyset$, with this set containing, say, the object n . Then there exists $v_1 \in r^{-1}(u_1)$ and $v_2 \in r^{-1}(u_2)$ such that $n \in \text{alter}_1(v_1)$ and $v \in \text{observe}_1(v_2)$. Thus $\text{alter}_1(v_1) \cap \text{observe}_1(v_2) \neq \emptyset$, and by WAC3 on AC_1 , we obtain that $v_1 \succrightarrow_1 v_2$. Since r is a refinement mapping, it follows that $u_1 = r(v_1) \succrightarrow_2 r(v_2) = u_2$, as required. \square

It is worth noting that we could also work with a more expressive notion of access control structure in which the functions alter and observe take as inputs the *actions* of a system M (rather than the domains.) Intuitively, this amounts to specifying constraints on the objects that each action is permitted to read and write. We will call an access control structure of this type an *action-based* access control structure, and refer to the former type as a *domain-based* access control structure.

Given $\text{AC} = (N, V, \text{contents}, \text{observe}, \text{alter})$, an action-based access control structure over actions A , and a domain mapping $\text{dom} : A \rightarrow D$, we may construct $\text{AC}/\text{dom} = (N, V, \text{contents}, \text{observe}', \text{alter}')$, a domain-based access control structure, by defining

$$\text{observe}'(u) = \cup\{\text{observe}(a) \mid a \in A, \text{dom}(a) = u\}$$

and

$$\text{alter}'(u) = \cup\{\text{alter}(a) \mid a \in A, \text{dom}(a) = u\}.$$

Given a system M , with domain u , let \sim_u^{oc} be the relation defined above, with respect to AC/dom , and for an *action* a define the relation \sim_a^{oc} on the states of M by $s \sim_a^{\text{oc}} t$ if $\text{contents}(s, n) = \text{contents}(t, n)$ for all $n \in \text{observe}(a)$.

The semantic conditions defining an action-based access control system can now be stated as a variant of WAC1-WAC3.

WAC1_{*a*}. For $u \in D$, if $s \sim_u^{\text{oc}} t$ then $\text{obs}_u(s) = \text{obs}_u(t)$.

WAC2_{*a*}. For all actions a , states s, t and names $n \in \text{alter}(a)$, if $s \sim_a^{\text{oc}} t$ and $\text{contents}(s, n) = \text{contents}(t, n)$ then $\text{contents}(s \cdot a, n) = \text{contents}(t \cdot a, n)$.

WAC3_{*a*}. If $\text{contents}(s \cdot a, n) \neq \text{contents}(s, n)$ then $n \in \text{alter}(a)$.

Intuitively, WAC1_{*a*} says that the observations of a domain depend only on the contents of objects that actions in that domain may observe. The remaining conditions are similar to the domain-based versions, except that we work at the level of actions rather than domains.

Proposition 1 *If M satisfies WAC1_{*a*}-WAC3_{*a*} with respect to the action-based access control structure AC and \mapsto , then M satisfies WAC1-WAC3 with respect to AC/dom and \mapsto .*

Example 5: To illustrate the interaction of architectural refinement and the implementation of architectures using access control structure, we reconsider the refinement of Example 2. As a further step of refinement towards the implementation of the system, we choose to implement architecture \mathcal{B} using an action-based access control structure $(N, V, \text{contents}, \text{observe}, \text{alter})$. The set N consists of the following objects:

1. local data $l_1, l_2, h_1, h_2, d, hdb$ for $L_1, L_2, H_1, H_2, D, HDB$, respectively,
2. high level files f_1, f_2
3. input buffers hin, din_h, din_l, lin for messages to H, D and L . (In the case of D , we have separate buffers din_h and din_l to receive communications from H and L respectively — this allows the sender to receive acknowledgements without creating a covert channel from H to L .)

The table in Figure 4 gives the actions associated to each domain, and the functions observe and alter .

Finally, we define observations in the system using the pairs $(L_i, \{l_i\}), (D, \{d\}), (H_i, \{h_i\}), (HDB, \{hdb, f_1\})$. Here the first component of a pair gives a domain u , and the observation $O_u(s)$ at state s is defined to be the sequence of values $\text{contents}(s, n)$ where n is an element of the second component. Note that we have not made f_2 observable - this might represent, e.g., that f_2 is used for internal data structures of the database and is not visible at the interface of the database.

Domain	Actions	observe	alter	Purpose
$L_i, i = 1, 2$	request(L_i) send(L_i, H) get(L_i) internal(L_i)	l_i l_i lin l_i	din_l, l_i hin, l_i lin, l_i l_i	request information from D send information to H read L input buffer local computation
D	get _{h} (D) get _{l} (D) query(D) respond(D) internal(D)	din_h din_l d d d	din_h, d din_l, d d, hdb d, lin d	read D input buffer from H read D input buffer from L send query to H database send response to L request local computation
$H_i, i = 1, 2$	request(H_i) internal(H_i)	h_i h_i	h_i, hdb h_i	send query/update to HDB local computation
HDB	get(HDB) respond(HDB, H_i) respond(HDB, D) internal(HDB)	hin hdb, f_1, f_2 hdb, f_1, f_2 hdb, f_1, f_2	hin, hdb h_i, hdb d, hdb hdb, f_1, f_2	read H input buffer respond to H_i request respond to D request local computation

Figure 4: Actions of an access control system

It is straightforward to check that this action-based access control structure induces a domain-based access control structure that is consistent with the policy of \mathcal{B} . Thus, by Theorem 6, any system M for this access control structure that satisfies the conditions $WAC1_a$ - $WAC3_a$ is TA-compliant with \mathcal{B} . Further, using either the preservation of TA-compliance under refinement (Theorem 4), or the preservation of access control structure under refinement (Theorem 7) and then Theorem 6, it also follows that $r(M)$ is TA-compliant with \mathcal{A} . Here $r(M)$ is the system where the observations are defined by the pairs $(L, \{l_1, l_2\})$, $(D, \{d\})$ and $(H, \{h_1, h_2, hdb, f_1\})$.

Note that this result applies to a *class* of systems, as we have not yet defined a unique system. To do so we would need to also give the values V , and define the effect that actions have on states. Once this is done, one way to ensure $WAC2_a$ and $WAC3_a$ might be by using static analysis to verify that the code implementing an action a reads only from $\mathbf{observe}(a)$ and writes only to $\mathbf{alter}(a)$. \square

6 System Refinement

A plausible intuition concerning security is that reducing the amount of information that domains can observe will make the system *more* secure. Intuitively, this is because, with less information, domains are less likely to be able to learn secrets that they are not supposed to know. In this section we consider refinement from this perspective. Not all the definitions we consider in this paper, it turns out, support the intuition.

Suppose that $M = \langle S, s_0, A, \mathbf{dom}, \mathbf{step}, \mathbf{obs} \rangle$ and $M' = \langle S', s'_0, A, \mathbf{dom}, \mathbf{step}', \mathbf{obs}' \rangle$

are two systems with the same set of domains D , actions A and domain assignment dom . Write $M \leq M'$ if for all sequences $\alpha, \beta \in A^*$, and all domains $u \in D$ we have $\text{obs}'_u(s'_0 \cdot \alpha) = \text{obs}'_u(s'_0 \cdot \beta)$ implies $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \beta)$. That is, the observations in system M contain less information than those in system M' .

For reasons that will become clear below, it turns out to be useful to have a more general version of this notion. Let X be a function mapping domains $u \in D$ and sequences $\alpha \in A^*$ to some value. We write $M \leq^X M'$ if for all sequences $\alpha, \beta \in A^*$, and all domains $u \in D$, if $X_u(\alpha) = X_u(\beta)$ then $\text{obs}'_u(s'_0 \cdot \alpha) = \text{obs}'_u(s'_0 \cdot \beta)$ implies $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \beta)$. Note that $M \leq M'$ implies $M \leq^X M'$ for all X , so this is a weaker notion.

Say that the system M is X -secure if for all $\alpha, \alpha' \in A^*$ and $u \in D$, if $X_u(\alpha) = X_u(\alpha')$ then $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$. Plainly, the notions of P-security, TO-security, ITO-security, TA-security and IP-security of a system M with respect to a policy \rightsquigarrow correspond to X -security for appropriate choices of the function X , viz., `purge`, `to`, `ito`, `ta` and `ipurge` as defined with respect to M and \rightsquigarrow .

The following result formalises the intuition that reducing the information in observations makes a system more secure.

Proposition 2 *If $M \leq^X M'$ and M' is X -secure then M is X -secure.*

Proof: Suppose $M \leq^X M'$ and M' is X -secure. Consider $\alpha, \alpha' \in A^*$ and $u \in D$ such that $X_u(\alpha) = X_u(\alpha')$. Then $\text{obs}'_u(s'_0 \cdot \alpha) = \text{obs}'_u(s'_0 \cdot \alpha')$ by X -security of M' . It follows using $M \leq^X M'$ that $\text{obs}_u(s_0 \cdot \alpha) = \text{obs}_u(s_0 \cdot \alpha')$. This is exactly what we need to show that M is X -secure. \square

Corollary 1 *For X any of P, TA, or IP, if $M \leq^X M'$ and M' is X -secure with respect to \rightsquigarrow then M is X -secure with respect to \rightsquigarrow .*

We note that we need to take some more care with the notions of ITO-security and TO-security, since the functions `ito` and `to` depend on the observations in the system to which they are applied, whereas Proposition 2 concerns the *same* function X in the systems M and M' . Indeed, it is *not* the case that if $M \leq M'$ and M' is TO-secure (ITO-secure), then M is TO-secure (ITO-secure).

Example 6: For this, consider the system M' in Figure 5, where there are domains H, D, L with policy $H \rightsquigarrow D \rightsquigarrow L$. Intuitively, H is a high security domain, D is a downgrader and the downgrader action d releases the information that the H action h has been performed. This system is TO-secure (hence ITO-secure also). For, suppose $\text{obs}_L(s_0 \cdot \alpha) = 1$ and $\text{obs}_L(s_0 \cdot \beta) = 0$. Then β does not contain an h followed later by a d , so must be of the form $\beta = d^k h^l$ with $k, l \geq 0$. Thus, $\text{to}_L(\beta) = \text{to}_L(d^k)$, which contains no D view containing observation 1. On the other hand, α contains an h and later d , so $\text{to}_L(\alpha)$ contains a D view containing observation 1. Thus, $\text{to}_L(\alpha) \neq \text{to}_L(\beta)$, and there




Figure 5: ITO-security and TO-security not preserved under system refinement

can be no violation of the condition for TO-security involving L . The argument for D is similar, and there is nothing to prove for L .

On the other hand, let M be the system obtained from M' by reducing the information observable to D , taking $\text{obs}_D(s) = \perp$ for all states s . Then in M we have

$$\begin{aligned}
 \text{ito}_L(hd) &= (\text{to}_L(h), \text{view}_D(hd), d) \\
 &= (\text{ito}_L(\epsilon), \perp d \perp, d) \\
 &= (\text{ito}_L(\epsilon), \text{view}_D(d), d) \\
 &= \text{ito}_L(d)
 \end{aligned}$$

but $\text{obs}_L(hd) = 1$ and $\text{obs}_L(d) = 1$. Hence M is not ITO-secure (and therefore also not TO-secure). \square

Intuitively, TO-security and ITO-security do more than place upper bounds on domains' information: they also add a type of lower bound constraint: information observable to a recipient must also have been observed by its sender. Thus, a further constraint on the systems M and M' is required in order for these properties to be preserved by the systems refinement. The following result identifies a sufficient condition for preservation of these properties.

For a sequence of actions α , and a domain u define $\text{res}_u(\alpha)$ to be $\alpha \upharpoonright \{u\}$, i.e. the subsequence of actions of domain u .

Proposition 3 *Suppose that $M \leq^{\text{res}} M'$ and for all sequences of actions α, α' and domains u , if $\text{view}_u^M(\alpha) = \text{view}_u^M(\beta)$ then $\text{view}_u^{M'}(\alpha) = \text{view}_u^{M'}(\beta)$. Then if M' is TO-secure (ITO-secure) with respect to \mapsto then M is TO-secure (ITO-secure) with respect to \mapsto .*

Proof: An easy induction using the property on the views shows that $\text{to}_u^M(\alpha) = \text{to}_u^M(\beta)$ implies $\text{to}_u^{M'}(\alpha) = \text{to}_u^{M'}(\beta)$. This yields the result. For, we get as an immediate consequence of this and the TO-security of M' that $\text{to}_u^M(\alpha) = \text{to}_u^M(\beta)$ implies $\text{obs}^{M'}(s_0 \cdot \alpha) = \text{obs}^{M'}(s_0 \cdot \beta)$. Also, by an easy induction, $\text{to}_u^M(\alpha) = \text{to}_u^M(\beta)$ implies $\text{res}_u(\alpha) = \text{res}_u(\beta)$. Thus, using $M \leq^{\text{res}} M'$, we conclude that $\text{to}_u^M(\alpha) = \text{to}_u^M(\beta)$ implies $\text{obs}^M(s_0 \cdot \alpha) = \text{obs}^M(s_0 \cdot \beta)$, as required for TO-security. The argument in the case of ITO-security is similar. \square

We remark that under the condition $M \leq^{\text{res}} M'$, we in fact have the determination of views in the other direction as well.

Proposition 4 *Suppose that $M \leq^{\text{res}} M'$. Then for all sequences of actions α, β , if $\text{view}_u^{M'}(\alpha) = \text{view}_u^{M'}(\beta)$ then $\text{view}_u^M(\alpha) = \text{view}_u^M(\beta)$.*

Proof: Since $M \leq^{\text{res}} M'$, there exists for each domain u a function $f_u : A_u^* \times O' \rightarrow O$ mapping sequences of actions of domain u and observations in M' to observations in M , such that for all sequences $\alpha \in A_u^*$, we have

$$f_u(\text{res}_u(\alpha), O^{M'}(s'_0 \cdot \alpha)) = O^M(s_0 \cdot \alpha).$$

For each domain u , define the mapping F_u from views in M' to views in M inductively, by

1. $F_u(o) = f_u(\epsilon, o)$, when o is an observation,
2. $F_u(\sigma a) = F_u(\sigma)a$ when $a \in A$, and
3. $F_u(\sigma o) = F_u(\sigma) \circ f_u(\sigma \upharpoonright A_u, o)$ when o is an observation and $\sigma \neq \epsilon$.

We claim that for all sequences $\alpha \in A^*$, we have $F_u(\text{view}_u^{M'}(\alpha)) = \text{view}_u^M(\alpha)$. The base case of $\alpha = \epsilon$ is immediate from the definitions. For the inductive case of a sequence αa where $a \in A$, consider three cases:

1. Case 1: $\text{dom}(a) \neq u$ and $O_u^{M'}(s'_0 \cdot \alpha) = O_u^{M'}(s'_0 \cdot \alpha a)$. Note that in this case we have $\text{res}_u(\alpha a) = \text{res}_u(\alpha)$, hence $O_u^M(s'_0 \cdot \alpha) = O_u^M(s'_0 \cdot \alpha a)$ by the fact that $M \leq^{\text{res}} M'$. Hence

$$\begin{aligned} F_u(\text{view}_u^{M'}(\alpha a)) &= F_u(\text{view}_u^{M'}(\alpha)) \\ &= \text{view}_u^M(\alpha) \quad (\text{by induction}) \\ &= \text{view}_u^M(\alpha a) \end{aligned}$$

2. Case 2: $\text{dom}(a) \neq u$ and $O_u^{M'}(s'_0 \cdot \alpha) \neq O_u^{M'}(s'_0 \cdot \alpha a)$. Note that in this case we have $\text{res}_u(\alpha a) = \text{res}_u(\alpha) = \text{view}_u^{M'}(\alpha) \upharpoonright A_u$. Hence

$$\begin{aligned} F_u(\text{view}_u^{M'}(\alpha a)) &= F_u(\text{view}_u^{M'}(\alpha) O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= F_u(\text{view}_u^{M'}(\alpha)) \circ f_u(\text{view}_u^{M'}(\alpha) \upharpoonright A_u, O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= F_u(\text{view}_u^{M'}(\alpha)) \circ f_u(\text{res}_u(\alpha a), O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= \text{view}_u^M(\alpha) \circ O_u^M(s_0 \cdot \alpha a) \\ &= \text{view}_u^M(\alpha a) \end{aligned}$$

3. Case 3: $\text{dom}(a) = u$. Then

$$\begin{aligned} F_u(\text{view}_u^{M'}(\alpha a)) &= F_u(\text{view}_u^{M'}(\alpha) a O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= F_u(\text{view}_u^{M'}(\alpha)) a f_u(\text{view}_u^{M'}(\alpha) a \upharpoonright A_u, O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= F_u(\text{view}_u^{M'}(\alpha)) a f_u(\text{res}_u(\alpha a), O_u^{M'}(s'_0 \cdot \alpha a)) \\ &= \text{view}_u^M(\alpha) a O_u^M(s_0 \cdot \alpha a) \\ &= \text{view}_u^M(\alpha a) \end{aligned}$$

□

Thus, under the conditions of Proposition 3, we have $\mathbf{view}_u^{M'}(\alpha) = \mathbf{view}_u^{M'}(\beta)$ iff $\mathbf{view}_u^M(\alpha) = \mathbf{view}_u^M(\beta)$. This may seem to make Proposition 3 uninterestingly weak. Below, we show that Proposition 3 is not as weak as it may seem, by giving a non-trivial application of Proposition 3 in which information is spread out through the views in M and M' rather different ways.

7 Action Observed Systems

The definitions of the previous sections are concerned with a machine model in which observations are made on states. A variant model has been considered in the literature, in which observations are associated to actions instead [Rus92]. A comparison between the two models for the set of semantics of intransitive information flow policies discussed above was carried out in [Mey07]. In this section we consider how the results on refinement in state-observed systems carry over to this model.

An *action-observed* machine is a tuple $\langle S, s_0, A, \mathbf{step}, \mathbf{out}, \mathbf{dom} \rangle$, where all the components are as in the state observed system model, except that the observation function \mathbf{obs} is replaced by a function $\mathbf{out} : S \times A \rightarrow O$. Intuitively, if s is a state and a is an action, then $\mathbf{out}(s, a)$ is the output, or return value, observed in domain $\mathbf{dom}(a)$ when action a is performed.

Each of the definitions of security for the state-observed system model has the form

M is secure with respect to \succrightarrow if for all sequences α, α' and domains u , if $X_u(\alpha) = X_u(\alpha')$ then $\mathbf{obs}_u(s_0 \cdot \alpha) = \mathbf{obs}_u(s_0 \cdot \alpha')$.

where $X_u(\alpha)$ is a function of $\alpha, u, \succrightarrow$ and M . We may obtain corresponding action-observed versions that have the form

M is secure with respect to \succrightarrow if for all sequences α, α' , domains u and actions a with $\mathbf{dom}(a) = u$, if $X_u(\alpha) = X_u(\alpha')$ then $\mathbf{out}(s_0 \cdot \alpha, a) = \mathbf{out}(s_0 \cdot \alpha', a)$.

In the cases of P-security, TA-security and IP-security, the corresponding functions \mathbf{purge} , \mathbf{ta} and \mathbf{ipurge} in fact depend only on u, α and \succrightarrow , and we may use exactly the same functions as X to obtain the action-observed definitions of security.

In the case of TO-security and ITO-security, there is also a dependence on observations in the model, which become outputs in the action-observed case. Here we need to reformulate the definitions somewhat. This is done as follows. First, the notion of view is adapted to the action-observed system model by defining $\mathbf{view}_u^a : A^* \rightarrow (A \cup O)^*$ for $u \in D$ inductively by $\mathbf{view}_u^a(\epsilon) = \epsilon$, and

$$\mathbf{view}_u^a(\alpha a) = \begin{cases} \mathbf{view}_u^a(\alpha) a \mathbf{out}(s_0 \cdot \alpha, a) & \text{if } \mathbf{dom}(a) = u \\ \mathbf{view}_u^a(\alpha) & \text{otherwise.} \end{cases}$$

That is, the view of a domain is just the sequence of actions that the domain has performed, together with the outputs obtained from those actions.

We now define an action-observed variant \mathbf{to}_u^a of \mathbf{to}_u , by $\mathbf{to}_u^a(\epsilon) = \epsilon$ and

$$\mathbf{to}_u^a(\alpha a) = \begin{cases} \mathbf{to}_u^a(\alpha) & \text{if } \mathbf{dom}(a) \not\rightarrow u, \\ (\mathbf{to}_u^a(\alpha), \mathbf{view}_{\mathbf{dom}(a)}^a(\alpha a), a) & \text{if } \mathbf{dom}(a) = u, \\ (\mathbf{to}_u^a(\alpha), \mathbf{view}_{\mathbf{dom}(a)}^a(\alpha), a) & \text{if } u \neq \mathbf{dom}(a) \rightarrow u. \end{cases}$$

Similarly, \mathbf{ito}_u^a is defined by $\mathbf{ito}_u^a(\epsilon) = \epsilon$ and

$$\mathbf{ito}_u^a(\alpha a) = \begin{cases} \mathbf{ito}_u^a(\alpha) & \text{if } \mathbf{dom}(a) \not\rightarrow u, \\ (\mathbf{ito}_u^a(\alpha), \mathbf{view}_{\mathbf{dom}(a)}^a(\alpha a), a) & \text{otherwise.} \end{cases}$$

Taking these functions as X in the above pattern for action-observed security yields the definitions of TO-security and ITO-security in the action-observed case. (The reader may note some subtle differences between these definitions in the state- and action-observed cases. We refer to [Mey07] for an explanation of these differences.)

These definitions of security on action-observed systems may be shown to be related to the similarly named definitions on state-observed systems, by means of mapping from the action-observed to the state-observed domain. Given an action-observed machine $M = \langle S, s_0, A, \mathbf{step}, \mathbf{out}, \mathbf{dom} \rangle$ with domains D and outputs O , define the state observed machine $F_{as}(M) = \langle S', s'_0, A, \mathbf{step}', \mathbf{obs}, \mathbf{dom} \rangle$ by

1. $S' = S \times (D \rightarrow O \cup \{\perp\})$,
2. $s'_0 = (s_0, f_0)$, where $f_0(d) = \perp$ for all $d \in D$,
3. $\mathbf{step}'((s, f), a) = (\mathbf{step}(s, a), f[\mathbf{dom}(a) \rightarrow \mathbf{out}(s, a)])$,
4. $\mathbf{obs}_u((s, f)) = f(u)$.

Here, we write $f[u \mapsto x]$ for the function f' that is identical to f except that $f'(u) = x$. Intuitively, in a state (s, f) , the value $f(d)$ for a domain d represents the observation most recently obtained in domain d , and is \perp if there has been no observation in domain d .

The following result states relationships between the definitions of security on the two types of model under this mapping.

Theorem 8 [Mey07] *Let X be any of P , TO , ITO , TA , or IP . Then an action-observed machine M is X -secure with respect to a policy \rightarrow (using the action-observed definitions) iff $F_{as}(M)$ is X -secure with respect to \rightarrow (using the state-observed definitions).*

We will use this result to derive a result on the soundness of architectural refinement that similar to Theorem 4, but for action-observed systems.

First, we define the result of applying a refinement mapping to an action-observed system. Let $M = (S, S_0, A, \mathbf{step}, \mathbf{out}, \mathbf{dom})$ be an action-observed system with set of domains D_1 . Let $r : D_1 \rightarrow D_2$. Then we define the system $r(M) = (S, S_0, A, \mathbf{step}, \mathbf{out}, \mathbf{dom}')$ simply by defining $\mathbf{dom}'(a) = r(\mathbf{dom}(a))$ for all $a \in A$; all other components are exactly as in M .

The following result characterizes how this operation relates to the mapping from action-observed to state-observed systems.

Lemma 3 *Let $r : D_1 \rightarrow D_2$, and let M be an action-observed system with set of domains D_1 . Then $F_{as}(r(M)) \leq^{\mathbf{res}} r(F_{as}(M))$, and for all sequences $\alpha, \beta \in A^*$, and domains $u \in D_2$ we have that $\mathbf{view}_u^{F_{as}(r(M))}(\alpha) = \mathbf{view}_u^{F_{as}(r(M))}(\beta)$ implies $\mathbf{view}_u^{r(F_{as}(M))}(\alpha) = \mathbf{view}_u^{r(F_{as}(M))}(\beta)$.*

Proof: We first characterize the observations in the systems $F_{as}(r(M))$ and $r(F_{as}(M))$. Let $\mathbf{dom} : A \rightarrow D_1$ be the domain function of M .

Note that in $r(M)$, the output of an action a is observed by domain $r(\mathbf{dom}(a))$. Since the observation made by domain u in the state reached after a sequence α in $F_{as}(r(M))$ is the output obtained by the latest action of domain u , this observation is the output obtained from the latest action a in α with $\mathbf{dom}(a) \in r^{-1}(u)$.

On the other hand, in $F_{as}(M)$ the observation of domain v is the output obtained from the latest v action, so in $r(F_{as}(M))$, the observation of domain u in the state reached after a sequence α is the mapping taking $v \in r^{-1}(u)$ to the output of the latest action a in α with $\mathbf{dom}(a) = v$.

To see that $F_{as}(r(M)) \leq^{\mathbf{res}} r(F_{as}(M))$, suppose that $\mathbf{res}_u(\alpha) = \mathbf{res}_u(\beta)$ and $\mathbf{obs}_u^{r(F_{as}(M))}(s_0 \cdot \alpha) = \mathbf{obs}_u^{r(F_{as}(M))}(s_0 \cdot \beta)$. Then the most recent action of domain v with $r(v) = u$ is the same in α and β . Thus,

$$\begin{aligned} \mathbf{obs}_u^{F_{as}(r(M))}(s_0 \cdot \alpha) &= \mathbf{obs}_u^{r(F_{as}(M))}(s_0 \cdot \alpha)(v) \\ &= \mathbf{obs}_u^{r(F_{as}(M))}(s_0 \cdot \beta)(v) \\ &= \mathbf{obs}_u^{F_{as}(r(M))}(\beta), \end{aligned}$$

as required.

To see that $\mathbf{view}_u^{F_{as}(r(M))}(\alpha) = \mathbf{view}_u^{F_{as}(r(M))}(\beta)$ implies $\mathbf{view}_u^{r(F_{as}(M))}(\alpha) = \mathbf{view}_u^{r(F_{as}(M))}(\beta)$, we define for each $u \in D_2$ a mapping G_u taking views of $r(F_{as}(M))$ to views of $F_{as}(r(M))$, by the following induction. Note that, by construction, the views of $r(F_{as}(M))$ and $F_{as}(r(M))$ have the property that they alternate actions and observations, i.e., there are no adjacent observations. (This is because actions a with $r(\mathbf{dom}(a)) \neq u$ do not produce output in any domains observable to u , hence u observations are invariant under such actions, and repeated observations are removed by the absorptive concatenation operator.) The base case is given by $G_u(\perp) = \perp^{r^{-1}(u)}$, where we write \perp^S for the function \perp with domain S and constant value \perp . For σ a view, $a \in A$ and o an observation in M , the inductive case is given by $G_u(\sigma a o) = G_u(\sigma) a O'$, where $O' = O[\mathbf{dom}(a) \mapsto o]$ and O is the final observation in $G_u(\sigma)$, which is a mapping from $r^{-1}(u)$ to observations in M .

We claim that $G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha)) = \mathbf{view}_u^{r(F_{as}(M))}(\alpha)$ for all $\alpha \in A^*$, from which it follows that $\mathbf{view}_u^{F_{as}(r(M))}(\alpha) = \mathbf{view}_u^{F_{as}(r(M))}(\beta)$ implies that $\mathbf{view}_u^{r(F_{as}(M))}(\alpha) = \mathbf{view}_u^{r(F_{as}(M))}(\beta)$. The proof of the claim is by induction. The base case is trivial. For the inductive step, we consider two cases. First, suppose that $\text{dom}(a) \neq u$. Then

$$\begin{aligned} G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha a)) &= G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha)) \\ &= \mathbf{view}_u^{r(F_{as}(M))}(\alpha) && \text{(by induction)} \\ &= \mathbf{view}_u^{r(F_{as}(M))}(\alpha a). \end{aligned}$$

Alternately, if $\text{dom}(a) = u$, then

$$G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha a)) = G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha) a o)$$

where $o = \text{out}(s_0 \cdot \alpha, a)$ (in M). Let O be the final observation in $G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha))$, and $O' = O[\text{dom}(a) \mapsto o]$. Then, by induction,

$$\begin{aligned} G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha) a o) &= G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha)) a O' \\ &= \mathbf{view}_u^{F_{as}(r(M))}(\alpha) a O'. \end{aligned}$$

In particular, since O is also the final observation in $\mathbf{view}_u^{F_{as}(r(M))}(\alpha)$, we have that $O = \text{obs}_u^{F_{as}(r(M))}(s_0 \cdot \alpha)$, and, by construction, $O' = \text{obs}_u^{F_{as}(r(M))}(s_0 \cdot \alpha a)$. We conclude that $G_u(\mathbf{view}_u^{F_{as}(r(M))}(\alpha a)) = \mathbf{view}_u^{r(F_{as}(M))}(\alpha a)$, as required. \square

We now obtain the following:

Corollary 2 *Let M be an action-observed system with domains D_1 , let $r : D_1 \rightarrow D_2$, and let \mapsto be a policy on D_2 . For X any of P , TO , ITO , TA or IP , if $r(F_{as}(M))$ is X -secure with respect to \mapsto then $F_{as}(r(M))$ is X -secure with respect to \mapsto .*

Proof: It is easily checked that for X equal to any of **purge**, **ta** or **ipurge**, if $X_u(\alpha) = X_u(\beta)$ then $\text{res}_u(\alpha) = \text{res}_u(\beta)$. Hence from $F_{as}(r(M)) \leq^{\text{res}} r(F_{as}(M))$ we obtain that $F_{as}(r(M)) \leq^X r(F_{as}(M))$ for these values of X . Hence by Lemma 3 and Corollary 1, we obtain that if $r(F_{as}(M))$ is X -secure with respect to \mapsto then $F_{as}(r(M))$ is X -secure with respect to \mapsto .

For the case of X equal to **to** or **ito**, we obtain this conclusion from Proposition 3 and Lemma 3. \square

We may now conclude that each of our definitions of security is preserved under refinement of action-observed systems.

Corollary 3 *Let $\mathcal{A}_1 \leq_r \mathcal{A}_2$, and let X any of P , TO , ITO , TA or IP . If M is an action-observed system and M is X -compliant with \mathcal{A}_1 then $r(M)$ is X -compliant with \mathcal{A}_2 .*

Proof: We have the following chain of implications:

1. M is X -secure with respect to \rightsquigarrow_1
2. implies $F_{as}(M)$ is X -secure with respect to \rightsquigarrow_1 (by Theorem 4)
3. implies $r(F_{as}(M))$ is X -secure with respect to \rightsquigarrow_2 (by Theorem 8)
4. implies $F_{as}(r(M))$ is X -secure with respect to \rightsquigarrow_2 (by Corollary 2)
5. implies $r(M)$ is X -secure with respect to \rightsquigarrow_2 (by Theorem 4)

□

8 Related Work

To the best of our knowledge, our work is the first consideration of the relationship between architectural refinement and intransitive noninterference. However, both formal theories of architecture refinement and refinement of noninterference security properties have been presented in the past.

In general, work on architectural refinement [PR97, Bar05] is concerned with behavioural notions of refinement, and has not taken security into account. An interesting exception is a sequence of papers by Moriconi *et al.*, [MQR95, MQ94], who develop a very abstract formal account of architecture refinement in which architectural designs are represented as logical theories and refinement is treated as a mapping of the symbols of the abstract theory to those of the concrete theory that must satisfy the logical condition of being a *faithful interpretation*. In order to apply this account to a particular type of architectural design notation, it is necessary to concretize the abstract theory by giving both a syntax for the architectural elements in the notation, and to develop a logical theory that represents the semantics of this notation. (E.g. this is done in [MQ94] for dataflow and shared-memory architectural styles.) In [MQRG97], the framework is applied to establish security properties of a number of secure variants of the X/Open Distributed Transaction Processing architecture. The security policy considered here is the Bell La-Padula policy [BP76], which lacks the kind of information flow semantics that we have studied here, although it can be related to noninterference for transitive policies [Rus92]. It is not clear whether a concretization of the Moriconi *et al.* theory could be developed that would enable it to represent the content of our results, but this would be an interesting topic for further study. Zhou and Alves-Foss [ZAF06] have also proposed a number of architecture refinement patterns for Multi-Level Secure systems development, but do not provide any formal semantics for their work.

The other area of related work, dealing with preservation of noninterference properties under notions of refinement, has typically been concerned with just the (transitive) two-domain policy stating that High level information may not flow to Low, rather than with the more general intransitive policies that we have

considered in this paper. (All of the literature discussed below differs from our work in this regard.)

That we have obtained positive results concerning refinement of security properties may be surprising to those familiar with the literature on formal security properties, where it is folklore that such properties are *not* preserved under refinement [Jac89], a fact known as the “refinement paradox”. However, our notions of refinement differ from the notions of refinement usually studied. Refinement is usually understood as a reduction in the set of possible behaviours of the system, which would be contrary to our assumptions that systems are *input-enabled* (actions always enabled) and *deterministic*.

It is possible to identify conditions under which reduction of the possible behaviours of a system preserves information flow security properties. Jacob [Jac89] presents a method in which an insecure system is first developed using a standard refinement methodology for functional properties, then made secure by a further deletion of behaviors in a fixpoint calculation. It is not guaranteed that this last step terminates, nor that it produces a useful system. Mantel [Man01] defines refinement operators that take as input a secure system, a set of transitions to be disabled, and a type of unwinding relation on the system that establishes the security property. The operators produce as output a refined system, as well as a new unwinding relation that establishes the security of the refined system. This is achieved by either disabling transitions other than those requested, or by maintaining some transitions whose disablement was requested. He considers a richer notion of information flow policy than we have treated, but with respect to a semantics that seems appropriate only for transitive policies. The practicality of these approaches has not been established.

A number of authors have also identified sufficient conditions under which data-refinement preserves transitive information flow policies [GCS91, O’H92]. Bossi *et al.* [BFPR03] develop conditions under which refinement of process algebra terms preserves bisimulation-based information flow security properties using a simulation-based notion of refinement. Roscoe [Ros95] defines *Low-determinism*, a very strong notion of security, which is always preserved under refinement, but at the cost of a significantly restricted range of applicability. Some recent works have also sought to overcome the refinement paradox by drawing a distinction between specification-level non-determinism and non-determinism that is inherent in a system, with the latter preserved under refinement [SS06, Jür05, Bib06].

We note that the notions of refinement we have considered in this paper are somewhat different from the standard notion of refinement considered in the literature, which reduces the set of possible behaviours of a system. Information flow properties are not generally preserved under such refinements [Jac89], although some authors have tried to obtain such results by modification of the refined system [Jac89, Man01]. Others have identified sufficient conditions for behavioral refinement to preserve information flow properties [GCS91, O’H92, BFPR03, Ros95]. Some recent works have also sought to overcome the refinement paradox by drawing a distinction between specification-level non-determinism and non-determinism that is inherent in a system, with

the latter preserved under refinement [SS06, Jür05, Bib06].

Another area in which refinement has been considered in the context of information flow security concerns refinement at the level of extended program notations [Mor09]. This approach aims to preserve the ignorance of a given agent during refinement, and has the advantage of providing an expressive framework for representing what an agent does not know. (We remark that [Mor09] contains a result related to the concerns of Section 6, showing that security is preserved by reducing the content of observations.) On the other hand, this line of work does not deal with the concerns of causal structure in a multi-agent setting that are our focus in the present paper.

9 Conclusion

This work is a contribution towards a formal design theory for information flow secure systems. Much remains to be done to realise such a theory. The notions of security studied here are based on an asynchronous modelling of systems - they do not take into account issues such as timing attacks and scheduling. Probabilistic reasoning, which is critical in practical security settings, is also ignored. Suitable variants of our definitions for the semantics of intransitive noninterference policies that take these concerns into account remain to be developed. We have considered only refinement of systems as a whole; it would be desirable to develop also an account of composition of architectural designs and policies, and to study how these interact with refinement, so that refinement can be carried out at the component level. Integrating our approach with approaches to refinement operating at lower levels of system description would also be desirable. We hope to address issues such as these in future work. Ultimately, one would like to have a theoretically sound and tool supported methodology that enables a system to be developed from the very abstract architectural level we have considered in this paper, all the way through to code running on particular hardware configurations.

References

- [AFHOT06] J. Alves-Foss, W.S. Harrison, P. Oman, and C. Taylor. The MILS architecture for high-assurance embedded systems. *International Journal of Embedded Systems*, 2(3/4):239–47, Feb 2006.
- [Bar05] M. A. Barbosa. A refinement calculus for software components and architectures. *ACM SIGSOFT Software Engineering Notes*, 30(5), September 2005.
- [BFPR03] A. Bossi, R. Focardi, C. Piazza, and S. Rossi. Refinement operators and information flow security. In *Proc. Int. Conf. on Software Engineering and Formal Methods*, pages 44–53, 2003.

- [Bib06] David Bibighaus. *Applying the doubly labeled transition system to the refinement paradox*. PhD thesis, Naval Postgraduate School, Monterey, 2006.
- [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.
- [GCS91] J. Graham-Cunning and J. Sanders. On the refinement of noninterference. In *Proc. IEEE Computer Security Foundations Workshop*, pages 35–42, 1991.
- [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.
- [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.
- [Jac89] J. Jacob. On the derivation of secure components. In *Proc. IEEE Symp. on Security and Privacy*, pages 242–247, 1989.
- [Jür05] J. Jürjens. *Secure Systems Development with UML*. Springer, 2005.
- [Man01] H. Mantel. Preserving information flow properties under refinement. In *Proc. IEEE Symp. Security and Privacy*, pages 78–91, 2001.
- [Mey07] R. van der Meyden. A comparison of semantic models of intransitive noninterference. submitted for publication, copy at <http://www.cse.unsw.edu.au/~meyden>, Dec 2007.
- [Mey08] R. van der Meyden. What, indeed, is intransitive noninterference? (submitted for publication, copy at <http://www.cse.unsw.edu.au/~meyden> – an extended abstract of this paper appears in Proc. ESORICS 2007), Jan 2008.
- [Mor09] Carroll Morgan. The shadow knows: Refinement and security in sequential programs. *Sci. Comput. Program.*, 74(8):629–653, 2009.
- [MQ94] M. Moriconi and X. Qian. Correctness and composition of software architectures. In *Proc. 2nd ACM SIGSOFT Symposium on Foundations of Software Engineering*, pages 164–174, 1994.

- [MQR95] M. Moriconi, X. Qian, and R.A. Riemenschneider. Correct architecture refinement. *IEEE Transactions on Software Engineering*, 21(4):356–372, April 1995.
- [MQRG97] M. Moriconi, X. Qian, R. A. Riemenschneider, and L. Gong. Secure software architectures. In *Proc. IEEE Symp. on Security and Privacy*, pages 84–893, 1997.
- [O’H92] C. O’Halloran. Refinement and confidentiality. In *Fifth Refinement Workshop*, pages 119–139. British Computer Society, 1992.
- [PR97] J. Philipps and B. Rumpe. Refinement of information flow architectures. In *Proc. 1st IEEE Int. Conf. on Formal Engineering Methods*, pages 203 – 212, 1997.
- [RG99] A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *IEEE Computer Security Foundations Workshop*, pages 228–238, 1999.
- [Ros95] A.W. Roscoe. CSP and determinism in security modelling. In *Proc. IEEE Symp. on Security and Privacy*, pages 114–221, 1995.
- [RR83] J.M. Rushby and R. Randell. A distributed secure system. *IEEE Computer*, 16(7):55–67, 1983.
- [Rus92] J. Rushby. Noninterference, transitivity, and channel-control security policies. Technical Report CSL-92-02, SRI International, Dec 1992.
- [SS06] F. Seehusen and K. Stolen. Information flow property preserving transformation of UML interaction diagrams. In *Proc. ACM symposium on access control models and technologies*, pages 150 – 159, 2006.
- [VBC⁺05] W.M. Vanfleet, R.W. Beckworth, B. Calloni, J.A. Luke, C. Taylor, and G. Uchenick. MILS:architecture for high assurance embedded computing. *Crosstalk: The Journal of Defence Engineering*, pages 12–16, Aug 2005.
- [ZAF06] J. Zhou and J. Alves-Foss. Architecture-based refinements for secure computer system design. In *Proc. Policy, Security and Trust*, Nov 2006.