The Generic Unwinding Theorem for CSP Noninterference Security

Pasquale Noce
Security Certification Specialist at Arjo Systems - Gep S.p.A.
pasquale dot noce dot lavoro at gmail dot com
pasquale dot noce at arjowiggins-it dot com

February 23, 2021

Abstract

The classical definition of noninterference security for a deterministic state machine with outputs requires to consider the outputs produced by machine actions after any trace, i.e. any indefinitely long sequence of actions, of the machine. In order to render the verification of the security of such a machine more straightforward, there is a need of some sufficient condition for security such that just individual actions, rather than unbounded sequences of actions, have to be considered.

By extending previous results applying to transitive noninterference policies, Rushby has proven an unwinding theorem that provides a sufficient condition of this kind in the general case of a possibly intransitive policy. This condition has to be satisfied by a generic function mapping security domains into equivalence relations over machine states.

An analogous problem arises for CSP noninterference security, whose definition requires to consider any possible future, i.e. any indefinitely long sequence of subsequent events and any indefinitely large set of refused events associated to that sequence, for each process trace.

This paper provides a sufficient condition for CSP noninterference security, which indeed requires to just consider individual accepted and refused events and applies to the general case of a possibly intransitive policy. This condition follows Rushby's one for classical noninterference security, and has to be satisfied by a generic function mapping security domains into equivalence relations over process traces; hence its name, Generic Unwinding Theorem. Variants of this theorem applying to deterministic processes and trace set processes are also proven. Finally, the sufficient condition for security expressed by the theorem is shown not to be a necessary condition as well, viz. there exists a secure process such that no domain-relation map satisfying the condition exists.

Contents

1 The Generic Unwinding Theorem

theory *GenericUnwinding*
imports *Noninterference-Ipurge-Unwinding.DeterministicProcesses*
begin

The classical definition of noninterference security for a deterministic state machine with outputs requires to consider the outputs produced by machine actions after any trace, i.e. any indefinitely long sequence of actions, of the machine. In order to render the verification of the security of such a machine more straightforward, there is a need of some sufficient condition for security such that just individual actions, rather than unbounded sequences of actions, have to be taken into consideration.

By extending previous results applying to transitive noninterference policies, Rushby [8] has proven an unwinding theorem that provides a sufficient condition of this kind in the general case of a possibly intransitive policy. This condition consists of a combination of predicates, which have to be satisfied by a generic function mapping security domains into equivalence relations over machine states.

An analogous problem arises for CSP noninterference security, whose definition given in [6] requires to consider any possible future, i.e. any indefinitely long sequence of subsequent events and any indefinitely large set of refused events associated to that sequence, for each process trace.

This paper provides a sufficient condition for CSP noninterference security, which indeed requires to just consider individual accepted and refused events and applies to the general case of a possibly intransitive policy. This condition follows Rushby's one for classical noninterference security; in some detail, it consists of a combination of predicates, which are the translations of Rushby's ones into Hoare's Communicating Sequential Processes model of computation [1]. These predicates have to be satisfied by a generic function mapping security domains into equivalence relations over process traces; hence the name given to the condition, *Generic Unwinding Theorem*. Variants of this theorem applying to deterministic processes and trace set processes (cf. [7]) are also proven.

The sufficient condition for security expressed by the Generic Unwinding

Theorem would be even more valuable if it also provided a necessary condition, viz. if for any secure process, there existed some domain-relation map satisfying the condition. Particularly, a constructive proof of such proposition, showing that some specified domain-relation map satisfies the condition whatever secure process is given, would permit to determine whether a process is secure or not by verifying whether the condition is satisfied by that map or not. However, this paper proves by counterexample that the Generic Unwinding Theorem does not express a necessary condition for security as well, viz. a process and a noninterference policy for that process are constructed such that the process is secure with respect to the policy, but no domain-relation map satisfying the condition exists.

The contents of this paper are based on those of [6] and [7]. The salient points of definitions and proofs are commented; for additional information, cf. Isabelle documentation, particularly [5], [4], [3], and [2].

For the sake of brevity, given a function F of type $'a_1 \Rightarrow \ldots \Rightarrow 'a_m \Rightarrow 'a_{m+1} \Rightarrow \ldots \Rightarrow 'a_n \Rightarrow 'b$, the explanatory text may discuss of F using attributes that would more exactly apply to a term of type $'a_{m+1} \Rightarrow \ldots \Rightarrow 'a_n \Rightarrow 'b$. In this case, it shall be understood that strictly speaking, such attributes apply to a term matching pattern $F\ a_1\ \ldots\ a_m$.

1.1 Propaedeutic definitions and lemmas

Here below are the translations of Rushby's predicates *weakly step consistent* and *locally respects* [8], applying to deterministic state machines, into Hoare's Communicating Sequential Processes model of computation [1].

The differences with respect to Rushby's original predicates are the following ones:

- The relations in the range of the domain-relation map hold between event lists rather than machine states.
- The domains appearing as inputs of the domain-relation map do not unnecessarily encompass all the possible values of the data type of domains, but just the domains in the range of the event-domain map.
- While every machine action is accepted in a machine state, not every process event is generally accepted after a process trace. Thus, whenever an event is appended to an event list in the consequent of an implication, the antecedent of the implication constrains the event list to be a trace, and the event to be accepted after that trace. In this way, the predicates do not unnecessarily impose that the relations in the range of the domain-relation map hold between event lists not being process traces.

definition *weakly-step-consistent* ::
 'a process ⇒ (*'a* ⇒ *'d*) ⇒ (*'a*, *'d*) *dom-rel-map* ⇒ *bool* **where**
weakly-step-consistent P D R ≡ ∀ *u* ∈ *range D*. ∀ *xs ys x*.
  (*xs*, *ys*) ∈ *R u* ∩ *R* (*D x*) ∧ *x* ∈ *next-events P xs* ∩ *next-events P ys* ⟶
  (*xs* @ [*x*], *ys* @ [*x*]) ∈ *R u*

definition *locally-respects* ::
 'a process ⇒ (*'d* × *'d*) *set* ⇒ (*'a* ⇒ *'d*) ⇒ (*'a*, *'d*) *dom-rel-map* ⇒ *bool* **where**
locally-respects P I D R ≡ ∀ *u* ∈ *range D*. ∀ *xs x*.
  (*D x*, *u*) ∉ *I* ∧ *x* ∈ *next-events P xs* ⟶ (*xs*, *xs* @ [*x*]) ∈ *R u*

In what follows, some lemmas propaedeutic for the proof of the Generic Unwinding Theorem are demonstrated.

lemma *ipurge-tr-aux-single-event*:
 ipurge-tr-aux I D U [*x*] = (*if* ∃ *v* ∈ *U*. (*v*, *D x*) ∈ *I*
   then []
   else [*x*])
⟨*proof*⟩

lemma *ipurge-tr-aux-cons*:
 ipurge-tr-aux I D U (*x* # *xs*) = (*if* ∃ *v* ∈ *U*. (*v*, *D x*) ∈ *I*
   then ipurge-tr-aux I D (*insert* (*D x*) *U*) *xs*
   else x # *ipurge-tr-aux I D U xs*)
⟨*proof*⟩

lemma *unaffected-domains-subset*:
  assumes
    A: *U* ⊆ *range D* **and**
    B: *U* ≠ {}
  shows *unaffected-domains I D U xs* ⊆ *range D* ∩ (−*I*) `` *range D*
⟨*proof*⟩

1.2 The Generic Unwinding Theorem: proof of condition sufficiency

Rushby's *Unwinding Theorem for Intransitive Policies* [8] states that a sufficient condition for a deterministic state machine with outputs to be secure is the existence of some domain-relation map R such that:

1. R is a *view partition*, i.e. the relations over machine states in its range are equivalence relations;

2. R is *output consistent*, i.e. states equivalent with respect to the domain of an action produce the same output as a result of that action;

3. R is weakly step consistent;

4. R locally respects the policy.

The idea behind the theorem is that a machine is secure if its states can be partitioned, for each domain u, into equivalence classes (1), such that the states in any such class C are indistinguishable with respect to the actions in u (2), transition into the same equivalence class C' as a result of an action (3), and transition remaining inside C as a result of an action not allowed to affect u (4).

This idea can simply be translated into the realm of Communicating Sequential Processes [1] by replacing the words "machine", "state", "action" with "process", "trace", "event", respectively, as long as a clarification is provided of what it precisely means for a pair of traces to be "indistinguishable" with respect to the events in a given domain. Intuitively, this happens just in case the events in that domain being accepted or refused after either trace are the same, thus the simplest choice would be to replace output consistency with *future consistency* as defined in [7]. However, indistinguishability between traces in the same equivalence class is not required in the case of a domain allowed to be affected by any domain, since the policy puts no restriction on the differences in process histories that may be detected by such a domain. Hence, it is sufficient to replace output consistency with *weak future consistency* [7].

Furthermore, indistinguishability with respect to individual refused events does not imply indistinguishability with respect to sets of refused events, i.e. refusals, unless for each trace, the corresponding refusals set is closed under set union. Therefore, for the condition to be sufficient for process security, the *refusals union closure* of the process [7] is also required. As remarked in [7], this property holds for any process admitting a meaningful interpretation, so that taking it as an additional assumption does not give rise to any actual limitation on the applicability of the theorem.

As a result of these considerations, the Generic Unwinding Theorem, formalized in what follows as theorem *generic-unwinding*, states that a sufficient condition for the CSP noninterference security [6] of a process being refusals union closed [7] is the existence of some domain-relation map R such that:

1. R is a view partition [7];

2. R is weakly future consistent [7];

3. R is weakly step consistent;

4. R locally respects the policy.

lemma *ruc-wfc-failures*:
assumes
RUC: *ref-union-closed P* **and**
WFC: *weakly-future-consistent P I D R* **and**
A: $U \subseteq$ *range D* $\cap (-I)$ ‘‘ *range D* **and**
B: $U \neq \{\}$ **and**
C: $\forall\, u \in U.\ (xs,\ xs') \in R\ u$ **and**
D: $(xs,\ X) \in$ *failures P*
shows $(xs',\ X \cap D -‘ U) \in$ *failures P*
⟨*proof*⟩

lemma *ruc-wfc-lr-failures-1*:
assumes
RUC: *ref-union-closed P* **and**
WFC: *weakly-future-consistent P I D R* **and**
LR: *locally-respects P I D R* **and**
A: $(xs\ @\ [y],\ Y) \in$ *failures P*
shows $(xs,\ \{x \in Y.\ (D\ y,\ D\ x) \notin I\}) \in$ *failures P*
⟨*proof*⟩

lemma *ruc-wfc-lr-failures-2*:
assumes
RUC: *ref-union-closed P* **and**
WFC: *weakly-future-consistent P I D R* **and**
LR: *locally-respects P I D R* **and**
A: $(xs,\ Z) \in$ *failures P* **and**
Y: $xs\ @\ [y] \in$ *traces P*
shows $(xs\ @\ [y],\ \{x \in Z.\ (D\ y,\ D\ x) \notin I\}) \in$ *failures P*
⟨*proof*⟩

lemma *gu-condition-imply-secure-aux* [*rule-format*]:
assumes
VP: *view-partition P D R* **and**
WFC: *weakly-future-consistent P I D R* **and**
WSC: *weakly-step-consistent P D R* **and**
LR: *locally-respects P I D R*
shows $U \subseteq$ *range D* $\longrightarrow$ $U \neq \{\} \longrightarrow$ $xs\ @\ ys \in$ *traces P* $\longrightarrow$
$(\forall\, u \in$ *unaffected-domains I D U* $[].\ (xs,\ xs') \in R\ u) \longrightarrow$
$(\forall\, u \in$ *unaffected-domains I D U ys*.
$(xs\ @\ ys,\ xs'\ @$ *ipurge-tr-aux I D U ys*$) \in R\ u)$
⟨*proof*⟩

lemma *gu-condition-imply-secure-1* [*rule-format*]:
assumes
RUC: *ref-union-closed P* **and**
VP: *view-partition P D R* **and**
WFC: *weakly-future-consistent P I D R* **and**
WSC: *weakly-step-consistent P D R* **and**
LR: *locally-respects P I D R*

shows $(xs \mathbin{@} y \# ys,\ Y) \in \mathit{failures}\ P \longrightarrow$
$(xs \mathbin{@} \mathit{ipurge\text{-}tr}\ I\ D\ (D\ y)\ ys,\ \mathit{ipurge\text{-}ref}\ I\ D\ (D\ y)\ ys\ Y) \in \mathit{failures}\ P$
⟨*proof*⟩

lemma *gu-condition-imply-secure-2* [*rule-format*]:
assumes
RUC: *ref-union-closed P* **and**
VP: *view-partition P D R* **and**
WFC: *weakly-future-consistent P I D R* **and**
WSC: *weakly-step-consistent P D R* **and**
LR: *locally-respects P I D R* **and**
Y: $xs \mathbin{@} [y] \in \mathit{traces}\ P$
shows $(xs \mathbin{@} zs,\ Z) \in \mathit{failures}\ P \longrightarrow$
$(xs \mathbin{@} y \# \mathit{ipurge\text{-}tr}\ I\ D\ (D\ y)\ zs,\ \mathit{ipurge\text{-}ref}\ I\ D\ (D\ y)\ zs\ Z) \in \mathit{failures}\ P$
⟨*proof*⟩

theorem *generic-unwinding*:
assumes
RUC: *ref-union-closed P* **and**
VP: *view-partition P D R* **and**
WFC: *weakly-future-consistent P I D R* **and**
WSC: *weakly-step-consistent P D R* **and**
LR: *locally-respects P I D R*
shows *secure P I D*
⟨*proof*⟩

It is interesting to observe that unlike symmetry and transitivity, the assumed reflexivity of the relations in the range of the domain-relation map is never used in the proof of the Generic Unwinding Theorem. Nonetheless, by assuming that such relations be equivalence relations over process traces rather than just symmetric and transitive ones, reflexivity has been kept among assumptions for both historical reasons – Rushby's Unwinding Theorem for deterministic state machines deals with equivalence relations – and practical reasons – predicate *refl-on* (*traces P*) may only be verified by a relation included in *traces P* × *traces P*, thus ensuring that traces be not correlated with non-trace event lists, which is a necessary condition for weak future consistency (cf. [7]).

Here below are convenient variants of the Generic Unwinding Theorem applying to deterministic processes and trace set processes (cf. [7]).

theorem *d-generic-unwinding*:
deterministic P ⟹
view-partition P D R ⟹
d-weakly-future-consistent P I D R ⟹
weakly-step-consistent P D R ⟹
locally-respects P I D R ⟹

secure P I D
⟨*proof*⟩

theorem *ts-generic-unwinding*:
trace-set T ⟹
view-partition (*ts-process T*) *D R* ⟹
d-weakly-future-consistent (*ts-process T*) *I D R* ⟹
weakly-step-consistent (*ts-process T*) *D R* ⟹
locally-respects (*ts-process T*) *I D R* ⟹
secure (*ts-process T*) *I D*
⟨*proof*⟩

1.3 The Generic Unwinding Theorem: counterexample to condition necessity

At a first glance, it seems reasonable to hypothesize that the Generic Unwinding Theorem expresses a necessary, as well as sufficient, condition for security, viz. that whenever a process is secure with respect to a policy, there should exist a set of "views" of process traces, one per domain, satisfying the apparently simple assumptions of the theorem.

It can thus be surprising to discover that this hypothesis is false, as proven in what follows by constructing a counterexample. The key observation for attaining this result is that symmetry, transitivity, weak step consistency, and local policy respect permit to infer the correlation of pairs of traces, and can then be given the form of introduction rules in the inductive definition of a set. In this way, a "minimum" domain-relation map *rel-induct* is obtained, viz. a map such that, for each domain u, the image of u under this map is included in the image of u under any map which has the aforesaid properties – particularly, which satisfies the assumptions of the Generic Unwinding Theorem.

Although reflexivity can be given the form of an introduction rule, too, it has been omitted from the inductive definition. This has been done in order to ensure that the "minimum" domain-relation map, and consequently the counterexample as well, still remain such even if reflexivity, being unnecessary (cf. above), were removed from the assumptions of the Generic Unwinding Theorem.

inductive-set *rel-induct-aux* ::
'a process ⇒ (*'d* × *'d*) *set* ⇒ (*'a* ⇒ *'d*) ⇒ (*'d* × *'a list* × *'a list*) *set*
for *P* :: *'a process* **and** *I* :: (*'d* × *'d*) *set* **and** *D* :: *'a* ⇒ *'d* **where**
rule-sym: (*u*, *xs*, *ys*) ∈ *rel-induct-aux P I D* ⟹
(*u*, *ys*, *xs*) ∈ *rel-induct-aux P I D* |
rule-trans: ⟦(*u*, *xs*, *ys*) ∈ *rel-induct-aux P I D*;
(*u*, *ys*, *zs*) ∈ *rel-induct-aux P I D*⟧ ⟹
(*u*, *xs*, *zs*) ∈ *rel-induct-aux P I D* |

rule-WSC: $[\![(u,\ xs,\ ys) \in \textit{rel-induct-aux}\ P\ I\ D;$
$(D\ x,\ xs,\ ys) \in \textit{rel-induct-aux}\ P\ I\ D;$
$x \in \textit{next-events}\ P\ xs \cap \textit{next-events}\ P\ ys]\!] \Longrightarrow$
$(u,\ xs\ @\ [x],\ ys\ @\ [x]) \in \textit{rel-induct-aux}\ P\ I\ D\ |$
rule-LR: $[\![u \in \textit{range}\ D;\ (D\ x,\ u) \notin I;\ x \in \textit{next-events}\ P\ xs]\!] \Longrightarrow$
$(u,\ xs,\ xs\ @\ [x]) \in \textit{rel-induct-aux}\ P\ I\ D$

definition *rel-induct* ::
$'a\ \textit{process} \Rightarrow ('d \times 'd)\ \textit{set} \Rightarrow ('a \Rightarrow 'd) \Rightarrow ('a, 'd)\ \textit{dom-rel-map}$ **where**
$\textit{rel-induct}\ P\ I\ D\ u \equiv \textit{rel-induct-aux}\ P\ I\ D\ ``\ \{u\}$

lemma *rel-induct-subset*:
assumes
VP: $\textit{view-partition}\ P\ D\ R$ **and**
WSC: $\textit{weakly-step-consistent}\ P\ D\ R$ **and**
LR: $\textit{locally-respects}\ P\ I\ D\ R$
shows $\textit{rel-induct}\ P\ I\ D\ u \subseteq R\ u$
⟨*proof*⟩

The next step consists of the definition of a trace set T_c, the corresponding trace set process P_c (cf. [7]), and a reflexive, intransitive noninterference policy I_c for this process, where subscript "c" stands for "counterexample". As event-domain map, the identity function is used, which explains why the policy is defined over events themselves.

datatype $event_c = a_c \mid b_c \mid c_c$

definition $T_c :: event_c\ list\ set$ **where**
$T_c \equiv \{[],$
$[a_c],\ [a_c,\ b_c],\ [a_c,\ b_c,\ c_c],\ [a_c,\ b_c,\ c_c,\ a_c],$
$[b_c],\ [b_c,\ a_c],\ [b_c,\ c_c],\ [b_c,\ a_c,\ c_c]\}$

definition $P_c :: event_c\ process$ **where**
$P_c \equiv \textit{ts-process}\ T_c$

definition $I_c :: (event_c \times event_c)\ set$ **where**
$I_c \equiv \{(a_c,\ a_c),\ (b_c,\ b_c),\ (b_c,\ c_c),\ (c_c,\ c_c),\ (c_c,\ a_c)\}$

Process P_c can be shown to be secure with respect to policy I_c. This result can be obtained by applying the Ipurge Unwinding Theorem, in the version for trace set processes [7], and then performing an exhaustive case distinction over all traces and domains, which obviously is possible by virtue of their finiteness.

Nevertheless, P_c and I_c are such that there exists no domain-relation map satisfying the assumptions of the Generic Unwinding Theorem. A proof *ad absurdum* is given, based on the fact that the pair of traces ($[a_c,\ b_c,\ c_c]$,

$[b_c,\ a_c,\ c_c]$) can be shown to be contained in the image of a_c under the "minimum" domain-relation map *rel-induct*. Therefore, it would also be contained in the image of a_c under a map satisfying the assumptions of the Generic Unwinding Theorem, so that according to weak future consistency, a_c should be a possible subsequent event for trace $[a_c,\ b_c,\ c_c]$ just in case it were such for trace $[b_c,\ a_c,\ c_c]$. However, this conclusion contradicts the fact that a_c is a possible subsequent event for the former trace only.

lemma *counterexample-trace-set*:
trace-set T_c
⟨*proof*⟩

lemma *counterexample-next-events-1*:
$(x \in$ *next-events* $($*ts-process* $T_c)$ *xs*$) = ($*xs* @ $[x] \in T_c)$
⟨*proof*⟩

lemma *counterexample-next-events-2*:
$(x \in$ *next-events* P_c *xs*$) = ($*xs* @ $[x] \in T_c)$
⟨*proof*⟩

lemma *counterexample-secure*:
secure P_c I_c *id*
⟨*proof*⟩

lemma *counterexample-not-gu-condition-aux*:
$([a_c,\ b_c,\ c_c],\ [b_c,\ a_c,\ c_c]) \in$ *rel-induct* P_c I_c *id* a_c
⟨*proof*⟩

lemma *counterexample-not-gu-condition*:
$\neg$ ($\exists$ *R*. *view-partition* P_c *id R* $\wedge$
weakly-future-consistent P_c I_c *id R* $\wedge$
weakly-step-consistent P_c *id R* $\wedge$
locally-respects P_c I_c *id R*)
⟨*proof*⟩

theorem *not-secure-implies-gu-condition*:
$\neg$ (*secure* P_c I_c *id* $\longrightarrow$
($\exists$ *R*. *view-partition* P_c *id R* $\wedge$
weakly-future-consistent P_c I_c *id R* $\wedge$
weakly-step-consistent P_c *id R* $\wedge$
locally-respects P_c I_c *id R*))
⟨*proof*⟩

end

References

[1] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, Inc., 1985.

[2] A. Krauss. *Defining Recursive Functions in Isabelle/HOL*. http://isabelle.in.tum.de/website-Isabelle2015/dist/Isabelle2015/doc/functions.pdf.

[3] T. Nipkow. *A Tutorial Introduction to Structured Isar Proofs*. http://isabelle.in.tum.de/website-Isabelle2011/dist/Isabelle2011/doc/isar-overview.pdf.

[4] T. Nipkow. *Programming and Proving in Isabelle/HOL*, May 2015. http://isabelle.in.tum.de/website-Isabelle2015/dist/Isabelle2015/doc/prog-prove.pdf.

[5] T. Nipkow, L. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, May 2015. http://isabelle.in.tum.de/website-Isabelle2015/dist/Isabelle2015/doc/tutorial.pdf.

[6] P. Noce. Noninterference security in communicating sequential processes. *Archive of Formal Proofs*, May 2014. http://isa-afp.org/entries/Noninterference_CSP.shtml, Formal proof development.

[7] P. Noce. The ipurge unwinding theorem for csp noninterference security. *Archive of Formal Proofs*, June 2015. http://isa-afp.org/entries/Noninterference_Ipurge_Unwinding.shtml, Formal proof development.

[8] J. Rushby. Noninterference, transitivity, and channel-control security policies. Technical report, SRI International, 1992.