

Using Architecture to Reason about Information Security

Ron van der Meyden

**(University of New South Wales
Sydney, Australia)**

**Joint work with
Stephen Chong
(Harvard)**

December 3, 2012


- ▶ Motivation: MILS Security
- ▶ Recap of Intransitive noninterference theory
- ▶ Extended theory for architectural specifications
- ▶ Using architecture to reason about information flow properties
- ▶ Connections to Access Control

Introduction: Rushby view of MILS


A two level design process comprised of

- ▶ **Policy Level:** an architectural design identifying components and their connections/permitted causal relationships.
- ▶ **Resource Sharing Level:** components implemented so as to share resources (processors, memory, network) with enforcement of architectural causality constraints using a variety of mechanisms (e.g., separation kernels, periods processing, crypto)

Policy Level:



Resource Sharing Level:



Further Objectives for MILS

- ▶ isolation of safety/security critical functionality in (small, formally verifiable) *trusted components*
- ▶ (formal) compositional derivation of global properties from architecture + local properties of the trusted components
- ▶ These global properties preserved by the resource sharing implementation

Overview of this talk

Questions concerning this vision: (& our answers, in this talk)

- ▶ What is the formal syntax and semantics of architectural designs?
 - ▶ Abstract syntax based on an extension of intransitive noninterference policies
 - ▶ A new semantics based on a knowledge-based approach to intransitive noninterference of van der Meyden (ESORICS 2007)
- ▶ Is it really possible to prove interesting security properties in this architecture + trusted component, local to global, pattern?
- ▶ How does the theory ground out in concrete resource sharing mechanisms?

Overview of this talk

Questions concerning this vision: (& our answers, in this talk)

- ▶ What is the formal syntax and semantics of architectural designs?
 - ▶ Abstract syntax based on an extension of intransitive noninterference policies
 - ▶ A new semantics based on a knowledge-based approach to intransitive noninterference of van der Meyden (ESORICS 2007)
- ▶ Is it really possible to prove interesting security properties in this architecture + trusted component, local to global, pattern?
 - ▶ Examples indicating the answer is 'Yes'.
- ▶ How does the theory ground out in concrete resource sharing mechanisms?

Overview of this talk

Questions concerning this vision: (& our answers, in this talk)

- ▶ What is the formal syntax and semantics of architectural designs?
 - ▶ Abstract syntax based on an extension of intransitive noninterference policies
 - ▶ A new semantics based on a knowledge-based approach to intransitive noninterference of van der Meyden (ESORICS 2007)
- ▶ Is it really possible to prove interesting security properties in this architecture + trusted component, local to global, pattern?
 - ▶ Examples indicating the answer is 'Yes'.
- ▶ How does the theory ground out in concrete resource sharing mechanisms?
 - ▶ Sufficient condition for architectural compliance in access control systems.

Noninterference policies (Goguen and Meseguer 1982)

Let D be a set of security domains/components/agents.

A noninterference policy is a reflexive relation $\succrightarrow \subseteq D \times D$

$u \succrightarrow v$ means


“actions of u are permitted to interfere with v ”, or

“actions of u are permitted to have effects observable to v ”, or

“information is permitted to flow from u to v ”

Hinke-Schaeffer Architecture

One of the proposed architectures for multi-level secure databases, as an intransitive noninterference policy:



(Strictly speaking, Hinke-Schaeffer = this policy level architecture, enforced at resource sharing level by the operating system.)

Deterministic System Model

Machines have the form $M = \langle S, s_0, Actions, D, dom, step, O \rangle$
where

- ▶ S is a set of states,

Deterministic System Model

Machines have the form $M = \langle S, s_0, \text{Actions}, D, \text{dom}, \text{step}, O \rangle$
where

- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,

Deterministic System Model

Machines have the form $M = \langle S, s_0, \text{Actions}, D, \text{dom}, \text{step}, O \rangle$
where

- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,
- ▶ Actions is a set of actions,

Deterministic System Model

Machines have the form $M = \langle S, s_0, \text{Actions}, D, \text{dom}, \text{step}, O \rangle$
where

- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,
- ▶ *Actions* is a set of actions,
- ▶ D is a set of *domains*

Deterministic System Model

Machines have the form $M = \langle S, s_0, Actions, D, \text{dom}, \text{step}, O \rangle$
where

- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,
- ▶ $Actions$ is a set of actions,
- ▶ D is a set of *domains*
- ▶ $\text{dom} : Actions \rightarrow D$ associates each action to a domain in D ,

Deterministic System Model

Machines have the form $M = \langle S, s_0, Actions, D, \text{dom}, \text{step}, O \rangle$
where



- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,
- ▶ $Actions$ is a set of actions,
- ▶ D is a set of *domains*
- ▶ $\text{dom} : Actions \rightarrow D$ associates each action to a domain in D ,
- ▶ $\text{step} : S \times Actions \rightarrow S$ is a deterministic transition function, and

Deterministic System Model

Machines have the form $M = \langle S, s_0, Actions, D, \text{dom}, \text{step}, O \rangle$
where

- ▶ S is a set of states,
- ▶ $s_0 \in S$ is the *initial state*,
- ▶ $Actions$ is a set of actions,
- ▶ D is a set of *domains*
- ▶ $\text{dom} : Actions \rightarrow D$ associates each action to a domain in D ,
- ▶ $\text{step} : S \times Actions \rightarrow S$ is a deterministic transition function, and
- ▶ $O_u : S \rightarrow Obs$ — the observation of/output to domain $u \in D$ at a state

Haigh and Young's Intransitive Purge Function



ipurge E

a b c d

Haigh and Young's definition: IP-security

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

Alternate definition: van der Meyden - ESORICS'07

Given a policy \succrightarrow , define, for each domain $u \in D$, the function ta_u , with domain Actions^* , inductively by $\text{ta}_u(\epsilon) = \epsilon$, and, for $\alpha \in \text{Actions}^*$ and $a \in \text{Actions}$,

$$\text{ta}_u(\alpha a) = \begin{cases} \text{ta}_u(\alpha) & \text{if } \text{dom}(a) \not\succrightarrow u \\ (\text{ta}_u(\alpha), \text{ta}_{\text{dom}(a)}(\alpha), a) & \text{if } \text{dom}(a) \succrightarrow u \end{cases}$$

Alternate definition: van der Meyden - ESORICS'07

Given a policy \rightsquigarrow , define, for each domain $u \in D$, the function ta_u , with domain $Actions^*$, inductively by $\text{ta}_u(\epsilon) = \epsilon$, and, for $\alpha \in Actions^*$ and $a \in Actions$,

$$\text{ta}_u(\alpha a) = \begin{cases} \text{ta}_u(\alpha) & \text{if } \text{dom}(a) \not\rightsquigarrow u \\ (\text{ta}_u(\alpha), \text{ta}_{\text{dom}(a)}(\alpha), a) & \text{if } \text{dom}(a) \rightsquigarrow u \end{cases}$$

Define a system M to be **TA-secure** with respect to a policy \rightsquigarrow if for all domains $u \in D$, and all $\alpha, \alpha' \in Actions^*$ such that $\text{ta}_u(\alpha) = \text{ta}_u(\alpha')$, we have $O_u(s_0 \cdot \alpha) = O_u(s_0 \cdot \alpha')$.

Reasons to believe this definition is better

Results from ESORICS-07:

- ▶ It does not admit a disturbing example from ESORICS-07
- ▶ TA-security \Rightarrow IP-security
- ▶ TA-security \equiv IP-security for transitive policies
- ▶ Rushby unwinding conditions for IP-security \Rightarrow TA-security
- ▶ A system bisimilar to M satisfies Rushby unwinding conditions $\Rightarrow M$ is TA-secure
- ▶ (A similar equivalence for a variant of Rushby's access control results.)

AND: It leads to the generalization of the present paper ...

Reasoning about Knowledge

Given a system M , define the view of domain u of a sequence $\alpha \in \text{Actions}^*$ to be the sequence $\text{view}_u(\alpha)$ of all actions and observations of that domain, with stuttering of observations eliminated (to model asynchrony).

E.g. if $\alpha = hhh$ generates (Low observations only):

$O_1 h O_1 h O_1 I O_2 h O_2$

then $\text{view}_{\text{Low}}(\alpha) = O_1 I O_2$

Propositions

A **proposition** ϕ is a fact about sequences of actions.

Formally $\phi \subseteq Actions^*$, and we say ϕ holds at $\alpha \in Actions^*$ if $\alpha \in \phi$.

ϕ is **non-trivial** if it is not \emptyset or $Actions^*$.


ϕ is **G-local**, for $G \subseteq D$, if it depends only on actions of the domains G

Formally, $\alpha|G = \beta|G$ implies $\alpha \in \phi$ iff $\beta \in \phi$.

Reasoning about Knowledge

Say domain u **knows** a proposition ϕ after a sequence $\alpha \in \text{Actions}^*$ in a system M if ϕ holds at β for all sequences $\beta \in \text{Actions}^*$ such that $\text{view}_u(\alpha) = \text{view}_u(\beta)$.

Notation: $M, \alpha \models \text{Knows}_u(\phi)$



Theorem: Suppose that M is TA-secure with respect to the Hinke-Schaeffer policy. Then for all $\{H_{user}, H_{DBMS}, H_F\}$ -local propositions ϕ , and $\alpha \in Actions^*$

$$M, \alpha \models \neg \text{Knows}_{L_{user}}(\phi)$$

Extended Architectures


(To model restrictions on the behaviour of trusted components)

An **extended architecture** is a pair (D, \rightsquigarrow) where $\rightsquigarrow \subseteq D \times D \times \mathcal{L}$, where

- ▶ D is a set of domains
- ▶ \mathcal{L} is a set of function names, including the special name \top
- ▶ $(u, v, f) \in \rightsquigarrow$ means “information is permitted to flow from u to v , but must be *filtered* through the function denoted by f ”.
- ▶ \top means “no constraints on information flow across this edge”

Example: Starlight Interactive Link

(Anderson et al. – A switch that allows a user to alternate their keyboard between High and Low level windows.)



Architectural Interpretations

An *interpretation* of an extended architecture consists of

- ▶ A set of actions *Actions*
- ▶ A domain assignment $\text{dom} : \text{Actions} \rightarrow D$
- ▶ An interpretation function I , such that for each $f \in \mathcal{L} \setminus \{\top\}$, $I(f)$ is a function with domain Actions^*

Intuitively, if $(u, v, f) \in \rightsquigarrow$ and $\alpha \in \text{Actions}^*$ and $a \in \text{Actions}$ with $\text{dom}(a) = u$, then $I(f)(\alpha a)$ is “the information permitted to flow from u to v when u does a after α .”

Given an extended architecture $\mathcal{A} = (D, \succrightarrow)$ and an architectural interpretation $\mathcal{I} = (Actions, \text{dom}, I)$, we can define for each $u \in D$ the function tff_u with domain $Actions^*$ by $\text{tff}_u(\epsilon) = \epsilon$ and

$$\text{tff}_u(\alpha a) = \begin{cases} \text{tff}_u(\alpha) & \text{if } \text{dom}(a) \not\succrightarrow u \\ \text{tff}_u(\alpha) (\text{tff}_{\text{dom}(a)}(\alpha), a) & \text{if } \text{dom}(a) \xrightarrow{\top} u \\ \text{tff}_u(\alpha) I(f)(\alpha a) & \text{if } \text{dom}(a) \xrightarrow{f} u \end{cases}$$

Given an extended architecture $\mathcal{A} = (D, \succrightarrow)$ and an architectural interpretation $\mathcal{I} = (\text{Actions}, \text{dom}, \text{I})$, we can define for each $u \in D$ the function tff_u with domain Actions^* by $\text{tff}_u(\epsilon) = \epsilon$ and

$$\text{tff}_u(\alpha a) = \begin{cases} \text{tff}_u(\alpha) & \text{if } \text{dom}(a) \not\succrightarrow u \\ \text{tff}_u(\alpha) (\text{tff}_{\text{dom}(a)}(\alpha), a) & \text{if } \text{dom}(a) \xrightarrow{\top} u \\ \text{tff}_u(\alpha) \text{I}(f)(\alpha a) & \text{if } \text{dom}(a) \xrightarrow{f} u \end{cases}$$

M **complies** with the interpreted architecture $(\mathcal{A}, \mathcal{I})$ if for all $u \in D$ and $\alpha, \beta \in \text{Actions}^*$, if $\text{tff}_u(\alpha) = \text{tff}_u(\beta)$ then $O_u(s_0 \cdot \alpha) = O_u(s_0 \cdot \beta)$.

Architectural Specifications


An architectural specification consists of

- ▶ An extended architecture \mathcal{A}
- ▶ A set \mathcal{C} of interpretations for this architecture

This captures architecture + constraints on the behaviour of trusted components.

M complies with an architectural specification $(\mathcal{A}, \mathcal{C})$ if it complies with $(\mathcal{A}, \mathcal{I})$ for some $\mathcal{I} \in \mathcal{C}$.

Example: Starlight Interactive Link



+ \mathcal{C} = all interpretations such that

- ▶ *Actions* contains a toggle action t with $\text{dom}(t) = S$
- ▶ $I(sf)(\alpha a) = a$ if $a = t$ or $\text{dom}(a) = S$ and α contains an odd number of t 's,
otherwise $I(sf)(\alpha a) = \epsilon$ (no information flow)

Reasoning based on Architectural Specifications


Say that a proposition ϕ is *toggle-High dependent* if it depends only on the subsequence of α consisting of

- ▶ all actions a with $\text{dom}(a) = H$
- ▶ all occurrences of actions a with $\text{dom}(a) = S$ that occur between an even numbered occurrence of t and any subsequent occurrence of t .

Theorem: Suppose that M complies with the Starlight architectural specification. Let ϕ be toggle-High dependent and non-trivial. Then

$$M, \alpha \models \neg \text{Knows}_L(\phi)$$


Electronic Election



Specification:

- ▶ *Actions* consists of actions of ElecAuth plus actions a_v for v a voter and $a \in \text{VoterActions}$; with $\text{dom}(a_v) = v$.
- ▶ For a permutation P of $\{v_1, \dots, v_n\}$ and $\alpha \in \text{Actions}^*$, let $P(\alpha)$ be the result of replacing each a_v in α by $a_{P(v)}$.
- ▶ For all α , $\text{results}(\alpha) = \text{results}(P(\alpha))$.

Electronic Election



Specification:

- ▶ *Actions* consists of actions of ElecAuth plus actions a_v for v a voter and $a \in \text{VoterActions}$; with $\text{dom}(a_v) = v$.
- ▶ For a permutation P of $\{v_1, \dots, v_n\}$ and $\alpha \in \text{Actions}^*$, let $P(\alpha)$ be the result of replacing each a_v in α by $a_{P(v)}$.
- ▶ For all α , $\text{results}(\alpha) = \text{results}(P(\alpha))$. (Examples: $\text{results}(\alpha)$ = number of votes cast for each candidate
 $\text{results}(\alpha)$ = the candidate(s) with the most votes.)

Theorem: If M complies with this specification, P is a permutation of voters and v a particular voter such that $P(v) = v$ and ϕ is a proposition, if

$$M, \alpha \models \neg \text{Knows}_v(\neg \phi)$$

then

$$M, \alpha \models \neg \text{Knows}_v(\neg P(\phi))$$

Example: “if Alice considers it possible that

Bob voted for Obama and Charlie voted for McCain,
then Alice considers it possible that

Charlie voted for Obama and Bob voted for McCain.”

Access Control

Following Rushby 92, a *system with structured state* is a machine $\langle S, s_0, \text{Actions}, \text{step}, O, \text{dom} \rangle$ together with

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

Reference Monitor Conditions

Define a binary relation \sim_u^{oc} of *observable content equivalence* on S for each domain $u \in D$, by $s \sim_u^{oc} t$ if $\text{contents}(s, n) = \text{contents}(t, n)$ for all $n \in \text{observe}(u)$.

RM1. If $s \sim_u^{oc} t$ then $O_u(s) = O_u(t)$.

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

RM3. If $\text{contents}(s \cdot a, n) \neq \text{contents}(s, n)$ then $n \in \text{alter}(\text{dom}(a))$.

(RM2' a variant, from van der Meyden - ESORICS 2007, of Rushby's RM2)

Consistency of access control with a noninterference policy:

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

Proposition

If M is a system with structured state satisfying RM1-RM3 and AOI with respect to noninterference policy \succrightarrow then M is TA-secure (hence IP-secure) for \succrightarrow .

Adapting the Conditions to Extended Architectures

AOI'. If $\text{alter}(u) \cap \text{observe}(v) \neq \emptyset$ then $u \xrightarrow{f} v$ for some f .

Extra conditions for filtered edges:

1. If $\text{dom}(a) \xrightarrow{f} u$ for $f \neq \top$ and $I(f)(\alpha, a) = \epsilon$ and $x \in \text{observe}(u) \cap \text{alter}(\text{dom}(a))$ then $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \alpha)(x)$.
12. If $\text{dom}(a) \xrightarrow{f} u$ with $f \neq \top$ and $\text{dom}(b) \xrightarrow{g} u$ with $f \neq \top$ and $I(f)(\alpha, a) = I(g)(\beta, b) \neq \epsilon$ and $x \in \text{observe}(u) \cap (\text{alter}(\text{dom}(a)) \cup \text{alter}(\text{dom}(b)))$ and $(s_0 \cdot \alpha)(x) = (s_0 \cdot \beta)(x)$ then $(s_0 \cdot \alpha a)(x) = (s_0 \cdot \beta b)(x)$.

Theorem

Let \mathcal{AI} be an interpreted architecture. Suppose that M is a system with structured state satisfying RM1-RM3, AOI' and I1-I2. Then M is TFF-compliant with \mathcal{AI} .

Conclusion/Future Work

The examples demonstrate cases where it is feasible to formally derive global properties from an abstract level of specification of architecture + properties of trusted components

Many issues remain:

- ▶ Are there classes of specifications that can be straightforwardly implemented?
- ▶ Connections to other implementation patterns: e.g., periods processing, network partitioning.
- ▶ Richer semantics of architectures, e.g., for timing, probabilistic attacks.
- ▶ Syntax for architectural specifications, efficiently automatable cases of verification.