Flots d'information, opacité – et quelques remarques sur le diagnostic

Béatrice Bérard

Sorbonne Université, LIP6

Based on joint work with:
K. Chatterjee, S. Haar, S. Haddad, O. Kouchnarenko, E. Lefaucheux, J. Mullins, M. Sassolas, S. Schmitz, S. Schwoon, N. Sznajder

MSR, 14 novembre 2019

General context: Security Properties

Information flow:

Transmission of information from a high level user to a low level user, in a possibly illegal and/or indirect way.

A class of Security Properties:

Avoid information flow to preserve secret data during communications. [Mantel 2000, Focardi, Gorrieri 2001, Bryans, Koutny, Mazaré, Ryan 2008].

Goals:

Check whether a system satisfies such properties.
[BKMR 2008, D'Souza, Holla, Raghavendra, Sprick 2011, Best, Darondeau, Gorrieri 2011, Best, Darondeau 2012, Cassez, Dubreil, Marchand 2012, Dimitrova, Finkbeiner, Kovács, Rabe, Seidl 2012, Clarkson, Finkbeiner, Koleini, Micinski, Rabe, Sanchez 2014]
Enforce those properties.
[Lafortune 12-19, with many co-authors], [Marchand 11-15, with many co-authors], [Tong, Ma, Li, Seatzu, Giua 16].

Partially Observable System

visit to a red state is
hidden from observer

observing ad^* discloses a visit
acd^* is ambiguous

$\mathcal{A}$:

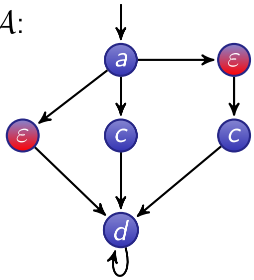

Partially Observable System

visit to a red state is
hidden from observer

observing *ad** discloses a visit
acd* is ambiguous

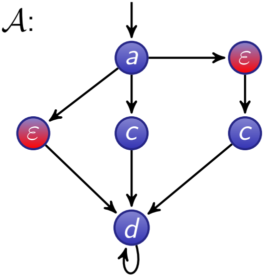

Goals: hide or detect information

- Opacity: the visit is a secret which must be kept
  [Bryans et al. 08]
- Diagnosis: the visit is a faulty event which must be detected
  [Sampath et al. 95]
  No black box: Observer knows the system

Outline

A common framework

A system with set L of behaviours

- A subset $M \subseteq L$ with $\overline{M} = L \setminus M$,
- An external agent observing the system via a function $\mathcal{O}$ on L.

Requirements: ordering ambiguity

$$
\begin{tikzcd}
 & \mathcal{O}(M) = \mathcal{O}(\overline{M}) & & \text{M symmetrically opaque} \\
\mathcal{O}(M) \subseteq \mathcal{O}(\overline{M}) \arrow[ur] & & \mathcal{O}(\overline{M}) \subseteq \mathcal{O}(M) \arrow[ul] & \text{M opaque $\mid$ $\overline{M}$ opaque} \\
 & \mathcal{O}(M) \cap \mathcal{O}(\overline{M}) \neq \emptyset \arrow[ul] \arrow[ur] & & \text{M weakly opaque or not diagnosable} \\
 & \mathcal{O}(M) \cap \mathcal{O}(\overline{M}) = \emptyset \arrow[u] & & \text{M diagnosable}
\end{tikzcd}
$$

Illustration

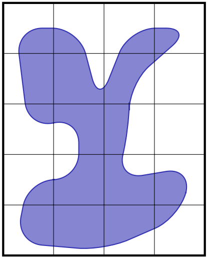

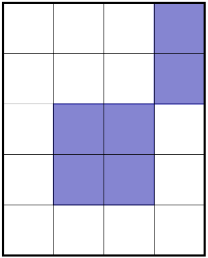

 M 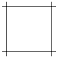 $\mathcal{O}^{-1}(o)$

$\mathcal{O}(M) \subseteq \mathcal{O}(\overline{M})$

$\mathcal{O}(M) \cap \mathcal{O}(\overline{M}) = \emptyset$

Verification

Model	Diagnosability	Opacity
finite LTS	NL-c.	PSPACE-c. [Cassez et al. 09]
safe (WF-)PN weak-fairness	PSPACE-c.	ESPACE-c. det. space $2^{O(n)}$
PN	EXPSPACE-c. +[Yin et al. 17]	undecidable +[Bryans et al. 08]
strict WF-PN no fair faults	Reach $\leq^{P}$ ¬Diag $\leq^{EXP}$ Reach	

[B., Haar, Schmitz, Schwoon 17]

Weak Fairness

A WF-Petri net

is a PN $\mathcal{N} = (P, T, w, m_0)$ with a subset $W \subseteq T$ of weakly fair transitions.
Trace $\sigma = t_1 t_2 \ldots \in Tr^\omega$ with markings $m_0 m_1 \ldots$ is weakly fair if $\forall t \in W$:

WF1 either there are infinitely many i with $t = t_i$

WF2 or there are infinitely many i where t_i conflicts with t in m_{i-1}:
$m_{i-1}(p) - w(p, t_i) < w(p, t)$ for some place p.

For safe PNS, equivalent to:

For each i, if t is enabled in m_{i-1}, there is a $j \geq i$ with ${}^\bullet t \cap {}^\bullet t_j \neq \emptyset$.

WF Diagnosability and WF Opacity

WF Diagnosability = Finite Diagnosability restricted to WF traces

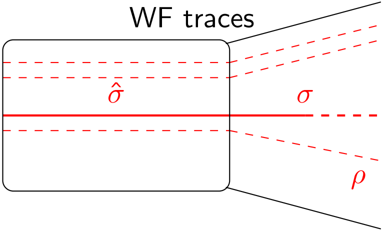

for any $\sigma \in Fty^{\omega}_{WF}(\mathcal{A})$
there is a prefix $\hat{\sigma}$ s.t.
any $\rho \in Tr^{\omega}_{WF}(\mathcal{A})$
with $\mathcal{O}(\hat{\sigma}) < \mathcal{O}(\rho)$
is also faulty

WF Diagnosability and WF Opacity

WF Diagnosability = Finite Diagnosability restricted to WF traces

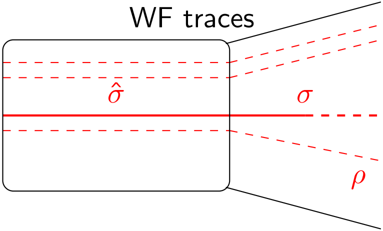

for any $\sigma \in Fty^{\omega}_{WF}(\mathcal{A})$
there is a prefix $\hat{\sigma}$ s.t.
any $\rho \in Tr^{\omega}_{WF}(\mathcal{A})$
with $\mathcal{O}(\hat{\sigma}) < \mathcal{O}(\rho)$
is also faulty

WF Opacity = Finite Opacity restricted to WF traces

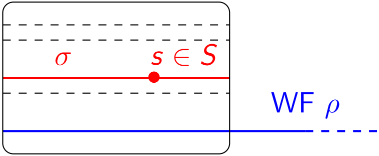

for each $\sigma \in Sec^{*}(\mathcal{A})$
there is $\rho \in Pub^{\omega}_{WF}(\mathcal{A})$
such that $\mathcal{O}(\sigma) \leq \mathcal{O}(\rho)$

Properties

$W = \emptyset$ corresponds to the standard notion

For a convergent WF PN $(\mathcal{N}, W)$:

- $(\mathcal{N}, \emptyset)$ is WF diagnosable iff $\mathcal{N}$ is diagnosable.
- Secret is WF opaque in $(\mathcal{N}, \emptyset)$ iff it is opaque in $\mathcal{N}$.

Properties

$W = \emptyset$ corresponds to the standard notion

For a convergent WF PN $(\mathcal{N}, W)$:

- $(\mathcal{N}, \emptyset)$ is WF diagnosable iff $\mathcal{N}$ is diagnosable.
- Secret is WF opaque in $(\mathcal{N}, \emptyset)$ iff it is opaque in $\mathcal{N}$.

WF Opacity is more discriminating than Opacity

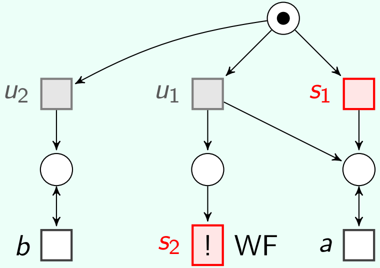

Opaque because
$\mathcal{O}(Sec^*) = a^*$
$\mathcal{O}(\overline{Sec^*}) = a^* + b^*$
but not WF-opaque

For secret trace $\sigma = s_1 a$, any infinite WF trace ρ such that $\mathcal{O}(\sigma) < \mathcal{O}(\rho)$ belongs to $s_1 a^\omega + u_1 a^* s_2 a^\omega$ hence contains a secret transition.

Properties

$W = \emptyset$ corresponds to the standard notion

For a convergent WF PN $(\mathcal{N}, W)$:

- $(\mathcal{N}, \emptyset)$ is WF diagnosable iff $\mathcal{N}$ is diagnosable.
- Secret is WF opaque in $(\mathcal{N}, \emptyset)$ iff it is opaque in $\mathcal{N}$.

Weak fairness increases diagnosability

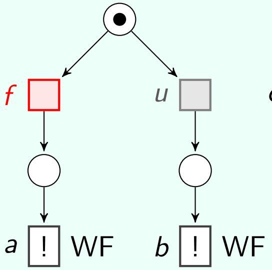

not Diag because
$\mathcal{O}(fc^\omega) = \mathcal{O}(uc^\omega)$
but Diag without c
and WF-diag

For WF faulty traces in fc^*ac^ω, finite prefixes containing a have observations in c^*ac^*, hence all infinite WF extensions are faulty.

Summary

Good news:

- Weak fairness for Diagnosability and Opacity comes at no additional cost in safe Petri nets;
- Standard Diagnosability is EXPSPACE-complete for Petri nets.

Bad news:

- Other strong undecidability results for the verification of Opacity in Petri nets;
- Non Diagnosability is equivalent to reachability when faults are not weakly fair.

Open problem: the complexity of verifying WF Diagnosability

Outline

Qualitative properties of Diagnosability and Opacity

Rational Information Flow Properties

Probabilistic Disclosure

- Probabilistic disclosure for Markov Chains
- Disclosing a secret under strategies

Examples

Given actions in A and set of traces $L \subseteq A^*$

- $A = V \uplus C \uplus N$ a partition into visible, confidential and neutral actions.

Removal of confidential actions:

An observer cannot see if the confidential actions are erased: for any $w \in L$, erasing all confidential actions in w results in a behaviour still in L.

Examples

Given actions in A and set of traces $L \subseteq A^*$

- $A = V \uplus C \uplus N$ a partition into visible, confidential and neutral actions.

Removal of confidential actions:
An observer cannot see if the confidential actions are erased: for any $w \in L$, erasing all confidential actions in w results in a behaviour still in L.

Insertion of X-admissible confidential actions, with $X \subseteq A$:
for any $w = w_1 w_2 \in L$ such that w_2 contains no confidential action and there exists $w_3 \in A^*$ and $c \in C$ with $w_3 c \in L$ and the X-letters in w_1 and w_3 are the same, then $w_1 c w_2$ also belongs to L.

Examples

Given actions in A and set of traces $L \subseteq A^*$

- $A = V \uplus C \uplus N$ a partition into visible, confidential and neutral actions.

Removal of confidential actions:

An observer cannot see if the confidential actions are erased: for any $w \in L$, erasing all confidential actions in w results in a behaviour still in L.

Insertion of X-admissible confidential actions, with $X \subseteq A$:

for any $w = w_1 w_2 \in L$ such that w_2 contains no confidential action and there exists $w_3 \in A^*$ and $c \in C$ with $w_3 c \in L$ and the X-letters in w_1 and w_3 are the same, then $w_1 c w_2$ also belongs to L.

- $A = V \uplus P$ a partition into visible actions and participant actions.

Strong anonymity of participants:

for any $w \in L$, replacing in w an action $a \in P$ by any other action in P produces a behaviour still in L.

Rational observers

- An automaton is a finite Labelled Transition System over a set of labels Lab. With final states and Lab is alphabet A, it accepts a regular language in A^*.

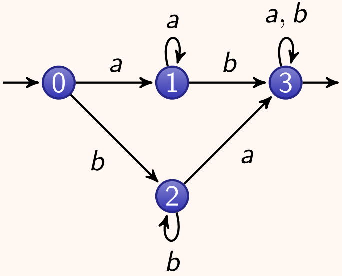

$L = a^+ b\{a,b\}^* \cup b^+ a\{a,b\}^*$

Rational observers

- An automaton is a finite Labelled Transition System over a set of labels Lab. With final states and Lab is alphabet A, it accepts a regular language in A^*.
- A transducer is an automaton with set of labels $Lab \subseteq A^* \times B^*$. With final states, it accepts a rational relation in $A^* \times B^*$.

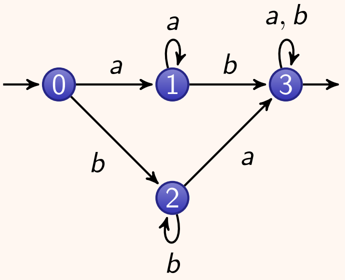

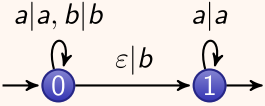

$L = a^+ b\{a,b\}^* \cup b^+ a\{a,b\}^*$

$R = \{(a,a),(b,b)\}^*(\varepsilon,b)(a,a)^*$

Rational observers

- An automaton is a finite Labelled Transition System over a set of labels Lab. With final states and Lab is alphabet A, it accepts a regular language in A^*.
- A transducer is an automaton with set of labels $Lab \subseteq A^* \times B^*$. With final states, it accepts a rational relation in $A^* \times B^*$.

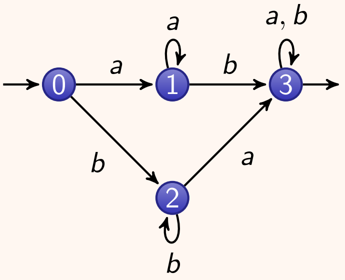

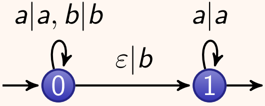

$abaaa \rightarrow abbaaa$

$L = a^+b\{a,b\}^* \cup b^+a\{a,b\}^*$

$R = \{(a,a),(b,b)\}^*(\varepsilon,b)(a,a)^*$

Rational observers

- An automaton is a finite Labelled Transition System over a set of labels *Lab*. With final states and *Lab* is alphabet A, it accepts a regular language in A^*.
- A transducer is an automaton with set of labels $Lab \subseteq A^* \times B^*$. With final states, it accepts a rational relation in $A^* \times B^*$.

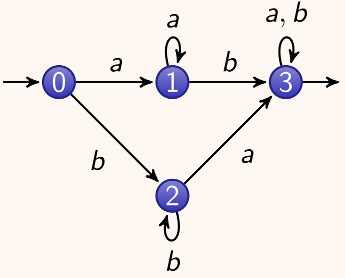

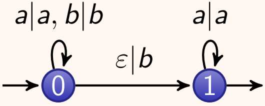

$$
\begin{aligned}
abaaa \rightarrow\; & ab\textcolor{red}{b}aaa \\
& aba\textcolor{red}{b}aa \\
& abaa\textcolor{red}{b}a \\
& abaaa\textcolor{red}{b}
\end{aligned}
$$

$L = a^+ b\{a, b\}^* \cup b^+ a\{a, b\}^*$

$R = \{(a, a), (b, b)\}^*(\varepsilon, b)(a, a)^*$

Rational observers

A rational observer

is a rational relation $\mathcal{O} \subseteq A^* \times B^*$.
Observation of $w \in A^*$: $\mathcal{O}(w) = \{w' \in B^* \mid (w, w') \in \mathcal{O}\}$.
Observation of $L \subseteq A^*$: $\mathcal{O}(L) = \cup_{w \in L} \mathcal{O}(w)$

Example: an Orwellian observer

Over $A = \{a, b\}$: $\mathcal{O}(\varepsilon) = \varepsilon$ and $\mathcal{O}(w) = \begin{cases} \pi_{\{b\}}(w) \text{ if } w \text{ ends with } a \\ \pi_{\{a\}}(w) \text{ if } w \text{ ends with } b \end{cases}$

Then $\mathcal{O} = \mathcal{O}_a \uplus \mathcal{O}_b \uplus \mathcal{O}_\varepsilon$ with:

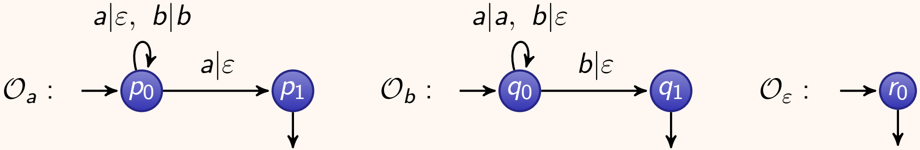

In $L = (a+b)(a^* + b^*)(a+b)$, the subset $M = a(a^* + b^*)(a+b)$ is opaque.

Rational Information Flow Properties

A rational information flow property (RIFP) for L

is any relation $L_1 \subseteq L_2$, where L_1 and L_2 are given by:

$$L_1, L_2 ::= \; L \mid \mathcal{O}(L_1) \mid L_1 \cup L_2 \mid L_1 \cap L_2$$

where $\mathcal{O}$ is a rational observer.

$RIF(\mathcal{L})$ for a class of languages $\mathcal{L}$

is the set of rational information flow properties for languages $L \in \mathcal{L}$.

Example 1: Removal of confidential actions

$A = V \uplus C \uplus N$ a partition into visible, confidential and neutral actions.

An observer cannot see if the confidential actions are erased: for any behaviour $w \in L$, erasing all confidential actions in w results in a behaviour still in L.

Translates as

$$\pi_{\overline{C}}(L) \subseteq L$$

where $\pi_{\overline{C}}$ is the projection from A^* onto $(A \setminus C)^*$:

$$\pi_{\overline{C}}(a) = \begin{cases} \varepsilon \text{ if } a \in C \\ a \text{ otherwise} \end{cases}$$

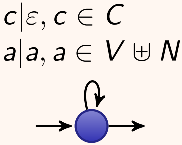

Example 1: Removal of confidential actions

$A = V \uplus C \uplus N$ a partition into visible, confidential and neutral actions.

An observer cannot see if the confidential actions are erased: for any behaviour $w \in L$, erasing all confidential actions in w results in a behaviour still in L.

Translates as

$$\pi_{\overline{C}}(L) \subseteq L$$

where $\pi_{\overline{C}}$ is the projection from A^* onto $(A \setminus C)^*$:

$$\pi_{\overline{C}}(a) = \begin{cases} \varepsilon \text{ if } a \in C \\ a \text{ otherwise} \end{cases}$$

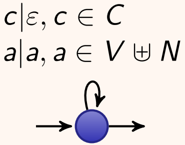

Proposition

Since $\pi_{\overline{C}}$ is a rational observer, removal of confidential actions is an RIFP.

Example 2: Insertion of confidential actions

$A = V \uplus C \uplus N$ and $X \subseteq A$.

For any $w = w_1 w_2 \in L$ such that w_2 contains no confidential event and there exists $w_3 \in A^*$ and $c \in C$ with $w_3 c \in L$ and the X-letters in w_1 and w_3 are the same, then $w_1 c w_2$ also belongs to L.

Translates as

$$\bigcup_{c \in C} (\textit{l-ins}_c(L) \cap \mathcal{O}_c^X(L)) \subseteq L$$

where for each $c \in C$,

- $\textit{l-ins}_c$ is the rational relation inserting c after the last confidential action,
- $\mathcal{O}_c^X$ is defined by $\mathcal{O}_c^X(u) = \pi_X^{-1}(\pi_X(c^{-1}u)).c.(V \uplus N)^*$ for $u \in A^*$.

Example 2: Insertion of confidential actions

$A = V \uplus C \uplus N$ and $X \subseteq A$.

For any $w = w_1 w_2 \in L$ such that w_2 contains no confidential event and there exists $w_3 \in A^*$ and $c \in C$ with $w_3 c \in L$ and the X-letters in w_1 and w_3 are the same, then $w_1 c w_2$ also belongs to L.

Translates as

$$\bigcup_{c \in C} (\mathit{l\text{-}ins}_c(L) \cap \mathcal{O}_c^X(L)) \subseteq L$$

where for each $c \in C$,

- $\mathit{l\text{-}ins}_c$ is the rational relation inserting c after the last confidential action,
- $\mathcal{O}_c^X$ is defined by $\mathcal{O}_c^X(u) = \pi_X^{-1}(\pi_X(c^{-1}u)).c.(V \uplus N)^*$ for $u \in A^*$.

Proposition

All operations are rational observers, hence insertion of X-admissible confidential actions is an RIFP.

Example 3: Strong anonymity

$A = V \uplus P$.

For any $w \in L$, replacing in w an action in P by another produces a behaviour in L.

Translates as $\mathcal{O}_{SA}^{P}(L) \subseteq L$
where $\mathcal{O}_{SA}^{P}$ is a substitution:

$$\mathcal{O}_{SA}^{P}(a) = \begin{cases} P \text{ if } a \in P \\ \{a\} \text{ otherwise} \end{cases}$$

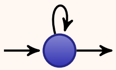

Example 3: Strong anonymity

$A = V \uplus P$.

For any $w \in L$, replacing in w an action in P by another produces a behaviour in L.

Translates as $\mathcal{O}_{SA}^{P}(L) \subseteq L$
where $\mathcal{O}_{SA}^{P}$ is a substitution:

$$\mathcal{O}_{SA}^{P}(a) = \begin{cases} P \text{ if } a \in P \\ \{a\} \text{ otherwise} \end{cases}$$

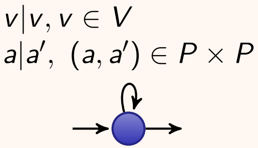

Proposition

A substitution is a rational observer, hence strong anonymity is an RIFP.

Verification of RIFPs

For a class of languages $\mathcal{L}$:

If $\mathcal{L}$ is closed under union, intersection, and rational transductions, and if the inclusion is decidable in $\mathcal{L}$, then any property in $RIF(\mathcal{L})$ is decidable.

Verification of RIFPs

For a class of languages $\mathcal{L}$:

If $\mathcal{L}$ is closed under union, intersection, and rational transductions, and if the inclusion is decidable in $\mathcal{L}$, then any property in $RIF(\mathcal{L})$ is decidable.

For the class $\mathcal{R}eg$ of regular languages:

The problem of deciding a property in $RIF(\mathcal{R}eg)$ is PSPACE-complete.

Because regular languages have all the required closure properties and inclusion is decidable in PSPACE in $\mathcal{R}eg$.

PSAPCE-hardness comes from the fact that $\mathcal{O}_K(w) = \{w\} \cap K$ is a rational relation if and only if K is a regular language.

Verification of RIFPs

For a class of languages $\mathcal{L}$:

If $\mathcal{L}$ is closed under union, intersection, and rational transductions, and if the inclusion is decidable in $\mathcal{L}$, then any property in $RIF(\mathcal{L})$ is decidable.

For the class $\mathcal{R}eg$ of regular languages:

The problem of deciding a property in $RIF(\mathcal{R}eg)$ is PSPACE-complete.

Because regular languages have all the required closure properties and inclusion is decidable in PSPACE in $\mathcal{R}eg$.
PSAPCE-hardness comes from the fact that $\mathcal{O}_K(w) = \{w\} \cap K$ is a rational relation if and only if K is a regular language.

Consequence:

Strong (and weak) anonymity [BKMR 2008], as well as all Basic Security Predicates [Mantel 2000], are decidable (in PSPACE) for regular languages. We retrieve results from [D'Souza et al., 2011].

The case of Opacity

For $M \subseteq L$ a regular subset of secret behaviours
and $\mathcal{O}$ a functional rational observer

From $\mathcal{O}(M) \subseteq \mathcal{O}(\overline{M})$:

Rational opacity for regular secrets is an RIFP.

Consequence:

We recover the decidability result (in PSPACE) for rational opacity with regular languages and regular secrets [Cassez et al., 2009].

Remark: Strong Anonymity translates as Opacity [BKMR08]

- $\mathcal{O}$ is the morphism into $(\Sigma \cup \{\sharp\})^*$ defined by:
  $\mathcal{O}(a) = \sharp$ if $a \in P$ and $\mathcal{O}(a) = a$ otherwise
- π_P the projection on P^*

L is strongly anonymous w.r.t. P iff for any $u \in P^*$,

$$Sec_u = \{w \in L \mid \pi_P(w) \neq u \wedge |\pi_P(w)| = |u|\}$$

is opaque for $\mathcal{O}$.

The case of Weak Non Inference

From [D'Souza et al., 2011]

With $A = V \uplus C \uplus N$,

WNI

L satisfies WNI if for all $w \in L$ there exists $w' \in L$ such that if w contains confidential actions, then $\pi_V(w) = \pi_V(w')$ and $\pi_C(w) \neq \pi_C(w')$.

WNI is undecidable on regular languages.

Consequence:

WNI is NOT an RIFP.

Summary

Good news:

- Many security properties from the literature are RIFPs;
- The complexity of verification is always in PSPACE for regular languages, whatever the (rational) observation.

To do:

- Find other classes satisfying the closure properties leading to decidability;
- Find links with model checking extensions of LTL like SecLTL [DFKRS12] or even CTL* like HyperCTL* [CFKMRS14].
- What about Opacity with a general rational observer ?

Outline

Qualitative properties of Diagnosability and Opacity

Rational Information Flow Properties

Probabilistic Disclosure

- Probabilistic disclosure for Markov Chains
- Disclosing a secret under strategies

A quantitative problem for opacity...

 Sec 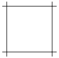 $\mathcal{O}^{-1}(o)$ 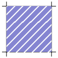 Classes leaking their inclusion into Sec

No disclosing path iff
$V = Sec \setminus \mathcal{O}^{-1}(\mathcal{O}(\overline{Sec}))$ is empty

Measuring the disclosure set V

...under uncertainty

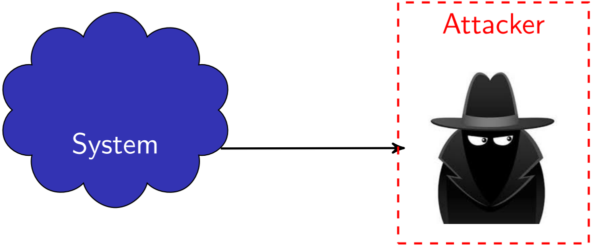

- Probabilistic choice: Markov Chains
  [B., Mullins, Sassolas 10,15] [Saboori, Hadjicostis 14]

...under uncertainty

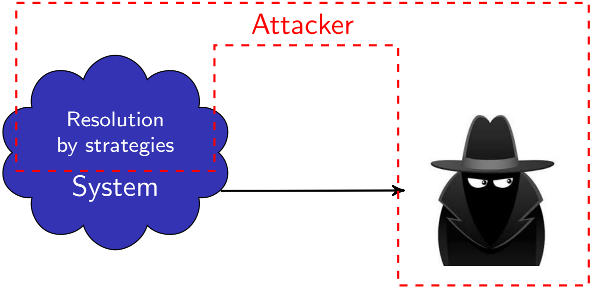

- Probabilistic choice: Markov Chains
  [B., Mullins, Sassolas 10,15] [Saboori, Hadjicostis 14]
- Combined with nondeterministic choice:
  [B., Chatterjee, Sznajder 15] for MDPs and POMDPs,
  [B., Haddad, Lefaucheux 17] for MDPs,
- Underspecification: [B., Kouchnarenko, Mullins, Sassolas 16] for IMCs.

A toy example

Access control to a database inspired from [Biondi et al. 13]

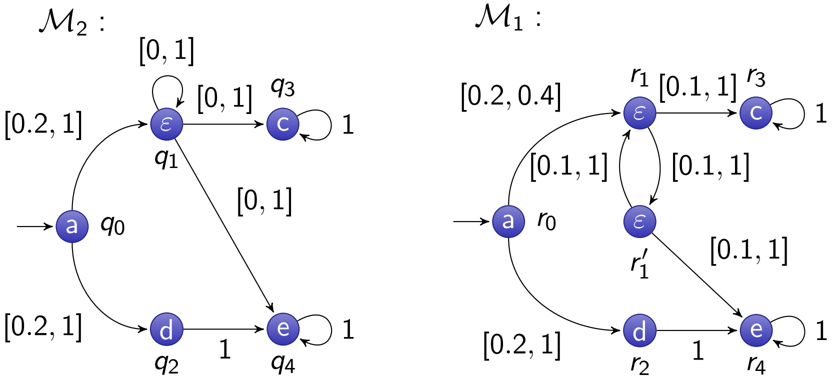

0: input user name, 1: input password, 3: access granted if correct
2: not on the list of authorized users, 4: reject
$Sec = \{0.1.3^\omega\}$; All states except 1 and $1'$ are observable.

Observable Markov chains

Example with *Sec*: visiting s_1 or s_2, hidden by $\mathcal{O}$

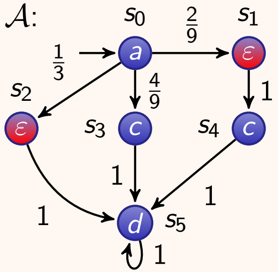

A Markov Chain $\mathcal{A} = (S, \Delta, \mathcal{O})$ over Σ:

- countable set S of states,
- $\Delta : S \to \mathcal{D}ist(S)$,
- $\mathcal{O} : S \to \Sigma \cup \{\varepsilon\}$ observation function.

equipped with an initial distribution μ_0.

Disclosure for MCs

ω-Disclosure of Sec in $(\mathcal{A}, \mu_0)$:

$$Disc_\omega(\mathcal{A}, \mu_0, Sec) = \mathbf{P}_{\mathcal{A},\mu_0}(V) \text{ for } V = Sec \setminus \mathcal{O}^{-1}(\mathcal{O}(\overline{Sec})).$$

Example with Sec: presence of s_1 or s_2, hidden by $\mathcal{O}$

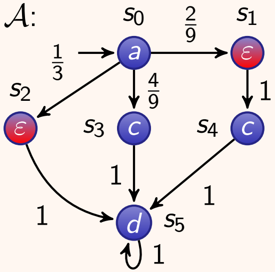

$Path(\mathcal{A})$	$\mathcal{O}$	Sec?	V?	$\mathbf{P}_\mathcal{A}$
$s_0 s_2 s_5^\omega$	ad^ω	✓	✓	$1/3$
$s_0 s_3 s_5^\omega$	acd^ω	✗	✗	$4/9$
$s_0 s_1 s_4 s_5^\omega$	acd^ω	✓	✗	$2/9$

$$Disc_\omega(\mathcal{A}, \mathbf{1}_{s_0}, Sec) = \frac{1}{3}$$

Finite disclosure

Restricting *Sec* to the set of pathes visiting states from a given subset

assuming a path remains secret once a secret state has been visited.

Observation sequence w in Σ^* is:
disclosing if all pathes in $\mathcal{O}^{-1}(w)$ are secret,
minimal disclosing if disclosing with no strict disclosing prefix.

$Disc(\mathcal{A}, \mu_0, Sec)$: probability of minimal disclosing observations

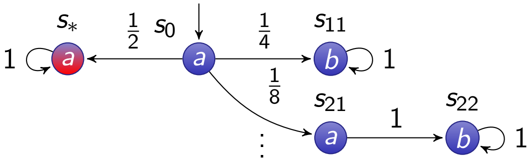

$Disc_\omega = \frac{1}{2}$

$Disc = 0$

***Disc* $\leq$ *Disc*$_\omega$**

equality if $\mathcal{A}$ is convergent and finitely branching.

Interactions with the system

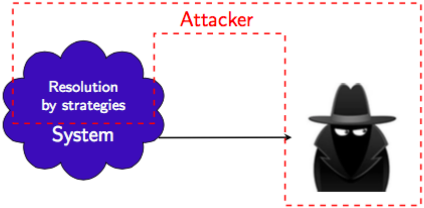

Active attacker

The attacker consists of two components:

- The passive external observer,
- Some piece of code inside the system.

Worst case corresponds to maximal disclosure.

System designer

The designer has provided a first version with the required functionalities.
He must develop the access policy...

... to obtain minimal disclosure.

Constraint Markov Chains

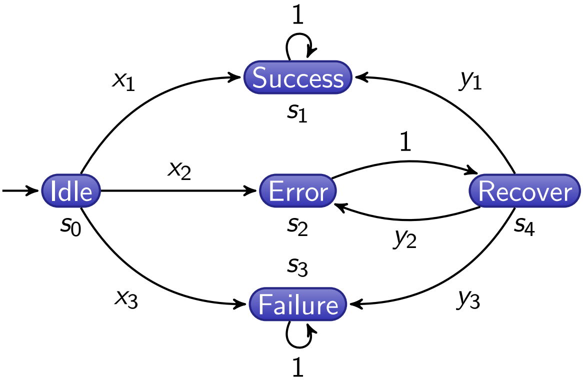

$\mathcal{M}_1 = (S, T_1, \mathcal{O})$:

$T_1(s_0)$ subset of:
$0 \leq x_1, x_2, x_3 \leq 1$
$x_1 + x_2 + x_3 = 1$

$T_1(s_4)$ subset of:
$0 \leq y_1, y_2, y_3 \leq 1$
$y_1 + y_2 + y_3 = 1$

A CMC over Σ: [Jonsson, Larsen 1991] [Caillaud et al., 2011]

$\mathcal{M} = (S, T, \mathcal{O})$ is like an OMC with

- finite set of states S,
- $T : S \to 2^{\mathcal{D}ist(S)}$.

Subclasses of CMCs

MDP: Markov Decision Processes

For each $s \in S$, $T(s)$ is a finite set.

LCMC: Linear CMCs

For each $s \in S$, $T(s)$ is the set of distributions that are solutions of a linear system.

IMC: Interval MC

For each s, $T(s)$ is described by a family of intervals $(I(s, s'))_{s' \in S}$.

Relations

- IMC is a strict subclass of LCMC,
- Any LCMC can be transformed in an exponentially larger MDP.

Examples

LCMC $\mathcal{M}_2$:

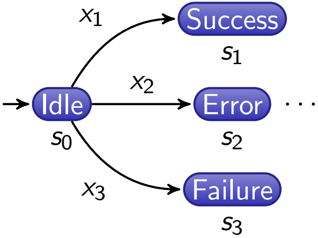

IMC $\mathcal{M}_3$:

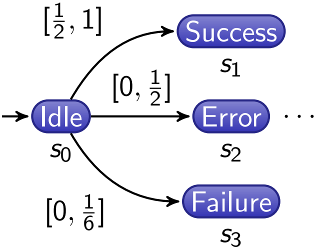

$$0 \le x_1, x_2, x_3 \le 1$$
$$x_1 + x_2 + x_3 = 1$$

$$x_2 \ge 2x_3$$
$$x_2 + x_3 \le \frac{1}{2}$$

$$\frac{1}{2} \le x_1 \le 1$$
$$0 \le x_2 \le \frac{1}{2}$$
$$0 \le x_3 \le \frac{1}{6}$$

$$\mu_1 = (1, 0, 0)$$
$$\mu_2 = (\tfrac{1}{2}, \tfrac{1}{2}, 0)$$
$$\mu_3 = (\tfrac{1}{2}, \tfrac{1}{3}, \tfrac{1}{6})$$

$$\mu_4 = (\tfrac{5}{6}, 0, \tfrac{1}{6}) \in T_3(s_0)$$
$$\mu_4 \notin T_2(s_0)$$

From LCMCs to MDPs

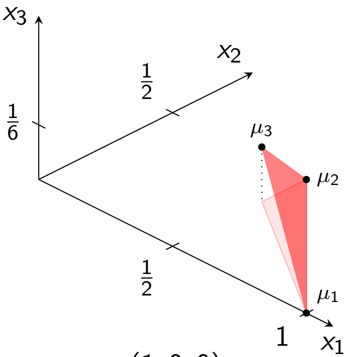

$$\mu_1 = (1, 0, 0)$$

$$\mu_2 = (\tfrac{1}{2}, \tfrac{1}{2}, 0)$$

$$\mu_3 = (\tfrac{1}{2}, \tfrac{1}{3}, \tfrac{1}{6})$$

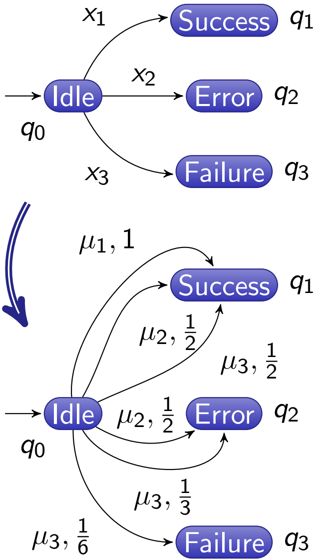

Strategies on CMCs

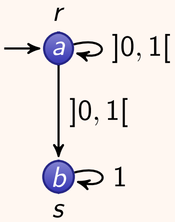

Strategies on CMCs

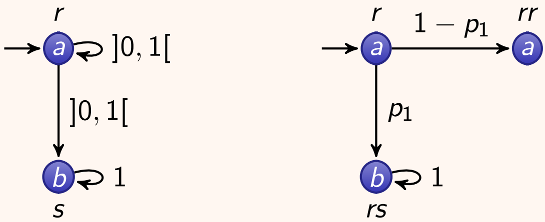

A strategy for $\mathcal{M} = (S, T, \mathcal{O})$ with initial distribution μ_0:

$\sigma : FRuns(\mathcal{M}) \to \mathcal{D}ist(S)$

For $\rho = s_0 \xrightarrow{\mu_1} s_1 \ldots \xrightarrow{\mu_n} s_n, \ \ \sigma(\rho) \in T(s_n)$.

Scheduling $\mathcal{M}$ with σ produces a (possibly infinite) MC $\mathcal{M}_\sigma$.

Strategies on CMCs

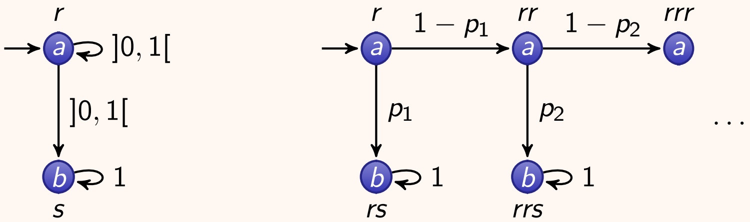

A strategy for $\mathcal{M} = (S, T, \mathcal{O})$ with initial distribution μ_0:

$\sigma : FRuns(\mathcal{M}) \to \mathcal{D}ist(S)$

For $\rho = s_0 \xrightarrow{\mu_1} s_1 \ldots \xrightarrow{\mu_n} s_n,\ \ \sigma(\rho) \in T(s_n)$.

Scheduling $\mathcal{M}$ with σ produces a (possibly infinite) MC $\mathcal{M}_\sigma$.

Randomized strategies on MDPs

An MDP with distributions μ_1 and μ_2 for s_0 and secret states $\{s_2, s_3\}$

$Disc = \frac{1}{2}$ with the two strategies choosing μ_1 or μ_2 in s_0
if they are known by the observer.

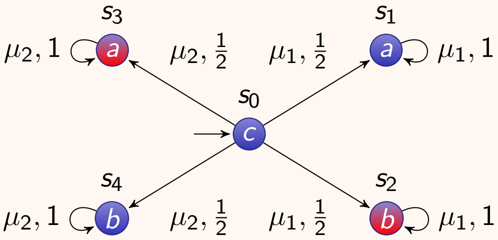

But $Disc = 0$ with randomized strategies σ_p such that
$\sigma_p(s_0) = p\mu_1 + (1-p)\mu_2$ with $0 < p < 1$. Necessary for minimisation.

A randomized strategy associates $\sigma(\rho) \in \mathcal{D}ist(T(s_n))$

with $\rho = s_0 \xrightarrow{\mu_1} s_1 \ldots \xrightarrow{\mu_n} s_n$ (instead of $\sigma(\rho)$ in $T(s_n)$).

Maximal and minimal disclosure

For *Sec* in $\mathcal{M}$ with initial distribution μ_0:

- $Disc_{\max}(\mathcal{M}, \mu_0, Sec) = sup_{\sigma \in Strat(\mathcal{M})} Disc(\mathcal{M}_\sigma, \mu_0, Sec)$
- $Disc_{\min}(\mathcal{M}, \mu_0, Sec) = inf_{\sigma \in Strat(\mathcal{M})} Disc(\mathcal{M}_\sigma, \mu_0, Sec)$

Several disclosure problems for a given $\mathcal{M}$

- Value problem: compute the disclosure $Disc_{\max}$ or $Disc_{\min}$.
- Quantitative decision problems: Given a threshold $\theta \in [0, 1]$, is $Disc_{\max} \geq \theta$? is $Disc_{\min} \leq \theta$?
- Qualitative decision problems:
  Limit-sure disclosure: the quantitative problem
  with $\theta = 1$ for maximisation and $\theta = 0$ for minimisation.

Maximal Disclosure

[BCS15] On MDPs, if observer ignores the strategies

or if no edge can be blocked by a strategy,

- The value can be computed in polynomial time;
- All problems are decidable.

[BHL17] On MDPs, if observer knows the strategies:

- Deterministic strategies are sufficient;
- The problem asking whether there exists a strategy producing value 1 is EXPTIME-complete;
- But the quantitative and limit-sure problems are undecidable.

Consequence:

The quantitative problem is undecidable for general LCMCs.

Minimal Disclosure

[BHL17] On MDPs, if observer knows the strategies:

- Families of randomized strategies are necessary;
- The value can be computed in EXPTIME;
- All problems are decidable.

Summary

Linear CMCs form a good class for compact specifications of probabilistic systems with:

- nice closure properties;
- an increased security criterion with schedulers as adversaries;
- But the quantitative problem is undecidable unless the structure is fixed.

Minimisation on MDPs

- requires randomized strategies;
- and all quantitative problems are decidable.

Conclusion

A lot of work to be done... on qualitative and quantitative aspects

Conclusion

A lot of work to be done... on qualitative and quantitative aspects

Thank you