

# Security properties and CSP\*

Steve Schneider

Royal Holloway, University of London  
Egham, Surrey, TW20 0EX

## Abstract

*Security properties such as confidentiality and authenticity may be considered in terms of the flow of messages within a network. To the extent that this characterisation is justified, the use of a process algebra such as Communicating Sequential Processes (CSP) seems appropriate to describe and analyse them. This paper explores ways in which security properties may be described as CSP specifications, how security mechanisms may be captured, and how particular protocols designed to provide these properties may be analysed within the CSP framework. The paper is concerned with the theoretical basis for such analysis. A sketch verification of a simple example is carried out as an illustration.*

## 1 Introduction

Security protocols are designed to provide properties such as authentication, key exchanges, key distribution, non-repudiation, proof of origin, integrity, confidentiality and anonymity, for users who wish to exchange messages over a medium over which they have little control. These properties are often difficult to characterise formally (or even informally). The protocols themselves often contain a great deal of combinatorial complexity, making their verification extremely difficult and prone to error. This paper promotes the view that process algebra can provide a single framework both for modelling protocols and for capturing security properties, facilitating verification and debugging. It is a discussion paper, proposing possible approaches rather than providing definitive answers. It considers only confidentiality and authenticity here; other properties have been considered in [7].

It has been argued that security properties should be considered as properties concerning the flow of messages within a network. To the extent that this characterisation is justified, the use of a process algebra

such as CSP [3] seems appropriate to describe and analyse them. This paper considers ways in which security properties may be described using the notation of CSP, how security mechanisms may be captured, and how particular protocols designed to provide these properties may be analysed within the CSP framework.

The approach presented is rather general, and it is clear that the modelling of particular properties and analysis of particular protocols will require tailoring of the model presented here. But this paper aims at exploring a general approach rather than trying to construct a universal model suitable for handling all possible security issues, which is probably an unrealistic goal.

Security properties are generally properties requiring that something bad should not occur (though they are not exclusively of this form); these tend to be considered as safety properties. Of course, particular communication protocols will also aim to be live (something good should occur), in that they will be designed to achieve goals such as delivery of messages. But there is a distinction to be drawn between the security requirements implemented by such a protocol, and its liveness requirements which are important for communication but which are generally independent of security. It is possible that there are some security properties which can be expressed only as liveness properties; these are outside the scope of this paper. Hence the traces model for CSP will be adequate for our present needs: to analyse properties of the form ‘something bad should not happen’ it is sufficient to focus on what systems may do, rather than what they must do. All equivalences and refinements expressed in this paper are therefore grounded in the traces model.

This paper is structured as follows: The relevant CSP notation is introduced in Section 2; Section 3 contains a discussion of how the properties of confidentiality and authentication may be captured within CSP independently of security protocols; Section 4 discusses the modelling of security protocols and of the networks on which they are implemented, and contains an exam-

ple verification of a toy confidentiality property to illustrate the material; finally Section 5 discusses the approach and its potential.

## 2 CSP notation

CSP is an abstract language designed specifically for the description of communication patterns of concurrent system components that interact through message passing. It is underpinned by a theory which supports analysis of systems described in CSP. It is therefore well suited to the description and analysis of network protocols: protocols can be described within CSP, as can the relevant aspects of the network. Their interactions can be investigated, and certain aspects of their behaviour can be verified through use of the theory. This section introduces the notation and ideas used in this paper. In particular, only the traces model for CSP is used here. For a fuller introduction to the language the reader is referred to [3].

### Events

Systems are modelled in terms of the events that they can perform. The set of all possible events (fixed at the beginning of the analysis) is denoted  $\Sigma$ . Events may be atomic in structure or may consist of a number of distinct components. For example, an event  $put.5$  consists of two parts: a channel name  $put$ , and a data value 5. An example of events used in this paper are those of the form  $c.i.j.m$  consisting of a channel  $c$ , a source  $i$ , a destination  $j$  and a message  $m$ . If  $M$  and  $N$  are sets of messages, then  $M.N$  will be the set of messages  $\{m.n \mid m \in M \wedge n \in N\}$ . If  $m$  is a single message then we elide the set brackets and define  $m.N$  to be  $\{m\}.N$ . Thus for example the set of events  $i.N.m = \{i.n.m \mid n \in N\}$ . A channel  $c$  is said to be of type  $M$  if any message  $c.m \in \Sigma$  has that  $m \in M$ .

### Processes

Processes are the components of systems. They are the entities that are described using CSP, and they are described in terms of the possible events that they may engage in. The process  $STOP$  is the process that can engage in no events at all; it is equivalent to deadlock. If  $P$  is a process then the process  $a \rightarrow P$  is only able to initially perform  $a$ , following which it will behave in the way described by  $P$ . The process  $P \square Q$  (pronounced ‘ $P$  choice  $Q$ ’) can behave either as  $P$  or as  $Q$ : its possible communications are those of  $P$  and those of  $Q$ . An indexed form of choice  $\square_{i \in I} P_i$  is able to behave as any of its arguments  $P_i$ .

Processes may also be composed in parallel. If  $A$  is a set of events then the process  $P \parallel[A] Q$  behaves as  $P$  and  $Q$  acting concurrently, with the requirement that they have to synchronise on any event in the synchronisation set  $A$ ; events not in  $A$  may be performed by either process independently of the other. A special form of parallel operator in which the two components do not interact on any events is  $P \parallel\!\!\! \parallel Q$  which is equivalent to  $P \parallel[\{\}] Q$ .

Processes may be recursively defined by means of equational definitions. Process names must appear on the left hand side of such definitions, and CSP expressions which may include those names appear on the right hand side. For example, the definition

$$LIGHT = on \rightarrow off \rightarrow LIGHT$$

defines a process  $LIGHT$  whose only possible behaviour is to perform  $on$  and  $off$  alternately.

### Traces

The semantics of a process  $P$  is defined to be the set of sequences of events ( $traces(P)$ ) that it may possibly perform. Examples of traces include  $\langle \rangle$  (the empty trace, which is possible for any process) and  $\langle on, off, on \rangle$  which is a possible trace of  $LIGHT$ .

A useful operator on traces is projection: If  $A$  is a set of events then the trace  $tr \upharpoonright A$  is defined to be the maximal subsequence of  $tr$  all of whose events are drawn from  $A$ . If  $A$  is a singleton set  $\{a\}$  then we overload notation and write  $tr \upharpoonright a$  for  $tr \upharpoonright \{a\}$ .

### Analysing processes

Specifications are given as predicates on traces, and a process  $P$  satisfies a specification  $S$  if all of its traces satisfy  $S$ :

$$P \text{ sat } S \Leftrightarrow \forall tr \in traces(P). S$$

A process  $P$  is refined by a process  $Q$  (written  $P \sqsubseteq Q$ ) if  $traces(Q) \subseteq traces(P)$ . This means that if  $P$  meets a specification then  $Q$  will also meet it. Model-checking techniques allow the refinement relation  $P \sqsubseteq Q$  to be checked mechanically (for finite-state processes) using the tool FDR [1]. Since two processes are equal if each refines the other, this means that equality of processes can also be checked using FDR.

## 3 Security properties

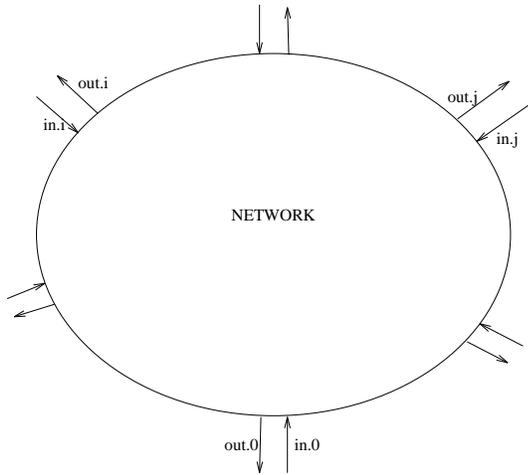
A network provides a means for users, such as people or applications programs, to communicate by sending

and receiving messages. This situation may be modelled at a high level of abstraction in CSP as a process *NET* which provides to each user two ways of interacting with it: sending messages to other parties, and receiving messages from other parties.

We will assume a universal set *MESSAGE* of all messages that might be sent by any party, and we will consider the users to be numbered up to *n*:

$$USER = \{0, 1, \dots, n\}$$

The channel employed by user *i* to input messages to the network will be the input channel *in.i*, of type *USER.MESSAGE*. An input of the form *in.i.j.m* is considered an instruction from user *i* to transmit message *m* to user *j*.



**Figure 1. High level view of the network**

The channel employed by user *i* to receive messages output from the network will be the output channel *out.i*, of type *USER.MESSAGE*. An output of the form *out.i.j.m* is considered to be receipt by user *i* of message *m* sent by user *j*.

Users' requirements on the network are expressed in terms of the behaviour of the network as a whole, and CSP has been used successfully for some years in the description and analysis of communications protocols. Common safety and liveness properties are readily expressed in terms of the possible behaviour of the network with respect to the users. For example, the property that no spurious messages are generated is captured as a safety property that requires any output message to have previously been input: if *out.i.j.m* appears in a trace, then *in.j.i.m* must have already occurred. The liveness requirement that no message is lost can be formalised as follows: for any input message *in.i.j.m* the corresponding output *out.j.i.m* must

eventually become available. Generally these properties are expressed precisely and formally in terms of the semantic models of CSP.

Although it is necessary to know the internal structure of the network in order to demonstrate that it provides particular services, the services or properties themselves should be expressible simply in terms of the interactions the network offers its users. This is the case for common communications protocols, and in this paper we take the view that security properties can be captured in the same way. We therefore examine and offer definitions of these properties before considering the network at any finer level of detail.

There are two views from which security properties can be considered:

- from the viewpoint of the users of the network, who do not know which other parties are to be trusted. Properties expressed from this viewpoint will generally include assumptions (implicitly or explicitly) that a user's communication partner will not act contrary to the aims of the protocol. For example, that any shared secrets should not be disclosed to third parties.
- from a high-level 'God's eye view' which identifies those nodes which follow their protocols faithfully, and also identifies those which are engaging in more general activity, perhaps in attempting to find a flaw in a protocol. If this view is taken, then care should be taken to ensure that this privileged information is not accidentally used in the protocol description: the responses of a node should not be dependent on information which is available only at the high-level view. In some circumstances, a node may not have knowledge concerning its communication partner; in other cases, a protocol may be invoked only when communicating with particular known and trusted users (how this knowledge and trust is obtained is outside the scope of this report).

In this report we will follow the high-level view. This means we can postulate the existence of an enemy whose identity is known and can be used in the formulation of security properties. We will use  $0 \in USER$  as the name of this enemy process. Later we will justify the decision to use only a single enemy, by arguing that further enemies do not increase the vulnerability of protocols: the single enemy in a sense encapsulates the behaviour of all enemies.

## Confidentiality

Confidentiality for a particular set of messages  $M$  is achieved when users may communicate any message drawn from the set  $M$  without the possibility of any user other than the intended recipient receiving it. In other words, if an input  $in.i.j.m$  occurs, then any (subsequent) output  $out.h.l.m$  must be for user  $j$ : i.e.  $h = j$ . Thus given user  $j$  and message  $m$ , if an output  $out.j.i.m$  occurs (for some  $i$ ) then some user  $l$  (not necessarily  $i$ ) must have sent that message to  $j$ : there must be some previous input of the form  $in.l.j.m$ . Thus  $j$  cannot obtain messages that were intended for some other user.

For analysis purposes, we will consider the system from the God's-eye view: the only user which might obtain messages intended for some other user will be user 0. Hence confidentiality will be captured as a specification requiring that any message output to user 0 must have actually been sent to user 0. We restrict attention to the message set  $M$  as being those messages which are intended to remain confidential. We also assume they cannot be generated by user 0 (which would be true for example for signed messages), though this is a simplifying assumption that is not justified in all circumstances. This assumption is implicit in the definition, since otherwise 0 could simply guess any confidential message. Other messages (such as encrypted messages or control messages) will in general be available to eavesdroppers, but confidentiality is not concerned with protecting these.

These considerations may be captured as a trace specification

**Definition 3.1**  $NET$  provides confidentiality for the set of messages  $M$  if and only if

$$NET \text{ sat } \forall m : M \bullet tr \upharpoonright out.0.USER.m \neq \langle \rangle \\ \Rightarrow tr \upharpoonright in.USER.0.m \neq \langle \rangle$$

□

This definition states that any message  $m$  which is output to user 0 (i.e. evidenced by a message  $out.0.i.m$  in the trace  $tr$ ) must have initially been sent to user 0 (i.e. there must be some  $j$  for which  $in.j.0.m$  appears in the trace).

This may also be expressed within the CSP process algebra:

**Theorem 3.2** A process  $NET$  provides confidentiality in the sense of Definition 3.1 if and only if

$$\forall m : M \bullet NET \parallel [in.USER.0.m] \parallel STOP = \\ NET \parallel [out.0.USER.m] \parallel STOP \\ in.USER.0.m$$

□

This states that if input of message  $m$  from any user to user 0 is blocked, then so is output of message  $m$  to user 0.

Observe that this is not equivalent to

$$NET \parallel [in.USER.0.M] \parallel STOP = \\ NET \parallel [out.0.USER.M] \parallel STOP \\ in.USER.0.M$$

For example,  $NET = in.1.0.m1 \rightarrow out.0.1.m1 \rightarrow in.1.2.m2 \rightarrow out.0.1.m2 \rightarrow STOP$  meets the latter equivalence but not the former: a message  $m2$  from user 1 to user 2 has been output by user 0, and this breaches confidentiality (assuming  $m1, m2 \in M$ ).

This property may also be captured in the traces model, as the property  $CONF$

$$CONF(tr) \hat{=} messages(tr \upharpoonright out.0.USER.M) \\ \subseteq messages(tr \upharpoonright in.USER.0.M)$$

This states that the messages (from message set  $M$ ) output from user 0 must be a subset of those that were sent to it. In other words, user 0 cannot obtain any messages from  $M$  that are not sent to it.

The fact that this property is a **sat** specification means that it is preserved by refinement.

### A simplification

Observe that if no messages are ever sent to 0 (perhaps if users are communicating with users they know and trust to be honest) then the characterisation of confidentiality may be simplified, since no messages will ever be sent to user 0.

The definition simplifies to

$$NET \text{ sat } \forall m : M \bullet tr \upharpoonright out.0.USER.m = \langle \rangle$$

which is equivalent to the simpler form

$$NET \text{ sat } tr \upharpoonright out.0.USER.M = \langle \rangle$$

This property may be expressed entirely within the process algebra in a number of different ways. The first way captures the idea that if attention is focussed entirely upon events from  $out.0.USER.M$ , then nothing should be observed:

$$NET \setminus (\Sigma \setminus out.0.USER.M) = STOP$$

The process  $STOP$  is a refinement of  $NET \setminus (\Sigma \setminus out.0.USER.M)$  (since in the traces model  $STOP$  is a refinement of every process), so achieving equality is equivalent to obtaining refinement in the other direction:

$$STOP \sqsubseteq NET \setminus (\Sigma \setminus out.0.USER.M)$$

An alternative characterisation is obtained by considering the effects of preventing  $NET$  from performing any events in  $out.0.USER.M$ . A system providing confidentiality should not be affected by this restriction:

$$NET = NET \parallel [out.0.USER.M] \parallel STOP$$

Since restricting the behaviour of  $NET$  can only reduce its behaviours, it follows automatically that the restriction is a refinement of  $NET$ . Hence the processes are equivalent precisely when there is a refinement in the other direction:

$$NET \parallel [out.0.USER.M] \parallel STOP \sqsubseteq NET$$

A final characterisation regards the system as acceptable if every event it can perform is in the set  $\Sigma \setminus out.0.USER.M$ . In other words, everything it can perform is also possible for a process which can always perform any of those events:

$$RUN_{\Sigma \setminus out.0.USER.M} \sqsubseteq NET$$

All of these characterisations are provably equivalent to the assertion  $NET \text{ sat } tr \upharpoonright out.0.USER.M = \langle \rangle$

It is straightforward using the process algebra to show that if a system provides confidentiality for two sets  $M1$  and  $M2$  separately, then it provides confidentiality for both sets simultaneously: if

$$NET \parallel [out.0.USER.M1] \parallel STOP = NET$$

and

$$NET \parallel [out.0.USER.M2] \parallel STOP = NET$$

then

$$\begin{aligned} NET \parallel [out.0.USER.(M1 \cup M2)] \parallel STOP &= ((NET \parallel [out.0.USER.M1] \parallel STOP) \parallel [out.0.USER.M2] \parallel STOP) \\ &= NET \parallel [out.0.USER.M2] \parallel STOP \\ &= NET \end{aligned}$$

## Message Authentication

This property requires that messages can be guaranteed to be ‘authentic’, in the sense that a particular message purporting to have come from a particular source really did come from that source. Authentication requires that messages cannot be forged.

In abstract terms, event  $b$  ‘authenticates’ event  $a$  if the observation of  $b$  is possible only if  $a$  occurred previously: the observation of  $b$  provides ‘evidence’ of  $a$ ’s previous occurrence.

**Definition 3.3** Event  $b$  authenticates event  $a$  in process  $P$  if and only if  $P \text{ sat } AUTH(tr)$ , where

$$AUTH(tr) = tr \upharpoonright b \neq \langle \rangle \Rightarrow tr \upharpoonright a \neq \langle \rangle$$

□

Observe that this specification does not restrict the number of occurrences of event  $b$  to each occurrence of event  $a$ .

The expression of this property in terms of a **sat** specification demonstrates that it is preserved by refinement.

This specification can also be captured as a process algebraic equation.

$$P \parallel [a, b] \parallel STOP = P \parallel [a] \parallel STOP$$

And since it is always the case that

$$P \parallel [a] \parallel STOP \sqsubseteq P \parallel [a, b] \parallel STOP$$

the condition is equivalent to

$$P \parallel [a, b] \parallel STOP \sqsubseteq P \parallel [a] \parallel STOP$$

For example, the process

$$\begin{aligned} P &= a \rightarrow b \rightarrow STOP \\ &\quad \square \\ &\quad b \rightarrow c \rightarrow STOP \end{aligned}$$

has  $c$  authenticating  $b$ :

$$\begin{aligned} P \parallel [b, c] \parallel STOP &= a \rightarrow STOP \\ &= P \parallel [b] \parallel STOP \end{aligned}$$

but it does not have  $b$  authenticating  $a$ :

$$\begin{aligned} P \parallel [a, b] \parallel STOP &= STOP \\ &\neq b \rightarrow c \rightarrow STOP \\ &= P \parallel [a] \parallel STOP \end{aligned}$$

In other words, a  $b$  event can occur even if  $a$  did not occur previously.

In the context of sending and receiving messages, we would require a received message  $out.j.i.m$  to authenticate a sent message  $in.i.j.m$ . In other words, receipt of the message  $i.m$  (‘ $m$  from  $i$ ’) by node  $j$  is possible only if that message was sent by node  $i$ . Thus on a system  $NET$  consisting of the medium, enemy and nodes, the property to check would be

$$\begin{aligned} NET \parallel [in.i.j.m] \parallel STOP &= \\ &\quad out.j.i.m \\ NET \parallel [in.i.j.m] \parallel STOP \end{aligned}$$

For example, a buffer process

$$COPY = in?x \rightarrow out!x \rightarrow COPY$$

has  $out.x$  authenticating  $in.x$  for any  $x$ : no message can be output unless it has previously been input.

This characterisation of authentication can be promoted to sets of events. The set  $B$  authenticates the set  $A$  in  $P$  if any message in  $B$  authenticates any of the messages in  $A$ . In other words, if any of the messages in  $B$  is seen, then one of the messages in  $A$  must previously have occurred. This is captured as follows

**Definition 3.4**  $B$  authenticates  $A$  in  $P$  if and only if

$$P \parallel [A \cup B] \parallel STOP = P \parallel [A] \parallel STOP$$

□

This form might be useful when we wish to check that a message is genuine even when its originator is unknown. This could be captured as the set  $out.j.USER.m$  authenticating the set  $in.USER.j.m$ . The authenticating message indicates that some honest node generated the original message.

The buffer process  $COPY$  has the weaker property of  $out.M$  authenticating  $in.M$ : no output can occur before input. This property is strictly weaker than the previous property, in which every output authenticates a corresponding input. For example, a random bit generator which allows any bit to be output following any input bit

$$RAND = in?x \rightarrow ( out!0 \rightarrow RAND \quad \square \quad out!1 \rightarrow RAND )$$

also has  $out.\{0, 1\}$  authenticating  $in.\{0, 1\}$ —no bit can be output if there wasn't previously an input—but does not have  $out.x$  authenticating  $in.x$  for any particular  $x$ .

The definition provides a straightforward proof of transitivity of authentication: if  $C$  authenticates  $B$ , and  $B$  authenticates  $A$ , then  $C$  authenticates  $A$ :

$$\begin{aligned} & P \parallel [A \cup C] \parallel STOP \\ &= P \parallel [A] \parallel STOP \parallel [C] \parallel STOP \\ &= P \parallel [A \cup B] \parallel STOP \parallel [C] \parallel STOP \\ &= P \parallel [A \cup B \cup C] \parallel STOP \\ &= P \parallel [B \cup C] \parallel STOP \parallel [A] \parallel STOP \\ &= P \parallel [B] \parallel STOP \parallel [A] \parallel STOP \\ &= P \parallel [A \cup B] \parallel STOP \\ &= P \parallel [A] \parallel STOP \end{aligned}$$

## 4 The network

### Architecture

A common architecture for which security protocols are designed consists of a network of *nodes* (typically workstations) which are able to communicate asynchronously by sending messages to each other over a *medium*, which acts as a postal service. The need for security arises from the fact that the users of this service (such as people, and applications programs) do not have control over the medium, and so it is possible for malicious agents to intercept or interfere with network traffic. The need for confidentiality in the face of an insecure medium creates the need for some form of encryption, and the need for authenticity when message forgery is possible also raises the need for some form of security mechanism.

A common approach to modelling this situation is to consider a set of nodes as connected to the medium, which is modelled as a single process. Although the medium will in general consist of a network of processes, this network may be considered at a higher level of abstraction as a single process. The only interactions the nodes may have with other nodes must be through this medium. As discussed earlier, we will find it convenient to model malicious interference by means of a separate enemy process node 0 which manipulates the essentially passive medium.

Thus the service provided to the users is modelled as

$$\left( \parallel_{i \in USER \setminus 0} NODE_i \right) \parallel [trans, rec] \parallel MEDIUM$$

The nodes are unable to interact directly, so their operation is entirely interleaved. They all communicate with the medium by means of channels *trans* and *rec*, used by the nodes to transmit and receive messages respectively.

Quite often the distinction between a user and the node by which the user communicates with the network is blurred when addressing security properties. An authentication check that a particular server remains up requires a response directly from that server rather than from the network operator responsible for it. Hence in some cases it is appropriate to think of the user and the node as being the same entity. However, for the purposes of this paper we find it convenient to treat them as distinct.

All forms of interference will be modelled by an intruder process  $ENEMY = NODE_0$  that is able to alter the condition of the medium via certain channels not available to the nodes. The entire system will be de-



events to signal particular modes of interference (in preference to having them occur nondeterministically) was originally taken in [6]. The advantage of this approach is that it allows greater control over the level and type of interference that may occur.

However, the enemy is not capable of producing all messages: for example, it cannot generate a message encrypted with a key it does not have (though of course it could reproduce such a message if it had previously received it).

In fact, the messages the intruder is able to generate will depend on the messages it has already seen pass as network traffic, the messages it is already able to generate, and the keys it has seen or which it owns.

We will use an information system [5] to define which messages can be generated by the enemy. It will have a trivial consistency relation: any set of messages is consistent. The definition of the relation  $\vdash$  of the information system will be dependent on, and should encapsulate the encryption mechanism. An information system defines a relation  $\vdash$  between finite sets of *tokens* and single tokens, indicating when the token can be generated from the set. In this case, we will use the relationship to indicate when the enemy, or indeed any other agent, can generate a particular message given the messages it has already seen.

Consider an example in which messages may be encrypted by means of either secret keys or public keys.

There will be the set *PUBLIC* of all the nodes' public keys—for simplicity we assume one for each node

$$PUBLIC = \{p_i \mid i \in USER\} \subseteq KEY$$

There will also be the set *S* of all the nodes' secret keys—one for each node

$$SECRET = \{s_i \mid i \in USER\} \subseteq KEY$$

This set is distinct from the set of public keys:

$$SECRET \cap PUBLIC = \emptyset$$

Finally there will be a set of shared keys *SHARED*, distinct both from public and secret keys:

$$SECRET \cap SHARED = \emptyset$$

$$PUBLIC \cap SHARED = \emptyset$$

The entailment relation  $\vdash: \mathbb{P}_{fn}(MESSAGE) \times MESSAGE$  will be a relation between a finite set of messages (that we think of as the enemy having seen) and messages that the enemy can generate. The relation is closed under the axioms for an information system:

A1. If  $m \in B$  then  $B \vdash m$

A2. If  $B \vdash m$  and  $B \subseteq B'$  then  $B' \vdash m$

A3. If  $B \vdash m_i$  for each  $m_i \in B'$  and  $B' \vdash m$  then  $B \vdash m$

We will abuse notation and allow the relation between possibly infinite sets and messages:

$$S \vdash m \Leftrightarrow \exists T \subseteq^{fn} S \bullet T \vdash m$$

We encapsulate the way in which messages can be generated by considering the possible structures for a message:

M1.  $B \vdash m \wedge B \vdash k \Rightarrow B \vdash k(m)$

M2.  $B \vdash m_1 \wedge B \vdash m_2 \Leftrightarrow B \vdash m_1.m_2$

where  $m, m_1$  and  $m_2$  are messages, and  $k$  is a key.

Certain properties of particular encoding mechanisms may also be captured by providing additional inference rules. For example, the relationship between secret and public keys may be captured by the following pair of rules:

K1.  $\{p_i(s_i(m))\} \vdash m$

K2.  $\{s_i(p_i(m))\} \vdash m$

where  $p_i \in P$  and  $s_i \in S$ .

For example, these rules allow us to deduce the obvious result that possession of a message encrypted with a secret key ( $s_i(m)$ , say), together with possession of the public key, allows the original message to be retrieved:

- |    |                                      |           |
|----|--------------------------------------|-----------|
| 1. | $\{p_i, s_i(m)\} \vdash s_i(m)$      | A1        |
| 2. | $\{p_i, s_i(m)\} \vdash p_i$         | A1        |
| 3. | $\{p_i, s_i(m)\} \vdash p_i(s_i(m))$ | M1, 1, 2  |
| 4. | $\{p_i, s_i(m)\} \vdash m$           | 3, K1, A3 |

The appropriate rule for shared keys is that possession of a shared key together with a message encrypted with that key allows generation of the original message:

K3.  $\{k, k(m)\} \vdash m$  if  $k \in SHARED$

It is also possible to encode various other deductions we might wish to include in the capability of the enemy, for example deducing a key from observing both an encoded message and that message in plaintext:

K4.  $\{m, k(m)\} \vdash k$

The rules that we give model different encryption and decryption capabilities of the enemy.

The rules can also be used to encapsulate properties of encryption. For example, if encryption were commutative, then we could include the rule

K5.  $\{k_1, k_2(m)\} \vdash k_2(k_1(m))$

## Medium

The description of *MEDIUM* involves a number of decisions about the best way to model the network medium.

We must allow the possibility of an intruder, who is able to manipulate the medium in particular ways. This could be done by building the intruder into the medium (so the medium itself has the capability of interfering with message traffic in particular ways); but we prefer to follow Roscoe's approach [6] of including a separate model of the intruder. This second approach gives greater separation between the medium itself, which would then be considered as essentially a passive service provided to the various nodes, and a malicious agent who has particular capabilities to manipulate the medium in particular ways. The capabilities of this agent will be made more explicit, and manipulation of the medium will be associated with particular events, which will make attacks on protocols easier to follow and understand.

The medium (containing the set of messages  $B$ ) may be described initially as  $MEDIUM(\emptyset)$ , where:

$$\begin{aligned}
 MEDIUM(B) &= INPUT(B) \\
 &\square \\
 &OUTPUT(B) \\
 &\square \\
 &IA(B)
 \end{aligned}$$

The process  $INPUT(B)$  permits input to the medium. We must decide on the type of messages that the medium will accept and offer. For the purposes of this paper, we will separate out the destination and source from the body of the message. Again there are other possibilities, for example, if the message is to be broadcast to all users then no explicit destination field is required.

$$\begin{aligned}
 INPUT(B) &= \\
 \square_i \text{ trans. } i?j?x &\rightarrow MEDIUM(B \cup \{i.j.x\})
 \end{aligned}$$

Here the channel *trans* is of type *USER.USER.MESSAGE*. A message *trans.i.j.m* should be thought of as node  $i$  sending an input  $j.m$  to the medium, indicating the wish that message  $m$  be delivered to node  $j$ . Thus  $i$  is the source,  $j$  the destination, and  $m$  the message.

We have abstracted away refusals, in the sense that input can never be refused, which amounts to making the assumption that nothing can be deduced from how or when the messages are accepted. This is a reasonable assumption, since there are protocols currently in

use to perform tasks such as masking network traffic. Hence at this level of abstraction we can assume that messages are always accepted by the network. (If this is later felt to be unrealistic, the definition of *INPUT* can be altered accordingly, so that messages may not be input after the number of messages in the network reaches some capacity threshold)

The process *OUTPUT* allows output from the medium:

$$\begin{aligned}
 OUTPUT(B) &= \\
 \square_{i.j.x \in B} \text{ rec. } j!i!x &\rightarrow MEDIUM(B \setminus \{i.j.x\})
 \end{aligned}$$

Here the channel *rec* is of type *USER.USER.MESSAGE*. A message *rec.j.i.m* corresponds to the receipt of a message  $m$  by node  $j$  which is labelled as coming from source node  $i$ .

Note that an empty external choice is simply equivalent to *STOP*, so when the set  $B$  is empty (i.e. the medium contains no messages) there is no possibility of output.

Finally, the process  $IA(B)$  describes the possible interactions with the medium due to Intruder Actions. A perfectly secure medium would treat this part of the process description as *STOP*. In cases we are considering, we model the ways in which the medium is susceptible to interference. Here, the medium is vulnerable to having messages removed, added, or leaked.

$$\begin{aligned}
 IA(B) &= \\
 \text{kill} &\rightarrow \square_{b \in B} MEDIUM(B \setminus \{b\}) \quad \text{if } B \neq \emptyset \\
 &MEDIUM(\emptyset) \quad \text{if } B = \emptyset \\
 &\square \\
 \text{add? } i?j?x &\rightarrow MEDIUM(B \cup \{i.j.x\}) \\
 &\square \\
 \square_{i.j.x \in B} \text{ leak! } i!j!x &\rightarrow MEDIUM(B)
 \end{aligned}$$

## Enemy action

In modelling the enemy we are concerned with messages that the enemy is able to generate. These may be used to disrupt a protocol, or may correspond to information about what the enemy has discovered concerning supposedly confidential messages.

Certain assumptions may be made concerning the enemy, depending on the property that is under analysis. When checking confidentiality, it is often assumed that the enemy is unable to generate those messages  $M$  (which can be generated by the users) that should be kept confidential. On the other hand, when checking authentication it should be assumed that the enemy (as well as the honest users) is capable of generating

those messages whose authenticity is ensured by the protocol, since if the enemy is unable to generate them then there is no need for an authentication protocol. For integrity, it is assumed that the enemy is capable of generating messages from  $M$ .

These assumptions may be incorporated into the description of the enemy, which may then be parameterised by a set of messages  $S$  that it has seen, and a set of messages  $INIT$  that it is initially able to generate. The assumptions can be expressed as conditions on  $INIT$  and on the set of messages  $M$  which the particular security property is concerned with.

The question also arises as to whether it is sufficient to model enemy actions using a single  $ENEMY$  process. In principle it is possible that a number of malicious agents acting together might effect an attack where a single agent would be unable to do so. Whether or not this is possible in the model being developed here depends on the actual description of the process  $ENEMY$ . In fact, the description to be presented enjoys the property that

$$ENEMY = ENEMY \parallel ENEMY$$

Hence for all analysis done at the level of traces we see that any number of enemies acting together are encapsulated within the description of a single enemy.

In addition to the messages that can be generated from those messages already seen, the enemy is able to generate particular plaintext messages. Furthermore, the enemy should be considered to be in possession of all of the nodes' public keys, and all of the users' names. We therefore use a set  $INIT$  to model all of the information initially in the possession of the enemy. Thus we have  $PUBLIC \subseteq INIT$  and  $USER \subseteq INIT$ . The relation  $\vdash$  gives the capability of the enemy to generate messages from messages already in its possession. The CSP description of the enemy will use this relation. The set  $S$  records those messages that have been read from the medium. This is initialised to  $\emptyset$ , so  $ENEMY$  is defined to be  $ENEMY(\emptyset)$ , where

$$\begin{aligned} ENEMY(S) &= KILL(S) \\ &\square \\ &ADD(S) \\ &\square \\ &LEAK(S) \\ &\square \\ &KNOWS(S) \end{aligned}$$

The first option allows the enemy to kill a message—to remove it from the medium. It is described simply as

$$KILL(S) = kill \rightarrow ENEMY(S)$$

In fact, when dealing with trace properties of communication protocols, the ability of the enemy to kill messages is entirely irrelevant. Although the possible removal of messages from the medium can interfere with liveness properties of communications protocols, it cannot compromise properties expressed only in terms of traces. This is because the medium allows the reordering of messages, so any particular message could always be ignored and remain in the medium without being killed. Any protocol that guarantees a security property if the enemy is unable to kill messages will therefore guarantee it in any case. An equally useful definition of  $KILL(S)$  would be  $STOP$  (which would be equivalent to omitting this option entirely).

The second course of action available to the enemy is to insert any message that it can generate onto the medium. These are any messages that can be generated from its initial set  $INIT$  together with the messages  $S$  that have since come into its possession.

$$ADD(S) = \square_{INIT \cup S \vdash x} add!i!j!x \rightarrow ENEMY(S)$$

Observe that this description incorporates the ability of the enemy to manipulate message address fields, thus giving the impression that a message comes from a source other than the genuine source.

If the enemy is considered to be simply an eavesdropper with no power to add messages to the medium, then the  $ADD$  component would simply be modelled as  $STOP$  (or omitted entirely).

The third option allows it to observe any message currently on the medium:

$$LEAK(S) = leak?i?j?x \rightarrow ENEMY(S \cup \{x\})$$

The final option is included to model the enemy's knowledge of particular messages. This is accomplished by allowing the enemy to output any message that can in fact be generated.

$$KNOWS(S) = \square_{INIT \cup S \vdash x} out.0.0!x \rightarrow ENEMY(S)$$

The channel  $out.0$  is used to indicate those messages that the enemy can deduce from what has already been seen together with what was known initially.

The argument  $S$  represents the set of messages that the enemy has already seen. Normally this will be the empty set at the beginning of a protocol run, but it is possible to model the effect that possession of a particular key might have on the vulnerability of a protocol, by including such a key, or some other message, in the set  $S$ .

Observe that we have allowed the insertion of any message into the medium, so in particular false sources can be attached to messages. Rerouting of a message

can also be modelled, by having the enemy read it via *leak*, kill it (this is cleaner though not essential, as already discussed), and then add the same message with a different destination field back to the medium.

Now the assumption that is made in the case of confidentiality can be formalised. We are assuming that none of the messages that we wish to keep confidential are in fact in the kernel of the messages that can initially be generated by the enemy:

$$M \cap \text{kernel}(\text{INIT}) = \emptyset$$

On the other hand, for integrity and authenticity, we are (implicitly) assuming that

$$M \cap \text{INIT} \neq \emptyset$$

in the sense that protocols designed to provide these services are intended to deal with messages that can be generated by an enemy.

When checking a confidentiality protocol, strong use is made of this assumption, since if the enemy can output a message that is supposed to be confidential then the protocol is considered to be insecure. However, there are situations such as key-exchange where a protocol is designed to provide both confidentiality and authenticity, in which case it is reasonable to begin analysis with  $M \cap \text{INIT} \neq \emptyset$ . In such situations, the above modelling of what the enemy knows is not adequate, and it would be necessary to construct a more sophisticated, complex model of the enemy which keeps track of incoming and outgoing messages and outputs on *out.0* only those messages it deduces have been generated by the legitimate users, in particular ignoring those messages in  $M$  that it puts onto the medium and then reads back via *leak*. We will not pursue this further in this paper, but will observe that it is a situation to bear in mind.

## Nodes

We must consider the nodes—which are the link between the user and the medium—to be under the control of the user. It is the nodes that will provide the security facilities required by the users, such as encrypting and deciphering messages.

The (finite) set of all nodes will be labelled using the set  $USER = \{0, 1, \dots, n\}$ .

The nodes provide the means by which users send messages over the network. A user communicating with the network is in fact communicating with the corresponding node. Nodes interact with users by inputting plaintext messages with intended destinations, and outputting such messages together with

their source. The process  $NODE_i$  thus communicates with its user via channels *in.i* and *out.i* of type  $USER.PLAINTEXT$ . An input *in.i.j.m* to node  $NODE_i$  is interpreted as a request from user *i* to send message *m* to user *j*. Similarly, an output communication *out.i.j.m* is interpreted as delivery to user *i* of message *m* purporting to come from *j*.

Nodes (with the exception of node 0) interact with the medium by transmitting (possibly enciphered) messages together with other control messages, intended recipients, and any other messages employed by the protocol being used. The channels used are *trans.i* for transmission, and *rec.i* for receipt of messages. These channels are of type  $USER.MESSAGE$ , where the set  $MESSAGE$  contains both plain and encrypted messages (as discussed later). A communication *trans.i.j.m* corresponds to  $NODE_i$  placing message *m* with destination *j* onto the medium. A communication *rec.i.j.m* corresponds to  $NODE_i$  receiving message *m* from the medium, with source purporting to be *j*.

The description of a  $NODE_i$  process will depend on the security property we are aiming to verify of the network. For confidentiality, authenticity, anonymity and integrity its description will consist of a CSP implementation of the particular protocol under analysis. For example, an extremely simple protocol to provide confidentiality of messages sent from user 1 to user 2 will be implemented using  $NODE_2$ 's public key  $p_2$  and secret key  $s_2$  as follows:

$$\begin{aligned} NODE_1 &= in.1.2?x \rightarrow trans.1!2!p_2(x) \rightarrow NODE_1 \\ NODE_2 &= rec.2?j?y \rightarrow \\ &\quad \begin{array}{ll} STOP & \text{if } s_2(y) \notin PLAINTEXT \\ out.2!j!s_2(y) & \\ \rightarrow NODE_2 & \text{otherwise} \end{array} \end{aligned}$$

where  $s_2(p_2(x)) = x$  for any message *x*. Observe that this protocol does not ensure authenticity.

The situation is different in the case of non-repudiation. In this case, verification is from the judge's viewpoint, and the judge does not have control over the nodes used in a non-repudiation protocol. In fact, from the judge's viewpoint, the parties could each be dishonest. Indeed, it is this possibility that generates the need for a non-repudiation protocol in the first place.

The judge has to allow for the possibility that each node has the capabilities of node 0. Thus non-repudiation has to be established in the context of nodes which can *kill*, *add*, and *leak* messages as well as interact with the medium in the usual ways. The nodes for which non-repudiation should be established

are therefore

$$\begin{aligned}
NODE_i(M) = & \\
& in.i?j?x \rightarrow NODE_i(M \cup \{x\}) \\
& \square \square_{INIT \cup M \vdash x} out.i!j!x \rightarrow NODE_i(M) \\
& \square \square_{INIT \cup M \vdash x} trans.i!j!x \rightarrow NODE_i(M) \\
& \square rec.i?j?x \rightarrow NODE_i(M \cup \{x\}) \\
& \square kill \rightarrow NODE_i(M) \\
& \square \square_{INIT \cup I \vdash x} add.i!j!x \rightarrow NODE_i(M) \\
& \square leak?l?j?x \rightarrow NODE_i(M \cup \{x\})
\end{aligned}$$

From a modelling point of view the interfaces of these processes with the network must be expanded to include the channels *add*, *leak* and *kill*.

Since the node is able to generate its own plaintext messages, the *in.i* channel is perhaps redundant, but is retained as a source of messages so that particular non-repudiation protocols will refine this node.

The set of messages *M* corresponding to the initial state of the node will contain all of the keys which the node may use to encrypt and decrypt messages.

### Meadows' example

In order to illustrate the above material, we will present a simple example used in [4] and [2]. It is not even a protocol, but is instead a simple example designed purely for illustrative purposes. In fact, it is not the kind of example that best illustrates the benefits of the process algebra approach, since process algebra would be of more use in exploring subtle patterns of interactions between different parties; here the interactions are fairly simple. However, it illustrates the approach. Although the proof of such an obvious property seems unduly long, it is also lengthier than might be expected in [4] and [2]. This is because a significant amount of formalisation needs to be done before the proof can actually proceed,

The example consists of a legitimate user who encrypts received messages with a particular key, and returns them to the medium. This could be described as a legitimate node (number 1 for definiteness) which receives messages on *rec.1*, encrypts them, and returns them on *trans.1*. The process algebra is as follows:

$$NODE_1 = rec.1?j?x \rightarrow trans.1!j!k(x) \rightarrow NODE_1$$

where *k* is a key possessed by *NODE*<sub>1</sub>.

The aim is to establish that the enemy cannot obtain a particular message *a* that it does not already possess. This is expressed as confidentiality with respect to *a*:

$$NET = NET \llbracket out.0.0.a \rrbracket STOP$$

or alternatively as

$$NET \text{ sat } (tr \upharpoonright out.0.0.a = \langle \rangle)$$

We make the standard assumption for confidentiality, that the enemy is not in possession of any messages containing *a*:

$$KERNEL(a) \not\subseteq kernel(INIT)$$

We may take the description of *NET* to consist of the node *NODE*<sub>1</sub>, the initially empty medium *MEDIUM*( $\emptyset$ ), and the enemy who has initially learned nothing: *ENEMY*( $\emptyset$ ).

It will prove useful to extract certain sets of messages from traces of the system:

$$\begin{aligned}
LEAK(tr) &= \{m \mid \exists i, j \bullet tr \upharpoonright leak.i.j.m \neq \langle \rangle\} \\
ADD(tr) &= \{m \mid \exists i, j \bullet tr \upharpoonright add.i.j.m \neq \langle \rangle\} \\
TRANS_k(tr) &= \{m \mid \exists j \bullet tr \upharpoonright trans.k.j.m \neq \langle \rangle\} \\
REC_k(tr) &= \{m \mid \exists j \bullet tr \upharpoonright rec.k.j.m \neq \langle \rangle\} \\
OUT_0(tr) &= \{m \mid tr \upharpoonright out.0.0.m \neq \langle \rangle\} \\
MESS(tr) &= LEAK(tr) \cup ADD(tr) \cup \\
&\quad TRANS_1(tr) \cup REC_1(tr) \cup \\
&\quad OUT_0(tr)
\end{aligned}$$

**Lemma 4.1** The kernel function is closed under the generates relation, i.e.

$$B \vdash m \Rightarrow kernel(m) \subseteq kernel(B)$$

□

**Proof** By considering all of the clauses that define the relation  $\vdash$ : A1, A2, A3, M1, M2, K1, K2, K3, and K4. The result follows for each clause, so it is true for the relation. □

In order to prove confidentiality of *NET* with respect to *a* we will use certain properties of its components. The required properties are described in the following lemma. They combine information about the state of the components (as maintained in *S* and *B*) and events that have occurred (extracted from the trace). A combination of information from both these sources is often required in establishing this kind of result. State-based approaches commonly include a 'history' variable as a component of the state in order to record trace information. The approach taken here is closer to event-based approaches which provide some way of extracting the state of the system from its trace.

**Lemma 4.2** The component processes meet the following specifications:

$$\begin{aligned}
ENEMY(S) & \text{ sat } E1_S(tr) \\
ENEMY(S) & \text{ sat } E2_S(tr) \\
MEDIUM(B) & \text{ sat } M1_S(tr) \\
MEDIUM(B) & \text{ sat } M2_S(tr) \\
NODE_1 & \text{ sat } N1_S(tr)
\end{aligned}$$

where

$$\begin{aligned}
E1_S(tr) & = \\
& \text{kernel}(ADD(tr)) \subseteq \\
& \text{kernel}(S) \cup \text{kernel}(INIT) \cup \text{kernel}(LEAK(tr)) \\
E2_S(tr) & = \\
& \text{kernel}(OUT_0(tr)) \subseteq \\
& \text{kernel}(S) \cup \text{kernel}(INIT) \cup \text{kernel}(LEAK(tr)) \\
M1_S(tr) & = \\
& \text{kernel}(LEAK(tr)) \subseteq \\
& \text{kernel}(B) \cup \text{kernel}(TRANS_1(tr)) \cup \text{kernel}(ADD(tr)) \\
M2_S(tr) & = \\
& \text{kernel}(REC(tr)) \subseteq \\
& \text{kernel}(B) \cup \text{kernel}(TRANS_1(tr)) \cup \text{kernel}(ADD(tr)) \\
N1_S(tr) & = \\
& \text{kernel}(TRANS_1(tr)) \subseteq \text{kernel}(REC_1(tr))
\end{aligned}$$

□

**Proof** This is a standard mutual recursion proof in CSP. The full proof has been omitted for reasons of space. □

Some further process algebra manipulation (details in the full paper) finally yields the required result.

**Theorem 4.3** The network is secure:

$$\begin{aligned}
(NODE_1 & \\
& |[trans.1, rec.1]| \\
MEDIUM & \\
& |[leak, add, kill]| \\
ENEMY) & \text{ sat } tr \upharpoonright out.0.0.a = \langle \rangle
\end{aligned}$$

□

**Proof** The strategy of the proof is as follows: we will use the properties established in Lemma 4.2 to prove that the kernels of all messages passed around the system must be contained in the kernel of *INIT*. Since it is given that  $a \notin \text{kernel}(INIT)$ , it follows that  $a$  can never be passed along channel *out.0*. □

The proof could easily be adapted to take other nodes into account. In fact, descriptions of the other nodes are not even necessary, all that is required of them is that they meet some particular specification: for example, that they do not transmit any messages whose kernel intersects with that of  $a$ :

$$\begin{aligned}
NODE_j & \text{ sat} \\
& \text{kernel}(TRANS_j(tr)) \cap \text{kernel}(a) = \emptyset
\end{aligned}$$

Of course, more complex specifications might be more appropriate, for example that the messages added to the network by the nodes do not intersect with  $a$  (though those that were passed to the node may be passed back):

$$\begin{aligned}
NODE_j & \text{ sat} \\
& \text{kernel}(TRANS_j(tr)) \\
& \cap (\text{kernel}(a) \setminus \text{kernel}(REC_j(tr))) \\
& = \emptyset
\end{aligned}$$

This latter specification is in fact met by *NODE<sub>1</sub>*.

## 5 Discussion

### Model-checking

This paper has been concerned with the expression of particular security properties and protocols within the framework of CSP, in order to provide a foundation for analysis and verification. This approach is motivated in part by the availability of model-checking tools such as FDR[1], and the work has always proceeded with an eye on applicability of these tools. However, it is inevitable that there will be some practical difficulties, and it may be necessary to adapt some of the properties. When this becomes necessary, we will have to establish that the properties we are checking do indeed correspond to the properties presented in this paper, or at least that results obtained by application of the tools allow us make the inferences we require.

For example, the sets *PLAINTEXT* and *MESSAGES* will generally be infinite, even when the base sets are very small. This makes them unsuitable for direct analysis by means of model-checking using current (February 1996) state-of-the-art technology, though the situation will improve as value-passing is introduced. Techniques such as Skolemisation (deducing results concerning all messages from verifications on place-holders) might also be appropriate here. In any case some simplifications will have to be made (perhaps concerning the maximum number of encryptions) in order to regain finitude of the message

space; and some additional justification of these assumptions will then be required to derive general correctness. This should not present any problems, since the protocols themselves will only perform encryptions to a certain level (generally no more than two) and so any interference involving deeper levels will be detected in any case. But nevertheless it will be necessary to prove that the imposition of a bound does not rule out any attacks on a protocol, in order to have confidence in the results of the analysis.

## Additional modelling issues

The modelling of the enemy as a separate process allows for the possibility of introducing tactics in the state space exploration when model-checking, for example by restricting the number of messages that the enemy will place on the medium. By accompanying enemy interference with the performance of events, we may introduce tactics by introducing further constraints in parallel, refining the system. This may prove useful when attempting to detect flaws, since a flaw in a refinement will be a flaw in the original system, but correctness of refined protocols does not imply that the original one is correct unless it can be demonstrated that the introduction of the tactic does not rule out any possible attacks.

Non-repudiation seems to be a completely different kind of property. Each party in a non-repudiating exchange of messages is concerned that the other might not be honest. Furthermore, it is not enough for each party to be satisfied that the other party received the required messages; each party aims to obtain evidence sufficient to convince an outside party that the exchange took place.

## Directions

Meadows' example appears to be particularly straightforward (which is what makes it a good example for comparing different approaches) because the proof rests on the fact that at no stage is the information required to generate the message  $a$  ever introduced into the system; the invariant for the system is therefore fairly simple, and does not rely particularly on encryption and decryption properties, but simply on the property that no generation of messages by the  $\vdash$  relation can introduce new information. It will be harder to find suitable invariants for scenarios where information is present in encrypted form (such as communication between two users via a shared key), where it will be necessary to prove that at no stage could it

ever be decrypted. More subtle properties of encryption and decryption will be required.

It seems disappointing that such a simple example as Meadows' still requires a lengthy proof. However, part of the point of doing such a proof is to explore the relationship between the language-theoretic ideas underpinning it and the invariant of the CSP recursive description. It seems likely that this relationship will be similar in many proof of this type, and we would hope to obtain theorems which allow results concerning the language of messages that can be generated to be translated immediately to the CSP setting without the need for a laborious manual translation. A close relationship between CSP protocol descriptions and rules for generating messages would allow more natural proofs. Once this is achieved we would expect the result that a particular set of rules cannot generate any message containing  $a$  to translate immediately into the result that the corresponding CSP description has the required confidentiality property.

## Acknowledgements

Thanks are due to Peter Ryan, Richard Moore, Irfan Zakiuddin, Paul Gardiner, Bill Roscoe, Gavin Lowe, Michael Goldsmith and Abraham Sidiropoulos for much lively discussion and perceptive comments on earlier versions of this work. Thanks also to the anonymous referees for their careful reviewing of the paper, and for their useful and constructive suggestions.

Thanks also to DRA Malvern for providing funding for this research.

## References

- [1] Formal Systems (Europe) Ltd. Failures divergences refinement user manual and tutorial, 1994.
- [2] J. W. Gray and J. McLean. Using temporal logic to specify and verify cryptographic protocols (progress report). In *Proceedings of the eighth IEEE Computer Security Foundations Workshop*, 1995.
- [3] C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [4] C. Meadows. Applying formal methods to the analysis of a key management protocol. *Journal of Computer Security*, 1(1), 1992.
- [5] A. W. Roscoe. Lecture notes on domain theory, 1986. Oxford University.
- [6] A. W. Roscoe. Prospects for describing, specifying and verifying key-exchange protocols in CSP and FDR, 1994. Formal Systems (Europe) Ltd.
- [7] S. A. Schneider. Modelling security properties with csp, 1996. Royal Holloway Technical Report CSD-TR-96-04.