

Non-interference through determinism*

A.W. Roscoe, J.C.P. Woodcock and L. Wulf

Oxford University Computing Laboratory, Parks Road, Wolfson Building, Oxford OX1 3QD, UK

The standard approach to the specification of a secure system is to present a (usually state-based) abstract security model separately from the specification of the system's functional requirements, and establishing a correspondence between the two specifications. This complex treatment has resulted in development methods distinct from those usually advocated for general applications.

We provide a novel and intellectually satisfying formulation of security properties in a process algebraic framework, and show that these are preserved under refinement. We relate the results to a more familiar state-based (Z) specification methodology. There are efficient algorithms for verifying our security properties using model checking.

Keywords: Security, non-interference, formal methods, process algebra, determinism, automatic verification

1. Introduction

Security requirements of a computer system are regarded as critical properties that demand the availability of mechanisms which control or protect programs and data. Three issues in particular are related to the area of computer security: (i) *confidentiality (secrecy)*, the problem of protecting information from unauthorised disclosure; (ii) *integrity*, the protection of information from unauthorised modification or destruction; and (iii) *denial of service*, the avoidance of major reduction in system performance.

It is possible to regard these security concerns as properties of *information flow* within the system and base a specification of security on the absence of undesired flows. The notion of *non-interference* [7] captures the idea that no information can flow from one user to another if the system view of the second is completely unaffected by actions of the first. We introduce a novel characterisation of non-interference based on the notion of deterministic views. This elegant formulation of non-interference has, unlike others described in the literature, the property of preserving security requirements under refinement.

The development of a secure system entails the construction of an abstract security model in addition to the specification of the system's functional requirements.

*This paper is a revised version of [16]. The notation has been slightly modified to make it compatible with that used in related publications [13, 14].

The model is intended to capture abstractly the complete set of security requirements, which are derived from the system's (possibly informal) security policy, and which form part of the total system requirements. Depending on the level of rigour required during development, it is necessary to either informally establish or formally prove a correspondence between the functional specification and the abstract model.

The construction of the security model has been attempted [4, 10] with the same methods as functional specifications, such as the Z notation [18]. We suggest that there are good reasons to employ a process algebraic notation for this purpose. Firstly, it is not the individual operations of the system, but the system as a whole that is to satisfy critical properties. Secondly, insecurity is introduced not by a single operation in isolation but by certain sequences of operations. And thirdly, it turns out to be possible to express non-interference constraints directly on a process representation of the system, thus eliminating the need for constructing a separate abstract model.

We therefore propose a process-algebraic approach (based on CSP [8]) to the specification of security properties. In particular the property of a process being *deterministic* is fundamental to the conditions we introduce for non-interference. This property can be verified using standard algorithms on finite-state systems, such as those implemented in the CSP model checker FDR¹ [12].

This paper is organised as follows. The following section defines the non-interference conditions and illustrates some of their properties. The conditions are generalised to systems with multiple users. Section 3 presents a functional specification of a file system that is intended to maintain confidential information. A systematic way of mapping this specification into process algebra is given in Section 4, and the particular conditions for non-interference in the process model are clearly stated. The security flaw of the system is detected by automatic verification in Section 5, and it is shown how the system can be made secure. Finally, we present our conclusions in Section 6.

2. Non-interference and determinism

There have been a number of CSP formulations of non-interference, such as Jacob's use of inference functions [9]. None of these approaches is based on the notion of determinism, which has only recently been recognised as the fundamental concept underlying the various definitions on non-interference [13]. This section will introduce some formal definitions of non-interference and analyse their properties.

¹FDR (Failures Divergence Refinement) is a product of Formal Systems (Europe) Ltd., 3 Alfred St., Oxford OX1 3EH, UK.

2.1. Notation and conventions

We will employ the failures-divergences model of CSP, in which a process is characterised by its failures and its divergences. We use the following notation to refer to various observations of a process P .

$\alpha(P)$	alphabet	set of events process P can engage in
TRACES(P)	traces	set of finite sequences of events P can engage in
FAILS(P)	failures	set of pairs (s, X) such that P can refuse events X after trace s
DIVS(P)	divergences	set of traces after which P may behave chaotically

The semantics of the *failures-divergences model* of CSP is detailed in, e.g., [8]. Of particular relevance below will be the concealment and interleaving operators whose formal semantics are given in the Appendix. Informally, $P \setminus A$ is a process that behaves like P except that occurrences of events in set A are concealed. A concealed event occurs automatically and instantaneously as soon as it can, without being observed or controlled by the environment of the process. The \parallel interleaving operator models asynchronous composition of processes: $P \parallel Q$ is a process whose trace forms an arbitrary interleaving of events from processes P and Q . An event can be refused by the composition only if both component processes refuse it.

We will interpret some processes U_i as users interacting with another process P called the system. A user U of P is defined by its interface to the system. For the moment, it is assumed that the system has exactly two users U_H and U_L , with $\alpha(P) = H \cup L$, $L \neq \emptyset$, $H \neq \emptyset$, and $H \cap L = \emptyset$. This latter condition of disjoint set of actions available to the users is convenient since it prohibits direct communication between users by synchronisation.

These simplifying assumptions will be relaxed in Subsection 2.6 where the non-interference conditions will be generalised to multi-user systems.

2.2. Abstracting events

In a system with two users U_H and U_L we will typically want one user (U_L) to be completely unaware of what the other (U_H) does. In other words, the system view of U_L should be unaffected by the presence or absence of events user U_H might engage in. If this is the case, we say that there is no flow of information from U_H to U_L , or that U_H is non-interfering with U_L .

In a sense, it is necessary to abstract away from the actual or potential behaviour of U_H and ensure that this abstraction cannot affect how the system appears to U_L . There are several ways this abstraction may be captured, e.g., by concealing or obscuring U_H 's actions. In CSP, the concealment of events is expressed using the \setminus hiding operator, and the obscuring of events may be achieved using \parallel interleaving.

It is well-known that concealment and interleaving of events may introduce non-determinism [8, pp. 113 and 120]. A non-deterministic system may, under the same conditions, behave differently towards its environment, due to some internal, uncontrollable choice. Though this choice cannot be observed directly, its external effects can, and thus provide clues on abstracted activities. The result of this abstraction will be that U_H 's actions become choices which, though not visibly directly to U_L , may resolve non-determinism that is. The absence of non-determinism under abstraction of U_H 's behaviour guarantees the absence of undesired information flow towards U_L .

The notion of determinism is formally defined as follows. A process P is *deterministic* if it is free of divergence, and if it never has a choice whether to refuse an event it can engage in.

$$P \text{ det} \Leftrightarrow \text{DIVS}(P) = \emptyset \wedge (tr \langle a \rangle \in \text{TRACES}(P) \Rightarrow (tr, \{a\}) \notin \text{FAILS}(P)).$$

A process lacking this property is *non-deterministic*; under identical environmental conditions it may behave differently in an unpredictable fashion.

2.3. Non-interference conditions

The conditions we propose are all based on the absence of non-determinism after the abstraction of "high-security" events, and are justified in detail in [13]. Concealment is the simplest method of abstracting from events in CSP which can serve as a first attempt to define the notion of non-interference.

Definition 2.1. A system P is said to be *eagerly secure* in L if concealment of H events does not introduce non-determinism, i.e.,

$$\mathcal{E}\text{Sec}_L(P) \Leftrightarrow (P \setminus H) \text{ det}.$$

The terminology will become clear later on. Another way of abstracting events in CSP is not by concealing but by obscuring their occurrence. This can be achieved used another process

$$\text{RUN}_H = x: H \rightarrow \text{RUN}_H$$

and combining it with the original system P by interleaving as $P \parallel \text{RUN}_H$.

This process can never refuse an H event since RUN_H is always prepared to contribute one in arbitrary places. An outside observer will not be able to tell whether such an action came originally from P or from RUN_H . As above, we postulate that abstraction by interleaving does not introduce non-determinism.

Definition 2.2. A system P is said to be *lazily secure* in L if obscuring H events by interleaving does not introduce non-determinism, i.e.,

$$\mathcal{L}\text{Sec}_L(P) \Leftrightarrow (P \parallel \text{RUN}_H) \text{ det}.$$

The terminology of the conditions reflects the semantics of the operators involved. The \setminus concealment operator is defined in a way such that hidden (internal) actions occur instantly. Abstraction of events by concealment is eager in the sense that the events cannot be prevented or delayed by the environment.

This situation contrasts with the usual interpretation of communications between interacting processes. The standard interpretation of the occurrence of an event is that the process and its environment have agreed on the action; it cannot occur without mutual consent. The agreement of U_H to engage in H events cannot be assumed to be immediately forthcoming. Abstraction by interleaving $P \parallel \text{RUN}_H$ does not force events from P to happen, it simply prevents an observer from knowing whether they came from P or from RUN_H . This lack of urgency explains why this is lazy abstraction. $P \parallel \text{RUN}_H$ can only be deterministic if the set of L events available before and after any H event of P are the same, since if the same event is communicated by RUN_H the state of P does not change. Lazy abstraction is thus sensitive not only to the effects of different actions by U_H , but also to the choice between action and inaction.

The possibility of infinite sequences of H actions gives rise to the danger that a system implementation will prefer them forever, thereby denying U_L the opportunity to communicate – which would be a clear breach of security. The eager security condition, which entails the assumption that H actions are never delayed, is necessarily sensitive to this possibility as $P \setminus H$ introduces divergence.² This is illustrated with Example 2.6 below.

In some practical examples of system models, the applicability of the lazy condition depends on the nature of the events concerned. If these are events which occur instantaneously – such as a system message appearing on the user's screen – then they are indistinguishable to U_L from the ordinary internal actions of Q . As long as these "signal" events are guaranteed to occur instantaneously there will be no refusal of a request by U_L at the interface to the system.

The H events can therefore be divided into two categories: signal events S which are guaranteed to occur instantly, and events D which cannot occur without the agreement of U_H and may thus be delayed. In many systems, delayable events take the form of inputs whereas the signals appear as output communications to the environment (including users).

Since S events resemble internal system actions we can abstract from them by hiding while we still use interleaving for ordinary events. The combination of the two forms of abstraction results in a *mixed* non-interference condition.

Definition 2.3. A system P whose H events can be partitioned into delay events D and signal events S satisfies $\mathcal{M}\text{Sec}_L^{(D,S)}(P)$ if $(P \setminus S) \parallel \text{RUN}_D$ is deterministic.

²The lazy condition (where H actions may be subject to delay) assumes that the implementation is sufficiently fair to avoid this insecurity.

2.4. Examples

Example 2.4. Consider the system P with $H = \{h_1, h_2\}$ and $L = \{l\}$ defined

$$P = (h_1 \rightarrow l \rightarrow P) \square (h_2 \rightarrow l \rightarrow P).$$

The system repeatedly offers a choice of a single H event followed by action l . Concealing H permits U_L to engage in l whenever desired independently of the (hidden) choice between h_1 and h_2 . Hence U_L 's view of $P \setminus H$ is deterministic and $\mathcal{E}Sec_L(P)$ holds.

We may doubt, however, whether the system should really be regarded as secure because the availability of l depends on the previous occurrence of an H action. We have

$$\langle h_1, l \rangle \in \text{TRACES}(P \parallel \text{RUN}_H), \text{ but } (\langle h_1, \{l\} \rangle \in \text{FAILS}(P \parallel \text{RUN}_H))$$

because the h_1 event from RUN_H has been interleaved with the empty trace from process P . This means that P has in fact made no progress and will therefore refuse to participate in l .

The lazy condition does not make the assumption that H actions occur so quickly such that no refusal to communicate l may be recognised by U_L . System P therefore is eagerly but not lazily secure, reflecting a dependence of U_L 's system view on activity of the other user.

Example 2.5. Let $H = \{h\}$ and $L = \{l_1, l_2\}$. In the system

$$Q = (h \rightarrow Q) \square (l_1 \rightarrow l_2 \rightarrow Q) \square (l_1 \rightarrow Q)$$

there is the possibility of an infinite sequence of h actions. This potentially endless delay of U_L 's actions is flagged by the eager condition since, e.g.,

$$\langle \rangle \in \text{DIVS}(Q \setminus H),$$

so Q is not eagerly secure. However $Q \parallel \text{RUN}_H$ is deterministic so that process Q does satisfy the lazy condition.

Example 2.6. Let $H = \{d_1, d_2, s_1, s_2\}$ and $L = \{l_1, l_2\}$. In the system

$$\begin{aligned} R &= (l_1 \rightarrow l_2 \rightarrow R) \\ &\square (d_1 \rightarrow s_1 \rightarrow R) \\ &\square (d_2 \rightarrow s_2 \rightarrow R) \end{aligned}$$

there is again the possibility of infinite pre-emption of L activity, so eager security cannot hold. R also fails the lazy condition since

$$\langle d_1, l_1 \rangle \in \text{TRACES}(R \parallel \text{RUN}_H), \text{ but } (\langle d_1, \{l_1\} \rangle \in \text{FAILS}(R \parallel \text{RUN}_H)).$$

The event l_1 will be removed from the interface when U_H engages in either d_1 or d_2 . Thus U_H will delay the system by communicating d_1 or d_2 until a further s action is taken. User U_L will recognise that the system refuses a request l_1 before U_H 's request is complete.

If we identify s_1 and s_2 as signal events, then we can partition H in the obvious way and apply the mixed security criterion, and find that $\mathcal{M}Sec_L^{(D,S)}(P)$ does indeed hold.

2.5. Properties of conditions

From the eager and lazy conditions based on the notion of determinism it is possible to derive conditions involving only observations of the process concerned. Eager security can be paraphrased as stating that nothing which is observed in L after trace tr will allow the H events which happened during tr to be inferred.

Proposition 2.7. *If system P satisfies $\mathcal{E}Sec_L(P)$, then $P \setminus H$ is free of divergence, and for any two traces $tr, tr' \in \text{TRACES}(P)$,*

$$tr \upharpoonright L = tr' \upharpoonright L \Rightarrow (P/tr) \setminus H =_{FD} (P/tr') \setminus H.$$

A corresponding consequence can be derived from the definition of lazy security.

Proposition 2.8. *If system P satisfies $\mathcal{L}Sec_L(P)$, then P is free of divergence, and for any two traces $tr, tr' \in \text{TRACES}(P)$,*

$$tr \upharpoonright L = tr' \upharpoonright L \Rightarrow (P/tr) \parallel \text{RUN}_H =_{FD} (P/tr') \parallel \text{RUN}_H.$$

The approach of postulating determinism after abstraction of high-security events can be generalized by analysing various models of U_H . The framework in which this can be done is provided by the condition

$$(P \parallel [H]U) \setminus H \text{ det} \tag{1}$$

for a suitability chosen process U which has to synchronise with P on every event in H . Process U can be regarded as a model of user U_H .

The user with the widest range of behaviour is one whose actions are unpredictable and uncontrollable. Such activity is represented in CSP by a process CHAOS defined as

$$\text{CHAOS}_H = \text{STOP} \square (x: H \rightarrow \text{CHAOS}_H)$$

displaying the most non-deterministic behaviour which is free of divergence. A system with such a non-deterministic user will lack interference only in the case of both eager and lazy security.

Proposition 2.9. *A system P satisfies $\mathcal{E}\text{Sec}_L(P)$ and $\mathcal{L}\text{Sec}_L(P)$ if, and only if, the process $(P \parallel [H] \text{CHAOS}_H) \setminus H$ is deterministic.*

It is shown in [13] that all three non-interference conditions (eager, lazy, and mixed) can in fact be expressed in the form of (1). For eager security, the model for user U_H is simply identical to RUN_H since $P \parallel [H] \text{RUN}_H = P$ for all processes P .

Corresponding formulations for lazy and mixed security require a more powerful model of CSP. In the infinite traces model [11] the failures-divergences representation of a process is augmented with its set of infinite traces. The model of user U_H required for lazy security is a process FINITE_H which behaves just like CHAOS_H but without ever engaging in an infinite trace. This restriction prohibits the occurrence of infinite H sequences resulting in divergence under concealment.

Proposition 2.10. *Eager, lazy, and mixed security can all be expressed in the general form of (1) as follows.*

$$\mathcal{E}\text{Sec}_L(P) \Leftrightarrow (P \parallel [H] \text{RUN}_H) \setminus H \text{ det},$$

$$\mathcal{L}\text{Sec}_L(P) \Leftrightarrow (P \parallel [H] \text{FINITE}_H) \setminus H \text{ det},$$

$$\mathcal{M}\text{Sec}_L^{(D,S)}(P) \Leftrightarrow (P \parallel [H] (\text{RUN}_S \parallel \text{FINITE}_D)) \setminus H \text{ det}.$$

These various ‘users’ suggest a more general approach to security specification: for a particular context, choose a process U which characterises all possible behaviours of U_H under which it is expected that confidentiality will be maintained. Usually this will be all its behaviours, but it is possible to imagine other circumstances, for example if the system P represents a mail system where it is allowable for a high-security user to send a message to a low-security one, we might expect to maintain confidentiality so long as no such messages are sent. (This type of property is known as *conditional* non-interference [7].)

The more non-deterministic the abstract model U the stronger is the equivalent security condition. When a more deterministic process is substituted for U , the properties of CSP refinement guarantee the preservation on non-interference.

More precisely, if P is a system component in context C , then refinement of P – replacing it with a less non-deterministic component – preserves determinism of the original system:

$$C(P) \text{ det} \wedge P \sqsubseteq P' \Rightarrow C(P') \text{ det}.$$

It is equally a consequence of this fact that refining P preserves the determinism of $(P \parallel [H] U) \setminus H$, and that therefore each of our non-interference properties is preserved under refinement. This is a result which may be exploited in system development or maintenance.

Proposition 2.11. *Eager, lazy, and mixed security are preserved under refinement:*

$$\mathcal{E}\text{Sec}_L(C(P)) \wedge P \sqsubseteq P' \Rightarrow \mathcal{E}\text{Sec}_L(C(P')),$$

$$\mathcal{L}\text{Sec}_L(C(P)) \wedge P \sqsubseteq P' \Rightarrow \mathcal{L}\text{Sec}_L(C(P')),$$

$$\mathcal{M}\text{Sec}_L^{(D,S)}(C(P)) \wedge P \sqsubseteq P' \Rightarrow \mathcal{M}\text{Sec}_L^{(D,S)}(C(P')).$$

A number of additional results concerning the composition and decomposition of secure systems may be derived [13, 14]. One such result, proved in detail in [14] is that a system P may be decomposed into two non-interacting parts if it is lazily secure with respect to two disjoint alphabets.

Proposition 2.12. *Let A, B be disjoint alphabets. $\mathcal{L}\text{Sec}_A(P)$ and $\mathcal{L}\text{Sec}_B(P)$ hold of a system if, and only if, there are two deterministic processes P_A with $\alpha(P_A) = A$ and P_B with $\alpha(P_B) = B$ such that $P = P_A \parallel [\emptyset] P_B$.*

Further properties of our non-interference conditions as well as the proofs of the propositions in this section may be found in the appendix and in [13, 14].

2.6. Generalisation

We will now generalise the determinism conditions for multi-user systems. If F is the system whose non-interference properties we attempt to establish, the system model can be described as

$$\text{SYSTEM} = \text{Users} \parallel [S] F, \quad \text{where } \text{Users} = \parallel_{i>0} U_i$$

and S is the users’ interface (common events) to the system. It is assumed that there is a security classification associated with each user. Let CLASS be the partially ordered set of these classifications. The total function $cl: \text{Users} \rightarrow \text{CLASS}$ assigns one classification to each user process. A further assumption is $\alpha(U_i) \cap \alpha(U_j) = \emptyset$ whenever $cl(U_i) \neq cl(U_j)$.

The function $low: \text{CLASS} \rightarrow \mathbb{P}\alpha(\text{SYSTEM})$ yields the set of events that are regarded as ‘low-level’ on a particular level of security classification. It is defined

$$low(c) = \bigcup \{ \alpha(U_i) \mid cl(U_i) \leq c \}.$$

The non-interference conditions of Subsection 2.3 hold for a multi-user system if they hold on each security level of the system.

Definition 2.13. A multi-user system P eagerly and (respectively) lazily-secure if,

$$\forall c_i \in \text{CLASS} \bullet \mathcal{E}\text{Sec}_{low(c_i)}(P), \text{ and}$$

$$\forall c_i \in \text{CLASS} \bullet \mathcal{L}\text{Sec}_{low(c_i)}(P).$$

In a realistic system it is typically the mixed non-interference condition that requires verification on each security level, as will be illustrated in the following case study.

3. A "secure" file system

This and the following section will illustrate the framework in which our non-interference conditions can be applied. The example is that of a file system in which confidential data is to be maintained.

It is widely accepted that a formal specification can increase the level of assurance that a system will meet its security requirements [5]. In fact governmental standards for the development of secure systems mandate the use of formal methods and proof. The Z notation [18] is particularly suited for this task since (i) it has a well-defined semantics; (ii) it has been successfully employed in industrial scale software development; and (iii) it has become increasingly popular for the specification and verification of secure systems [3, 4, 10].

We begin the specification of the file system by introducing some basic types. The set of users of the system is represented by type *USER*, each of which holds an associated security classification from the set *CLASS*. *FID* represents the set of file identifiers, and *DATA* refers to the set of possible data that may be stored in a file. This type contains a special value *NULL* representing invalid data. These are the basic types we will use

$$[USER, CLASS, FID, DATA]$$

There is a security classification associated with each user. We use a global function *cl* to obtain the appropriate class by supplying it with a user identification. It is declared as a *total* function; there cannot be users without classification.

$$| cl: USER \rightarrow CLASS$$

3.1. File model

Each file has the structure³

<i>File</i> <i>class</i> : <i>CLASS</i> <i>data</i> : <i>DATA</i>

³In Z, formal notation is separated from informal descriptions by so-called schema boxes. A schema contains a number of declarations and, if there are any constraints on these declarations, a separating line followed by appropriate predicates. Schemas are used to represent structured state as well as operations on structures.

where the *class* component relates to the level of security of the stored *data*. Each file initialised with the level of clearance at which the file is created.

<i>Init</i> <i>File'</i> <i>clear?</i> : <i>CLASS</i> <hr/> <i>class'</i> = <i>clear?</i> <i>data'</i> = <i>NULL</i>
--

We follow a standard convention of decorating inputs with ?, outputs with !, and states after completion of the operation with a prime '. Unprimed variables or schemas refer to states before the operation.

Two operations are provided on files: reading stored data, and writing new data, provided the file access is carried out with the appropriate clearance. Reading is permitted only when the operation is carried out with appropriate access permission $clear? \geq class$, in which case stored data is output as *data!*. The notation $\exists File$ indicates that reading a file does not change its state.

<i>Rd0</i> $\exists File$ <i>clear?</i> : <i>CLASS</i> <i>data!</i> : <i>DATA</i> <hr/> <i>clear?</i> \geq <i>class</i> <i>data!</i> = <i>data</i>

Storing new data in a file is carried out with a *Wr0* operation which is permitted only if the user clearance is equal to the file classification. The $\Delta File$ schema component indicates that writing data changes the file state; the input data *new?* is stored in the *data* component of *File*.

<i>Wr0</i> $\Delta File$ <i>clear?</i> : <i>CLASS</i> <i>new?</i> : <i>DATA</i> <hr/> <i>clear?</i> = <i>class</i> = <i>class'</i> <i>data'</i> = <i>new?</i>
--

To indicate the success or failure of an operation, we define the system's response as type

$$RESP ::= ok \mid fail.$$

Each operation on a file is accompanied by an indication of whether it has succeeded. The output message is defined by the (horizontal) schemas

$$\text{Success} \hat{=} [\text{resp!} : \text{RESP} \mid \text{resp!} = \text{ok}],$$

$$\text{Failure} \hat{=} [\text{resp!} : \text{RESP} \mid \text{resp!} = \text{fail}].$$

We do not give the user any indication of whether a failure was caused by a functional error or a security breach, in order to avoid a potential channel of information flow.

If a request for file access is carried out without valid clearance the operation fails, and the file status remains unchanged ($\exists \text{File}$). The case of invalid read access is described as

$\begin{array}{l} \text{NoRdAccess} \\ \exists \text{File} \\ \text{Failure} \\ \text{clear?} : \text{CLASS} \\ \text{data!} : \text{DATA} \\ \hline \text{clear?} < \text{class} \\ \text{data!} = \text{NULL} \end{array}$
--

The corresponding error condition for writing is

$\begin{array}{l} \text{NoWrAccess} \\ \exists \text{File} \\ \text{Failure} \\ \text{clear?} : \text{CLASS} \\ \hline \text{clear?} \neq \text{class} \end{array}$

The total read and write operations are Rd and Wr specified as

$$Rd \hat{=} (Rd0 \wedge \text{Success}) \vee \text{NoRdAccess},$$

$$Wr \hat{=} (Wr0 \wedge \text{Success}) \vee \text{NoWrAccess}.$$

If the request is carried out with appropriate clearance the system reports with *ok*, otherwise the user just receives a *fail* message and the file remains unaltered.

3.2. File system

Our file system is given by

$\begin{array}{l} \text{FileSystem} \\ \text{files} : \text{FID} \rightarrow \text{File} \end{array}$

Component *files* is declared as a partial function from file identifiers to files. This means that no two files can have the same name. The system initially contains no files:

$$\text{FInit} \hat{=} [\text{FileSystem}' \mid \text{files}' = \emptyset].$$

In addition to the initialisation occurring when a file is created at the system level, we want the operations of reading and writing a file to be available at the system interface. This is achieved by *promoting* the schemas *Init*, *Rd*, and *Wr* with the aid of two “framing” schemas:

$\begin{array}{l} \Phi1 \\ \Delta \text{FileSystem} \\ \text{file?} : \text{FID} \\ \text{user?} : \text{USER} \\ \hline \text{clear?} = \text{cl}(\text{user?}) \\ \text{files}' = \text{files} \oplus \{\text{file?} \mapsto \theta \text{File}'\} \end{array}$

$\begin{array}{l} \Phi2 \\ \Phi1 \\ \hline \text{file?} \in \text{dom files} \\ \theta \text{File} = \text{files}(\text{file?}) \end{array}$
--

The promoted operations will require both a file name (*file?*) and a user identification (*user?*) as input. The user's classification is then the clearance at which the file operation is carried out. The three operations available at the interface are

$$\text{Create0} \hat{=} \exists \text{File}' \bullet (\Phi1 \wedge \text{Init}),$$

$$\text{Read0} \hat{=} \exists \Delta \text{File} \bullet (\Phi2 \wedge \text{Rd}),$$

$$\text{Write0} \hat{=} \exists \Delta \text{File} \bullet (\Phi2 \wedge \text{Wr}).$$

It is necessary to ensure that no operation is carried out on files which do not exist. This error condition can occur if the user supplies an invalid file identifier.

$\begin{array}{l} \text{UnknownFile} \\ \exists \text{FileSystem} \\ \text{Failure} \\ \text{file?} : \text{FID} \\ \hline \text{file?} \notin \text{dom files} \end{array}$
--

Similarly, a request for file creation cannot succeed if the suggested name has already been used for another file.

<i>FileExists</i>
$\exists FileSystem$
<i>Failure</i>
<i>file?</i> : <i>FID</i>
$file? \in \text{dom files}$

The total operations available at the file system interface are then given by

$$Create \hat{=} (Create0 \wedge Success) \vee FileExists,$$

$$Read \hat{=} Read0 \vee UnknownFile,$$

$$Write \hat{=} Write0 \vee UnknownFile.$$

We suggest that a security analysis is best carried out on a process algebraic representation of the system. This representation may be regarded as a security model [5] which can in fact be derived by translation. It is therefore unnecessary to engage in the formidable task to prove a correspondence between model and specification. In the coming section we map the functional specification of the file system into CSP and state the non-interference conditions that require verification.

4. Z into CSP

The Z specification may be translated into CSP according to the technique described in [20]. The theoretical basis for this work may be found in [19].

First we interpret the Z specification as an action system [2] whose state is specified by *File*. It has two actions corresponding to the operations *Rd* and *Wr*. However, each of these operations also has an output, and we must be careful to separate the two parts of the operation and associate an action with each, since we cannot regard input and output as happening simultaneously. When a user has invoked an operation, but has not consumed its output, then the system will do nothing else while that output is pending. When no output is pending, all operation actions are enabled.

This interpretation of a Z specification is informal (albeit systematic), but it does correspond to the intuitive meaning given to Z specifications (see [18], for example).

Consider the *Wr* operation. We must separate it into two parts: the first part consumes the input and then stores its output in the state; the second part waits for the opportunity of delivering its output. Define a new free type that is either a response or nothing:

$$RESP_+ ::= nullresp \mid outresp(\langle RESP \rangle)$$

and augment the state of a file with a component that contains the pending output (if it exists)

<i>File</i> ₊
<i>File</i>
<i>wrpend</i> : <i>RESP</i> ₊

The first part of the operation is as follows:

<i>Wr</i> ₊
$\Delta File_+$
<i>clear?</i> : <i>CLASS</i>
<i>new?</i> : <i>DATA</i>
$wrpend = nullresp$
$\exists resp! : RESP \mid wrpend' = outresp(resp!) \bullet Wr$

and the second part is

<i>Wr</i> ₋
$\Delta File_+$
$\exists File$
<i>resp!</i> : <i>RESP</i>
$wrpend \neq nullresp$
$resp! = outresp^{\sim}(wrpend)$
$wrpend' = nullresp$

We can prove that the only change we are making to *Wr* by splitting into two is to delay its output:

$$\vdash Wr = \exists wrpend, wrpend' : RESP_+ \bullet Wr_+ \circ_w Wr_-$$

According to [20], we can now translate our specification of the write operation into two actions.

$$wr?clear?new \wedge wrpend = nullresp \rightarrow Wr_+$$

$$wrout!(outresp^{\sim}(wrpend)) \wedge wrpend \neq nullresp \rightarrow wrpend := nullresp$$

Thus, upon receipt of the communication of a clearance and some new data, then, providing that there is no write-output pending, the *Wr*₊ operation is performed. Output may be transmitted whenever it is pending.

We can make similar transformations for the other operations.

The actions may now be embedded in a CSP-framework process. We now have a CSP process which is formally equivalent to the *File* abstract data type.

$$File = init?class \rightarrow File(class, NULL, (nullresp, nulldata), nullresp)$$

```

File(class, data, rdpend, wrpend) =
  if rdpend = nullresp  $\wedge$  wrpend = nullresp
  then
    rd?clear  $\rightarrow$ 
      if clear  $\geq$  class
      then File(class, data, (outresp(ok), outdata(data)), wrpend)
      else File(class, data, (outresp(fail), outdata(NULL)), wrpend)
     $\square$  wr?clear?new  $\rightarrow$ 
      if clear = class
      then File(class, new, rdpend, outresp(ok))
      else File(class, data, rdpend, outresp(fail))
  else
    if rdpend  $\neq$  nullresp
    then rdout!outresp $\sim$ (rdpend)
       $\rightarrow$  File(class, data, (nullresp, nulldata), wrpend)
    else wrout!outresp $\sim$ (wrpend)
       $\rightarrow$  File(class, data, rdpend, nullresp)

```

The *File* process may now be transformed using the laws of CSP, and, if desired, the state variables containing the pending outputs elided.

In [20], the connection is made between the technique of promotion in Z, and the use of subordination or the means of sharing through interleaving. In this way, the file system can be created as a system of CSP processes.

The structure of the resulting CSP implementation is illustrated in Fig. 1. *FILES* will be a shared pool of files accessible through the interface *FSYS*. Each

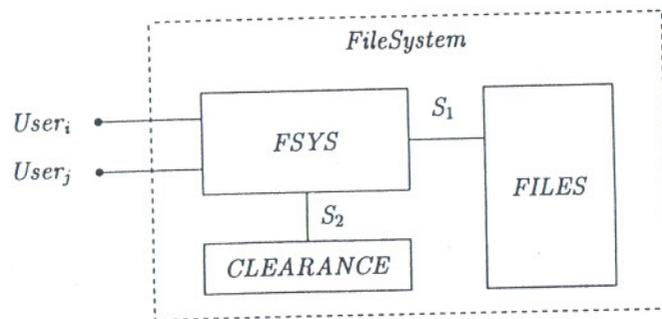


Fig. 1. The file system implemented by communicating processes.

file has an associated name and classification, and may contain arbitrary data. The process *File* models a file waiting to be initialised.

$$File = init?file?class \rightarrow File(file, class, NULL)$$

A file after initialisation may be read or written to.

$$File(file, class, data) = Rd(file, class, data) \square Wr(file, class, data)$$

The read operation is implemented by process *Rd* as

```

Rd(file, class, data) = rd.file?clear  $\rightarrow$ 
  if clear  $\geq$  class
  then rdout.file!ok!data  $\rightarrow$  File(file, class, data)
  else rdout.file!fail!NULL  $\rightarrow$  File(file, class, data)

```

Storing new data in a file is realised with process

```

Wr(file, class, data) = wr.file?clear?new  $\rightarrow$ 
  if clear = class
  then wrout.file!ok  $\rightarrow$  File(file, class, new)
  else wrout.file!fail  $\rightarrow$  File(file, class, data)

```

The total pool of files is given by

$$FILES = \parallel_{0 \leq i < n} File.$$

4.1. The system interface

It is not possible to conjoin *FILES* with the set of user processes directly because users must be protected from a number of functional errors, such as reading a file which does not exist. To this purpose, we will provide a system interface process *FSYS* which manages access to the individual files.

$$\alpha(FSYS) = S_0 \cup S_1 \cup S_2$$

$$S_0 = \{create, createout, read, readout, write, writeout\}$$

$$S_1 = \{init, rd, rdout, wr, wrout\}$$

$$S_2 = \{clear\}$$

FSYS holds state variable *files*, the set of current file names

$$FSYS(files) = Create(files) \square Read(files) \square Write(files)$$

The three services available at the interface are implemented with processes *Create*, *Read*, and *Write*, respectively.

```

Create(files) = create?user?file →
  if file ∉ files
  then clear.user?class → init!file!class → createout.user!ok
    → FSYS(files ∪ {file})
  else createout.user!fail → FSYS(files)

```

```

Read(files) = read?user?file →
  if file ∈ files
  then clear.user?class → rd.file!class → rdout.file?resp?result
    → readout.user.file!resp!result → FSYS(files)
  else readout.user.file!fail!NULL → FSYS(files)

```

```

Write(files) = write?user?file?new →
  if file ∈ files
  then clear.user?class → wr.file!class!new → wrout.file?resp
    → writeout.user.file!resp → FSYS(files)
  else writeout.user.file!fail → FSYS(files)

```

Process *CLEARANCE* provides the appropriate classification of a user when required.

$$CLEARANCE = (\square clear.u!(cl(u)) \rightarrow CLEARANCE)$$

for all $u \in USER$

The complete file system is given by parallel composition of the interface process, the file pool, and the clearance process, with intermediate channels concealed.

$$FileSystem = ((FSYS(\emptyset) || [S_1] FILES) || [S_2] CLEARANCE) \setminus S_1 \cup S_2$$

4.2. Security specification

Any particular instance of the file system can be subjected to the security conditions presented in Section 2. We will consider the case of three users with the following classifications.

USER	CLASS
Lisa	3 (highest)
Mari	2
Nina	1 (lowest)

It is convenient to partition the events at the system interface into “delay” and “signal” events on each level of user classification (except the top level).

$$H2d = \{ create.user.file, read.user.file, write.user.file.data \mid user \in \{Lisa\} \}$$

$$H2s = \{ createout.user.resp, readout.user.file.resp.data, writeout.user.file.resp \mid user \in \{Lisa\} \}$$

$$H1d = \{ create.user.file, read.user.file, write.user.file.data \mid user \in \{Lisa, Mari\} \}$$

$$H1s = \{ createout.user.resp, readout.user.file.resp.data, writeout.user.file.resp \mid user \in \{Lisa, Mari\} \}$$

The file system satisfies eager security if

$$(FileSystem \setminus (H2d \cup H2s)) \det \wedge (FileSystem \setminus (H1d \cup H1s)) \det.$$

The file system is lazily secure if

$$(FileSystem ||| RUN_{(H2d \cup H2s)}) \det \wedge (FileSystem ||| RUN_{(H1d \cup H1s)}) \det.$$

The file system satisfies the mixed security property if

$$((FileSystem \setminus H2s) ||| RUN_{H2d}) \det \wedge ((FileSystem \setminus H1s) ||| RUN_{H1d}) \det.$$

It turns out that none of these conditions is met – i.e., that the system contains undesired information flows. Since it may not be obvious that the conditions fail to hold (and why not), we employ a verification tool.

5. Automated verification

The effort of formulating the eager/lazy/mixed non-interference conditions would be futile without a method of verifying them. Luckily, the absence of non-determinism on which the conditions are based can be automatically verified using standard algorithms on finite-state systems. We show that the CSP proof tool FDR can be used to complete the security analysis.

5.1. FDR

The FDR tool [12] has been originally designed to verify behavioural CSP specifications, in particular refinement relations between processes. These refinement checking capabilities are employed to decide whether a given process is deterministic using the following algorithm:

1. Search through the state space of P , resolving all non-determinism that is encountered. In a "stable" state (in which internal progress is impossible) a single representative for each available action is selected, whereas in a state where internal actions are possible we chose one of these arbitrarily. This search either finds a divergence of P (in which case it is clearly non-deterministic) or yields a deterministic process Q that refines the original P .
2. Use the refinement checker to confirm whether $Q \sqsubseteq P$. The check succeeds if, and only if, P is deterministic.

The algorithm is justified by the fact that the deterministic processes are maximal in the failures-divergences model of CSP, and are therefore incomparable. Thus, for some arbitrary deterministic refinement Q of P ,

$$P \text{ det} \Leftrightarrow P =_{FD} Q.$$

5.2. Making the file system secure

Checking the security specification of Subsection 4.2 using FDR confirms that the file system is neither eagerly nor lazily secure. There is no eager security since there is the danger of infinite Lisa and Lisa/Mari activity, thus pre-empting any progress of the low-level users. This danger is detected by recognising divergence of both processes whose determinism was to be verified. The reason for lack of lazy security lies in the basic structure of the system interface: a menu of services is offered to users with various classifications, and a choice of service by a particular user is followed by a system response on the same security level.

This structure resembles that of the (much simpler) process R of Example 2.6 which was already observed to be insecure under the eager and lazy conditions. As was motivated there, these "pure" conditions are inappropriate for a system whose interface has a request-response type structure (like R or *FileSystem*), and it becomes necessary to partition events into delay and signal events in order to apply the mixed condition.

However the mixed condition fails to hold as well, which must be of serious concern to the system designers. FDR flags the security breach as follows:

$$\langle \text{create.lisa.id}, \text{create.nina.id}, \text{createout.nina.ok} \rangle$$

$$\in \text{TRACES}((\text{FileSystem} \setminus H1s) \parallel \parallel \text{RUN}_{H1d})$$

since event *create.lisa.id* has been contributed by RUN_{H1d} . We also have

$$\langle \langle \text{create.lisa.id}, \text{create.nina.id} \rangle, \{ \text{createout.nina.ok} \} \rangle$$

$$\in \text{FAILS}((\text{FileSystem} \setminus H1s) \parallel \parallel \text{RUN}_{H1d})$$

since *create.lisa.id* could in fact have been an actual user request for file creation. The file system, however, was specified to prohibit the existence of two files with

the same name; this feature is a security flaw since a user who attempts to create a file (with identifier *id* say) and fails has learned that a file named *id* of higher classification exists. This clear breach of non-interference is reflected in the failure of the mixed condition.

The question remains how the flaw can be overcome. One idea may be to change the *Create* operation so that a request of file creation always succeeds. This approach is probably unsatisfactory if creation of a file which already exists results in stored data to be lost. A more promising approach is to somehow associate classifications with file identifiers in order to guarantee that files on different security levels have different names.

A simple way of implementing this is to provide pairwise disjoint sets of identifiers for the different levels. For the system in Section 4 one might consider partitioning the set *FID* into three sets (say)

$$FID_1 = \{a, b\}, FID_2 = \{c, d\}, FID_3 = \{e, f\},$$

so that for all $i \in \text{CLASS}$

$$FID = \bigcup FID_i.$$

Doing so entails the re-definition of the *Create* operation which now needs to confirm whether the use of a particular identifier is valid with regard to the user's classification:

```

Create(files) = create?user?file →
  if file ∉ files
  then clear.user?class →
    if file ∈ FIDclass
    then init!file!class → createout.user!ok → FSYS(files ∪ {file})
    else createout.user!fail → FSYS(files)
  else createout.user!fail → FSYS(files)

```

Verification of the mixed security condition now shows that the mixed security condition is satisfied, *provided* that the number of files available through the system is equal to or exceeds the combined total of identifiers for all levels of classification. So if

$$FILES = \parallel_{0 \leq i < n} File$$

we require $n \geq \text{size}(FID)$. Without the proviso the file system does not pass the mixed condition, again because an attempt of the file creation may fail. This time the security breach is caused by the potential exhaustion of the pool of available files.

6. Conclusion

This paper presents process algebraic specifications as a practical framework for the development of systems with security constraints. The approach is illustrated with an example of a file system intended to maintain secret data, but in fact our results apply equally to systems with security concerns other than confidentiality. This is a consequence of defining general non-interference conditions which require the system view of particular users to be unaffected by the actions taken by others.

Our non-interference conditions are based on the notion of deterministic views. This elegant characterisation of secure systems has only recently been recognised as the fundamental concept underlying various definitions of non-interference, such as those surveyed in [6]. Although these alternative definitions are cast in rather different notation without employing determinism, Roscoe [13] demonstrates that many are either straightforward consequence of, or closely related to, the conditions for eager, lazy, and mixed security. For example, our lazy property $\mathcal{L}Sec_L(P)$ corresponds precisely both to Graham-Cumming's own non-interference property and those of Allen [1] and Ryan [17] for systems whose overall behaviour is deterministic (as opposed to the abstractions used in formulating our properties). A significant advantage of our conditions in comparison to others is the preservation of non-interference under refinement, thus eliminating the potential compromise of security during development. A detailed discussion of this phenomenon, and an explanation of why it is desirable, may be found in [13].

The general framework envisaged for the development of secure systems falls into two parts: functional specifications of the system using state-based notations as for general applications, followed by an analysis of non-interference properties of a process-algebraic representation of the system. The main advantage of this method is in avoiding the complex treatment of establishing a correspondence between the specification and a separate generic security model. In contrast, the mapping of the specification into process algebra can in many cases be carried out by direct translation (tool support for this task is, however, at present not available). Process algebras such as CSP based on possible sequences (traces) of events provide an ideal notation for non-interference analysis since they naturally incorporate the notion of (non-)determinism, thus permitting the application of the conditions of Section 2. These conditions can then be automatically verified using a currently available proof tool.

Initial experience with the CSP model checker FDR [12] shows that a security analysis as illustrated in Section 5 can be carried out within minutes. This result propounds the hope that the verification approach will scale up to systems of realistic size. The size of problem we can deal with will benefit from the proposed development [12] of FDR to incorporate *implicit* model-checking techniques such as the hierarchical compression of intermediate state-spaces. Verification speed will further increase by the exploitation of behavioural independence of processes from particular values of data communicated. This property of *data-independence*

[15] has already shown promise in significant reduction of state spaces as well as the induction of properties of arbitrary data types based on finite checks.

Future work is required to formalise the mapping of state-based specifications to process descriptions. The techniques of [19, 20] still have to be extended to be applicable to specifications with complex semantics, and utilised to provide tool support for the translation into process algebra. We intend to apply these techniques and the framework outlined in this paper in a case study of a large-scale secure system. A further avenue of research is to explore potential applications of our determinism-based conditions for non-interference on systems with critical requirements other than security, such as in the areas of safety-critical systems, fault tolerance, and feature independence.

References

- [1] P.G. Allen, A comparison of non-interference and non-deducibility using CSP, in: *Proc. 1991 IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 1991, pp. 43–54.
- [2] R.-J. Back and R. Kurki-Suonio, Decentralization of process nets with centralized control, in: *Proc. 2nd Annual Symposium on Principles of Distributed Computing*, Montreal, 1983.
- [3] T. Boswell, Specification and validation of a security policy model, in: *Proc. FME '93 Symposium*, LNCS 670, Springer-Verlag, April, 1993.
- [4] R. Collinson, A critical look at functional specifications, in: *Proc. FME '94 Symposium*, LNCS 873, Springer-Verlag, October, 1994.
- [5] M. Gasser, *Building a Secure Computer System*, Van Nostrand Reinhold, 1988.
- [6] J. Graham-Cumming, *The Formal Development of Secure Systems*, Oxford University, DPhil Thesis, 1992.
- [7] J. Goguen and J. Meseguer, Security policy and security models, in: *Proc. 1982 IEEE Symposium on Security and Privacy*, IEEE Computer Society Press, 1982.
- [8] C.A.R. Hoare, *Communicating Sequential Processes*, Prentice-Hall, 1985.
- [9] J.L. Jacob, Specifying security properties, in: *Developments in Concurrency and Communication*, C.A.R. Hoare, ed., ACM Press, 1990.
- [10] R.B. Jones, Methods and tools for the verification of critical properties, in: *Proc. 5th Refinement Workshop*, C.B. Jones, R.C. Shaw and T. Denvir, eds, Springer-Verlag, London, 1992.
- [11] A.W. Roscoe, Unbounded non-determinism in CSP, *Journal of Logic and Computation* 3 (1993).
- [12] A.W. Roscoe, Model checking CSP, in: *A Classical Mind*, A.W. Roscoe, ed., Prentice-Hall, 1994.
- [13] A.W. Roscoe, CSP and determinism in security modelling, in: *Proc. 1995 IEEE Symposium on Security and Privacy*, IEEE Computer Society Press, 1995.
- [14] A.W. Roscoe and L. Wulf, Composing and decomposing systems under security properties, in: *Proc. 1995 IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 1995.
- [15] A.W. Roscoe and H. MacCarthy, Verifying a replicated database: A case study in model-checking CSP, submitted for publication.
- [16] A.W. Roscoe, J.C.P. Woodcock and L. Wulf, Non-interference through determinism, in: *Proc. Third European Symposium on Research in Computer Security (ESORICS'94)*, LNCS 875, Springer-Verlag.
- [17] P.Y.A. Ryan, A CSP formulation of non-interference, *Cipher*, IEEE Computer Society Press, 1991, 19–27.
- [18] J.M. Spivey, *The Z Notation: A Reference Manual* (2nd ed.), Prentice-Hall International, 1992.
- [19] J.C.P. Woodcock and C. Morgan, Refinement of state-based concurrent systems, *Proc. VDM Symposium 1990*, LNCS 428, Springer-Verlag.
- [20] J.C.P. Woodcock, CSP interpretations of Z specifications, in preparation.

Appendix A: CSP reference

A.1. Processes

αP	alphabet (possible events) of process P
P/t	process P after it has engaged in trace t
$P \parallel [A] Q$	parallel composition with synchronisation (only) on events A
$P \parallel\!\!\parallel Q$	parallel interleaving (without synchronisation)
$P \sqcap Q$	internal (non-deterministic) choice between P and Q
$P \square Q$	external (deterministic) choice between P and Q
$a: A \rightarrow P$	prefixing of P with an event from A

$STOP$, RUN , and $CHAOS$ are special processes. Process $STOP$ never engages in any event. RUN_A is always willing to contribute an event from set A :

$$RUN_A = a: A \rightarrow RUN_A.$$

$CHAOS_A$ is the most non-deterministic, but divergence-free, process communicating in A :

$$CHAOS_A = STOP \sqcap (a: A \rightarrow CHAOS_A).$$

A.2. Semantics

In the failures-divergences model, each process P is determined by its failure set, $FAILS(P)$, and its divergence set, $DIVS(P)$.

- each failure is a pair (s, X) where s is a finite trace of the process and X is a set of events which it may refuse after s , and
- each divergence is a finite trace on which the process can perform an infinite sequence of internal actions.

In the failures-divergences model of CSP, two processes are regarded as equal if they agree in their failures and their divergences:

$$P =_{FD} Q \Leftrightarrow FAILS(P) = FAILS(Q) \wedge DIVS(P) = DIVS(Q)$$

When a process Q is more deterministic than another process P we say that P is refined by Q . This relation is written $P \sqsubseteq Q$ and formally defined by

$$P \sqsubseteq Q \Leftrightarrow FAILS(P) \supseteq FAILS(Q) \wedge DIVS(P) \supseteq DIVS(Q)$$

The semantics of the hiding operator in the failures-divergences model is given by

$$FAILS(P \setminus A) = \{(s \setminus A, X) \mid (s, X \cup A) \in FAILS(P)\} \cup \{(s, X) \mid s \in DIVS(P \setminus A)\}$$

$$\begin{aligned} DIVS(P \setminus A) &= \{(s \setminus A) \hat{\ } t \mid s \in DIVS(P)\} \cup \\ &\quad \{(s \setminus A) \hat{\ } t \mid (\forall n \in \mathbb{N} \bullet (\exists u \in A^* \bullet \\ &\quad \#u > n \wedge s \hat{\ } u \in TRACES(P)))\} \end{aligned}$$

The set of all possible interleavings of traces s and t is $s \parallel\!\!\parallel t$, defined as

$$s \parallel\!\!\parallel t = t \parallel\!\!\parallel s$$

$$s \parallel\!\!\parallel \langle \rangle = \{s\}$$

$$\langle a \rangle \hat{\ } s \parallel\!\!\parallel \langle b \rangle \hat{\ } t = \{\langle a \rangle \hat{\ } r \mid r \in s \parallel\!\!\parallel \langle b \rangle \hat{\ } t\} \cup \{\langle b \rangle \hat{\ } r \mid r \in \langle a \rangle \hat{\ } s \parallel\!\!\parallel t\}$$

The semantics of process interleaving is defined as

$$FAILS(P \parallel\!\!\parallel Q) = \{(r, X) \mid \exists s, t \bullet r \in s \parallel\!\!\parallel t \wedge$$

$$(s, X) \in FAILS(P) \wedge (t, X) \in FAILS(Q)\} \cup$$

$$\{(u, X) \mid u \in DIVS(P \parallel\!\!\parallel Q) \wedge X \subseteq \Sigma\}$$

$$DIVS(P \parallel\!\!\parallel Q) = \{r \hat{\ } u \mid \exists s, t \bullet r \in s \parallel\!\!\parallel t \wedge$$

$$(s \in DIVS(P) \wedge t \in TRACES(Q)) \wedge$$

$$(s \in TRACES(P) \wedge t \in DIVS(Q))\}$$

The set of all possible interleavings of traces s and t which agree on events S is $s \parallel [S] t$, defined as

$$s \parallel [S] t = t \parallel [S] s$$

$$s \parallel [S] \langle \rangle = \{s\} \text{ (if } s \upharpoonright S = \langle \rangle)$$

$$\langle a \rangle \hat{\ } s \parallel [S] \langle b \rangle \hat{\ } t = \{\langle a \rangle \hat{\ } r \mid r \in s \parallel [S] t\} \text{ (if } a \in S)$$

$$\langle a \rangle \hat{\ } s \parallel [S] \langle b \rangle \hat{\ } t = \{\langle b \rangle \hat{\ } r \mid r \in \langle a \rangle \hat{\ } s \parallel [S] t\} \text{ (if } a \in S \wedge b \notin S)$$

$$\langle a \rangle \hat{\ } s \parallel [S] \langle b \rangle \hat{\ } t = \{\langle a \rangle \hat{\ } r \mid r \in s \parallel [S] \langle b \rangle \hat{\ } t\} \cup$$

$$\{\langle b \rangle \hat{\ } r \mid r \in \langle a \rangle \hat{\ } s \parallel [S] t\} \text{ (if } a, b \notin S)$$

The shared interface parallel operator between processes is defined

$$FAILS(P \parallel [S] Q) = \{(r, (X \cap Y) \cup (X \cap S) \cup (Y \cap S)) \mid$$

$$\exists s, t \bullet r \in s \parallel [S] t \wedge$$

$$(s, X) \in FAILS(P) \wedge (t, Y) \in FAILS(Q)\} \cup$$

$$\{(u, X) \mid u \in DIVS(P \parallel [S] Q) \wedge X \subseteq \Sigma\}$$

$$DIVS(P \parallel [S] Q) = \{r \hat{\ } u \mid \exists s, t \bullet r \in s \parallel [S] t \wedge$$

$$(s \in DIVS(P) \wedge t \in TRACES(Q)) \wedge$$

$$(s \in TRACES(P) \wedge t \in DIVS(Q))\}$$

Appendix B: Proofs

Proof of Proposition 2.8. If tr and tr' are two traces of P such that $tr \upharpoonright L = tr' \upharpoonright L$, then we can find a trace tr'' of $P \parallel\parallel RUN_H$ such that

- $tr'' \upharpoonright L = tr' \upharpoonright L = tr \upharpoonright L$, and
- tr and tr' are both subsequences of tr'' .

In other words, there exist $u, v \in H^*$ such that tr'' is both an interleaving of tr and u , and one of tr' and v . If $(P/tr) \parallel\parallel RUN_H \neq (P/tr') \parallel\parallel RUN_H$, with the difference showing up after trace s , then there would be nondeterminism in $P \parallel\parallel RUN_H$ after trace $tr'' \hat{\ } s$. \square

Proof of Proposition 2.9. (1) Suppose that $(P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H$ is deterministic. Since we have

$$P \setminus H = (P \parallel\parallel [H] \parallel\parallel RUN_H) \setminus H \supseteq (P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H,$$

by standard properties of CSP (as $RUN_H \supseteq CHAOS_H$), $P \setminus H$ must equal $(P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H$ as deterministic processes have no proper refinements. Thus P satisfies $\mathcal{E}Sec_L(P)$. If it fails $\mathcal{L}Sec_L(P)$ then $P \parallel\parallel RUN_H$ exhibits nondeterminism. Since certainly it is divergence-free, it follows that there must be $s \in TRACES(P \parallel\parallel RUN_H)$ and $a \in L$ such that

$$(s, \{a\}) \in FAILS(P \parallel\parallel RUN_H) \text{ and } s \hat{\ } (a) \in TRACES(P \parallel\parallel RUN_H).$$

(The fact that a is in L rather than H is a consequence of the fact that $P \parallel\parallel RUN_H$ never refuses an element of H .) As before, there are thus traces tr_1 and tr_2 of P , and $u, v \in H^*$ such that

- s is an interleaving of tr_1 and u , and one of tr_2 and v (from which it follows that $tr_1 \upharpoonright L = tr_2 \upharpoonright L$);
- $(tr_1, \{a\}) \in FAILS(P)$;
- $tr_2 \hat{\ } (a) \in TRACES(P)$.

It follows that

$$(tr_2 \hat{\ } (a)) \upharpoonright L \in TRACES(P \setminus H) = TRACES((P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H),$$

and also that $(tr_1 \upharpoonright L, \{a\}) \in FAILS((P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H)$ (the process $CHAOS_H$ contributes the refusal of H that $\setminus H$ requires). But since $tr_1 \upharpoonright L = tr_2 \upharpoonright L$ this contradicts the assumption that $(P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H$ is deterministic. This establishes implication one way.

(2) For the converse, suppose P satisfies $\mathcal{E}Sec_L(P)$ and $\mathcal{L}Sec_L(P)$. From the former we know that $(P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H$ is divergence-free, for the sets of traces and divergences of P are the same as those of $P \parallel\parallel [H] \parallel\parallel CHAOS_H$ and it is these that determine the divergences after hiding. If $(P \parallel\parallel [H] \parallel\parallel CHAOS_H) \setminus H$ were nondeterministic, there must therefore be a trace $t \hat{\ } (a)$ (necessarily in L^*) such that $(t, \{a\})$ is a failure. These imply the existence of traces tr and tr' of P such that

- $tr \upharpoonright L = tr' \upharpoonright L = t$;
- $(tr, \{a\}) \in FAILS(P)$;
- $tr' \hat{\ } (a) \in TRACES(P)$.

Reversing the argument that produced very similar circumstances in the first half of this proof quickly yields that $P \parallel\parallel RUN_H$ is nondeterministic, giving a contradiction. \square