

Events in Security Protocols

Federico Crazzolaria* Glynn Winskel
Computer Laboratory, University of Cambridge
England

{fc232,gw104}@cl.cam.ac.uk

ABSTRACT

The events of a security protocol and their causal dependency can play an important role in the analysis of security properties. This insight underlies both strand spaces and the inductive method. But neither of these approaches builds up the events of a protocol in a compositional way, so that there is an informal spring from the protocol to its model. By broadening the models to certain kinds of Petri nets, a restricted form of contextual nets, a compositional event-based semantics is given to an economical, but expressive, language for describing security protocols; so the events and dependency of a wide range of protocols are determined once and for all. The net semantics is formally related to a transition semantics, strand spaces and inductive rules, as well as trace languages and event structures, so unifying a range of approaches, as well as providing conditions under which particular, more limited, models are adequate for the analysis of protocols. The net semantics allows the derivation of general properties and proof principles which are demonstrated in establishing an authentication property, following a diagrammatic style of proof.

1. INTRODUCTION

The last few years have seen the emergence of successful intensional, event-based, approaches to reasoning about security protocols. The methods are concerned with reasoning about the events that a security protocol can perform, and make use of a causal dependency that exists between events. For example, to show secrecy in a protocol it is shown that there can be no earliest event violating a secrecy property; any such event is shown to depend on some earlier event which itself violates secrecy - because the behaviour of the protocol does not permit such an infinite regress, the secrecy property is established. In a similar way, dependency between events is used to establish forms of authentication by showing that a sequence of communication events of one

agent entails a corresponding sequence of events of the intended participant.

Both the method of strand spaces [12] and the inductive method of Paulson [9] have been designed to support such an intensional, event-based, style of reasoning. Strand spaces are based on an explicit causal dependency of events, whereas in Paulson's method the dependency is implicit in the inductive rules, which might express, for instance, that the input of a message depends on its previous output. Both methods have successfully tackled a number of protocols though in an *ad hoc* fashion. Both make an informal spring from a protocol to its representation, as either a strand space or a set of inductive rules. Both methods do not address how to build up their representation of a protocol in a compositional fashion.

We show that Petri nets, and specifically a restricted form of contextual nets [8], provide a common framework in which to understand both the strand-space and inductive methods, and it seems, although we understand it less well, the recent multiset rewriting method of [3]. But, more importantly, by moving to a broader class of models we can show how event-based models can be structured in a compositional way and so used to give a formal semantics to security protocols which supports proofs of their correctness. To make the case, and provide semantics to a whole range of protocols once and for all, we study the semantics of **SPL** (Security Protocol Language). We demonstrate the usefulness of the net semantics in deriving (in contrast to postulating) proof principles for security protocols and apply them to prove an authentication property - the diagrammatic style of proof may be of interest in itself. We establish precise relationships between the net semantics and transition semantics, strand spaces, inductive rules, and trace languages and event structures. The results formally back up the adequacy of strand-space and inductive-rule representations for broad classes of security protocols and properties, showing when nothing is lost in moving to these more restrictive models.

2. SECURITY PROTOCOLS

As a running example we consider the Needham-Schröder-Lowe (NSL) protocol:

- (1) $A \longrightarrow B : \{m, A\}_{Pub(B)}$
- (2) $B \longrightarrow A : \{m, n, B\}_{Pub(A)}$
- (3) $A \longrightarrow B : \{n\}_{Pub(B)}$

This protocol, like many others of its kind, has two roles: one for the initiator, here played by agent A (say Alice), and one

*BRICS, Centre of the Danish National Research Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CCS'01, November 5-8, 2001, Philadelphia, Pennsylvania, USA.
Copyright 2001 ACM 1-58113-385-5/01/0011 ...\$5.00.

for the responder, here B (Bob). It is a public-key protocol that assumes an underlying public-key infrastructure, such as RSA [11]. Both agents have their own, secret private key. Public keys in contrast are available to all participants in the protocol. The NSL protocol makes use of *nonces* which one can think of as newly generated, unguessable numbers whose purpose is to ensure the freshness of messages.

The protocol describes an interaction between A in the role of initiator and B as responder: A sends to B a new nonce m together with her own agent name A , both encrypted with B 's public key. When the message is received by B , he decrypts it with his secret private key. Once decrypted, B prepares an encrypted message for A that contains a new nonce together with the nonce received from A and his name B . Acting as responder, B sends it to A , who recovers the clear text using her private key. A convinces herself that this message really comes from B by checking whether she got back the same nonce sent out in the first message. If that is the case, she acknowledges B by returning his nonce. B does a similar test.

Although in this informal explanation only two agents in their respective roles are described, the protocol is really a shorthand for a situation in which a network of distributed agents are each able to participate in multiple concurrent sessions as both initiator and responder. There is no assurance that they all stick to the protocol, or indeed that communication goes to the intended agent. An attacker might disassemble and pretend to be one or several agents, taking advantage of any leaked keys it possesses in deciphering, and preparing the messages it sends. However in this paper we assume cryptography unbreakable and random numbers unguessable, following the lead of Dolev and Yao [4].

The NSL protocol aims at distributing nonces m and n in a secure way, allowing no one but the initiator and the responder to know them (*secrecy*). Another aim of the protocol is that, for example, Bob should be guaranteed that m is indeed the nonce sent by Alice (*authentication*). Lowe pointed out that the NSL protocol is prone to a “middle-man” attack, violating both these secrecy and authentication properties, if, as in the original protocol, the name B is not included in the second message [6].

3. SPL - A LANGUAGE FOR SECURITY PROTOCOLS

In order to be more explicit about the activities of participants in a protocol and those of a possible attacker, and to express these compositionally, we design an economical process language for the purpose. The language **SPL** (Security Protocol Language) is close to an asynchronous Pi-Calculus [7] and is similar to that adopted in [1], though in its treatment of new names its transition semantics will be closer to that in [10] (it separates concerns of freshness from concerns of scope which are combined in the Pi-Calculus restriction).

3.1 The syntax of SPL

We start by giving the syntactic sets of the language:

- An infinite set of names \mathbf{N} , with elements n, m, A, \dots . Names range over nonces as well as agent names, and can also include other values.
- Variables over names x, y, \dots, X, Y, \dots .
- Variables over messages $\psi, \psi', \psi_1, \dots$.

Name expr.	v	::=	$n, A, \dots \mid x, X, \dots$
Key expr.	k	::=	$Pub(v) \mid Priv(v) \mid Key(\vec{v})$
Messages	M	::=	$v \mid k \mid M_1, M_2 \mid \{M\}_k \mid \psi$
Processes	p	::=	$out\ new\vec{x}M.p \mid in\ pat\vec{x}\vec{\psi}M.p \mid \parallel_{i \in I} p_i$

Figure 1: Syntax of SPL

- Indices $i \in \mathbf{Indices}$ with which to index components of parallel compositions.

The other syntactic sets of the language are described by the grammar shown in Figure 1. Note we use “vector” notation; for example, the “vector” \vec{x} abbreviates some possibly empty list x_1, \dots, x_l .

We take $fv(M)$, the free variables of a message M , to be the set of variables which appear in M , and define the free variables of process terms by:

$$\begin{aligned} fv(out\ new\vec{x}M.p) &= (fv(p) \cup fv(M)) \setminus \{\vec{x}\} \\ fv(in\ pat\vec{x}\vec{\psi}M.p) &= (fv(p) \cup fv(M)) \setminus \{\vec{x}, \vec{\psi}\} \\ fv(\parallel_{i \in I} p_i) &= \bigcup_{i \in I} fv(p_i). \end{aligned}$$

As usual, we say that a process without free variables is closed, as is a message without variables. We shall use standard notation for substitution into the free variables of an expression, though we will only be concerned with the substitution of names or closed (variable-free) messages, obviating the problems of variable capture.

We use $Pub(v)$, $Priv(v)$ for the public, private keys of v , and we use $Key(\vec{v})$ for the symmetric key shared by agents with names in \vec{v} . Keys can be used in building up encrypted messages. Messages consist of a name or a key, the composition of two messages M_1, M_2 , or an encryption $\{M\}_k$ representing the message M encrypted using the key k .

An informal explanation of the language:

out new $\vec{x}M.p$ This process chooses fresh, distinct names $\vec{n} = n_1, \dots, n_l$ and binds them to the variables $\vec{x} = x_1, \dots, x_l$. The message $M[\vec{n}/\vec{x}]$ is output to the network and the process resumes as $p[\vec{n}/\vec{x}]$. The communication is *asynchronous* in the sense that the action of output does not await input. The *new* construct is like that of Pitts and Stark [10] and abstracts out an important property of a value chosen randomly from some large set: such a value is likely to be new.

in pat $\vec{x}\vec{\psi}M.p$ This process awaits an input that matches the pattern M for some binding of the pattern variables $\vec{x}\vec{\psi}$ and resumes as p under this binding. All the pattern variables $\vec{x}\vec{\psi}$ must appear in the pattern M .

$\parallel_{i \in I} p_i$ This process is the parallel composition of all components p_i for i in the indexing set I . The set I is a subset of **Indices**. Indices will help us distinguish in what agent, which role and what run a particular action occurs. The process, written *nil*, abbreviates the empty parallel composition (where the indexing set is empty).

CONVENTION 3.1. It simplifies the writing of process expressions if we adopt some conventions. Firstly, we simply write *out M.p* when the list of “*new*” variables is empty. Secondly, and more significantly, we allow ourselves to write $\dots in\ M.p \dots$ in an expression, to be understood as meaning the expression $\dots in\ pat\vec{x}\vec{\psi}M.p \dots$ where the pattern variables $\vec{x}\vec{\psi}$ are precisely those variables left free in M by the

$$\begin{aligned}
Init(A, B) &\equiv out\ new\ x\{x, A\}_{Pub(B)} \cdot \\
&\quad in\{x, y, B\}_{Pub(A)} \cdot \\
&\quad out\{y\}_{Pub(B)} \cdot \\
&\quad nil \\
Resp(B) &\equiv in\{x, Z\}_{Pub(B)} \cdot \\
&\quad out\ new\ y\{x, y, B\}_{Pub(Z)} \cdot \\
&\quad in\{y\}_{Pub(B)} \cdot \\
&\quad nil
\end{aligned}$$

Figure 2: Initiator and responder code

surrounding expression. For example, we can describe a responder in NSL as the process $Resp(B)$ in Figure 2. For the first input, the variables x, Z in $\{x, Z\}_{Pub(B)}$ are free in the whole expression, so by convention are pattern variables. On the other hand, in the second input the variable y in $\{y\}_{Pub(B)}$ is bound by the outer $out\ new\ y \dots$ and so by the convention is not a pattern variable, and has to be that value sent out earlier. *Replication* of a process, $!p$, abbreviates $\|_{i \in \omega} p$, consisting of countably infinite copies of p set in parallel.

An obvious structural induction defines the set of names of a process. We define $size(p)$ of a process term p to be an ordinal measuring the depth of process operations in the term.

DEFINITION 3.1. The *size* of a closed process term is an ordinal given by the structural induction:

$$\begin{aligned}
size(out\ new\ \vec{x}M.p) &= 1 + size(p) \\
size(in\ pat\ \vec{x}\vec{p}M.p) &= 1 + size(p) \\
size(\|_{i \in I} p_i) &= 1 + sup_{i \in I} size(p_i).
\end{aligned}$$

3.2 NSL as a process

As an illustration, we can program the NSL protocol in our language, and so formalise the introductory description given in the Section 2. We assume given a set of agent names, **Agents**, of agents participating in the protocol. The agents participating in the NSL protocol play two roles, as initiator and responder with any other agent. Abbreviate by $Init(A, B)$ the program of initiator $A \in \mathbf{Agents}$ communicating with $B \in \mathbf{Agents}$ and by $Resp(B)$ the program of responder $B \in \mathbf{Agents}$. The code of both an arbitrary initiator and an arbitrary responder is given in Figure 2. In the code we are forced to formalise aspects that are implicit in the informal description, such as the creation of new nonces, the decryption of messages and the matching of nonces.

We can model the attacker by directly programming it as a process. Figure 3, shows a general, active attacker or “spy”. The spy has the capability of composing eavesdropped messages, decomposing composite message, and using cryptography whenever the appropriate keys are available; the available keys are all the public keys and the leaked private keys. By choosing a different program for the spy we can restrict or augment its power, e.g., to passive eavesdropping or active falsification.

The whole system is obtained by putting all components in parallel. Components are replicated, to model multiple concurrent runs of the protocol. The system is described in Figure 4.

$$\begin{aligned}
Spy_1 &\equiv in\ \psi_1.in\ \psi_2.out\ \psi_1, \psi_2.nil && \text{(composing)} \\
Spy_2 &\equiv in\ \psi_1, \psi_2.out\ \psi_1.out\ \psi_2.nil && \text{(decomp'g)} \\
Spy_3 &\equiv in\ x.in\ \psi.out\ \{\psi\}_{Pub(x)}.nil && \text{(encrypting)} \\
Spy_4 &\equiv in\ Priv(x).in\ \{\psi\}_{Pub(x)}.out\ \psi.nil && \text{(decrypting)} \\
Spy &\equiv \|\|_{i \in \{1, \dots, 4\}} Spy_i
\end{aligned}$$

Figure 3: Attacker code

$$\begin{aligned}
P_{init} &\equiv \|\|_{A, B} ! Init(A, B) \\
P_{resp} &\equiv \|\|_A ! Resp(A) \\
P_{spy} &\equiv ! Spy \\
NSL &\equiv \|\|_{i \in \{resp, init, spy\}} P_i
\end{aligned}$$

Figure 4: The system

3.3 A transition semantics

We first give a, fairly traditional, transition semantics to **SPL**. It says how input and output actions affect configurations; a configuration expresses the state of execution of the process, the messages so far output to the network and the names currently in use.

A *configuration* consists of a triple $\langle p, s, t \rangle$ where p is a closed process term, s is a subset of the set of names \mathbf{N} , and t is a subset of closed (i.e., variable-free) messages. We say the configuration is *proper* iff the names in p and t are included in s . The idea is that a closed process p acts in the context of the set of names s that have been used so far, and the set of messages t which have been output, to input a message or to generate new names before outputting a message.

Actions α may be inputs or new-outputs, possibly tagged by indices to show at which parallel component they occur:

$$\alpha ::= out\ new\ \vec{n}.M \mid in\ M \mid i : \alpha$$

where M is a closed message, \vec{n} are names and i is an index drawn from **Indices**. We write $out\ M$ for an output action, outputting a message M , where no new names are generated.

The way configurations evolve is expressed by transitions $\langle p, s, t \rangle \xrightarrow{\alpha} \langle p', s', t' \rangle$, given by the rules shown in Figure 5.

The transition semantics allows us to state formally many security properties. However, it does not support directly local reasoning of the kind one might wish to apply in the analysis of security protocols. To give an idea of the difficulty, imagine we wished to establish that the nonce generated by B as responder in NSL was never revealed as an open message on the network. A reasonable way to prove such a property is to find a stronger invariant, a property which can be shown to be preserved by all the actions of the process. Equivalently, one can assume that there is an earliest action α_i in a run which violates the invariant, and derive a contradiction by showing that this action must depend on a previous action, which itself violates the invariant.

An action might depend on another action through being, for example, an input depending on a previous output, or simply through occurring at a later control point in a process. A problem with the transition semantics is that it masks such local dependency, and even the underlying process events on which the dependency rests. The wish to support arguments based on local dependency leads to a more refined semantics based on events.

(**output**) Provided the names \vec{n} are all distinct and not in s ,

$$\langle out\ new\ \vec{x}M.p, s, t \rangle \xrightarrow{out\ new\ \vec{n}M[\vec{n}/\vec{x}]} \langle p[\vec{n}/\vec{x}], s \cup \{\vec{n}\}, t \cup \{M[\vec{n}/\vec{x}]\} \rangle$$

(**input**) Provided $M[\vec{n}/\vec{x}, \vec{N}/\vec{\psi}] \in t$,

$$\langle in\ pat\ \vec{x}\vec{\psi}M.p, s, t \rangle \xrightarrow{in\ M[\vec{n}/\vec{x}, \vec{N}/\vec{\psi}]} \langle p[\vec{n}/\vec{x}, \vec{N}/\vec{\psi}], s, t \rangle$$

(**par**)

$$\frac{\langle p_j, s, t \rangle \xrightarrow{\alpha} \langle p'_j, s', t' \rangle}{\langle \parallel_{i \in I} p_i, s, t \rangle \xrightarrow{j:\alpha} \langle \parallel_{i \in I} p'_i, s', t' \rangle} \quad j \in I$$

where p'_i is p'_j for $i = j$, else p_i .

Figure 5: Transition semantics

4. THE EVENTS OF SPL

We must first address the issue of what constitutes an event of a security protocol. Here, we follow the lead from Petri nets (see Appendix), and define events in terms of how they affect conditions. Conditions are to represent some form of local state and we discern conditions of three kinds: *control*, *output* and *name* conditions.

The set of *control conditions* \mathbf{C} consists of output or input processes, perhaps tagged by indices, and is given by the grammar

$$b ::= out\ new\ \vec{x}M.p \mid in\ pat\ \vec{x}\vec{\psi}M.p \mid i : b$$

where $i \in \mathbf{Indices}$. A condition in \mathbf{C} stands for the point of control in a (single-thread) process. When C is a subset of control conditions we will write $i : C$ to mean $\{i : b \mid b \in C\}$.

The set of *output conditions* \mathbf{O} consists of closed message expressions. An individual condition M in \mathbf{O} stands for the message M having been output on the network. Output conditions are *persistent*; once they are made to hold they continue to hold forever. This squares with our understanding that once a message has been output to the network it can never be removed, and can be input repeatedly.

The set of *name conditions* is precisely the set of names \mathbf{N} . A condition n in \mathbf{N} stands for the name n being in use.

We define the *initial control conditions* of a closed process term p , to be the subset $Ic(p)$ of \mathbf{C} , given by the following structural induction:

$$Ic(out\ new\ \vec{x}M.p) = \{out\ new\ \vec{x}M.p\}$$

$$Ic(in\ pat\ \vec{x}\vec{\psi}M.p) = \{in\ pat\ \vec{x}\vec{\psi}M.p\}$$

$$Ic(\parallel_{i \in I} p_i) = \bigcup_{i \in I} i : Ic(p_i)$$

where the last case also includes the base case *nil*, when the indexing set is empty.

We shortly define the set of *events* \mathbf{Events} as a subset of $Pow(\mathbf{C}) \times Pow(\mathbf{O}) \times Pow(\mathbf{N}) \times Pow(\mathbf{C}) \times Pow(\mathbf{O}) \times Pow(\mathbf{N})$.

So an individual event $e \in \mathbf{Events}$ is a tuple

$$e = (\overset{c}{e}, \overset{o}{e}, \overset{n}{e}, \overset{c'}{e}, \overset{o'}{e}, \overset{n'}{e})$$

where $\overset{c}{e}$ is the set of \mathbf{C} -preconditions of e , $\overset{c'}{e}$ is the set of \mathbf{C} -postconditions of e , etc. Write $\overset{e}{e}$ for $\overset{c}{e} \cup \overset{o}{e} \cup \overset{n}{e}$, all preconditions of e , and $\overset{e'}{e}$ for all postconditions $\overset{c'}{e} \cup \overset{o'}{e} \cup \overset{n'}{e}$.

Earlier in the transition semantics we used actions α to specify the nature of transitions. An event e is associated with a unique action $act(e)$.

The set of events associated with **SPL** is given by an inductive definition. Define \mathbf{Events} to be the smallest set which includes all output, input and indexed events:


- *Output events* $\mathbf{Out}(out\ new\ \vec{x}M.p; \vec{n})$, where $\vec{n} = n_1, \dots, n_l$ are *distinct* names to match the variables $\vec{x} = x_1, \dots, x_l$, consist of an event e with these pre- and postconditions:

$$\overset{c}{e} = \{out\ new\ \vec{x}M.p\}, \quad \overset{o}{e} = \emptyset, \quad \overset{n}{e} = \emptyset,$$

$$\overset{c'}{e} = Ic(p[\vec{n}/\vec{x}]), \quad \overset{o'}{e} = \{M[\vec{n}/\vec{x}]\}, \quad \overset{n'}{e} = \{n_1, \dots, n_l\}.$$

The *action* of an output event is

$$act(\mathbf{Out}(out\ new\ \vec{x}M.p; \vec{n})) = out\ new\ \vec{n}.M[\vec{n}/\vec{x}].$$



An occurrence of the event $\mathbf{Out}(out\ new\ \vec{x}M.p; \vec{n})$ affects the control conditions and puts the new names n_1, \dots, n_l into use, necessarily for the first time as according to the token game the event occurrence must avoid contact with names already in use.

The definition includes the special case when \vec{x} and \vec{n} are empty lists, and we write $\mathbf{Out}(out\ M.p)$ for the output event with no name conditions and action *out M*.


- *Input events* $\mathbf{In}(in\ pat\ \vec{x}\vec{\psi}M.p; \vec{n}, \vec{L})$, where \vec{n} is a list of names to match \vec{x} and \vec{L} is a list of closed messages to match $\vec{\psi}$, consist of an event e with these pre- and postconditions:

$$\overset{c}{e} = \{in\ pat\ \vec{x}\vec{\psi}M.p\}, \quad \overset{o}{e} = \{M[\vec{n}/\vec{x}, \vec{L}/\vec{\psi}]\}, \quad \overset{n}{e} = \emptyset,$$

$$\overset{c'}{e} = Ic(p[\vec{n}/\vec{x}, \vec{L}/\vec{\psi}]), \quad \overset{o'}{e} = \emptyset, \quad \overset{n'}{e} = \emptyset.$$

The *action* of an input event is

$$act(\mathbf{In}(in\ pat\ \vec{x}\vec{\psi}M.p; \vec{n}, \vec{L})) = in\ M[\vec{n}/\vec{x}, \vec{L}/\vec{\psi}].$$



- *Indexed events* $i : e$, where $e \in \mathbf{Events}$, $i \in \mathbf{Indices}$, and

$$\overset{c}{(i : e)} = i : \overset{c}{e}, \quad \overset{o}{(i : e)} = \overset{o}{e}, \quad \overset{n}{(i : e)} = \overset{n}{e},$$

$$\overset{c'}{(i : e)} = i : \overset{c'}{e}, \quad \overset{o'}{(i : e)} = \overset{o'}{e}, \quad \overset{n'}{(i : e)} = \overset{n'}{e}.$$

The *action* of an indexed event is $act(i : e) = i : act(e)$. When E is a subset of events we will generally use $i : E$ to mean $\{i : e \mid e \in E\}$.

In defining the set of conditions and, inductively, the set of events, we have in fact defined a (rather large) net from the syntax of **SPL**. The **SPL**-net has conditions $\mathbf{C} \cup \mathbf{O} \cup \mathbf{N}$ and events **Events**. Its markings \mathcal{M} will be subsets of conditions and so of the form

$$\mathcal{M} = c \cup s \cup t$$

where $c \subseteq \mathbf{C}$, $s \subseteq \mathbf{N}$, and $t \subseteq \mathbf{O}$. By assumption the set of conditions \mathbf{O} are persistent so the net is a contextual net with the following token game - see Appendix A.3.

Letting $c \cup s \cup t$ and $c' \cup s' \cup t'$ be two markings,

$$c \cup s \cup t \xrightarrow{e} c' \cup s' \cup t' \text{ iff}$$

- i) $e \subseteq c \cup s \cup t$ & $e^c \cap c = \emptyset$ & $e^n \cap s = \emptyset$ and
- ii) $c' = (c \setminus e) \cup e^c$ & $s' = s \cup e^n$ & $t' = t \cup e^o$.

In particular, the occurrence of e begins the holding of its name postconditions e^n - these names have to be distinct from those already in use to avoid contact.

5. RELATING SEMANTICS

The behaviour of the **SPL**-net is closely related to the transition semantics given earlier.

THEOREM 5.1.

- i) If $\langle p, s, t \rangle \xrightarrow{\alpha} \langle p', s', t' \rangle$, then for some $e \in \mathbf{Events}$ with $act(e) = \alpha$, $Ic(p) \cup s \cup t \xrightarrow{e} Ic(p') \cup s' \cup t'$.
- ii) If $Ic(p) \cup s \cup t \xrightarrow{e} \mathcal{M}'$ in the **SPL**-net, then for some closed process term p' , for some $s' \subseteq \mathbf{N}$ and $t' \subseteq \mathbf{O}$, $\langle p, s, t \rangle \xrightarrow{act(e)} \langle p', s', t' \rangle$ and $\mathcal{M}' = Ic(p') \cup s' \cup t'$.

DEFINITION 5.1. Let $e \in \mathbf{Events}$. Let p be a closed process, $s \subseteq \mathbf{N}$, and $t \subseteq \mathbf{O}$. Write $\langle p, s, t \rangle \xrightarrow{e} \langle p', s', t' \rangle$ iff $Ic(p) \cup s \cup t \xrightarrow{e} Ic(p') \cup s' \cup t'$ in the **SPL**-net.

6. THE EVENTS OF A PROCESS

Generally for a process p only a small subset of the events **Events** can ever come into play. For this reason it's useful to restrict the events to those reachable in the behaviour of a process.

The set $Ev(p)$ of events of a closed process term p is defined by induction on size:

$$\begin{aligned} Ev(out\ new\ \vec{x}M.p) &= \\ &\{\mathbf{Out}(out\ new\ \vec{x}M.p; \vec{n}) \mid \vec{n} \text{ distinct names}\} \\ &\cup \bigcup \{Ev(p[\vec{n}/\vec{x}]) \mid \vec{n} \text{ distinct names}\} \\ Ev(in\ pat\ \vec{x}\vec{\psi}M.p) &= \\ &\{\mathbf{In}(in\ pat\ \vec{x}\vec{\psi}M.p; \vec{n}, \vec{L}) \mid \vec{n} \text{ names, } \vec{L} \text{ closed messages}\} \\ &\cup \bigcup \{Ev(p[\vec{n}/\vec{x}, \vec{L}/\vec{\psi}]) \mid \vec{n} \text{ names, } \vec{L} \text{ closed messages}\} \\ Ev(\|_{i \in I} p_i) &= \bigcup_{i \in I} i : Ev(p_i). \end{aligned}$$

As an example, the events $Ev(NSL)$ are shown in the Appendix.

A closed process term p denotes a net $Net(p)$ consisting of the global set of conditions $\mathbf{C} \cup \mathbf{O} \cup \mathbf{N}$ built from **SPL**, events $Ev(p)$ and initial control conditions $Ic(p)$. We can

define the token game on the net $Net(p)$ exactly as we did earlier for the **SPL**-net, but this time events are restricted to being in the set $Ev(p)$. It's clear that if an event transition is possible in the restricted net $Net(p)$ then so is it in the **SPL**-net. The converse also holds provided one starts from a marking whose control conditions either belong to $Ic(p)$ or are conditions of events in $Ev(p)$.

DEFINITION 6.1. Let p be a closed process term. Define its *control-conditions* by $p^c = Ic(p) \cup \bigcup \{e^c \mid e \in Ev(p)\}$.

LEMMA 6.1. Let $\mathcal{M} \cap \mathbf{C} \subseteq p^c$. Let $e \in \mathbf{Events}$. Then, $\mathcal{M} \xrightarrow{e} \mathcal{M}'$ in the **SPL**-net iff $e \in Ev(p)$ & $\mathcal{M} \xrightarrow{e} \mathcal{M}'$ in $Net(p)$.

Consequently, in analysing those sequences of event transitions a closed process p can perform it suffices to study the behaviour of $Net(p)$ with its restricted set of events $Ev(p)$. This simplification is especially useful in proving invariance properties because these amount to an argument by cases on the form of events a process can do.

PROPOSITION 6.2. Let $e \in \mathbf{Events}$. Suppose that $\langle p, s, t \rangle$ and $\langle p', s', t' \rangle$ are configurations, and that $\langle p, s, t \rangle$ is proper. If $\langle p, s, t \rangle \xrightarrow{e} \langle p', s', t' \rangle$, then $\langle p', s', t' \rangle$ is also proper.

Important convention: From now on we assume that all configurations $\langle p, s, t \rangle$ are proper.

7. PROVING SECURITY PROPERTIES

To demonstrate the viability of the net semantics as a tool in proving security properties, we use the semantics to derive general principles for proving secrecy and authentication. The principles capture the kind of dependency reasoning found in the strand spaces and inductive methods. To illustrate the principles in action, we apply them to establish an authentication guarantee for the responder part of the NSL protocol. We introduce a diagrammatic style of reasoning which we find helpful.

7.1 General proof principles

From the net semantics we can derive several principles useful in proving authentication and secrecy of security protocols. Write $M \sqsubset M'$ to mean message M in a subexpression of message M' , i.e., \sqsubset is the smallest binary relation on messages such that:

$$\begin{aligned} M &\sqsubset M \\ M &\sqsubset N \Rightarrow M \sqsubset N, N' \text{ and } M \sqsubset N', N \\ M &\sqsubset N \Rightarrow M \sqsubset \{N\}_k \end{aligned}$$

where M, N, N' are messages and k is a key expression. We also write $M \sqsubset t$ iff $\exists M'. M \sqsubset M' \wedge M' \in t$, for a set of messages t .

PROPOSITION 7.1. (*Well-foundedness*) Given a property \mathcal{P} on configurations, if a run

$$\langle p_0, s_0, t_0 \rangle \xrightarrow{e_1} \dots \xrightarrow{e_r} \langle p_r, s_r, t_r \rangle \xrightarrow{e_{r+1}} \dots,$$

contains a configurations s.t. $\mathcal{P}(p_0, s_0, t_0)$ and $\neg \mathcal{P}(p_j, s_j, t_j)$, then there is an event e_h , $0 < h \leq j$, such that $\mathcal{P}(p_i, s_i, t_i)$ for all $i \leq h$ and $\neg \mathcal{P}(p_h, s_h, t_h)$.

We say that a name $m \in \mathbf{N}$ is *fresh on an event* e if $m \in e^n$ and we write $Fresh(m, e)$.

PROPOSITION 7.2. (*Freshness*) Within a run

$$\langle p_0, s_0, t_0 \rangle \xrightarrow{e_1} \dots \xrightarrow{e_r} \langle p_r, s_r, t_r \rangle \xrightarrow{e_{r+1}} \dots,$$

the following properties hold:

- i) If $n \in s_i$ then either $n \in s_0$ or there is a previous event e_j such that $\text{Fresh}(n, e_j)$.
- ii) Given a name n there is at most one event e_i s.t. $\text{Fresh}(n, e_i)$.
- iii) If $\text{Fresh}(n, e_i)$ then for all $j < i$ the name n does not appear in $\langle p_j, s_j, t_j \rangle$.

PROPOSITION 7.3. (*Control precedence*) Within a run

$$\langle p_0, s_0, t_0 \rangle \xrightarrow{e_1} \dots \xrightarrow{e_r} \langle p_r, s_r, t_r \rangle \xrightarrow{e_{r+1}} \dots,$$

if $b \in {}^c e_i$ either $b \in \text{Ic}(p_0)$ or there is an earlier event e_j , $j < i$, such that $b \in e_j^c$.

PROPOSITION 7.4. (*Output-input precedence*) In a run

$$\langle p_0, s_0, t_0 \rangle \xrightarrow{e_1} \dots \xrightarrow{e_r} \langle p_r, s_r, t_r \rangle \xrightarrow{e_{r+1}} \dots,$$

if $M \in {}^o e_i$, then either $M \in t_0$ or there is an earlier event e_j , $j < i$, such that $M \in e_j^o$.

7.2 An example: authentication for NSL

We will prove authentication for a responder in an NSL protocol in the sense that: to any complete session of agent B_0 as responder, apparently with agent A_0 , there corresponds a complete session of agent A_0 as initiator. We refer to the Appendix for the events of NSL.

In the proof it's helpful to make use of a form of diagrammatic reasoning which captures the precedence of events. When the run is understood, we draw $e \longrightarrow e'$ when e precedes e' in the run, allowing $e = e'$.

THEOREM 7.5 (AUTHENTICATION). If a run of NSL

$$\langle \text{NSL}, s_0, t_0 \rangle \xrightarrow{e_1} \dots \xrightarrow{e_r} \langle p_r, s_r, t_r \rangle \xrightarrow{e_{r+1}} \dots,$$

contains the responder events b_1, b_2, b_3 , with actions

$$\begin{aligned} \text{act}(b_1) &= \text{resp} : B_0 : i : \text{in}\{m_0, A_0\}_{\text{Pub}(B_0)}, \\ \text{act}(b_2) &= \text{resp} : B_0 : i : \text{out new } n_0 \{m_0, n_0, B_0\}_{\text{Pub}(A_0)}, \\ \text{act}(b_3) &= \text{resp} : B_0 : i : \text{in}\{n_0\}_{\text{Pub}(B_0)}, \end{aligned}$$

for an index i , and $\text{Priv}(A_0) \not\sqsubseteq t_0$, then the run contains initiator events a_1, a_2, a_3 with $a_3 \longrightarrow b_3$, where, for some index j ,

$$\begin{aligned} \text{act}(a_1) &= \text{init} : (A_0, B_0) : j : \text{out new } m_0 \{m_0, A_0\}_{\text{Pub}(B_0)}, \\ \text{act}(a_2) &= \text{init} : (A_0, B_0) : j : \text{in}\{m_0, n_0, B_0\}_{\text{Pub}(A_0)}, \\ \text{act}(a_3) &= \text{init} : (A_0, B_0) : j : \text{out}\{n_0\}_{\text{Pub}(B_0)}. \end{aligned}$$

PROOF. By control precedence we obtain:


$$b_1 \longrightarrow b_2 \longrightarrow b_3.$$

Consider the property of configurations

$$Q(p, s, t) \Leftrightarrow \forall M \in t. n_0 \sqsubset M \Rightarrow \{m_0, n_0, B_0\}_{\text{Pub}(A_0)} \sqsubset M.$$

By freshness, the property Q holds immediately after b_2 , but clearly not immediately before b_3 . By well-foundedness

there is a earliest event following b_2 but preceding b_3 that violates Q . Let e be such an event.




Inspecting the events of the NSL protocol (see Appendix), using the assumption that $\text{Priv}(A_0) \not\sqsubseteq t_0$, one can show that e can only be an initiator event a'_3 with action


$$\text{act}(a'_3) = \text{init} : (A, B_0) : j : \text{out}\{n_0\}_{\text{Pub}(B_0)}$$

for some index j and agent A . There must also be preceding events a'_1, a'_2 with actions

$$\begin{aligned} \text{act}(a'_1) &= \text{init} : (A, B_0) : j : \text{out new } m \{m, A\}_{\text{Pub}(B_0)} \\ \text{act}(a'_2) &= \text{init} : (A, B_0) : j : \text{in}\{m, n_0, B_0\}_{\text{Pub}(A)} \end{aligned}$$



Since $\text{Fresh}(b_2, n_0)$, the event b_2 must precede a'_2 . The property Q holds on configurations up to a'_3 and, in particular, on the configuration immediately before a'_2 . From this we conclude that $m = m_0$ and $A = A_0$. Hence $a'_3 = a_3$, $a'_2 = a_2$, and $a'_1 = a_1$ as described below.



(Since $\text{Fresh}(a_1, m_0)$, the event a_1 precedes b_1 .) \square

8. RELATING SECURITY MODELS

We have related our net semantics of SPL to a transition semantics. Now we establish its relations to the security models of strand spaces, inductive rules, as well as other traditional models. In security protocols we are largely interested in safety properties, which reduce to a property holding of all finite behaviours. Thus it suffices to show how a finite behaviour in one model can be matched by the finite behaviour in another. In relating the net semantics to strand spaces and inductive rules we need to constrain process terms, to allow some repetition of actions, though this does not seem unduly restrictive in formalising security protocols.

8.1 Strand spaces

In relating the net semantics to strand spaces we must face the fact that strand spaces don't compose readily, not using traditional process operations at least. Their form doesn't allow prefixing by a single event. Nondeterminism only arises through the choice as to where input comes from, and there is not a recognisable nondeterministic sum of strand spaces. Even an easy definition of parallel composition by juxtaposition is thwarted if "unique origination" is handled as a global condition on the entire strand space. This complicates the relation between a compositional semantics and strand spaces.

We can however relate the net behaviour of a $!$ -par process to that of an associated strand space; a $!$ -par process is a closed process of the form $!!_{i \in I} p_i$ for which no subterm

p_i contains a parallel composition. In proving the relation (though unfortunately not in this short write-up) we find it useful to extend strand spaces in order to compose them, chiefly with conflict to permit their nondeterministic sum, and then finally to observe that for processes with replication the conflict can be eliminated, without upsetting the strand-space behaviour. (Strand spaces can be viewed as special forms of event structures - see below; so ideas, such as the use of a conflict relation, can be adapted from there.)

DEFINITION 8.1. *A strand space consists of $S = \langle S_i \rangle_{i \in I}$ an indexed set of strands. An individual strand S_i , where $i \in I$, is a finite sequence of output or input events carrying respectively output or input actions of the kind $out\ \bar{n}M$ or $in\ M$, where M is a closed message and \bar{n} a list of distinct names that are intended to be fresh (“uniquely originating”) at the event. We permit only strands on which any “new” names do not appear in previous actions of the strand. (A set of strands is canonically a strand space in which each strand has itself as index.)*

As usual, a strand space can be seen as a graph whose nodes are of the form (i, l) with $i \in I$ index of a strand and l position of an event in that strand ($1 \leq l \leq length(s_i)$). Each node uniquely identifies an event in a strand. Edges are of two different kinds: \Rightarrow between two nodes that identify two events of a same strand, one immediately preceding the other and \rightarrow between two nodes identifying respectively an output event and an input event with the same message. A bundle of a strand space S is a finite, acyclic subgraph s.t.

- if a node belongs to the bundle, so do all nodes that precede it on its strand, together with their incident \Rightarrow edges,
- each input node has exactly one incoming \rightarrow edge,
- two different strands that have a “new” name in common don’t both contribute to the same bundle.

We write \Rightarrow_C and \rightarrow_C for edges of a strand space graph belonging to a bundle C . Our definition is not quite standard. But the only significant difference is in the treatment of unique origination which is taken care of in the definition of bundle rather than being a condition on the entire strand space - the “parametric strand spaces” of [3] achieve the same effect and are closely related.

A strand space can be seen as a form of *event structure* [13]. A strand space determines a *stable event structure*, whose family of configurations is the same as the bundles of the strand space; the bundles of a strand space when ordered by inclusion form a *stable family* which ensures not only that each configuration of events in the family can be equipped with a local partial order of causal dependency, but that at the cost renaming events these local partial orders can be extended to a global partial order of causal dependency, yielding a *prime event structure*.

Often in strand spaces the precise identity of indices does not matter. A *re-indexing* of a strand space $S = \langle S_i \rangle_{i \in I}$ is a permutation π of I such that S_i and $S_{\pi(i)}$ are sequences of the same length with the same actions at corresponding events. A re-indexing of a strand space induces a re-indexing on its bundles; a bundle’s nodes and arcs are changed according to the correspondence given by π .

To relate the net behaviour of a process to its behaviour as a strand space we need to linearise bundles. More precisely:

DEFINITION 8.2. *Given a bundle C of a strand space S , a linearisation of C is a sequence of nodes $e_1 \dots e_k$ such that $\{e_1, \dots, e_k\}$ are all the nodes of C and for all e in C and all e_i in L , if $e \Rightarrow_C e_i$ or $e \rightarrow_C e_i$ then e precedes e_i in the sequence. An event-linearisation of a bundle is the sequence of strand-space events associated with the nodes of a linearisation.*

Let p be a l -par process and s a set of names containing all names in p . Take $Tr(p, s)$ to be the strand space with strands consisting of all the maximal sequences $e_1 \dots e_k$ of events in $Ev(p)$ such that:

- i) ${}^c e_1 \subseteq Ic(p)$ and
- ii) for all $i, 1 \leq i < k$, we have $e_i^n \cap (s \cup \{e_j^n \mid j < i\}) = \emptyset$ and $e_i^c = {}^c e_{i+1}$.

Sequences satisfying the above conditions are necessarily finite as the size of control conditions strictly decreases along the sequence. The events of the net are already associated with input and output actions. The net and strand space behaviour are closely related:

THEOREM 8.3. *Given p a l -par process and s set of names containing all names in p , we have that:*

- i) *The sequence of events in a finite run in $Net(p)$ from the initial configuration $\langle p, s, \emptyset \rangle$ is an event-linearisation of a bundle over $Tr(p, s)$.*
- ii) *Every bundle over $Tr(p, s)$ can be re-indexed so that any of its event-linearisations is a run in $Net(p)$.*

The only way a strand space can cope with there being a nonempty set of initial output messages is through the slight clumsiness of introducing extra output events; we avoid this above by assuming the initial set of output messages is empty.

8.2 Inductive rules

Paulson’s inductive rules for a security protocol capture the actions it and a spy can perform [9]. Through allowing persistent conditions, we can represent a collection of inductive rules as a net in which the events stand for rule instances and runs to sequences of rule instances which form a derivation from the rules. In particular, instances of inductive rules for security protocols can be represented as events in a net for which all but the name conditions are persistent. According to such a semantics, once a protocol can input it can do so repeatedly. Once it can output generating new names it can do so repeatedly, provided this doesn’t lead to clashes with names already in use. Paulson’s traces and the associated runs of the net will necessarily include such “stuttering.”

We define a net of rule instances from a closed process term. Take the set of “rule-conditions” to consist of name conditions and persistent output conditions, as before, but now with additional persistent conditions consisting of closed input and output process terms. Let r be the function from **SPL**-conditions to rule-conditions which removes the indices tagging control conditions and leaves output and name conditions unchanged. Extend r to **SPL**-events: let r replace all the control conditions of an **SPL**-event by their images under r - intuitively, an event is replaced by a rule instance. Define the “net of rule instances” $R(p)$ of a closed process term p to be the net with rule-conditions and events the image $r\ Ev(p)$.

For a closed process term p , let p^* be the process term obtained by inserting a replication before every input and output process subterm in p . Note that $R(p^*) = R(p)$ as R drops indices. Now, having restricted to a process with sufficient replication, we can establish a close relation between the behaviours of $Net(p^*)$ and $R(p^*)$.

THEOREM 8.4. *Let p be a closed process term. Let t be a subset of closed messages and s a subset of names including those of p and t . Let $\mathcal{M}_0 = Ic(p^*) \cup s \cup t$.*

- i) *A run $\mathcal{M}_0 \xrightarrow{e_1} \dots \xrightarrow{e_l} \mathcal{M}_l$ of $Net(p^*)$ yields a run $r\mathcal{M}_0 \xrightarrow{r(e_1)} \dots \xrightarrow{r(e_l)} r\mathcal{M}_l$ of $R(p^*)$.*
- ii) *To $\mathcal{M}'_0 \xrightarrow{e'_1} \dots \xrightarrow{e'_l} \mathcal{M}'_l$ a run of $R(p^*)$ with $\mathcal{M}'_0 = r\mathcal{M}_0$, there is a run $\mathcal{M}_0 \xrightarrow{e_1} \dots \xrightarrow{e_l} \mathcal{M}_l$ of $Net(p^*)$, with $r(e_i) = e'_i$ and $r(\mathcal{M}_i) = \mathcal{M}'_i$ for all i , $0 < i \leq l$.*

8.3 Basic nets, trace languages and event structures

Because strand spaces can be easily turned into event structures, Section 8.1 yields an event structure for each l -par process. But, without any restrictions, we can relate the net semantics to traditional independence models such as event structures and Mazurkiewicz trace languages. The crux of the construction is that of eliminating the persistent conditions from the net $Net(p)$, of a closed process term p , in an initial marking $Init(p) \cup s \cup t$, to produce a basic net. It's well-known how to "unfold" a basic net to a Mazurkiewicz trace language and event structure. Assume $Net(p)$ has input events In and output events Out . Then:

THEOREM 8.5. *There is a basic net \mathcal{N} with events*

$$E = Out \cup \{(*, e) \mid e \in In \ \& \ e \in t\} \cup \{(e_1, e) \mid e \in In \ \& \ e_1 \in Out \ \& \ e = e_1^o\}$$

such that letting the map $\sigma : E \rightarrow Out \cup In$ leave events in Out unchanged and project pairs $(, e)$, (e_1, e) to the component e :*

- i) *If \mathcal{N} has a run with events e'_1, \dots, e'_k , then there is a run with events $\sigma(e'_1), \dots, \sigma(e'_k)$ of $Net(p)$ from the initial marking $Init(p) \cup s \cup t$.*
- ii) *If $Net(p)$ has a run e_1, \dots, e_k from the initial marking $Init(p) \cup s \cup t$, then \mathcal{N} has a run e'_1, \dots, e'_k where $e_1, \dots, e_k = \sigma(e'_1), \dots, \sigma(e'_k)$.*

The construction used to obtain \mathcal{N} above is an example of the construction for eliminating colours from a coloured net - see [14]; first colours are introduced to the persistent conditions and input events of $Net(p)$ to distinguish the different ways in which they are made to occur, and then eliminated through splitting the conditions and events according to their colours. The result in this case is a basic net. Its runs form a Mazurkiewicz trace language from which we can then obtain an event structure - see [15].

9. CONCLUDING REMARKS

A role of the language **SPL** is that it can support, and so help relate, different semantics useful in the analysis of security protocols - we have seen several examples. Future goals are to relate to a probabilistic semantics, moving away from

the perfect cryptography assumption, and to study what equivalences and compositional reasoning fit with the rather intensional event-based methods dealt with here. We hope our work helps towards a more high-level analysis of security protocols. To this end, we see the net semantics of **SPL** as giving a potentially useful, concrete model theory for logics for security protocols. The net runs are histories on which to interpret security properties, perhaps expressed in the style of BAN logic. On another tack, Petri nets form models of linear logic [5], close it seems to the linear logic for security protocols based on multiset rewriting [2]. As suggested by an anonymous referee, the expressiveness of **SPL** could be useful in characterising natural classes of protocols with decidable security properties.

Acknowledgements. We are grateful to Stewart Lee and an anonymous referee for suggestions and encouragement. GW thanks Martın Abadi for a timely overview of approaches to security protocols while they were both visiting Microsoft Research, Cambridge.

10. REFERENCES

- [1] M. Abadi and A. Gordon. A calculus for cryptographic protocols: The Spi calculus. In *4th ACM CCS*, 1997.
- [2] I. Cervesato, N. A. Durgin, M. Kanovich, and A. Scedrov. Interpreting strands in linear logic. In *FMCS'00*, 2000.
- [3] I. Cervesato, N. A. Durgin, P. D. Lincoln, J. C. Mitchell, and A. Scedrov. Relating strands and multiset rewriting for security protocol analysis. In *13th IEEE CSFW*, 2000.
- [4] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Trans. on Inf. Theory*, 2(29), 1983.
- [5] U. Engberg and G. Winskel. Linear logic on Petri nets. In *REX, a decade of concurrency, 1993*, volume 803 of *LNCS*, 1994.
- [6] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In *TACAS'96*, volume 1055 of *LNCS*, 1996.
- [7] R. Milner. *Communicating and mobile systems: The π -calculus*. Cambridge University Press, 1999.
- [8] U. Montanari and F. Rossi. Contextual nets. *Acta Informatica*, (32), 1995.
- [9] L. C. Paulson. Proving security protocols correct. In *LICS, 14th Symp.*, July 1999.
- [10] A. M. Pitts and I. Stark. Observable properties of higher order functions that dynamically create local names, or: What's new? In *MFCS 18th Int. Symp.*, volume 711 of *LNCS*, 1993.
- [11] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Comm. of the ACM*, 21(2), 1978.
- [12] J. Thayer, J. Herzog, and J. Guttman. Strand spaces: Why is a security protocol correct? In *1998 IEEE Symposium on Security and Privacy*.
- [13] G. Winskel. Event structures. In *Advanced Course on Petri nets*, volume 255 of *LNCS*, 1986.
- [14] G. Winskel. Petri nets, algebras, morphisms, and compositionality. *Inf. & Comput.*, 72, 1987.
- [15] G. Winskel and M. Nielsen. Models for concurrency. In *Handbook of Logic in Comp. Sci.*, volume 4, 1995.

APPENDIX

A. PETRI NETS

The explanation of general Petri nets involves a little algebra of multisets (or bags), which are like sets but where multiplicities of elements matters. It's convenient to also allow infinite multiplicities, so we adjoin an extra element ∞ to the natural numbers, though care must be taken to avoid subtracting ∞ . ∞ -Multisets support addition $+$ and multiset inclusion \leq , and even multiset subtraction $X - Y$ provided $Y \leq X$ and Y has no infinite multiplicities, in which case we call Y simply a multiset.

A.1 General Petri nets

A *general Petri net* (*place-transition system*) consists of

- a set of *conditions* (or *places*), P ,
- a set of *events* (or *transitions*), T ,
- a *precondition map* pre , which to each $t \in T$ assigns a multiset $pre(t)$ over P - it is traditional to write $\cdot t$ for $pre(t)$,
- a *postcondition map* $post$ which to each $t \in T$ assigns an ∞ -multiset $post(t)$ over P , traditionally written $t \cdot$,
- a *capacity function* Cap which is an ∞ -multiset over P , assigning a nonnegative number or ∞ to each condition p , bounding the multiplicity to which the condition can hold; a capacity of ∞ means the capacity is unbounded.

A state of a Petri net consists of a *marking*, an ∞ -multiset \mathcal{M} over P bounded by the capacity function, *i.e.* $\mathcal{M} \leq Cap$. A marking captures a notion of distributed, global state.

Token game for general nets: Markings can change as events occur, precisely how being expressed by the transitions $\mathcal{M} \xrightarrow{t} \mathcal{M}'$ events t determine between markings \mathcal{M} and \mathcal{M}' . For markings $\mathcal{M}, \mathcal{M}'$ and $t \in T$, define

$$\mathcal{M} \xrightarrow{t} \mathcal{M}' \text{ iff } \cdot t \leq \mathcal{M} \text{ and } \mathcal{M}' = \mathcal{M} - \cdot t + t \cdot$$

An event t is said to have *concession* (or be *enabled*) at a marking \mathcal{M} iff its occurrence leads to a marking, *i.e.* iff

$$\cdot t \leq \mathcal{M} \text{ and } \mathcal{M} - \cdot t + t \cdot \leq Cap.$$

There is a widely-used graphical notation for nets in which events are represented by squares, conditions by circles and the pre- and postcondition maps by directed arcs carrying numbers or ∞ as weights. A marking is represented by the presence of tokens on a condition, the number of tokens representing the multiplicity to which the condition holds. When an event with concession occurs tokens are removed from its preconditions and put on its postconditions with multiplicities according to the pre- and postcondition maps. Because of this presentation, the transition relation on Petri nets is described as the “token game”.

A.2 Basic nets

We instantiate the definition of general Petri nets to an important case where in all the multisets the multiplicities are either 0 or 1, and so can be regarded as sets. In particular, we take the capacity function to assign 1 to every condition, so that markings become simply subsets of conditions. The general definition now specialises to the following.

A *basic Petri net* consists of

- a set of *conditions*, B ,

- a set of *events*, E , and
- two maps: a *precondition map* $pre : E \rightarrow \mathcal{P}ow(B)$, a *postcondition map* $post : E \rightarrow \mathcal{P}ow(B)$. We can still write $\cdot e$ for the preconditions and $e \cdot$ for the postconditions of $e \in E$ and we require $\cdot e \cup e \cdot \neq \emptyset$.

Now a *marking* consists of a subset of conditions, specifying those conditions which hold.

Token game for basic nets: For markings $\mathcal{M}, \mathcal{M}' \subseteq B$ and event $e \in E$, define

$$\begin{aligned} \mathcal{M} \xrightarrow{e} \mathcal{M}' \text{ iff } & (1) \cdot e \subseteq \mathcal{M} \ \& \ (\mathcal{M} \setminus \cdot e) \cap e \cdot = \emptyset \text{ and} \\ & (2) \mathcal{M}' = (\mathcal{M} \setminus \cdot e) \cup e \cdot. \end{aligned}$$

Property (1) expresses that the event e has *concession* at the marking \mathcal{M} . Returning to the definition of concession for general nets, of which it is an instance, it ensures that the event does not load another token on a condition that is already marked. Property (2) expresses in terms of sets the marking that results from the occurrence of an event. So, an occurrence of the event ends the holding of its preconditions and begins the holding of its postconditions. (It is possible for a condition to be both a precondition and a postcondition of the same event, in which case the event is imagined to end the precondition before immediately restarting it.)

There is *contact* at a marking \mathcal{M} when for some event e

$$\cdot e \subseteq \mathcal{M} \ \& \ (\mathcal{M} \setminus \cdot e) \cap e \cdot \neq \emptyset.$$

The occurrence of an event is blocked through conditions, which it should begin, holding already. Blocking through contact is consistent with the understanding that the occurrence of an event should end the holding of its preconditions and begin the holding of its postconditions; if the postconditions already hold, and are not also preconditions of the event, then they cannot begin to hold on the occurrence of the event. Avoiding contact ensures the freshness of names in the semantics of name creation.

Basic nets are important because they are related to many other models of concurrent computation, in particular, Mazurkiewicz trace languages (languages subject to trace equivalence determined by the independence of actions) and event structures (sets of events with extra relations of causality and conflict) - see [15].

A.3 Nets with persistent conditions

Sometimes we have use for conditions which once established continue to hold and can be used repeatedly. If we are to use net events to represent rules of the kind we find in inductive definitions, we need conditions that persist.

Persistent conditions can be understood as an abbreviation for conditions within general nets which once they hold, do so with infinite multiplicity. Consequently any number of events can make use of them as preconditions but without their ever ceasing to hold. Such conditions, having unbounded capacity, can be postconditions of several events without there being conflict.

To be more precise, we modify the definition of basic net given above by allowing certain conditions to be *persistent*. A net with persistent conditions will still consist of events and conditions related by pre- and postcondition maps which to an event will assign a set of preconditions and a set of postconditions. But, now amongst the conditions are the persistent conditions forming a subset P . A marking of a net with persistent conditions will be simply a subset of

conditions, of which some may be persistent. Nets with persistent conditions have arisen independently several times and have been studied for example in *contextual nets* [8].

A net with persistent conditions can be understood on its own terms, or as standing for a general net with the same sets of conditions and events. The general net's capacity function will be either 1 or ∞ on a condition, being ∞ precisely on the persistent conditions. When p is persistent, $p \in e'$ is interpreted in the general net as arc weight $(e')_p = \infty$, and $p \in e$ as $(e)_p = 1$. A marking of a net with persistent conditions will correspond to a marking in the general Petri net in which those persistent conditions which hold do so with infinite multiplicity. Graphically, we will distinguish persistent conditions by drawing them as double circles: \odot


Token game with persistent conditions: The token game is modified to account for the subset of conditions P being persistent. Let \mathcal{M} and \mathcal{M}' be markings (*i.e.* subsets of conditions), and e an event. Define

$$\begin{aligned} \mathcal{M} \xrightarrow{e} \mathcal{M}' \text{ iff } & (1) \ 'e \subseteq \mathcal{M} \ \& \ (\mathcal{M} \setminus (e \cup P)) \cap e' = \emptyset \text{ and} \\ & (2) \ \mathcal{M}' = (\mathcal{M} \setminus e) \cup e' \cup (\mathcal{M} \cap P) . \end{aligned}$$



B. THE EVENTS OF NSL

We classify the events $Ev(NSL)$ of the NSL protocol:

Initiator events:




Responder events:




Spy events: For compactness we overlap at shared conditions.


Composing, $Spy_1 \equiv in\ \psi_1.in\ \psi_2.out(\psi_1, \psi_2)$



Decomposing, $Spy_2 \equiv in(\psi_1, \psi_2).out\ \psi_1.out\ \psi_2$



Encrypting, $Spy_3 \equiv in\ x.in\ \psi.out\ \{\psi\}_{Pub(x)}$



Decrypting, $Spy_4 \equiv in\ Priv(x).in\ \{\psi\}_{Pub(x)}.out\ \psi$

