Math. Struct. in Comp. Science (2009), *vol.* 19, *pp.* 1065–1090. © Cambridge University Press 2009
doi:10.1017/S0960129509990120

Structural non-interference in elementary and trace nets

NADIA BUSI and ROBERTO GORRIERI

Università di Bologna, Dipartimento di Scienze dell'Informazione, Mura Anteo Zamboni 7,
40127 Bologna, Italy
Email: gorrieri@cs.unibo.it

Received 28 February 2008; revised 2 July 2008

In memory of Nadia Busi

Several notions of non-interference have been proposed in the literature for studying the problem of confidentiality in concurrent systems. The common feature of these non-interference properties is that they are all defined as extensional properties based on some notion of behavioural equivalence on systems. Here, instead, we address the problem of defining non-interference by looking at the structure of the systems under investigation. We use a simple class of Petri nets, namely, contact-free elementary net systems, as the system model and define *structural* non-interference properties based on the absence of particular places in the net: such places show that a suitable causality or conflict relation is present between a high-level transition and a low-level one. We characterise one structural property, called *PBNI+*, which we show to be equivalent to the well-known behavioural property *SBNDC*. It essentially captures all the *positive* information flows (that is, a low-level user can deduce that some high-level action has occurred). We start by providing a characterisation of *PBNI+* on contact-free elementary net systems, then extend the definition to cope with the richer class of trace nets.

1. Introduction

Non-interference has been defined in the literature as an extensional property based on some observational semantics: the high-level part of a system does not interfere with the low-level part if whatever is done at the high level produces no visible effect on the low-level part of the system. The original notion of non-interference in Goguen and Meseguer (1982) was defined, using trace semantics, for system programs that are deterministic. Generalised notions of non-interference were then designed to include (non-deterministic) labelled transition systems and finer notions of observational semantics such as bisimulation (see, for example, Ryan (2001), Focardi and Gorrieri (1995), Roscoe (1995), Ryan and Schneider (1999) and Focardi and Gorrieri (2001)). Relevant properties in this class are the trace-based properties *SNNI* and *NDC*, as well as the bisimulation-based properties *BSNNI*, *BNDC* and *SBNDC* proposed by Focardi and Gorrieri some years ago (Focardi and Gorrieri 1995; Focardi and Gorrieri 2001) on a CCS-like (Milner 1989) process algebra. In particular, *SNNI* states that a system R is secure if $R \setminus H$ (all the high-level actions are prevented) and R/H (all the high-level actions are permitted but are unobservable) are trace equivalent. *BSNNI* is like *SNNI*, but with trace equivalence replaced by weak bisimulation. Intuitively, *NDC* states that a system R is secure if $R \setminus H$ is trace equivalent to R in parallel with any high-level process Π where all the actions in H are restricted (hence cooperation on high-level actions is forced). *BNDC* is the same as *NDC*, but with trace equivalence replaced by weak bisimulation. And *SBNDC* says that a system R is secure if, whenever a high-level action h is performed, the two instances of the system before and after performing h are bisimilar from a low-level point of view.

In the first part of the paper we show that these non–interference properties can also be naturally defined on Petri nets: in particular, to keep the presentation as simple as possible, we use elementary nets (Engelfriet and Rozenberg 1998). The advantage of this proposal is the import into the Petri net theory of security notions, which makes possible the study of security problems. Technically, what we do is to introduce two operations on nets, namely parallel composition (with synchronisation in TCSP-like style (Brookes *et al.* 1984)) and restriction, and suitable notions of observational equivalences on the low-level part of the system (low trace equivalence and low bisimulation). Then, five security properties are defined and compared in a rather direct way. In particular, the two properties based on low trace semantics, namely *SNNI* and *NDC*, are shown to be equivalent. On the other hand, in the bisimulation case, *BSNNI* is strictly weaker than *BNDC*, which, surprisingly, turns out to be equivalent to *SBNDC*. In this approach, the security property is based on the dynamics of systems; these properties are all defined by means of one (or more) equivalence check(s). So non-interference checking is as difficult as equivalence checking, which is a well-studied hard problem in concurrency theory.

In the second part of the paper we address the problem of statically defining non-interference for elementary nets by looking at the structure of the net systems under investigation:

- in order to get a better understanding of the relationship between a flow of information and the causality (or conflict) relation between the activites originating such a flow, hence grounding more firmly the intuition about what is an interference; and
- in order to find more efficiently checkable non-interference properties that are sufficient (sometimes also necessary) conditions for those that have already received some support in the literature.

We define structural non-interference properties based on the absence of particular places in the net. We identify two special classes of places: *causal places*, that is, places for which there are both an incoming high-level transition and an outgoing low-level transition; and, *conflict places*, that is, places for which there are both low- and high-level outgoing transitions. Intuitively, causal places represent a potential source of interference (*hilo* flow for *high input – low output*), because the occurrence of the high-level transition is a prerequisite for the execution of the low-level transition. Similarly, conflict places represent potential source of interference (*holo* flow for *high output – low output*), because if the low-level event is not executable, we can deduce that a certain high-level transition has occurred. The absence of causal and conflict places can be easily checked by a simple inspection of the (finite) net structure; interestingly enough, this is a sufficient condition to ensure *SBNDC*.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

In order to characterise *SBNDC* more precisely, we have to refine the notions of causal place and conflict place slightly, yielding the so-called *active* causal place and *active* conflict place. These new definitions are also based on a limited exploration of the state-space of the net (that is, of its marking graph), so the absence of such places is not a purely structural property, but a hybrid property. Moreover, active causal places and active conflict places capture all the *positive* information flows, that is, the information a low-level user can deduce about the actions the high-level user has performed. We show that when active causal and active conflict places are absent, we get a property, called *Positive Place–Based Non–Interference* (*PBNI+* for short), which turns out to be equivalent to *SBNDC*. More precisely, the main result of the paper states that a net N has no active causal places and no active conflict places if and only if it satisfies *SBNDC*. This result shows that all the positive information flows are captured by *SBNDC*.

In the third part of the paper we extend the definition of *PBNI+* to cope with the richer class of trace nets (Badouel and Darondeau 1995). We provide an example showing how *PBNI+* can be used to capture the information flows arising in a shared variable that can be accessed and modified by both high- and low-level users.

The paper is organised as follows. In Section 2 we recall the basic definitions concerning transition systems and elementary net systems. In Section 3 we recast the behavioural approach to non-interference properties, which was originally defined in a process algebraic setting, on elementary nets. The original structural property *PBNI+* for elementary nets is introduced in Section 4, where the main technical result (*SBNDC* if and only if *PBNI+*) is given. In Section 5, after recalling the basic definitions for trace nets, we extend the definition of *PBNI+* to trace nets. Finally, some concluding remarks are given in Section 6.

2. Basic definitions

Here we recall the basic definitions for transition systems and elementary net systems that we will use in the rest of the paper.

2.1. *Transition systems*

Definition 2.1. A transition system is a triple $TS = (St, E, \rightarrow)$ where:

— St is the set of states

— E is the set of events

— $\rightarrow \subseteq St \times E \times St$ is the transition relation.

In the following we use $s \overset{e}{\rightarrow} s'$ to denote $(s, e, s') \in \rightarrow$. Given a transition $s \overset{e}{\rightarrow} s'$, s is said to be the *source*, s' the *target* and e the *label* of the transition. A *rooted* transition system is a pair (TS, s_0) where $TS = (St, E, \rightarrow)$ is a transition system and $s_0 \in St$ is the *initial state.*

2.2. *Elementary net systems*

Definition 2.2. An *elementary net* is a tuple $N = (S, T, F)$, where

— S and T are the (finite) sets of *places* and *transitions*, with $S \cap T = \varnothing$

— $F \subseteq (S \times T) \cup (T \times S)$ is the flow relation.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

A subset of S is called a *marking*. Given a marking m and a place s, if $s \in m$, we say that the place s contains a token, otherwise we say that s is empty.

Let $x \in S \cup T$. The *preset* of x is the set ${}^{\bullet}x = \{y \mid F(y, x)\}$. The *postset* of x is the set $x^{\bullet} = \{y \mid F(x, y)\}$. The preset and postset functions are generalised in the obvious way to a set of elements: if $X \subseteq S \cup T$, then ${}^{\bullet}X = \bigcup_{x \in X} {}^{\bullet}x$ and $X^{\bullet} = \bigcup_{x \in X} x^{\bullet}$. A transition t is enabled at marking m if ${}^{\bullet}t \subseteq m$ and $t^{\bullet} \cap m = \emptyset$. The firing (execution) of a transition t enabled at m produces the marking $m' = (m \setminus {}^{\bullet}t) \cup t^{\bullet}$. This is usually written as $m[t\rangle m'$. We use the notation $m[t\rangle$ to mean that there exists m' such that $m[t\rangle m'$.

An *elementary net system* is a pair (N, m_0), where N is an elementary net and m_0 is a marking of N, called the *initial marking*. With abuse of notation, we use (S, T, F, m_0) to denote the net system $((S, T, F), m_0)$.

The set of *markings reachable from* m, denoted by $[m\rangle$, is defined as the least set of markings such that:

— $m \in [m\rangle$
— if $m' \in [m\rangle$ and there exists a transition t such that $m'[t\rangle m''$, then $m'' \in [m\rangle$.

The set of *firing sequences* is defined inductively as follows:

— m_0 is a firing sequence;
— if $m_0[t_1\rangle m_1 \ldots [t_n\rangle m_n$ is a firing sequence and $m_n[t_{n+1}\rangle m_{n+1}$, then

$$m_0[t_1\rangle m_1 \ldots [t_n\rangle m_n[t_{n+1}\rangle m_{n+1}$$

is a firing sequence also.

Given a firing sequence $m_0[t_1\rangle m_1 \ldots [t_n\rangle m_n$, we say $t_1 \ldots t_n$ is a *transition sequence*. The set of transition sequences of a net N is denoted by $TS(N)$. We use σ to range over $TS(N)$. Let $\sigma = t_1 \ldots t_n$. We use $m[\sigma\rangle m_n$ as an abbreviation for $m[t_1\rangle m_1 \ldots [t_n\rangle m_n$ and also write $t_i \in \sigma$ to mean transtion t_i occurs in the transition sequence σ.

The *marking graph* of a net system N is the transition system

$$MG(N) = ([m_0\rangle, T, \{(m, t, m') \mid m \in [m_0\rangle \wedge t \in T \wedge m[t\rangle m'\}).$$

A net is *transition simple* if the following condition holds for all $x, y \in T$:

$$\text{if } {}^{\bullet}x = {}^{\bullet}y \text{ and } x^{\bullet} = y^{\bullet}, \text{ then } x = y.$$

A marking m contains a *contact* if there exists a transition $t \in T$ such that ${}^{\bullet}t \subseteq m$ and $\neg(m[t\rangle)$. A net system is *contact free* if no marking in $[m_0\rangle$ contains a contact. A net system is *reduced* if each transition can occur at least once: for all $t \in T$ there exists $m \in [m_0\rangle$ such that $m[t\rangle$. In the following we consider contact-free elementary net systems that are transition simple and reduced.

3. A behavioural approach to non-interference for Petri nets

In this section we recall from Busi and Gorrieri (2004b) some basic definitions of security properties defined on elementary nets, which were originally proposed in Focardi and Gorrieri (1995; 1997; 2001) in a process algebraic setting. Our aim is to analyse systems that can perform two kinds of actions: high-level actions, representing the interaction of

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

the system with high-level users, and low-level actions, representing the interaction with low-level users. We want to verify if the interplay between the high-level user and the high-level part of the system can affect the view of the system as observed by a low-level user. We assume that the low-level user knows the structure of the system, and we check if, in spite of this, he is unable to infer the behaviour of the high-level user by observing the low view of the execution of the system. Hence, we consider nets whose set of transitions is partitioned into two subsets: the set H of high-level transitions and the set L of low-level transitions. To emphasise this partition, we use the following notation. Let L and H be two disjoint sets: we use (S, L, H, F, m_0) to denote the net system $(S, L \cup H, F, m_0)$.

The non-interference properties we are going to introduce are based on some notion of *low* observability of a system, that is, what can be observed of a system from the point of view of low-level users. The low view of a transition sequence is just the subsequence where high-level transitions are discarded.

Definition 3.1. Let $N = (S, L, H, F, m_0)$ be an elementary net system. The *low view* of a transition sequence σ of N is defined as follows:

$$\begin{aligned}\Lambda_N(\varepsilon) &= \varepsilon\\ \Lambda_N(\sigma t) &= \begin{cases}\Lambda_N(\sigma)t & \text{if } t \in L\\ \Lambda_N(\sigma) & \text{otherwise.}\end{cases}\end{aligned}$$

The definition of Λ_N is extended in the obvious way to sets of transition sequences:

$$\Lambda_N(\Sigma) = \{\Lambda_N(\sigma) \mid \sigma \in \Sigma\} \quad \text{for} \quad \Sigma \subseteq (L \cup H)^*.$$

Definition 3.2. Let N_1 and N_2 be two elementary net systems. We say that N_1 is *low-view trace equivalent* to N_2 (denoted $N_1 \overset{\Lambda}{\approx}_{tr} N_2$) if and only if $\Lambda_{N_1}(TS(N_1)) = \Lambda_{N_2}(TS(N_2))$.

We define the operations of parallel composition (in TCSP-like style, see Brookes *et al.* (1984)) and restriction on nets, which will be useful for defining some non-interference properties.

Definition 3.3. Let $N_1 = (S_1, L_1, H_1, F_1, m_{0,1})$ and $N_2 = (S_2, L_2, H_2, F_2, m_{0,2})$ be two net systems such that $S_1 \cap S_2 = \varnothing$ and $(L_1 \cup L_2) \cap (H_1 \cup H_2) = \varnothing$. The parallel composition of N_1 and N_2 is the net system

$$N_1 \mid N_2 = (S_1 \cup S_2, L_1 \cup L_2, H_1 \cup H_2, F_1 \cup F_2, m_{0,1} \cup m_{0,2}).$$

Note that synchronisation occurs over those (low- or high-level) transitions that are shared by the two nets, that is, a transition t that occurs both in N_1 and N_2 has preset (postset), in $N_1 \mid N_2$, given by the union of the disjoint presets (postsets) in N_1 and N_2, respectively. Observe that if N_1 and N_2 are contact free, transition simple and reduced, then $N_1 \mid N_2$ is contact free, transition simple and reduced also.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

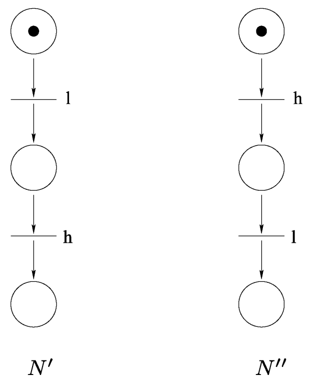

Fig. 1. The net system N' is *SNNI* while N'' is not *SNNI*.

Definition 3.4. Let $N = (S, L, H, F, m_0)$ be a safe net system and let U be a set of transitions. The restriction on U is defined as $N \backslash U = (S, L', H', F', m_0)$, where

$$\begin{aligned} L' &= L \setminus U \\ H' &= H \setminus U \\ F' &= F \setminus (S \times U \cup U \times S) \end{aligned}$$

We can see immediately that if N is contact free, transition simple and reduced, then so is $N \backslash U$.

We will now introduce the first behavioural information-flow security property. *Strong Non-deterministic Non-Interference* (*SNNI* for short) is a trace-based property, which says intuitively that a system is secure if what the low-level part can see does not depend on what the high-level part can do.

Definition 3.5. Let $N = (S, L, H, F, m_0)$ be a net system. We say that N is *SNNI* if and only if $N \overset{\Lambda}{\approx}_{tr} N \backslash H$.

The intuition is that, from the point of view of low-level users, the system where the high-level transitions are prevented should offer the same traces as the system where the high-level transitions can be freely performed. In essence, a low-level user cannot infer, by observing the low view of the system, that some high-level activity has occurred.

As a matter of fact, this non-interference property captures the information flows from high to low, while it admits flows from low to high. For instance, the net N' of Figure 1 is *SNNI*, while the net N'' is not *SNNI* because the low view of the transition sequence *hl* of N'' is l, but no transition sequence l is performable by $N'' \backslash H$.

An alternative notion of non-interference, called *Non-Deducibility on Composition* (or *NDC* for short), says that the low view of a system N in isolation must not be altered when considering each potential interaction of N with the high-level users of the external environment.

Definition 3.6. Let $N = (S, L, H, F, m_0)$ be a net system. We say that N is a high-level net if $L = \emptyset$.

Definition 3.7. Let $N = (S, L, H, F, m_0)$ be a net system. N is *NDC* if and only if for all high-level nets $K = (S_K, \emptyset, H_K, F_K, m_{0,K})$, we have $N \backslash H \overset{\Lambda}{\approx}_{tr} (N \mid K) \backslash (H \setminus H_K)$.

The left-hand term represents the system N when isolated from high-level users (hence, the low view of N in isolation), while the right-hand term expresses the low view of N interacting with the (common transitions of the) high-level environment K (note that the activities resulting from such interactions are invisible by the definition of low-view equivalence). *NDC* is a very intuitive property: whatever high-level system K is interacting with N, the low-level effect is unobservable. However, it is difficult to check this property because of the universal quantification over high-level systems. Luckily enough, we can prove that *SNNI* and *NDC* are actually the same non-interference property.

Theorem 3.8. Let $N = (S, L, H, F, m_0)$ be a net system. N is *SNNI* if and only if N is *NDC*.

Proof.

$NDC \Rightarrow SNNI$:

Take the contact-free high-level net

$$K = (\{s, s'\}, \emptyset, H \cup \{\tau\}, \{s\} \times H \cup H \times \{s'\} \cup \{(s', \tau), (\tau, s)\}, \{s\}),$$

where τ is an auxiliary high-level transition not in H. Then the implication follows by the observation that

$$\Lambda_N(TS((N \mid K) \backslash \emptyset)) = \Lambda_N(TS(N)).$$

$SNNI \Rightarrow NDC$:

This implication follows by the following two observations:

(i) $\Lambda_N(TS(N \backslash H)) \subseteq \Lambda_N(TS((N \mid K) \backslash (H \setminus H_K)))$ for all high-level nets K.

(ii) $\Lambda_N(TS((N \mid K) \backslash (H \setminus H_K))) \subseteq \Lambda_N(TS(N))$ for all high-level nets K. □

The two properties above are based on (low) trace semantics. It is well known (Focardi and Gorrieri 1995; 2001) that bisimulation semantics is more appropriate than trace semantics because it also captures some indirect information flows due to, for example, deadlocks. For this reason, we now consider non-interference properties based on bisimulation. To this end, we first need to introduce a notion of low-view bisimulation.

Definition 3.9. Let $N_1 = (S_1, L_1, H_1, F_1, m_{0,1})$ and $N_2 = (S_2, L_2, H_2, F_2, m_{0,2})$ be two net systems. A *low-view bisimulation* from N_1 to N_2 is a relation on $\mathscr{P}(S_1) \times \mathscr{P}(S_2)$[†] such that if $(m_1, m_2) \in R$, then for all $t \in \bigcup_{i=1,2} L_i \cup H_i$:

— If $m_1[t\rangle m_1'$, then there exist σ, m_2' such that we have $m_2[\sigma\rangle m_2'$, $\Lambda_{N_1}(t) = \Lambda_{N_2}(\sigma)$ and $(m_1', m_2') \in R$.

— If $m_2[t\rangle m_2'$, then there exist σ, m_1' such that we have $m_1[\sigma\rangle m_1'$, $\Lambda_{N_2}(t) = \Lambda_{N_1}(\sigma)$ and $(m_1', m_2') \in R$.

[†] We use $\mathscr{P}(S)$ to denote the powerset of S.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

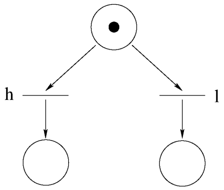

Fig. 2. A net system that is *SNNI* but not *BSNNI*.

If $N_1 = N_2$, we say that R is a low-view bisimulation on N_1.

We say that N_1 is *low-view bisimilar* to N_2, denoted $N_1 \overset{\Lambda}{\approx}_{bis} N_2$, if there exists a low-view bisimulation R from N_1 to N_2 such that $(m_{0,1}, m_{0,2}) \in R$.

We also say that two markings m and m' of the net N_1 are low-view bisimilar if there exists a low-view bisimulation R on N_1, that is, on $\mathscr{P}(S_1) \times \mathscr{P}(S_1)$, such that $(m, m') \in R$.

The first obvious variation on the theme is to define the bisimulation-based version of *SNNI*, yielding *BSNNI*.

Definition 3.10. Let $N = (S, L, H, F, m_0)$ be a net system. We say that N is *BSNNI* if and only if $N \overset{\Lambda}{\approx}_{bis} N \backslash H$.

Obviously, *BSNNI* $\subseteq$ *SNNI*. The converse is not true: the net N in Figure 2 is *SNNI* but not *BSNNI*. Note that *SNNI* fails to capture the indirect information flow present in this net: if the low-level transition l cannot be performed, the low-level user can infer that the high-level transition h has been performed, hence deducing one piece of (positive) high-level knowledge.

Similarly, *BNDC* can be defined from *NDC*, yielding a rather appealing security property, which is finer than *BSNNI*.

Definition 3.11. Let $N = (S, L, H, F, m_0)$ be a net system. N is *BNDC* if and only if for all high-level nets $K = (S_K, \emptyset, H_K, F_K, m_{0,K})$, we have $N \backslash H \overset{\Lambda}{\approx}_{bis} (N \mid K) \backslash (H \setminus H_K)$.

Theorem 3.12. Let $N = (S, L, H, F, m_0)$ be a net system. If N is *BNDC*, then N is *BSNNI*.

Proof. Take the contact-free high-level net K defined in the proof of Theorem 3.8. Then the implication follows by the obvious observation that $(N \mid K) \backslash \emptyset \overset{\Lambda}{\approx}_{bis} N$. □

Unfortunately, the converse is not true. Figure 3 shows a net that is *BSNNI* but not *BNDC* ; it is easy to see why this is the case by looking at the respective marking graphs in Figure 4.

BNDC is, intuitively, quite appealing, but it is difficult to check: one has to perform the bisimulation check against all possible high-level systems, and, in principle, there are infinitely many of them. The next property, called *Strong Bisimulation Non-Deducibility on Composition* (*SBNDC* for short), is actually an alternative characterisation of *BNDC* and is more easily checkable.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

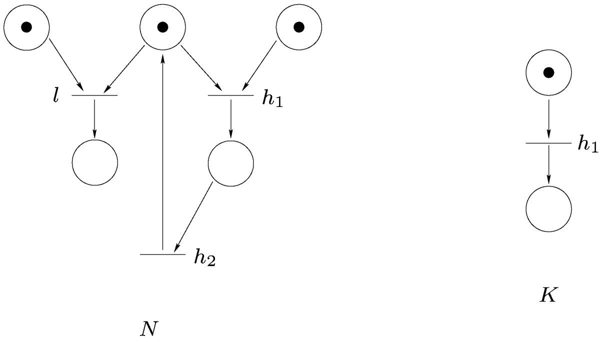

Fig. 3. A net system that is *BSNNI* but not *BNDC*.

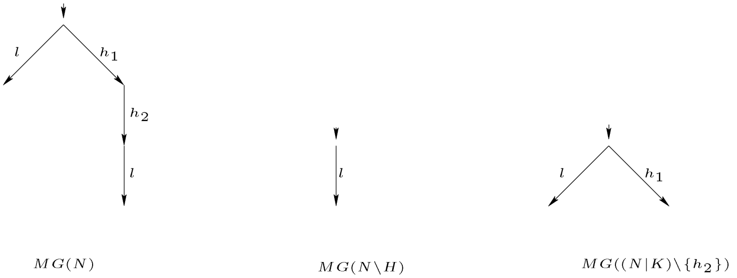

Fig. 4. The marking graphs of the net systems N, $N\backslash H$ and $(N \mid K)\backslash\{h_2\}$.

Definition 3.13. Let $N = (S, L, H, F, m_0)$ be a net system. N is *SBNDC* if and only if for all markings $m \in [m_0\rangle$ and for all $h \in H$, the following holds:
if $m[h\rangle m'$, there exists a low-view bisimulation R on $N\backslash H$ such that $(m, m') \in R$.

The intuition behind *SBNDC* is that, whenever a high-level transition h is performed, the markings before h and after h are observationally indistinguishable for a low-level observer. Note that *SBNDC* is clearly decidable for (finite) elementary net systems because the number of reachable markings is finite, as well as the set H of high-level transitions. We now prove in two steps that *SBNDC* is indeed an alternative characterisation of *BNDC*.

Theorem 3.14. Let $N = (S, L, H, F, m_0)$ be a net system. If N is *BNDC*, then N is *SBNDC*.

Proof. Suppose N is *BNDC*. Let $m \in [m_0\rangle$ and $h \in H$ be such that $m[h\rangle m'$. We show that there exists a low-view bisimulation on $N\backslash H$ that contains the pair (m, m').

We need an auxiliary function $\mathscr{H}_N$, called *high-view* and defined in a very similar way to the low-view function Λ_N, that extracts from a transition sequence σ the subsequence of its high-level transitions. Let σ be a transition sequence such that $m_0[\sigma\rangle m$. Let $\mathscr{H}_N(\sigma) = h_1 \dots h_n$. Note that the sequence $h_1 \dots h_n$ may contain repeated elements. We construct the high-level net $K_1 = (S_1, \emptyset, H_1, F_1, m_{01})$, shown in Figure 5, where:

— $S_1 = \bigcup_{i=1}^{n+1}\{preen_i\} \cup \bigcup_{i=1}^{n+1}\{posten_i\} \cup \{end\}$

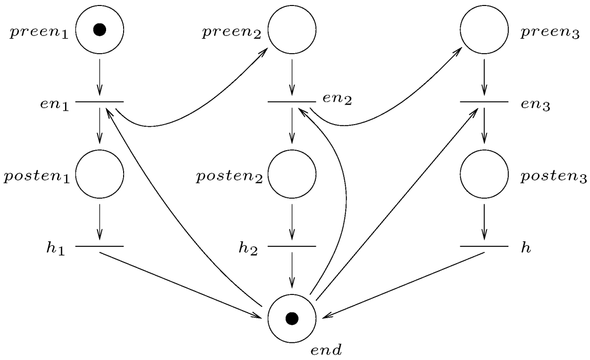

Fig. 5. The net K_1 used in the proof of Theorem 3.14, for the case $n = 2$.

— $H_1 = \bigcup_{i=1}^{n+1}\{en_i\} \cup \{h_1,\ldots,h_n\} \cup \{h\}$.

— $F_1 = \bigcup_{i=1}^{n+1}\{(preen_i, en_i)\} \cup \bigcup_{i=1}^{n+1}\{(end, en_i)\} \cup \bigcup_{i=1}^{n+1}\{(en_i, posten_i)\} \cup \bigcup_{i=1}^{n}\{(en_i, preen_{i+1})\} \cup \bigcup_{i=1}^{n}\{(posten_i, h_i)\} \cup \{(posten_{n+1}, h)\} \cup \bigcup_{i=1}^{n}\{(h_i, end)\} \cup \{(h, end)\}$

— $m_{01} = \{preen_1, end\}$.

As N is *BNDC*, we have that $(N\backslash H) \overset{\Lambda}{\approx}_{bis} ((N|K_1)\backslash(H \setminus H_1))$. Hence, there exists a low-view bisimulation R_1 from $(N\backslash H)$ to $((N|K_1)\backslash(H \setminus H_1))$. Hence, $(m_0, m_0 \cup m_{01}) \in R_1$.

Moreover, it is easy to see that the only (maximal) transition sequence that can be performed by K_1 is $en_1h_1en_2h_2\ldots en_{n+1}h$, leading to marking $\{end\}$. Take $\sigma_1,\ldots,\sigma_{n+1}$ such that $\sigma = \sigma_1h_1\sigma_2h_2\ldots\sigma_nh_n\sigma_{n+1}$. As $m_0[\sigma\rangle m[h\rangle m'$, we have

$$m_0 \cup m_{01}[\sigma'\rangle m \cup \{end, preen_{n+1}\}[en_{n+1}h\rangle m' \cup \{end\}$$

in $((N|K_1)\backslash(H \setminus H_1))$, where $\sigma' = \sigma_1en_1h_1\sigma_2en_2h_2\ldots\ en_nh_n\sigma_{n+1}$. Note that $\Lambda_N(\sigma') = \Lambda_N(\sigma' en_{n+1}h)$.

As $(m_0, m_0 \cup m_{01}) \in R_1$ and R_1 is a low-view bisimulation, we have that there exists $\bar{m}$ such that $m_0[\Lambda_N(\sigma')\rangle\bar{m}$, such that $(\bar{m}, m' \cup \{end\})$ belong to R_1. As the net K_1 with marking $\{end\}$ can perform no transition, it is easy to see that the relation $R_1' = \{(m_1, m_2) \mid (m_1, m_2 \cup \{end\}) \in R_1\}$ is a low-view bisimulation on $N\backslash H$. Moreover, as $(\bar{m}, m' \cup \{end\})$ belongs to R_1, we get that $(\bar{m}, m')$ belongs to R_1'.

Now we take the net $K_2 = (S_2, \emptyset, H_2, F_2, m_{02})$ obtained from K_1 by removing transitions h and en_{n+1} and the corresponding places, that is,

— $S_2 = S_1 \setminus \{preen_{n+1}, posten_{n+1}\}$

— $H_2 = H_1 \setminus \{en_{n+1}, h\}$

— $F_2 = F_1 \cap ((S_2 \times H_2) \cup (H_2 \times S_2))$

— $m_{02} = m_{01}$.

Following a reasoning similar to that performed for net K_1, we get that there exists a low-view bisimulation R_2 from $(N\backslash H)$ to $((N|K_2)\backslash(H\backslash H_2))$ with $(m_0, m_0\cup m_{02}) \in R_2$. As $m_0[\sigma\rangle m$, we have $m_0\cup m_{02}[\sigma'\rangle m\cup\{end\}$ in $((N|K_2)\backslash(H\backslash H_2))$, where $\sigma' = \sigma_1en_1h_1\sigma_2en_2h_2\ldots en_nh_n\sigma_{n+1}$.

As $(m_0, m_0 \cup m_{02}) \in R_2$, R_2 is a low-view bisimulation and we have seen above that the marking reached from m_0 after firing $\Lambda_N(\sigma')$ is $\bar{m}$, we have $(\bar{m}, m \cup \{end\}) \in R_2$. As the

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

net K_2 with marking $\{end\}$ can perform no transition, it is easy to see that the relation $R_2' = \{(m_1, m_2) \mid (m_1, m_2 \cup \{end\}) \in R_2\}$ is a low-view bisimulation on $N \backslash H$. Moreover, as $(\bar{m}, m \cup \{end\})$ belongs to R_2, we get that $(\bar{m}, m)$ belongs to R_2'. Thus, we have $(\bar{m}, m') \in R_1'$ and $(\bar{m}, m) \in R_2'$. It is easy to see that if R is a low-view bisimulation, then R^{-1} is a low-view bisimulation and that if R and S are low-view bisimulations, then $R \circ S$ is a low-view bisimulation. Let $R = (R_2')^{-1} \circ R_1'$. We have that R is a low-view bisimulation on $N \backslash H$ and $(m, m') \in R$. Hence, N is *SBNDC*. □

Theorem 3.15. Let $N = (S, L, H, F, m_0)$ be a net system. If N is *SBNDC*, then N is *BNDC*.

Proof. Suppose that N is *SBNDC*. Let $K = (S_k, \emptyset, H_k, F_k, m_{0k})$ be a high-level net. To show that N is *BNDC*, we need to provide a low-view bisimulation from $N \backslash H$ to $(N|K) \backslash (H \setminus H_k)$.

Let $R = \{(m, \bar{m} \cup m_k \mid m, \bar{m} \in [m_0\rangle, m_k \in [m_{0k}\rangle$ and let there exist a low-view bisimulation R' on $N \backslash H : (m, \bar{m}) \in R'\}$. We want to prove that R is the required low-view bisimulation. We have that $(m_0, m_0 \cup m_{0k}) \in R$.

Take $(m, \bar{m} \cup m_k) \in R$.

If $m[l\rangle m'$, as there exists a low-view bisimulation R' containing $(m, \bar{m})$, then there exists $\bar{m}'$ such that $\bar{m}[l\rangle \bar{m}'$ and $(m', \bar{m}') \in R'$. From $\bar{m}[l\rangle \bar{m}'$, we get $\bar{m} \cup m_k[l\rangle \bar{m}' \cup m_k$. From $(m', \bar{m}') \in R'$, we get $(m', \bar{m}' \cup m_k) \in R$. Conversely, if $\bar{m} \cup m_k[l\rangle m_1$, as K contains only high-level transitions, there exists $\bar{m}'$ such that $\bar{m}[l\rangle \bar{m}'$ and $m_1 = \bar{m}' \cup m_k$. As there exists R' such that $(m, \bar{m}) \in R'$, there exists m' such that $m[l\rangle m'$ and $(m', \bar{m}') \in R'$. Hence, following the same reasoning as above, we have $(m', \bar{m}' \cup m_k) \in R$.

If $\bar{m} \cup m_k[h\rangle m_1$, there exist $\bar{m}'$ and m_k' such that $\bar{m}[h\rangle \bar{m}'$, $m_k[h\rangle m_k'$ and $m_1 = \bar{m}' \cup m_k'$. By *SBNDC*, from $\bar{m}[h\rangle \bar{m}'$, we get that there exists a low-view bisimulation R'' such that $(\bar{m}, \bar{m}') \in R''$. As, by hypothesis, $(m, \bar{m}) \in R'$, we have $(m, \bar{m}') \in R' \circ R''$, which is also a low-view bisimulation. Hence, we get $(m, \bar{m}' \cup m_k') \in R$. (As $\Lambda_N(h) = \varepsilon$, we also have $m[\Lambda_N(h)\rangle m$.)

Hence, R is a low-view bisimulation from $N \backslash H$ to $(N|K) \backslash (H \setminus H_k)$. □

The two theorems above give the following corollary.

Corollary 3.16. Let $N = (S, L, H, F, m_0)$ be a net system. N is *BNDC* if and only if N is *SBNDC*.

The above corollary holds because we are in an unlabelled setting, that is, transitions are not labelled. Focardi and Gorrieri (1995; 2001) proved that, for the Security Process Algebra, *SBNDC* is strictly finer than *BNDC*. Nonetheless, it is an interesting observation that two properties that are defined in a radically different way actually coincide in this setting.

4. A structural approach to positive non-interference

We defined two notions of non-interference in Busi and Gorrieri (2004a; 2004b), namely, *PBNI* and *RBNI*, which were aimed at capturing any kind of information flow from high-level users to low-level users. Those notions capture both positive and negative

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

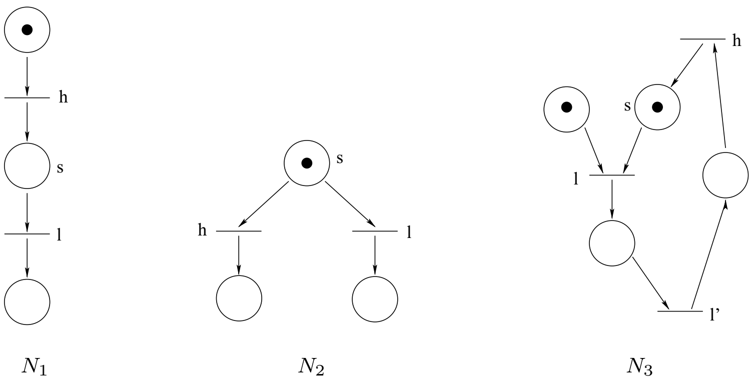

Fig. 6. Examples of net systems containing conflict and (potentially) causal places.

pieces of information on the high-level behaviour of the system. More precisely, a *positive* information flow arises when the occurrence of a high-level transition can be deduced from the low-level behaviour of the system, whereas a *negative* information flow is concerned with the deduction that a high-level transition has *not* occurred.

In this section we provide a characterisation of positive information flows, that is, we consider a system secure if it is not possible to deduce that some high-level action has been performed by observing the low-level behaviour. In this sense, the behavioural properties studied in the previous section are able to capture positive (and also some negative) information flows. However, we still need to determine if all the positive information flows are captured. With this end in mind, we propose the structural *PBNI+* property, which is based on the absence of some kinds of places in a net system and captures all the positive information flows in a rather intuitive way.

Consider a net system $N = (S, L, H, F, m_0)$. Consider a low-level transition l of the net: if l can fire, we know that the places in the preset of l are marked before the firing of l; moreover, we know that such places become unmarked after the firing of l. If there exists a high-level transition h that produces a token in a place s in the preset of l (see the system N_1 in Figure 6), then the low-level user can infer that h has occurred if he can perform the low-level transition l. We note that there exists a causal dependency between the transitions h and l because the firing of h produces a token that is consumed by l. Hence, the occurrence of l gives a piece of positive high-level information (the execution of h) to a low-level user.

Consider now the situation illustrated in the system N_2 of Figure 6: in this case, place s is in the preset of both l and h, that is, l and h are competing for the use of the resource represented by the token in s. Aware of the existence of such a place, a low-level user knows that the high-level action h has been performed if he is not able to perform the low-level action l, hence deducing a piece of positive information about the high-level behaviour. Place s represents a conflict between transitions l and h, because the firing of h prevents l from firing.

Our idea is to consider a net system secure if it does not contain places of the kinds illustrated above.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

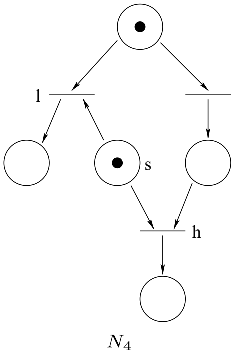

Fig. 7. A net system containing a potentially conflict place but no active conflict places.

In order to avoid the definition of a security notion that is too strong, and that rules out systems that do not reveal information about the high-level actions that have been performed, we need to refine the concepts illustrated above. Consider, for instance, the net system N_3 shown in Figure 6. Although s is a potentially causal place, the net system has to be considered secure, as the (unique) possible firing of l is at the initial marking, hence it is not caused by h. For s to be a source of information on the occurrence of h, there must exist a firing sequence where l consumes a token produced by h. In other words, s is an active causal place if there exists a path in $MG(N_3)$ connecting (an occurrence of) h to (an occurrence of) l, such that the transitions occurring after h and before l do not produce tokens in s.

As far as conflicts are concerned, consider the net system N_4 shown in Figure 7. At first sight, the net N_4 might appear to be insecure because of the presence of the conflict place s. However, we note that the occurrence of h has no effect on the low-level behaviour of the system, as the possibility of firing l has already been ruled out by the firing of transition l'. Hence, for s to be a source of information about the occurrence of h, there must exist a reachable marking where the firing of h rules out the possibility of firing (immediately or after some other transitions) l. In other words, s is an active conflict place if there exists a path in $MG(N_4)$ connecting the source of (an occurrence of) h to (an occurrence of) l such that the transitions occurring in the path do not produce tokens in s.

Definition 4.1. Let $N = (S, L, H, F, m_0)$ be an elementary net system. Let s be a place of N such that $s^\bullet \cap L \neq \emptyset$ (that is, a token in s can be consumed by a low-level transition).

The place $s \in S$ is a *potentially causal place* if ${}^\bullet s \cap H \neq \emptyset$ (that is, a token in s can be produced by a high-level transition). A potentially causal place s is an *active causal place* if the following condition holds: there exist $l \in s^\bullet \cap L$, $h \in {}^\bullet s \cap H$, $m \in [m_0\rangle$ and a transition sequence σ such that $m[h\sigma l\rangle$ and $s \notin t^\bullet$ for all $t \in \sigma$.

The place $s \in S$ is a *potentially conflict place* if $s^\bullet \cap H \neq \emptyset$ (that is, the token in s can also be consumed by a high-level transition). A potentially conflict place is an *active conflict place* if the following condition holds: there exist $l \in s^\bullet \cap L$, $h \in s^\bullet \cap H$, $m \in [m_0\rangle$ and a transition sequence σ such that $m[h\rangle$, $m[\sigma l\rangle$ and $s \notin t^\bullet$ for all $t \in \sigma$.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

It is quite convincing that the presence of an active causal place, such as place s in net N_1 of Figure 6, may determine the only *causal* positive information flow that exists from high to low: if l is performed, then h has also been performed. Note that it is not possible to define a negative causal information flow: in no way, by performing (or not) the low-level transition l, we can deduce the fact that the high-level transition h has *not* been performed.

It is quite convincing that the presence of an active conflict place, such as place s in net N_2 of Figure 6, may determine the only *conflict* positive information flow that exists from high to low: if l is not executable, then h has been performed. Observe that the definition of active conflict place is correct because it implicitly uses the assumption of contact-freeness: no transition in σ can generate a contact for h. This ensures that if l after σ is not executable, then h has been performed[†]. Note, however, that a negative information flow can also be derived: if l has been performed, then h cannot occur. Note also that if the role of h and l are exchanged in the definition of active conflict place (that is, $m[l\rangle$, $m[\sigma h\rangle$), we get a definition that may only reveal a negative information flow: if l is performed, then h cannot occur, because the non-occurrence of l says nothing about the possible firing of h (if σ is non-empty). Since we are interested in capturing positive information flows only, we will ignore this additional source of information flows.

Definition 4.2. Let $N = (S, L, H, F, m_0)$ be an elementary net system. We say that N is *PBNI+* (*positive Place-Based Non-Interference*) if, for all $s \in S$, s is neither an active causal place nor an active conflict place.

We have that the absence of both active causal and active conflict places is a necessary and sufficient condition for *SBNDC*, thereby showing that the behavioural property *SBNDC* captures all the positive information flows. This result, which represents the main technical achievement of this paper, is proved in two steps.

Theorem 4.3. Let $N = (S, L, H, F, m_0)$ be an elementary net system. If N is *PBNI+*, then N is *SBNDC*.

Proof. Let N be *PBNI+*. We will show that N is *SBNDC*. Take $m \in [m_0\rangle$ such that $m[h\rangle m'$ for some $h \in H$. We have to prove that there exists a low-view bisimulation R on $N \backslash H$ such that $(m, m') \in R$.

Let

$$\begin{aligned} R = \{(m_1, m_2) \mid & \forall l \in L\ \forall s \in {}^\bullet l : \\ & m_1(s) \neq m_2(s) \Rightarrow (\forall \sigma \forall i \in \{1,2\} : m_i[\sigma l\rangle \Rightarrow \exists l_1 \in \sigma : s \in l_1^\bullet)\} \end{aligned}$$

be the candidate relation.

— We show first that R is a low-view bisimulation on $N \backslash H$.
Let $(m_1, m_2) \in R$. Suppose $m_1[l\rangle m_1'$. We show that $m_2[l\rangle$ also. Suppose there exists $s \in {}^\bullet l$ such that $m_2(s) = 0$; hence, $m_1(s) \neq m_2(s)$. As $(m_1, m_2) \in R$ and $m_1[l\rangle$, by definition

[†] Note that transitions in σ may disable h simply by consuming other tokens needed for h to fire, but in such a case, after σ, l will be always executable.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

of R (with $\sigma = \varepsilon$), there must exist $t \in \varepsilon$, which gives a contradiction. Hence, $\forall s \in {}^\bullet l$ $m_2(s) \geqslant 1$, so there exists m_2' such that $m_2[l\rangle m_2'$.

Now we show that $(m_1', m_2') \in R$. Suppose $(m_1', m_2') \notin R$. Then there exist l', $s' \in {}^\bullet l'$ such that $m_1'(s') \neq m_2'(s')$, and there exists σ and i such that $m_i'[\sigma l'\rangle$ and $s' \in l_1^\bullet$ for no $l_1 \in \sigma$. As $m_i[l\rangle m_i'$ for $i = 1, 2$, we have $m_1(s') \neq m_2(s')$ and there exists $i \in \{1, 2\}$ such that $m_i[l\sigma l'\rangle$ and $s' \in l_1^\bullet$ for no $l_1 \in \sigma$. Since $m_1'(s') \neq m_2'(s')$, necessarily $s' \notin l^\bullet$, hence $s' \in l_1^\bullet$ for no $l_1 \in l\sigma$. Thus we obtain $(m_1, m_2) \notin R$, which gives a contradiction. Hence we have $(m_1', m_2') \in R$.

The symmetric case can be proved in the same way, which shows that R is a low-view bisimulation on $N \backslash H$.

— We now show that $(m, m') \in R$.

Suppose there exists s and $l \in s^\bullet$ such that $m(s) \neq m'(s)$. We show that $\forall \sigma : m[\sigma l\rangle \Rightarrow \exists t \in \sigma : s \in t^\bullet$ and $\forall \sigma : m'[\sigma l\rangle \Rightarrow \exists t \in \sigma : s \in t^\bullet$.

As $m[h\rangle m'$, we can deduce from $m(s) \neq m'(s)$ that one of the following holds:

– $s \in h^\bullet$.

Hence, s is a potentially causal place.

Take a sequence σ such that $m[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that $s \in t^\bullet$. There are two possible subcases:

(i) $\sigma = h\sigma'$.

As *PBNI+* holds, s is not an active causal place. Hence, for all $\bar{m} \in [m_0\rangle$ and for all $\bar{\sigma}$, we have if $\bar{m}[h\bar{\sigma} l\rangle$, there exists $t \in \bar{\sigma}$ such that $s \in t^\bullet$. As $m[\sigma l\rangle$ and $\sigma = h\sigma'$, there exists $t \in \sigma'$ such that $s \in t^\bullet$.

(ii) $\sigma = \varepsilon$ or $\sigma = t'\sigma'$ with $t' \neq h$.

As $s \in h^\bullet$, we have $m(s) = 0$. As $m[\sigma l\rangle$ and $s \in {}^\bullet l$, there must exist a transition $t \in \sigma$ that produces one token in s, that is, such that $s \in t^\bullet$. In particular, $\sigma \neq \varepsilon$.

Now consider a sequence σ such that $m'[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that $s \in t^\bullet$. As $m[h\rangle m'$, we have $m[h\sigma l\rangle$, hence, because *PBNI+* holds, there exists $t \in \sigma$ such that $s \in t^\bullet$.

– $s \in {}^\bullet h$.

So s is a potentially conflict place.

Take a sequence σ such that $m[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that $s \in t^\bullet$. As *PBNI+* holds, s cannot be an active conflict place. Hence, for all $\bar{m} \in [m_0\rangle$ and for all $\bar{\sigma}$, we have if $\bar{m}[h\rangle$ and $\bar{m}[\bar{\sigma} l\rangle$, there exists $t \in \bar{\sigma}$ such that $s \in t^\bullet$.

As $m[h\rangle m'$ and $m[\sigma l\rangle$, there exists $t \in \sigma$ such that $s \in t^\bullet$.

Now take a sequence σ such that $m'[\sigma l\rangle$. As $s \in {}^\bullet h$, we have $m'(s) = 0$. As $s \in {}^\bullet l$, from $m'[\sigma l\rangle$ we have that there must exist a transition $t \in \sigma$ producing one token in s, that is, $s \in t^\bullet$. □

Theorem 4.4. Let $N = (S, L, H, F, m_0)$ be an elementary net system. If N is *SBNDC*, then N is *PBNI+*.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

Proof. Suppose N is *SBNDC*. We show that no place in N can be an active causal place or an active conflict one:

— Suppose first that s is an active causal place.
Then, there exist $h \in {}^\bullet s$, $l \in s^\bullet$, $m \in [m_0\rangle$ and σ such that $m[h\sigma\ l\rangle$ and $\forall t \in \sigma : s \notin t^\bullet$. Among the markings and the transition sequences that satisfy the conditions above, take m and σ such that σ contains the minimum number of transitions in H.
There are two cases to consider:

  - All transitions in σ belong to L.
  We have $m[h\rangle m'$. By *SBNDC*, there exists a low-view bisimulation on $N\backslash H$ containing the pair (m, m'). As $m'[\sigma l\rangle$, we also have $m[\sigma l\rangle$. But from $h \in {}^\bullet s$ and $m[h\rangle$, we deduce that $s \notin m$. We also know that $\forall t \in \sigma : s \notin t^\bullet$. So, after the firing of σ, place s is still empty, contradicting the fact that $m[\sigma l\rangle$.
  - There exists a high-level transition in σ.
  Let h' be the last high-level transition in σ. Hence, there exist σ_1, σ_2 such that $\sigma = \sigma_1 h' \sigma_2$ and all transitions in σ_2 belong to L. Thus, there exist m_1, m_2 such that $m[h\sigma_1\rangle m_1[h'\rangle m_2[\sigma_2 l\rangle$.
  From $m_1[h'\rangle m_2$, by *SBNDC* there exists a low-view bisimulation on $N\backslash H$ containing the pair (m_1, m_2). From $m_2[\sigma_2 l\rangle$, we also get that $m_1[\sigma_2 l\rangle$, thus obtaining the firing sequence $m[h\sigma_1\sigma_2 l\rangle$, which contradicts the assumption that the chosen transition sequence was the one with the least number of high-level transitions.

— Now suppose s is an active conflict place.
The proof proceeds in an analogous way to the case above. □

Corollary 4.5. Let $N = (S, L, H, F, m_0)$ be an elementary net system. Then N is *PBNI+* if and only if N is *SBNDC*.

An obvious consequence is that if N has no *potentially causal* or *potentially conflict* places, then N is *SBNDC*. Hence, a simple strategy for checking if N is *SBNDC* is to first identify potential causal/conflict places, a procedure that is linear in the size of the net[†]. If no place of these sorts is found, then N is *PBNI+* . Otherwise, any such candidate place should be further studied to check if it is actually an *active* causal/conflict place, a procedure that requires a limited exploration of the marking graph[‡].

Moreover, we state that *PBNI+* is a compositional property with respect to parallel composition and restriction.

Theorem 4.6. Let $N_i = (S_i, L_i, H_i, F_i, m_{0,i})$ $(i = 1, 2)$ be two *PBNI+* net systems and let $U \subseteq L$ be a set of low-level transitions. Then, $N_1 \mid N_2$ is *PBNI+* as well as $N_1\backslash U$.

[†] In the worst case, the identification of the potentially causal or conflict places is $O(n \times m)$, where n is the number of places and m is the number of transitions. In the average case, this is $O(n \times k)$ where k is the maximum number of transitions connected to a place (hence k is typically rather small).

[‡] Efficient algorithms for checking if potential places are active are beyond the scope of this paper. However, they will have to cope with the nature of the marking graph, which is exponential in the number of places. Initial work in this direction has been reported in Frau (2008).

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

This compositionality principle can be helpful in proving the security of a large net N obtained as the composition of many subnets: $N = N_1 \mid \ldots \mid N_k$. As a matter of fact, if we prove that all the k subnets N_i are *PBNI+*, we can derive for free that N is *PBNI+* also. This may be of little use in checking potentially causal/conflict places, because of the linearity of this check. However, it actually is of great help when exploration of the marking graph is needed to check if a potential causal/conflict place is actually an active one, because in this case the problem of state space explosion can be kept under control. On the other hand, if any of the N_i is not *PBNI+*, we cannot conclude that N is not *PBNI+*, hence in this case we are forced to analyze N as a whole.

We conclude this section with an observation that gives a better indication of the different discriminating power of trace semantics and bisimulation semantics in the definition of security properties; this observation shows that the absence of active causal places is enough to ensure *NDC* (which is the same as *SNNI*), hence showing that information flows due to conflicts are completely ignored by the trace-based version of *BNDC* (which is the same as *SBNDC*).

Lemma 4.7. Let $N = (S, L, H, F, m_0)$ be a net system without active causal places. If $m_0[\sigma\rangle m_1$, then there exists m_2 such that $m_0[\Lambda_N(\sigma)\rangle m_2$ and $m_2(s) \geqslant m_1(s)$ for all $s \in {}^\bullet L$ such that there exist $l \in L$, $s \in {}^\bullet l$ and $\bar{\sigma} \in TS(N)$ such that $m_1[\bar{\sigma} l\rangle$ and $s \notin t^\bullet$ for all $t \in \bar{\sigma}$.

Proof. We use induction on the length of σ. For simplicity, instead of repeating the formal condition on s (namely, for all $s \in {}^\bullet L$ such that $\exists l \in L$, $s \in {}^\bullet l$, $m_1[\bar{\sigma} l\rangle$ and $s \notin t^\bullet$ for all $t \in \bar{\sigma}$), we will write 'for all s that can be consumed by a low-level transition l' (in some cases l will be assumed).

The case $\sigma = \varepsilon$ is trivial.

If $\sigma = \sigma' t$, there exists m_1' such that $m_0[\sigma'\rangle m_1'$ and $m_1'[t\rangle m_1$. By the induction hypothesis, there exists m_2' such that $m_0[\Lambda_N(\sigma')\rangle m_2'$ and $m_2'(s) \geqslant m_1'(s)$ for all the places s that can be consumed by a low-level transition l. There are two cases to consider:

— If $t \in L$, as $m_1'[t\rangle m_1$ by definition of firing (in a contact-free net), we have ${}^\bullet t \subseteq m_1'$. As, by the induction hypothesis, $m_2'(s) \geqslant m_1'(s)$ for all $s \in {}^\bullet t$, we have ${}^\bullet t \subseteq m_2'$. So $m_2'[t\rangle m_2$ and $m_2 = m_2' \setminus {}^\bullet t \cup t^\bullet$. Since $m_2'(s) \geqslant m_1'(s)$ for all s that can be consumed by a low-level transition, and since $m_2 = m_2' \setminus {}^\bullet t \cup t^\bullet$, and since $m_1 = m_1' \setminus {}^\bullet t \cup t^\bullet$, we have $m_2(s) \geqslant m_1(s)$ for all s that can be consumed by a low-level transition. Thus, $m_0[\Lambda_N(\sigma' t)\rangle m_2$ and $m_2(s) \geqslant m_1(s)$ for all s that can be consumed by a low-level transition.

— If $t \in H$, by $m_1'[t\rangle m_1$ and the definition of firing, we have $m_1(s) = m_1'(s) - {}^\bullet t(s) + t^\bullet(s)$ for all $s \in S$. As N has no active causal places and $t \in H$, we have $t^\bullet(s) = 0$ for all s that can be consumed by a low-level transition l. Hence, $m_1'(s) \geqslant m_1(s)$ for all such s. By the induction hypothesis, $m_2'(s) \geqslant m_1'(s)$ for all s that can be consumed by a low-level transition l, hence $m_2'(s) \geqslant m_1(s)$ for all such s. As $\Lambda_N(\sigma' t) = \Lambda_N(\sigma')$, we have $m_0[\Lambda_N(\sigma' t)\rangle m_2'$, with $m_2'(s) \geqslant m_1(s)$ for all s that can be consumed by a low-level transition l. □

Corollary 4.8. Let $N = (S, L, H, F, m_0)$ be a net system. If N has no active causal places, then N is *SNNI*.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

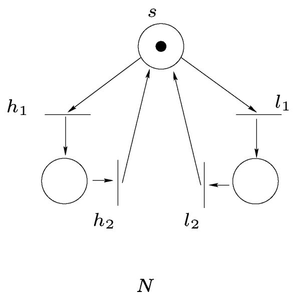

Fig. 8. A *BSNNI* net with an active causal place.

The converse of this corollary does not hold: the net N in Figure 8 is *SNNI*, in fact it is also *BSNNI*, but place s is an active causal place (and also a conflict place). The same argument applies also to the net N in Figure 3, which is *BSNNI* but the marked place in the middle is an active causal place (and also an active conflict place).

5. Non-interference in trace nets

In this section we extend the definition of *PBNI+* to cope with the richer class of trace nets (Badouel and Darondeau 1995) and show that the results presented in the previous section for elementary nets also hold in this setting. Finally, we provide an example to show how our property can be used to capture the information flows arising in a shared variable that can be accessed and modified by both high- and low-level users.

5.1. *Trace nets*

Trace nets (Badouel and Darondeau 1995) are an extension of elementary nets: in addition to the classical flow arcs, we add arcs permitting the testing for presence/absence of tokens in a place, and arcs permitting the filling/emptying of a place regardless of its previous contents.

Definition 5.1. A *trace net* is a tuple $N = (S, T, W)$, where:

— S and T are the (finite) sets of *places* and *transitions*, with $S \cap T = \emptyset$.
— $W : (S \times T) \to \{in, out, nop, read, inhib, set, reset\}$ is the flow function, with $\forall t \in T \exists s \in S : W(s,t) \neq nop$.

The arcs of kind *in* and *out* correspond to the flow arcs of elementary nets: more precisely, a flow arc from a place s to a transition t is represented in a trace net by setting $W(s,t) = in$, while a flow arc from t to s is represented by setting $W(s,t) = out$. The arcs of kind *read* and *inhib* permit us to test a condition on a place, without altering its contents. A read (respectively, inhibitor) arc from s to t requires that s contains a token (respectively, no tokens) for t to fire. The arcs of kind *set* and *reset* permit us to set the contents of the place to a given value, independently of the previous contents of the place. A set (respectively, reset) arc from s to t sets the number of tokens in place s to 1

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

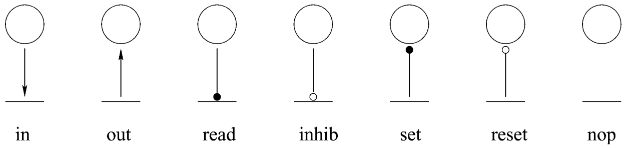

Fig. 9. Graphic conventions for drawing trace nets.

(respectively, 0) when t fires. Finally, an arc of kind *nop* denotes the absence of any kind of relation between the contents of s and the firing of t.

We adopt the graphical convention proposed in Badouel and Darondeau (1998), and depicted in Figure 9 to draw trace nets: input (respectively, output) arcs are represented as directed edges with an arrow on the transition (respectively, place) side; read (respectively, inhibitor) arcs are represented as edges with a small black (respectively, white) circle on the transition side; set (respectively, reset) arcs are represented as edges with a small black (respectively, white) circle on the place side.

To simplify the definitions of the enabling of a transition and of the firing rule, we introduce the following auxiliary relations. Intuitively, $t\ test1\ s$ (respectively, $t\ test0\ s$) holds if it is necessary that s contains one (respectively, zero) tokens for t to fire. On the other hand, $t\ set1\ s$ (respectively, $t\ set0\ s$) holds if, after the firing of t, s contains one (respectively, zero) tokens.

Given $s \in S$ and $t \in T$, we define the following relations:

$$\begin{array}{rll} t\ test1\ s & \text{if and ony if} & W(s,t) \in \{in, read\} \\ t\ test0\ s & \text{if and ony if} & W(s,t) \in \{out, inhib\} \\ t\ set1\ s & \text{if and ony if} & W(s,t) \in \{out, set\} \\ t\ set0\ s & \text{if and ony if} & W(s,t) \in \{in, reset\}. \end{array}$$

A transition t is enabled at marking m if $\{s \mid t\ test1\ s\} \subseteq m$ and $\{s \mid t\ test0\ s\} \cap m = \emptyset$.

The firing (execution) of a transition t enabled at m produces the marking $m' = \{s \in S \mid t\ set1\ s\} \cup \{s \in m \mid W(s,t) = nop\}$. This is usually written as $m[t\rangle m'$.

A *trace net system* is a pair (N, m_0), where N is a trace net and m_0 is a marking of N, called the *initial marking*. With abuse of notation, we use (S, T, W, m_0) to denote the trace net system $((S, T, W), m_0)$.

The definitions of reachable markings, firing sequences, marking graph and reduced system given in Section 2 for elementary nets apply also for trace nets.

In the following we consider trace net systems that are reduced.

5.2. *Positive Place-Based Non-Interference for trace nets*

As we have already done for elementary nets, we consider trace nets whose set of transitions is partitioned into two subsets: the set H of high-level transitions and the set L of low-level transitions. Also, for trace nets, given two disjoint sets L and H, we use (S, L, H, W, m_0) to denote the trace net system $(S, L \cup H, W, m_0)$.

We extend the definitions of causal and conflict places for elementary nets given in Section 4 to trace nets.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

An extension of the definition of potentially causal place for contact-free elementary nets to trace nets leads to the following: a place s is a potentially causal place if there exist a high-level transition h that puts a token in s (either through an output or a set arc) and a low-level transition l that needs place s to be full to fire (because s and l are connected by a read or an input arc). However, as contact freeness no longer holds for trace nets, we also have to take into account causal dependencies arising from the fact that a high-level transition h removes a token contained in a place s (either through an input or a reset arc) that is required to be empty for a low-level transition l to fire (because s and l are connected by an output or an inhibitor arc).

Also, for potentially conflict places, two kinds of conflicts can arise. In a similar way to the case for elementary nets, s is a potentially conflict place if there exists a high-level transition h that removes a token from s and a low-level transition l that needs place s to be full to fire. Morevoer, s also has to be considered a potentially conflict place if there exists a high-level transition h that produces a token in s and a low-level transition l that needs place s to be empty in order to fire.

Definition 5.2. Let $N = (S, L, H, W, m_0)$ be a trace net system.

The place $s \in S$ is a *potentially causal place* if there exist $h \in H$, $l \in L$ and $X \in \{0, 1\}$ such that $h \; setX \; s$ and $l \; testX \; s$.

A potentially causal place s is an *active causal place* if there exist $h \in H$, $l \in L$ and $X \in \{0, 1\}$ such that:

— $h \; setX \; s$

— $l \; testX \; s$

— there exists a marking $m \in [m_0\rangle$ and a transition sequence σ such that

  – $m[h\sigma l\rangle$

  – $s \in m$ if and only if $X = 0$

  – for all $t \in \sigma$: $\neg(t \; setX \; s)$.

The place $s \in S$ is a *potentially conflict place* if there exist $h \in H$, $l \in L$ and $X \in \{0, 1\}$ such that $h \; setX \; s$ and $l \; test(1 - X) \; s$.

A potentially conflict place is an *active conflict place* if there exist $h \in H$, $l \in L$ and $X \in \{0, 1\}$ such that:

— $h \; setX \; s$

— $l \; test(1 - X) \; s$

— there exists a marking $m \in [m_0\rangle$ and a transition sequence σ such that

  – $m[h\rangle$ and $m[\sigma l\rangle$

  – $s \in m$ if and only if $X = 0$

  – for all $t \in \sigma$: $\neg(t \; set(1 - X) \; s)$.

Definition 5.3. Let $N = (S, L, H, W, m_0)$ be a trace net system. We say that N is *PBNI+* (*positive Place-Based Non-Interference*) if, for all $s \in S$, s is neither an active causal place nor an active conflict place.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

The definitions of parallel composition and restriction presented in Section 3 are extended to trace nets in the obvious way as follows.

Definition 5.4. Let $N_1 = (S_1, L_1, H_1, W_1, m_{0,1})$ and $N_2 = (S_2, L_2, H_2, W_2, m_{0,2})$ be two trace net systems such that $S_1 \cap S_2 = \emptyset$ and $(L_1 \cup L_2) \cap (H_1 \cup H_2) = \emptyset$. The parallel composition of N_1 and N_2 is the trace net system

$$N_1 \mid N_2 = (S_1 \cup S_2, L_1 \cup L_2, H_1 \cup H_2, W_1 \cup W_2, m_{0,1} \cup m_{0,2}).$$

Definition 5.5. Let $N = (S, L, H, W, m_0)$ be a trace net system and let U be a set of transitions. The restriction on U is defined as $N \backslash U = (S, L', H', W', m_0)$, where

$$\begin{aligned} L' &= L \setminus U \\ H' &= H \setminus U \\ W' &= W\lceil(S \times ((L \cup H) \setminus U)). \end{aligned}$$

The results presented in Section 3 also hold for trace nets.

Theorem 5.6. Let $N = (S, L, H, W, m_0)$ be a trace net system. N is *BNDC* if and only if N is *SBNDC*.

Proof. The proofs of Theorems 3.14 and 3.15 can be used for trace nets without any significant change. □

Also, the theorems of Section 4, with slight modification, hold for trace nets too. Here we give full details for one of the proofs; the other can be adapted similarly.

Theorem 5.7. Let $N = (S, L, H, W, m_0)$ be a trace net system. If N is *PBNI+*, then N is *SBNDC*.

Proof. We follow the proof idea of the corresponding theorem, Theorem 4.3, on elementary net systems. Take $m \in [m_0\rangle$ such that $m[h\rangle m'$ for some $h \in H$. We have to prove that there exists a low-view bisimulation R on $N \backslash H$ and that $(m, m') \in R$.

Let $R = \{(m_1, m_2) \mid (\forall l \in L\ \forall s)\ l\ testX\ s\ \wedge\ m_1(s) \neq m_2(s) \Rightarrow ((\forall \sigma \forall i \in \{1,2\})\ m_i[\sigma l\rangle \Rightarrow \exists l_1 \in \sigma\ :\ l_1\ setX\ s\ )\}$ be the candidate relation.

— We show that R is a low-view bisimulation on $N \backslash H$.

Let $(m_1, m_2) \in R$. Suppose $m_1[l\rangle m_1'$. We show that $m_2[l\rangle$ also. Suppose there exists s such that $l\ testX\ s$ and $m_2(s) = (1 - X)$. So $m_1(s) \neq m_2(s)$. As $(m_1, m_2) \in R$ and $m_1[l\rangle$, by the definition of R (with $\sigma = \varepsilon$), there must exist $t \in \varepsilon$, which gives a contradiction. Hence, for all s such that $l\ testX\ s$, we have $m_2(s) = X$, so there exists m_2' such that $m_2[l\rangle m_2'$.

Now we show that $(m_1', m_2') \in R$. Suppose $(m_1', m_2') \notin R$. Then there exist l', s' such that, $l'\ testX\ s'$ and $m_1'(s') \neq m_2'(s')$, and there exist σ and i such that $m_i'[\sigma l'\rangle$ and $l_1\ setXs'$ for no $l_1 \in \sigma$.

As $m_i[l\rangle m_i'$ for $i = 1, 2$, we have $m_1(s') \neq m_2(s')$ and there exists $i \in \{1,2\}$ such that $m_i[l\sigma l'\rangle$ and $l_1\ setXs'$ for no $l_1 \in \sigma$.

Since $m_1'(s') \neq m_2'(s')$, we necessarily have $\neg(l\ setX\ s')$, so $l_1\ setXs'$ for no $l_1 \in l\sigma$. Thus we have $(m_1, m_2) \notin R$, which gives a contradiction. So we have $(m_1', m_2') \in R$.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

The symmetric case can be proved in the same way, so we have now shown that R is a low-view bisimulation on $N\backslash H$.

— We show that $(m, m') \in R$.
Suppose there exists s and l such that l *testX* s and $m(s) \neq m'(s)$. We show that $\forall\sigma : m[\sigma l\rangle \Rightarrow \exists t \in \sigma : t$ *setX* s and that $\forall\sigma : m'[\sigma l\rangle \Rightarrow \exists t \in \sigma : t$ *setX* s.
As $m[h\rangle m'$, we can deduce from $m(s) \neq m'(s)$ that one of the following holds:

- h *setX* s.
Hence, s is a potentially causal place. Moreover, $m'(s) = X$ and $m(s) = (1 - X)$ because $m(s) \neq m'(s)$.
Take a sequence σ such that $m[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that t *setX* s. There are two subcases:

  (i) $\sigma = h\sigma'$.
  As *PBNI+* holds, s is not an active causal place. Hence, for all $\bar{m} \in [m_0\rangle$ and for all $\bar{\sigma}$, we have if $\bar{m}[h\bar{\sigma}l\rangle$ and $\bar{m}(s) = (1 - X)$, then there exists $t \in \bar{\sigma}$ such that t *setX* s. As $m[\sigma l\rangle$, and $\sigma = h\sigma'$, there exists $t \in \sigma'$ such that t *setX* s.

  (ii) $\sigma = \varepsilon$ or $\sigma = t'\sigma'$ with $t' \neq h$.
  As we have already observed that $m(s) = (1 - X)$, and we know that $m[\sigma l\rangle$ and l *testX* s, there must exist a transition $t \in \sigma$ such that t *setX* s. (In particular, $\sigma \neq \varepsilon$.)

  Now consider a sequence σ such that $m'[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that t *setX* s.
As $m[h\rangle m'$, we have $m[h\sigma l\rangle$. As h *setX* s and $m(s) \neq m'(s)$, we have $m(s) = (1 - X)$. Hence, because *PBNI+* holds (s cannot be an active causal place), there exists $t \in \sigma$ such that t *setX* s.

- h *set*$(1 - X)$ s.
Hence, s is a potentially conflict place. Moreover, $m'(s) = (1 - X)$ and $m(s) = X$ because $m(s) \neq m'(s)$.
Take a sequence σ such that $m[\sigma l\rangle$. We show that there exists $t \in \sigma$ such that t *setX* s.
As *PBNI+* holds, s cannot be an active conflict place. Hence, for all $\bar{m} \in [m_0\rangle$ and for all $\bar{\sigma}$, we have if $\bar{m}[h\rangle$ and $\bar{m}[\bar{\sigma}l\rangle$ and $\bar{m}(s) = X$, then there exists $t \in \bar{\sigma}$ such that t *setX* s. As $m[h\rangle m'$ and $m[\sigma l\rangle$, there exists $t \in \sigma$ such that t *setX* s.
Now take a sequence σ such that $m'[\sigma l\rangle$. We have already observed that $m'(s) = (1 - X)$. As l *testX* s, we get from $m'[\sigma l\rangle$ that there must exist a transition $t \in \sigma$ such that t *setX* s. □

Theorem 5.8. Let $N = (S, L, H, W, m_0)$ be a trace net system. If N is *SBNDC*, then N is *PBNI+* .

Proof. It is enough to follow the proof of Theorem 4.4, with the proviso of interpreting membership of transition to pre/postsets of places in terms of *setX* or *testX*, for example, $h \in {}^{\bullet}s$ as h *setX* s. □

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

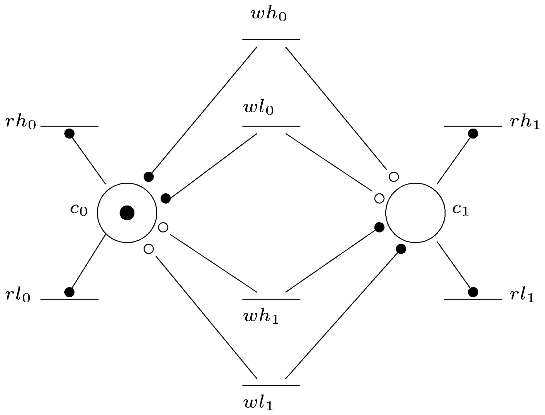

Fig. 10. The trace net modelling a binary memory cell.

5.3. *Example: binary memory cell*

In this section we recast the example of a binary memory cell proposed in Bossi *et al.* (2004) in our framework.

A binary memory cell can contain a binary value, that is, either 0 or 1. The memory cell is accessible to both high- and low-level users, who can read and write a value in the cell. The trace net representing the binary cell is shown in Figure 10.

A token in place c_0 (respectively, c_1) represents the fact that the current value contained in the cell is 0 (respectively, 1). Each operation is modelled by two transitions: one for each binary value that can be contained in the cell. For example, transition wh_0 (respectively, wh_1) is performed by a high-level user who writes the value 0 (respectively, 1) in the cell. A write operation of, for example, value 0 in the cell is represented by a transition that sets the contents of place c_0 (that is, puts one token in place c_0 regardless of its previous contents), and resets the contents of place c_1 (that is, removes the possible token present in place c_1). A read operation of, for example, value 0 is represented by a transition with a read arc on place c_0, that is, a transition that can happen only if place c_0 contains a token.

As has already been pointed out in Bossi *et al.* (2004), the binary memory cell depicted in Figure 10 is completely insecure since a high-level user can send confidential information to a low-level user through the binary cell. In fact, the binary cell is not *PBNI+* because of the existence of (at least) the active causal place c_1. Note that c_1 is a potentially causal place, because the high-level transition wh_1 has a set arc on c_1, and the low-level transition rl_1 has a read arc on c_1. Moreover, if we also consider the firing sequence $\{c_0\}[wh_1\rangle\{c_1\}[rl_1\rangle\{c_1\}$, the conditions for the potentially causal place c_1 to be an active causal place are fulfilled.

In order to avoid the flow of information from the high-level user to the low-level user, we can either forbid all the read operations performed by a low-level user or forbid all the write operations performed by a high-level user, thus obtaining the trace nets depicted in Figures 11 and 12.

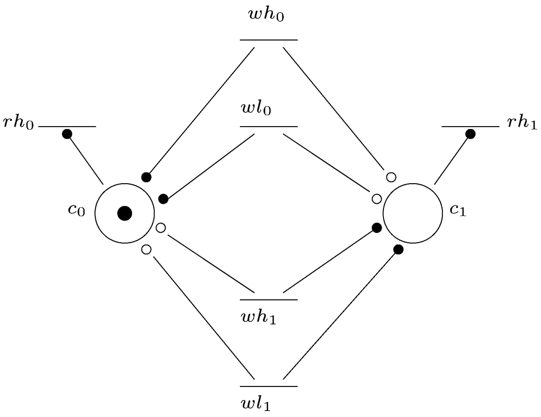

Fig. 11. The trace net obtained by removing the low-level read operations.

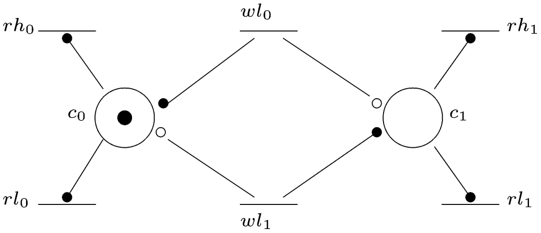

Fig. 12. The trace net obtained by removing the high-level write operations.

The trace nets obtained in this way do not contain any (potential) causal or conflict place.

6. Conclusion

We have proposed a structural non-interference property *PBNI+* on elementary nets to make explicit the essence of (positive) information flows. These can be of two kinds: *causal flows*, when a low-level transition consumes a token produced by a high-level transition, or *conflict flows*, when a high-level and a low-level transition compete for the same token. We have shown that *PBNI+* turns out to be equivalent to the behavioural property *SBNDC* (which is in turn equivalent to *BNDC*), which was originally proposed in a process algebraic setting (Focardi and Gorrieri 1995; Focardi and Gorrieri 2001). We argue that *PBNI+* is often more easily checkable than *SBNDC*, which is based on (a number of) equivalence checks. We have also shown that the basic definition can be easily extended to the richer class of trace net systems, allowing us to analyse one non-trivial example.

The property *PBNI+* is structural because no notion of observational equivalence is considered in its definition; however, to be precise, the definition of *PBNI+* requires a limited exploration of the state space (marking graph), hence it is in some sense a *hybrid* property.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

The main difference between *PBNI+* compared with the non-interference properties *PBNI* and *RBNI*, which were introduced in Busi and Gorrieri (2004a; 2004b) (the latter being based on a suitable exploitation of the theory of regions (Badouel and Darondeau 1998)), is that *PBNI+* only captures flows of positive information from high-level users to low-level users (together with some inevitable negative information flows), whereas *RBNI* (and in some cases *PBNI* also) also captures unnecessary negative flows. Consider, for example, the net system N_4 in Figure 7: the net N_4 does not satisfy *PBNI* and *RBNI*, because of the presence of the conflict place s. Indeed, there exists a negative information flow from high-level users to low-level users: if the low-level action l is performed, the low-level user knows that the high-level action h has not been performed (and will never be performed). On the other hand, N_4 is *PBNI+*, because s is not an active conflict place. Indeed, no positive information can flow from high-level users to low-level users.

In order to keep the presentation as simple as possible, the current investigation was carried out for two classes of 'safe' net systems, that is, for nets whose places can contain at most one token. A natural generalisation is to consider Place/Transition systems, where each place can contain more than one token. Such a class of nets is particularly interesting because the marking graph associated with a finite P/T net system may be infinite. We claim that *PBNI+* can also be defined on this richer class of nets with a minor change in the definition of potential causal and conflict places. Moreover, we claim that *PBNI+* is also the same as *SBNDC* for finite P/T net systems. We conjecture that *PBNI+* can be checked in a finite amount of time by inspecting the (finite) coverability tree of a finite P/T net, thereby possibly providing the first decidability result for a behavioural information flow security property, like *SBNDC*, on a class of infinite state systems.

Another interesting future research area is the extension of this approach to other, rather general, non-interference frameworks, such as, for example, the one discussed in Gruska (2007).

Finally, we would like to mention that the theory developed here has been implemented recently in a tool, called the *Petri Net Security Checker* (Frau 2008), where a user-friendly graphical interface allows the user to build nets with transitions of two different confidentiality levels and to check if they satisfy *PBNI+*.

Acknowledgements

This paper is an extended and revised version of Busi and Gorrieri (2004c), which was prepared by the second author after the first author's premature death. The second author would like to thank the anonymous referees for helpful suggestions and comments, and is very grateful to Philippe Darondeau for his careful proof-reading of the paper and for improvements in some proofs.

References

Badouel, E. and Darondeau, Ph. (1995) Trace nets and process automata. *Acta Informatica* **32** 647–679.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press

Badouel, E. and Darondeau, Ph. (1998) Theory of regions. In: Reisig, W. and Rozenberg, G. (eds.) Lectures on Petri Nets I: Basic Models. *Springer-Verlag Lecture Notes in Computer Science* **1491** 529–586.

Bossi, A., Focardi, R., Macedonio, D., Piazza, C. and Rossi, S. (2004) Unwinding in Information Flow Security. *Electronic Notes in Theoretical Computer Science* **99** 127–154.

Brookes S. D., Hoare C. A. R. and Roscoe A. W. (1984) A Theory of Communicating Sequential Processes. *Journal of the ACM* **31** (3) 560–599.

Busi, N. and Gorrieri, R. (2004a) Structural Non-Interference with Petri Nets. Workshop on Issues in the Theory of Security (WITS'04).

Busi, N. and Gorrieri, R. (2004b) A Survey on Non-Interference with Petri Nets. Advanced Course on Petri Nets 2003. *Springer-Verlag Lecture Notes in Computer Science* **3098** 328–344.

Busi, N. and Gorrieri, R. (2004c) Positive Non-Interference in Elementary and Trace Nets. Proceedings 25th International Conference on Application and Theory of Petri Nets. *Springer-Verlag Lecture Notes in Computer Science* **3099** 1–16.

Engelfriet J. and Rozenberg G. (1998) Elementary Net Systems. In: Reisig, W. and Rozenberg, G. (eds.) Lectures on Petri Nets I: Basic Models. *Springer-Verlag Lecture Notes in Computer Science* **1491** 12–121.

Focardi R. and Gorrieri R. (1995) A Classification of Security Properties. *Journal of Computer Security* **3** (1) 5–33.

Focardi R. and Gorrieri R. (1997) The Compositional Security Checker: A Tool for the Verification of Information Flow Security Properties. *IEEE Transactions on Software Engineering* **23** (9) 550–571.

Focardi R. and Gorrieri R. (2001) Classification of Security Properties (Part I: Information Flow). In: Focardi, R. and Gorrieri, R. (eds.) Foundations of Security Analysis and Design – Tutorial Lectures. *Springer-Verlag Lecture Notes in Computer Science* **2171** 331–396.

Frau S. (2008) Uno strumento sotware per l'analisi di proprietà di sicurezza su reti di Petri. Masters thesis (in Italian), University of Bologna, March 2008.

Goguen J. A. and Meseguer J. (1982) Security Policy and Security Models. In: *Proceedings of Symposium on Security and Privacy*, IEEE CS Press 11–20.

Gruska D. P. (2007) Observation Based System Security. *Fundamenta Informaticae* **79** (3-4) 335–346.

Milner R. (1989) *Communication and Concurrency*, Prentice-Hall.

Petri C. A. (1962) *Kommunikation mit Automaten*, Ph.D. Thesis, Institut für Instrumentelle Mathematik, Bonn, Germany.

Reisig W. (1985) *Petri Nets: An Introduction*, EATCS Monographs in Computer Science, Springer-Verlag.

Roscoe A. W. (1995) CSP and Determinism in Security Modelling. In: *Proceedings of IEEE Symposium on Security and Privacy*, IEEE CS Press 114–127.

Ryan P. Y. A. (2001) Mathematical Models of Computer Security. In: Focardi, R. and Gorrieri, R. (eds.) Foundations of Security Analysis and Design – Tutorial Lectures. *Springer-Verlag Lecture Notes in Computer Science* **2171** 1–62.

Ryan P. Y. A. and Schneider S. (1999) Process Algebra and Noninterference. In: *Proceedings of 12th Computer Security Foundations Workshop*, IEEE CS Press 214–227.

https://doi.org/10.1017/S0960129509990120 Published online by Cambridge University Press