

# On the Benefits of Using the Up–To Techniques for Bisimulation Verification

Daniel Hirschhoff\*

CERMICS - ENPC / INRIA, Marne la Vallée, France

**Abstract.** We advocate the use of the up–to techniques for bisimulation in the field of automatic verification. To this end, we develop a tool to perform proofs using the up to structural congruence, the up to restrictions and the up to parallel composition proof techniques for bisimulation between  $\pi$ -calculus terms. The latter technique is of particular interest because it allows one to reason on infinite state space processes. To use it in full effect, we adapt the “on the fly” bisimulation checking algorithm, leading to a form of computational completeness. The usefulness of these techniques in dealing with the expressive power of the  $\pi$ -calculus is illustrated on two non trivial examples, namely the treatment of persistent data structures and the alternating bit protocol. These examples are also good opportunities to study how well-known  $\pi$ -calculus encodings behave in the framework of automatic verification.

## Introduction

This paper studies the applicability of the so-called up–to techniques for bisimulation in the field of verification. Bisimilarity, and its correlated proof method bisimulation, have become popular notions of equivalence used to reason about concurrent systems. While having a simple and clear mathematical definition, they are far from being straightforward to handle in the framework of automatic verification, due to their richness and to the many subtle phenomena they are able to catch in the study of concurrency. The up–to techniques for bisimulation are presented in [San95]; they can simplify bisimulation proofs by reducing the size of the relations one has to consider. More precisely, an up–to technique is represented by a function  $\mathcal{F}$  from relations to relations, such that proving that processes related in a relation  $\mathcal{R}$  evolve to processes related in  $\mathcal{F}(\mathcal{R})$  is enough to show that  $\mathcal{R}$  is contained in bisimilarity. One can thus consider relations that are smaller than bisimulations, the gap being filled by the application of  $\mathcal{F}$ .

In this paper, we concentrate on bisimulation between  $\pi$ -calculus terms; the  $\pi$ -calculus [Mil91] has become a widely accepted algebra for modelling concurrency, and has demonstrated a great expressive power. It is our belief that, to use  $\pi$ -calculus as a specification language in the field of verification, one should

---

\* [dh@cermics.enpc.fr](mailto:dh@cermics.enpc.fr), CERMICS - ENPC/INRIA - 6-8 av. Blaise Pascal, Cité Descartes, Champs sur Marne, 77455 Marne la Vallée Cedex 2 - France, phone: + 33 (0) 1 64 15 35 29 - fax: + 33 (0) 1 64 15 35 86

be able to cope with the richness of this formalism, in particular by designing *specific* verification techniques for it, and not only transpose methods that are applied to less expressive formalisms. The Mobility WorkBench [VM94] is a good example of a verification tool developed specifically for the  $\pi$ -calculus (and more recently for its evolution *fusion calculus*), taking advantage in particular of the nice characterisations given by *open bisimulation* [San96]. With these ideas in mind, the up-to methods seem to be good candidates for this task: the bisimulation up to bisimilarity proof technique, and more importantly the up to restrictions and the up to parallel composition proof techniques, have been developed in the theoretical study of  $\pi$ -calculus, and typically allow one to deal with name extrusion (up to restrictions proof technique) and with replication (up to parallel composition proof technique).

Feasibility of the automation of the up-to techniques has been addressed in [Hir98], where some methods to decide if a pair of processes belongs to a relation up to structural congruence, up to restrictions, and up to parallel composition, are introduced. These methods lead to the definition of a general verification algorithm for up-to bisimulation between  $\pi$ -calculus terms, embedded in a prototype tool. In the present work, we extend the capabilities of the system, and, more importantly, try to evaluate the outcome of such an effort through the study of two non trivial examples that exploit the  $\pi$ -calculus' expressive power. With respect to [Hir98], the current version of the system is provided with a richer notion of structural congruence, and, moreover, the general bisimulation verification algorithm is modified in order to take into account the way the up to parallel composition proof technique works.

The plan of the paper is as follows: in Section 1, we introduce the formal background on  $\pi$ -calculus and the up-to techniques for bisimulation. Section 2 is devoted to the description of our implementation, and of the general bisimulation algorithm we use to handle the up to parallel composition proof technique in a computationally complete fashion. We then describe two case studies: persistent data structures, as introduced by Milner [Mil91], are examined in Section 3, while Section 4 is devoted to the study of the alternating bit protocol, in an encoding adapted from [Mam98]. We finally conclude by discussing future work.

## 1 Preliminaries

In this Section, we introduce the syntax and semantics of the language we use, a restricted though expressive subset of polyadic  $\pi$ -calculus where replication is allowed only on prefixed processes. We then define bisimulation, together with the up-to techniques for bisimulation.

**Definition 1 (Syntax)** *Given an infinite countable set of names  $\mathcal{N}$ , ranged over by  $a, b, c, \dots, p, q, \dots, x, y, \dots$ , we range over (possibly empty) name lists with  $\mathbf{a}, \mathbf{b}, \dots$ , and we define prefixes, ranged over by  $\alpha, \beta$ , as follows:*

$$\alpha \stackrel{def}{=} a(\mathbf{b}) \mid \bar{a}[\mathbf{b}];$$

Processes, ranged over with  $P, Q, R, \dots$ , are then defined as follows:

$$P \stackrel{\text{def}}{=} \mathbf{0} \mid \alpha.P \mid (\nu x)P \mid P_1 \mid P_2 \mid !\alpha.P.$$

Prefixed processes are either receptions (prefix  $a(\mathbf{b})$ ) or emissions ( $\bar{a}[\mathbf{b}]$ );  $\mathbf{0}$  is the inactive process; the restriction operator  $\nu$  makes a name private to the restricted process; parallel composition is written  $\mid$ , and  $!$  stands for replication, allowed only on prefixed processes; intuitively, a replicated process represents any number of copies of this process put in parallel.

**Conventions and notations:** In prefixes  $a(\mathbf{b})$  and  $\bar{a}[\mathbf{b}]$ ,  $a$  and  $\mathbf{b}$  are called respectively the *subject* and the *object* parts of the prefix. We shall omit the object part of a prefix when it is empty, and use a monadic notation for single name object parts (thus writing e.g.  $a(b)$  and  $\bar{a}b$ ). We shall as well omit the continuation of a prefix when it is the inactive process  $\mathbf{0}$ . Free and bound names are defined by saying that restriction and abstraction (embodied in the input prefix) are binding operators. As usual, we shall silently use  $\alpha$ -conversion to avoid capture of bound names.

**Definition 2 (Transition system)** *The operational semantics of  $\pi$ -calculus terms is defined as a labeled transition system. Actions, ranged over with  $\mu, \mu'$ , are given by the following syntax:*

$$\mu \stackrel{\text{def}}{=} a(\mathbf{b}) \mid (\nu \mathbf{b}') \bar{a}[\mathbf{b}]_{\mathbf{b}' \subseteq \mathbf{b}} \mid \tau,$$

which reads as follows: an action is either a reception, a (possibly bound) output, or the silent action  $\tau$ , denoting internal communication. Bound and free names of actions are defined as usual.

The judgment  $P \xrightarrow{\mu} P'$ , meaning that process  $P$  is liable to perform action  $\mu$  to become  $P'$ , is defined by the rules of Figure 1 (symmetrical versions of rules  $\text{PAR}_l$  and  $\text{CLOSE}_1$  are omitted;  $::$  is the constructor used to add an element to a list). Note that we adopt an “early” version of the operational semantics.

The notion of semantical equivalence on processes we are interested in is bisimilarity, defined as follows:

**Definition 3 (Bisimulation, bisimilarity)** *A relation  $\mathcal{R}$  over processes is a bisimulation iff, whenever  $P \mathcal{R} Q$  and  $P \xrightarrow{\mu} P'$ , there exists a process  $Q'$  s.t.  $Q \xrightarrow{\mu} Q'$  and  $P' \mathcal{R} Q'$ , and conversely for the transitions of  $Q$ . Bisimilarity, written  $\sim$ , is the greatest bisimulation.*

**Definition 4 (Structural congruence)** *Structural congruence, written  $\equiv$ , is the smallest equivalence relation on  $\pi$ -calculus terms that is a congruence generated by the following rules:*

$$\begin{aligned} P \mid \mathbf{0} &\equiv \mathbf{P} & P \mid Q &\equiv Q \mid P & P \mid (Q \mid R) &\equiv (P \mid Q) \mid R & (\nu x) \mathbf{0} &\equiv \mathbf{0} \\ !\alpha.P \mid \alpha.P &\equiv !\alpha.P & !\alpha.P \mid !\alpha.P &\equiv !\alpha.P & (\nu x)P \mid Q &\equiv (\nu x)(P \mid Q) & \text{if } x \notin \text{fn}(Q). \end{aligned}$$

(INP) $a(\mathbf{b}).P \xrightarrow{a(c)} P_{\mathbf{b}:=c}$	(OUT) $\bar{a}[\mathbf{b}].P \xrightarrow{\bar{a}[\mathbf{b}]} P$
(BANG) $\frac{\alpha.P \xrightarrow{\mu} P'}{!\alpha.P \xrightarrow{\mu} !\alpha.P P'}$	(RES) $\frac{P \xrightarrow{\mu} P'}{(\nu x)P \xrightarrow{\mu} (\nu x)P'} \quad x \notin n(\mu)$
(PAR <sub>l</sub> ) $\frac{P \xrightarrow{\mu} P'}{P Q \xrightarrow{\mu} P' Q} \quad bn(\mu) \cap fn(Q) = \emptyset$	
(OPEN) $\frac{P \xrightarrow{(\nu b')\bar{a}[\mathbf{b}]} P'}{(\nu t)P \xrightarrow{(\nu t::b')\bar{a}[\mathbf{b}]} P'} \quad t \neq a, t \in \mathbf{b} \setminus \mathbf{b}'$	
(CLOSE <sub>1</sub> ) $\frac{P \xrightarrow{a(\mathbf{b})} P' \quad Q \xrightarrow{(\nu b')\bar{a}[\mathbf{b}]} Q'}{P Q \xrightarrow{\tau} (\nu b')(P' Q')} \quad \mathbf{b}' \cap fn(P) = \emptyset$	

**Fig. 1.** Operational semantics

Note that the rule for processes of the form  $!\alpha.P|!\alpha.P$  is new with respect to the usual definition of structural congruence, and defines a straightforward extension of this relation. Note as well that  $\equiv$  is included in  $\sim$ .

We shall actually use a slightly enriched version of structural congruence, where we add the following law:

$$(x \text{ occurs in } \alpha \text{ in subject position}) \Rightarrow (\nu x)\alpha.P \equiv \mathbf{0}.$$

This rule is usually not included in the definition of structural congruence, because its signification is too “semantical”, i.e. it expresses a behavioural property of a process rather than a geometrical one, as structural congruence is intended to do. However, in the framework of the up to bisimilarity proof technique, it is preferable to equate as many terms as possible, hence an extended version of structural congruence is of interest; moreover, this particular law will turn out to be very useful in the example treated in Section 4. We extend the notation  $\equiv$  to our “enriched” structural congruence relation.

To introduce *up-to bisimulation* (see [San95]), we consider functions from relations to relations, ranged over by  $\mathcal{F}$ , and modify the definition of bisimulation:

**Definition 5 (Up-to bisimulation)** *Given a function  $\mathcal{F}$  from relations to relations, we say that a relation  $\mathcal{R}$  is a bisimulation up to  $\mathcal{F}$  iff the property of Definition 3 holds when we replace “ $P' \mathcal{R} Q'$ ” by “ $P' \mathcal{F}(\mathcal{R}) Q'$ ”.*

Intuitively, an up-to bisimulation relation  $\mathcal{R}$  is “smaller” than a bisimulation, the gap being filled by function  $\mathcal{F}$ , that helps building the “future” of processes related in  $\mathcal{R}$ . This can be useful, as it allows us to reduce the size of the relations we handle for the task of proving bisimilarity results. This is possible when we are dealing with a *correct* function  $\mathcal{F}$ : a function  $\mathcal{F}$  is said to be correct when ( $\mathcal{R}$  is a bisimulation up to  $\mathcal{F}$ ) implies ( $\mathcal{R} \subseteq \sim$ ). In the context of  $\pi$ -calculus, we shall use the up-to techniques given by the following Proposition:

**Proposition 6** ([San95]) *The up to structural congruence, up to restrictions and up to parallel composition proof techniques for  $\pi$ -calculus terms are defined by the following functions over relations, respectively called  $\mathcal{F}_1$ ,  $\mathcal{F}_2$  and  $\mathcal{F}_3$ :*

$$\begin{aligned}\mathcal{F}_1(\mathcal{R}) &\stackrel{\text{def}}{=} \{(P, Q); \exists P_0, Q_0 \text{ s.t. } P \equiv P_0, P_0 \mathcal{R} Q_0 \text{ and } Q_0 \equiv Q\}; \\ \mathcal{F}_2(\mathcal{R}) &\stackrel{\text{def}}{=} \{(P, Q); \exists P_0, Q_0, \mathbf{x} \text{ s.t. } P = (\nu \mathbf{x}) P_0, Q = (\nu \mathbf{x}) Q_0 \text{ and } P_0 \mathcal{R} Q_0\}; \\ \mathcal{F}_3(\mathcal{R}) &\stackrel{\text{def}}{=} \{(P, Q); \exists P_0, Q_0, T \text{ s.t. } P = P_0|T, Q = Q_0|T \text{ and } P_0 \mathcal{R} Q_0\}.\end{aligned}$$

Functions  $\mathcal{F}_1$ ,  $\mathcal{F}_2$  and  $\mathcal{F}_3$  define correct up-to proof techniques, i.e. proving that  $\mathcal{R}$  is a bisimulation up to  $\mathcal{F}$  (where  $\mathcal{F}$  is one of the aforementioned functions) is sufficient to prove that  $\mathcal{R}$  relates bisimilar processes. Moreover, these functions can be combined together, still yielding a correct technique.

Note that  $\mathcal{F}_1$  corresponds to a restriction of the general bisimulation up to bisimilarity technique (where  $\sim$  would replace  $\equiv$  in the definition of  $\mathcal{F}_1$ ).

## 2 Automatising Up-To Bisimulation

We discuss here the questions related to the automation of the up-to proof techniques for bisimulation. As stated before, the improvements with respect to [Hir98] are twofold: firstly, a narrower relation is handled in the bisimulation up to bisimilarity proof technique, and secondly, the general up-to bisimulation checking algorithm is modified in order to perform breadth-first search (and to handle *expansion* as well: see Section 4).

### 2.1 Deciding Up-To Closure

The cornerstone of our method is the ability to decide whether a pair of processes belongs to the closure of a relation up to some correct proof technique; in the sequel, we concentrate on our most powerful proof technique, namely the up to (extended) structural congruence, up to restrictions and up to parallel composition proof technique, embedded by function  $\mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_3$  (see Proposition 6 above), where  $\circ$  denotes function composition.

**Proposition 7** ([Hir98]) *Given a pair  $(P, Q)$  of processes and a relation  $\mathcal{R}$ , we can decide whether  $(P, Q) \in \mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_3(\mathcal{R})$ .*

The proof given in [Hir98] is straightforwardly adapted to handle extended structural congruence as defined above. Let us now explain how we exploit this result for up-to bisimulation verification.

### 2.2 The General Checking Method

As we front the task of bisimulation verification, an immediate though determining remark is that functions  $\mathcal{F}_1$ ,  $\mathcal{F}_2$  and  $\mathcal{F}_3$  are *syntactical* operators; this

compells us to use an “on the fly” algorithm for bisimulation checking, in order to have access to the actual term corresponding to a given state in the verification process (as opposed for example to the partition refinement algorithm [PT87], which works on *unfoldings* of processes). The “classical” on the fly checking method [FM91] can easily be adapted to the up-to methods, so as to define a semi-decision procedure for up-to bisimulation verification in the  $\pi$ -calculus, as is done in [Hir98]. However, to exploit the up-to methods in full effect, we need to modify the general bisimulation verification algorithm; this is the subject of the remainder of this Section.

Indeed, one of the crucial improvements brought by the up-to techniques is the ability, through the up to parallel composition proof technique, to handle in some cases *replicated terms*, i.e. terms possibly having an infinite behaviour. Thus, when applied, this proof method can allow one to cut infinite branches in the state space of the processes being compared. Using depth-first search, as the original algorithm of [FM91] does, it can be the case that we miss two matching branches which can be cut using the up to parallel composition technique, and instead enter an infinite loop (corresponding to an infinite growth of the state space). To avoid this, it is preferable to adopt a breadth-first strategy, in order to enter an infinite loop *only when compulsory*. Figure 2 presents the overall behaviour of Mounier’s original “on the fly” algorithm, that progressively builds a candidate bisimulation relation, by exploring the states which can be reached starting from two processes.

```

W:=  $\emptyset$ ;
(*)  $\mathcal{R}=\{(P, Q)\}$ , V:=  $\emptyset$ , R:=  $\emptyset$ , status:=true; insert (P, Q) in  $\mathbb{S}$ ;
while  $\mathbb{S}$  is not empty and status = true do
  choose a pair (P0, Q0) in  $\mathbb{S}$  and remove it from  $\mathbb{S}$ ;
  if (P0, Q0) succeeds then (add (P0, Q0) to V (or R, see below); propagate) else
  if (P0, Q0) fails then
    if (P0, Q0)  $\in$  R
      then (remove (P0, Q0) from V and insert it in W, status:=false; propagate)
      else (insert (P0, Q0) in W; propagate);
    else (* neither succeeds nor fails *)
      compute the successors of (P0, Q0) and insert them in  $\mathbb{S}$ ;
endwhile;
if P ~ Q then (if status then true else loop back to (*) ) else false

```

**Fig. 2.** General “On The Fly” Bisimulation Checking Algorithm

Let us explain informally how the verification method works (a more detailed description can be found in [FM91] and in the author’s forthcoming PhD thesis). The algorithm explores the state space given by the cartesian product of the transition systems induced by the two processes being checked. The data structures involved in the algorithm are: a structure  $\mathbb{S}$  (we shall return later on to the

nature of  $\mathbb{S}$ ), that contains the states that still have to be examined; three sets  $V$ ,  $W$ , and  $R$ , containing the pairs of processes that are respectively supposed to be bisimilar, known to be non-bisimilar, and known to be bisimilar. When two new processes turn out to be bisimilar, if the corresponding pair is in  $V$ , it is stored in  $R$ , and symmetrically for the case of non-bisimilarity (from  $V$  to  $W$ ). It can be the case, though, that during the exploration of the state space, an incorrect hypothesis is assumed regarding the bisimilarity of two processes (i.e. a pair in  $R$  is actually made of non-bisimilar processes): to handle this, a boolean flag called `status` records the “reliability” of the current computation. When a computation leads to an unreliable result, we start all over again, keeping track in  $W$  of the states that have been proved non bisimilar.

We still have to explain the behaviour of functions **succeeds**, **fails**, and **propagate**: function **fails** tests if a failure can be detected by inspecting the immediate transitions of the two processes to be examined (a failure occurs if one process performs an action that cannot be mimicked by the other term), or if the two processes are known to be non-bisimilar (set  $W$ ). Symmetrically, **succeeds** checks if the two processes are trivially bisimilar (i.e. if they are  $\alpha$ -convertible or both have no transition) or if there is an assumption stating bisimilarity between them (sets  $V$  and  $R$ ). Finally, function **propagate** propagates a newly found information along the state space, possibly discovering new (non-)bisimilarity properties, or cutting out branches that do not have to be explored anymore.

As expected, the nature of the structure  $\mathbb{S}$  determines the exploring strategy: with respect to Mounier’s algorithm, we replace the *stack* of states that have to be examined by a *queue*, thus performing breadth-first instead of depth-first search. Some extra information regarding the representation of the “vertical” structure of the state space (i.e. the relationships between states due to the transition relation), that comes for free in the case of a stack structure, has to be provided within the objects that are stored in  $\mathbb{S}^1$ . Using breadth-first search, we get a *computationally complete* checking algorithm:

**Theorem 8** *We say that a bisimulation checking algorithm is computationally complete with respect to a given up-to technique  $\mathcal{F}$  whenever it diverges if and only if no finite bisimulation up to  $\mathcal{F}$ , relating the two processes to be checked and derivatives of these processes, exists.*

*With this terminology, the breadth-first version of the algorithm of Figure 2 is computationally complete with respect to the function  $\mathcal{F}_1 \circ \mathcal{F}_2 \circ \mathcal{F}_3$  (see Proposition 6).*

The proof of this result follows immediately from the very definition of the breadth-first search. Note that the hypothesis about the relation containing only *derivatives* of the processes we examine is determinant for the proof. What is really important here is the form of completeness we get: while for finite-state processes, depth-first and breadth-first searches differ only “strategically”, in

<sup>1</sup> This structural information is even more intricate in the weak case, where some care has to be provided to avoid  $\tau$ -loops before an action actually fires (i.e. in the “first  $\Rightarrow$  part” of  $\Rightarrow \xrightarrow{\mu} \Rightarrow$ : see Section 4).

the case of replicated terms, the breadth-first version, in conjunction with the up to parallel composition proof technique, provides a real gain in expressiveness, as expressed by Theorem 8. This improvement of the system, together with the additional structural congruence law seen above, will be determinant for the treatment of the examples in Sections 3 and 4.

### 2.3 A Tool for Up-To Bisimulation Verification

The methods presented so far are implemented in a tool, written in O’Caml, and running under Unix. The system allows the user to define a pair of processes and to check whether they are bisimilar by choosing between three up-to proof techniques and between strong bisimulation and *expansion* (see below for the definition of this notion). Facilities for interactively simulating the behaviour of a process and for debugging (in the case of a check failure) are also provided. Internally, the system works systematically up to structural congruence, i.e. with *normalised* terms, as defined in [Hir98]. The tool is available at <http://cermics.enpc.fr/~dh/pi>.

## 3 First Case Study: Persistent Values and Sharing

One of the applications of the up to parallel composition proof technique is the proof of the so-called *replication theorems*, which express some properties of processes that model resources. Intuitively, a term of the form  $(\nu a) (!a(\mathbf{b}).P \mid Q)$  is viewed as an agent  $Q$  having access to a private resource  $P$ , located at channel  $a$ , with the possibility of instantiating the resource with some parameters  $\mathbf{b}$ . The resource is usually replicated, because it is meant to be “always available”. Resource processes of the form  $!a(\mathbf{b}).P$  are ubiquitous in the study of  $\pi$ -calculus: they can be found in particular in the encodings of the  $\lambda$ -calculus, where application is represented as (the translation of) the function having access to its argument as a resource. Resource processes arise also in the study of higher-order  $\pi$ -calculus, of object calculi, and of data structures. We focus here on the latter subject, by studying Milner’s encoding of lists in the  $\pi$ -calculus.

In [Mil91], lists are represented using two kinds of  $\pi$ -calculus terms, corresponding to each constructor: a process of the form  $l(c, n).\bar{c}[v; l']$  represents a *Cons* node, situated at  $l$ , and containing a value that can be accessed at channel  $v$  and a reference to the remainder of the list situated at  $l'$ . The term  $(\nu l', v) (l(c, n).\bar{c}[v; l'] \mid v.V \mid l'.L)$  is thus seen as the list  $Cons(V, L)$ , where  $V$  is the value located at  $v$  and  $L$  is the tail of the list, located at  $l'$ . A process of the form  $l(c, n).\bar{\pi}$  represents the empty list *Nil* situated at  $l$ . Lists can be interrogated by sending two names at their location, one for each possible constructor.

Such data structures, however, are *linear*: reading them destroys them; to make them persistent, one uses the replication operator. As said in [Mil91], there are *a priori* two ways to achieve this, whether one chooses to replicate the nodes and the value cells in the lists, or directly the subcomponents. Accordingly, list

$Cons(V, L)$  seen above would become

$$(\nu l', v) (!l(c, n).\bar{c}[v; l'] \mid !v.V \mid \dots) \quad \text{or} \quad !l(c, n).(\nu l', v) (\bar{c}[v; l'] \mid !v.V \mid \dots)$$

respectively (the dots “...” stand here for the continuation of the lists, whose shape depends on the nature of the topmost constructor of  $L$ ,  $Cons$  or  $Nil$ : what is relevant here is the expression of a  $Cons$  node in a persistent list). As we shall see, this design choice gets reflected both on the behaviour of the corresponding data structure and on questions related to bisimulation checking.

To illustrate this, we shall work in the following with a very simple list, consisting only in the element  $\mathbf{true}^2$ ; its representation as an ephemeral data structure, given by  $R_0 \stackrel{def}{=} (\nu b, l_1) (l_0(c.n)\bar{c}[b; l_1] \mid !b\mathbf{true} \mid l_1(c, n).\bar{\pi})$ , is simplified into  $R \stackrel{def}{=} (\nu b) (l_0(c).\bar{c}b \mid !b\mathbf{true})$ . This process can be viewed as a single constant cell holding the value  $\mathbf{true}$ , and is only reminiscent of the list  $[\mathbf{true}]$ ; however, it will be sufficient for our task, and the reasoning made below holds for “real” lists as well. As said above, there are two ways to transform  $R$  into a persistent structure, represented by the following two processes:

$$R_1 \stackrel{def}{=} (\nu b) (!l_0(c).\bar{c}b \mid !b\mathbf{true}) \quad \text{and} \quad R_2 \stackrel{def}{=} !l_0(c).(\nu b) (\bar{c}b \mid !b\mathbf{true}).$$

As stated in [Mil91], the difference between  $R_1$  and  $R_2$  is reflected on *sharing properties* of our data structures. The problem of value sharing can be expressed in terms of security: suppose we have an agent  $A$ , willing to interrogate the resource  $R$  located at  $l_0$  any number of times, and to send a signal on either one of two channels  $o$  and  $p$ , depending on the value read at  $l_0$ . This agent can be represented by the process

$$A \stackrel{def}{=} !\bar{l}_0 c_1.c_1(b).(\nu t, f) (\bar{b}[t, f] \mid t.\bar{o} \mid f.\bar{p}),$$

where  $c_1$  is a name that is specific to  $A$  for its communications with  $R$ . Suppose now that an evil agent  $E$  is willing to interfere with the communications of  $A$ , by interrogating  $R$  and trying to send the value  $\mathbf{false}$  as if this was the value of the cell located at  $l_0$ . To do this,  $E$  needs to interrogate once  $R$ , to get name  $b$  where the value is located, and then “mimick”  $\mathbf{false}$  along  $b$ . Hence:

$$E \stackrel{def}{=} \bar{l}_0 c_2.c_2(b).b(t, f)\bar{f}.$$

Our whole system is made of the parallel composition of the resource (being either  $R_1$  or  $R_2$ ), “innocent” agent  $A$ , and “evil” agent  $E$ :

$$\text{Sys}_i \stackrel{def}{=} (\nu l_0, c_1, c_2) (R_i \mid A \mid E) \quad i = 1, 2.$$

<sup>2</sup> Along the lines of the encoding of booleans in the  $\lambda$ -calculus, the value  $\mathbf{true}$  located at  $b$ , written  $b\mathbf{true}$ , is a process waiting for two names  $t$  and  $f$ , and returning a signal on  $t$  ( $\mathbf{false}$  would return a signal on  $f$ ): we write  $b\mathbf{true} \stackrel{def}{=} b(t, f).\bar{t}$  (see [Mil91]).

We verify that, depending on which representation we take for the persistent cell  $R$ ,  $E$  can either disturb the communications performed by  $A$ , or be innocuous. To this end, we compare the behaviour of the term  $\text{Sys}_i$  for  $i = 1, 2$  with the behaviour of the process where each agent  $A$  and  $E$  have their own private copy of the resource (and hence cannot interfere):

$$\text{Sys}'_i \stackrel{\text{def}}{=} (\nu l_0, c_1)(R_i \mid A) \mid (\nu l_0, c_2)(R_i \mid E) \quad i = 1, 2.$$

Indeed, if we run the tool with these definitions, we get the following results:

- **case  $\text{Sys}_1$ :** in this case, a location  $b$  for the value `true` is created once for all and transmitted to every process that interrogates the list at  $l_0$ . This way, even in the case where  $A$  performs only one request at  $l_0$ , agent  $E$  can send the wrong value `false` to  $A$ , and  $\text{Sys}_1$  and  $\text{Sys}'_1$  are not bisimilar. We check this on a simple session of our tool; we first define the two components of the pair of processes to be checked:

```
> Left  (~l0,c1,c2)((~b)(!l0(c).c[b]|b(t;f).t[]))
| l0[c1].c1(b).b(t;f).f[] | l0[c2].c2(b).(~t,f)(b[t;f]|t.o[]|f.p[])
> Right (~l0,c1)((~b)(!l0(c).c[b]|b(t;f).t[]))
| l0[c1].c1(b).b(t;f).f[] | (~l0,c2)((~b)(!l0(c).c[b]|b(t;f).t[]))
| l0[c2].c2(b).(~t,f)(b[t;f]|t.o[]|f.p[])
```

The syntax for processes is quite natural:  $\sim$  stands for abstraction,  $[]$  for output and  $()$  for input. We now check for bisimilarity:

```
> Check
The processes are not bisimilar.
```

- **case  $\text{Sys}_2$ :** in this case, a new name  $b$  is created each time the list is interrogated at  $l_0$ , and hence processes  $A$  and  $E$  cannot interfere. Even in the case where an infinite number of processes interrogate the list, we can prove that  $\text{Sys}_2$  and  $\text{Sys}'_2$  are bisimilar:

```
> Left  (~l0)(((!l0(c).(~b)c[b].b(t;f).t[])) | l0[c1].c1(b).b(t;f).f[]
| !l0[c2].c2(b).(~t)(~f)(b[t;f]|t.o[]|f.p[]))
> Right (~l0)(((!l0(c).(~b)c[b].b(t;f).t[]))
| l0[c1].c1(b).b(t;f).f[] | (~l0)(((!l0(c).(~b)c[b].b(t;f).t[]))
| !l0[c2].c2(b).(~t)(~f)(b[t;f]|t.o[]|f.p[]))
> Check
The processes are bisimilar.
```

When the tool performs this verification, it uses the up to parallel composition proof technique to erase copies of the cell each time they appear in both terms: an infinite proof is thus replaced by a finite one.

The latter example exploits the up to parallel composition proof technique to handle the representation of data structures in the  $\pi$ -calculus. In this context, our tool has been useful in dealing with non trivial properties of processes, and it

should be stressed that other existing systems cannot handle replicated processes like those we manipulate here.

Furthermore, it should be remarked that security properties like the one we have considered can be addressed using type systems for the  $\pi$ -calculus. As exposed for example in [San97], an easy way to prevent  $E$  from “pretending” to be  $R$  by locating the value `false` at  $b$  is to forbid processes that interact with  $R$  to use  $b$  in input subject position (as  $E$  does).

## 4 Second Case Study: the Alternating Bit Protocol

We now focus on another example: the well-known *alternating bit protocol*, which is probably the most widely used benchmark for verification systems in the field of concurrency. Our purpose here is not to study it *per se*, but instead to use it to shed light on issues related to the behaviour of encodings in the  $\pi$ -calculus. To this end, we shall need *expansion* between processes:

**Definition 9 (Expansion)** We write  $\Rightarrow$  for the reflexive transitive closure of  $\xrightarrow{\tau}$ ,  $\xrightarrow{\mu}$  for  $\Rightarrow \xrightarrow{\mu} \Rightarrow$ , and  $\xrightarrow{\hat{\mu}}$  for  $\xrightarrow{\mu}$  if  $\mu \neq \tau$ , = or  $\xrightarrow{\tau}$  otherwise.

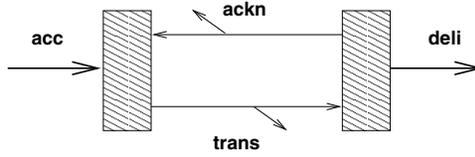
We say that a relation  $\mathcal{R}$  is an expansion iff, whenever  $PRQ$ ,  $P \xrightarrow{\mu} P'$  implies that there exists a  $Q'$  s.t.  $Q \xrightarrow{\hat{\mu}} Q'$  and  $P'\mathcal{R}Q'$ , and  $Q \xrightarrow{\mu} Q'$  implies that there exists  $P'$  s.t.  $P \xrightarrow{\hat{\mu}} P'$  and  $P'\mathcal{R}Q'$ . In this case, we write  $P \lesssim Q$  for  $PRQ$ .

$\lesssim$  is a preorder on processes that is more realistic than (strong) bisimulation for reasoning about larger case-studies such as protocols. It allows one to prove that a given term “respects” a behaviour, modulo some extra  $\tau$  steps (given by transitions  $\xrightarrow{\hat{\mu}}$ ): typically, the duality *specification vs. implementation* can be expressed using expansion. Our treatment of the up-to techniques is straightforwardly adapted to handle expansion; the user of our tool can choose between a *strong* and a *weak* mode, corresponding to relations  $\sim$  and  $\lesssim$  respectively.

We first informally introduce the protocol, then present an encoding adapted from [Mam98], and finally discuss the encodings we use in order to automatically perform the correctness proof of the protocol. Our purpose in the latter part is to show that the question of defining encodings in the framework of verification differs from the usual problem (i.e. the study of *expressiveness*), as exemplified on our case study.

### 4.1 The Protocol

Figure 3 presents the entities involved in the communication protocol: a first agent receives a message (whose content is not taken into account here) on channel `acc`. It then transmits this message on a channel called `trans` to a second agent. But channel `trans` is unreliable, and some messages transiting on it may get lost: to handle that, a boolean tag is associated to the message being sent on this channel, and the first agent keeps sending this information on `trans`. By reception, the second agent transmits the message on `deli` and is willing to send



**Fig. 3.** The Alternating Bit Protocol

back the boolean to the first agent on *ackn*, as an acknowledgment. Channel *ackn* being lossy as well, the second agent has to do this repeatedly. When the acknowledgment eventually arrives, a new cycle can start with the negated boolean, so that the first agent can tell when the message has actually been received.

The  $\pi$ -calculus processes<sup>3</sup> implementing this protocol are defined on Figure 4 (again, our encoding is adapted from [Mam98]): the first agent’s behaviour is given by processes *Send* and *Wait*, while the second agent corresponds to process *Receive*; message losses on *trans* and *ackn* are represented by process *Noise*, that sends a signal on a channel *loss* whenever a signal gets lost (one could see these synchronisations on *loss* as a modelisation of a timer mechanism to detect losses). Note that the agents are parametrised upon two booleans: this will be discussed below. Finally, process *Specif* is an ideal specification of the protocol’s behaviour: it receives a signal on *acc*, sends one on *deli*, and starts again by sending a signal on the trigger channel *c*.

$$\begin{aligned}
 \text{Send} &\stackrel{def}{=} !\text{send}(b_1, b_2).\text{acc}.\overline{\text{trans}}[b_2; b_1].\overline{\text{wait}}[b_2; b_1] \\
 \text{Wait} &\stackrel{def}{=} !\text{wait}(b_1, b_2).(\text{ackn}(b'_1, b'_2).\overline{\text{send}}[b'_1; b'_2] + \text{loss}.\overline{\text{trans}}[b_1; b_2].\overline{\text{wait}}[b_1; b_2]) \\
 \text{Receive} &\stackrel{def}{=} !\text{rec}(b_1, b_2).\text{trans}(b'_1, b'_2).\text{if } b_1 = b'_1 \\
 &\quad \underline{\text{then}} \overline{\text{ackn}}[b_1; b_2] \mid \overline{\text{rec}}[b_1; b_2] \quad \underline{\text{else}} \overline{\text{deli}}.(\overline{\text{ackn}}[b_2; b_1] \mid \overline{\text{rec}}[b_2; b_1]) \\
 \text{Noise} &\stackrel{def}{=} !\text{noise}.( \text{trans}(b, b').\overline{\text{loss}}.\overline{\text{noise}} + \text{ackn}(b, b').\overline{\text{loss}}.\overline{\text{noise}} ) \\
 \text{System} &\stackrel{def}{=} (\nu \text{ trans, ackn, send, wait, rec, loss, noise, } b_1, b_2) \\
 &\quad ( \text{Send} \mid \text{Wait} \mid \text{Receive} \mid \text{Noise} \mid !b_1 \text{ true} \mid !b_2 \text{ false} \mid \overline{\text{send}}[b_1; b_2] \mid \overline{\text{rec}}[b_1; b_2] \mid \overline{\text{noise}} ) \\
 \text{Specif} &\stackrel{def}{=} (\nu c) (!c.\overline{\text{acc}}.\overline{\text{deli}}.\overline{c} \mid \overline{c})
 \end{aligned}$$

**Fig. 4.** Modelling the Alternating Bit Protocol

<sup>3</sup> The language we use here is actually value-passing CCS; for the encodings discussed below, however, we shall need the full power of  $\pi$ -calculus.

## 4.2 The Encodings – Discussion

**Choice** – To model the alternating bit protocol in our system, we have to deal with two constructs that are not present in our language, namely operators on booleans and the choice operator. Well-known encodings of these constructs exist (see [Mil91,Nes97]); these encodings belong to theoretical studies, and have been mechanised e.g. in the programming language PICT [PT97]. We shall see here the problems that arise as we try to use them in the context of verification, and how they can be treated.

Various encodings of guarded choice are investigated in [NP96,Nes97]; the general idea is to represent each branch of the choice by a parallel component, and to implement a lock mechanism to prevent other branches to fire when one branch has committed. In general, once a choice has been taken, the “dead” branches remain, and, depending on the encoding we choose, are “more or less alive”; when it comes to automatic verification, we need to be able to erase these branches *syntactically*, to prevent our relation from growing. We have chosen to encode the binary choice of Figure 4 as follows ( $\llbracket \cdot \rrbracket$  represents the encoding function):

$$\llbracket a(\mathbf{x}).P + b(\mathbf{y}).Q \rrbracket \stackrel{def}{=} (\nu l) (\bar{l} \mid a(\mathbf{x}).l.\bar{b}(\mathbf{x}).P \mid b(\mathbf{y}).l.\bar{a}(\mathbf{y}).Q).$$

This encoding represents an adaptation of those described in [NP96]; here, a lock channel  $l$  is created and when one branch is chosen, it deactivates the other branch by consuming its head prefix, thus leading to a subterm of the form  $(\nu l)(l.T)$ , that can immediately be garbage collected using the additional law for structural congruence (see Section 1). Our simplified encoding is correct due to an important property of the terms we consider: each time a choice construct is encountered in a run of the protocol (agents *Wait* and *Noise*), at most one process can interact with the branches involved in the choice (that is, we cannot have simultaneous emissions on *ackn* and *loss*, nor on *ackn* and *trans*). Thus, it cannot be the case that both branches commit before the lock  $l$  is consumed, and the untaken branch can be safely deactivated before the chosen branch proceeds.

Furthermore, note as well that the simplest encoding of guarded choice, defined as follows

$$\llbracket a(\mathbf{x}).P + b(\mathbf{y}).Q \rrbracket \stackrel{def}{=} (\nu l) (\bar{l} \mid l.a(\mathbf{x}).P \mid l.b(\mathbf{y}).Q),$$

and called *internal choice* in [NP96], cannot be adopted here, as the choice of the committing branch depends on the context (since at most one branch can commit), and cannot be done internally *before* synchronisation.

The original encoding of [NP96] could not be chosen, since it allows the computation to proceed before the dead branches have been deactivated.

**Booleans** – We have already seen how booleans are encoded in the  $\pi$ -calculus. The encoding presented in [Mam98] uses a single boolean as parameter for the various agents, and the negation operator for the recursive calls in *Send* and

*Receive*. We could have encoded the latter operator as

$$\llbracket \neg b \rrbracket \stackrel{def}{=} !b'(t, f)(\nu t', f') (\bar{b}[t'; f'] \mid t'.\bar{f} \mid f'.\bar{t}); \quad (*)$$

this encoding is not convenient for the purpose of verification, because each time the protocol loops, it keeps adding new terms, induced by the negation operator to the system, making the relation grow *ad infinitum* (intuitively,  $\neg\neg b$  does not reduce to  $b$ ). We therefore use a trick, that consists in parametrising all the agents upon *two* booleans, that are exchanged as we want to make a recursive call with the negated bit (see Figure 4).

We also have to manage the testing construct in agent *Receive*: here again, we take advantage of the additional law for structural congruence to get rid of inactive terms once the test has taken place. We represent the boolean operator “=” (that can be seen as the negation of the XOR operator) by:

$$\llbracket b = b' \rrbracket \stackrel{def}{=} b_0(T, F).(\nu t, f, t', f')(\bar{b}[t; f] \mid \bar{b}'[t'; f'] \mid t.(t'.\bar{T}|f'.\bar{F}) \mid f.(t'.\bar{F}|f'.\bar{T})).$$

Once the test has been performed, non-chosen branches are automatically garbage collected in the normalisation process.

Having defined the encoding of the protocol into our simple language, we provide our tool with the corresponding definitions, and the correctness proof of the property  $Specif \lesssim System$  is performed automatically.

## Conclusion

We have seen how the up-to techniques for bisimulation can be adapted for the purpose of verification. The examples of persistent lists and of the proof of the alternating bit protocol suggest that the task of dealing with encodings in the  $\pi$ -calculus can be managed using these techniques. It is questionable, though, whether one should adopt a standpoint analogous to the design choices of PICT [PT97], where a full-size programming language is built by adding successive layers to a simple core language through encodings of higher-level constructs. Indeed, developing an implementation in a system built this way adds transitions to the terms and calls for clever use of garbage collection (as seen on the examples); such issues are critical in the field of verification. Nevertheless, the experiments shown above are interesting both from a theoretical and from a practical point of view.

From a theoretical point of view, we have seen how the application of theoreticians' techniques (as the up-to methods originally are) for the purpose of verification gives extra insight on results such as the encodings of language constructs or of data. A striking example is the straightforward definition of the negation operator for booleans (see equation marked with (\*)), which is catastrophic in terms of state space growth, as stated above.

From a practical point of view, this work seems encouraging for the use of the up-to techniques in the field of verification. Of course, constructs like booleans

or the choice operator should probably belong to a “real size” system. However, the task of specification in a verification tool usually involves some formalisations that, on a larger scale, are quite akin to the encodings we have seen, at least from a methodological point of view. In this approach, the up-to techniques have been shown to be helpful in terms both of efficiency and of expressiveness.

Regarding future work, it could be interesting to study whether the up-to techniques can be integrated in a tool like the Mobility WorkBench [VM94], that is a much richer system than our prototype, and that moreover supports various supplementary facilities. On its own, however, our tool can be enriched to perform some more experiments concerning the verification of  $\pi$ -calculus bisimilarity results; larger case studies, for example, could be considered, or other encodings. Work is also in progress to adapt the up-to bisimulation verification method to *open terms*, i.e. terms possibly containing process variables, in order to be able to prove not only bisimilarity results, but also general *bisimilarity laws*. Finally, the issue of the usefulness of the up-to techniques within other models of concurrency should also be investigated.

*Acknowledgments* We thank the anonymous referees for useful remarks about the presentation of this work, as well as Bruno Sudret for inspiring discussions.

## References

- FM91. J.-C. Fernandez and L. Mounier. “On the fly” verification of behavioural equivalences and preorders. In *Proceedings of CAV’91*, LNCS, 1991.
- Hir98. D. Hirschhoff. Automatically Proving Up To Bisimulation. In *Proceedings of MFCS’98 Workshop on Concurrency*, volume 18 of *ENTCS*, 1998.
- Mam98. B. Mammas. Une méthodologie de preuves orientée contraintes et basée sur les systèmes de transitions modales. Technical report, LIP6, UPMC, 1998.
- Mil91. R. Milner. The polyadic  $\pi$ -calculus: a tutorial. Technical Report ECS-LFCS-91-180, LFCS, October 1991.
- Nes97. U. Nestmann. What is a ‘good’ encoding of guarded choice? In *Proceedings of EXPRESS’97*, volume 7 of *ENTCS*, 1997.
- NP96. U. Nestmann and B. C. Pierce. Decoding choice encodings. In *Proceedings of CONCUR ’96*, number 1119. LNCS, Springer Verlag, August 1996.
- PT87. R. Paige and R. E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
- PT97. B. C. Pierce and D. N. Turner. Pict: A Programming Language Based on the Pi-Calculus. Technical Report CSCI 476, Computer Science Department, Indiana University, 1997.
- San95. D. Sangiorgi. On the bisimulation proof method. In *Proceedings of MFCS ’95*, volume 969 of *LNCS*, 1995.
- San96. D. Sangiorgi. A theory of bisimulation for the  $\pi$ -calculus. *Acta Informatica*, 33:69–97, 1996.
- San97. D. Sangiorgi. The name discipline of uniform receptiveness. In *Proceedings of ICALP ’97*, volume 1256 of *LNCS*, pages 303–313, 1997.
- VM94. B. Victor and F. Moller. The Mobility Workbench — a tool for the  $\pi$ -calculus. In D. Dill, editor, *Proceedings of CAV’94*, volume 818 of *LNCS*, pages 428–440. Springer-Verlag, 1994.

