



Basic Research in Computer Science

BRICS RS-95-53

Klarlund et al.: Automated Logical Verification based on Trace Abstractions

# Automated Logical Verification based on Trace Abstractions

Nils Klarlund  
Mogens Nielsen  
Kim Sunesen

BRICS Report Series

RS-95-53

ISSN 0909-0878

November 1995

**Copyright © 1995, BRICS, Department of Computer Science  
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work  
is permitted for educational or research use  
on condition that this copyright notice is  
included in any copy.**

**See back inner page for a list of recent publications in the BRICS  
Report Series. Copies may be obtained by contacting:**

**BRICS  
Department of Computer Science  
University of Aarhus  
Ny Munkegade, building 540  
DK - 8000 Aarhus C  
Denmark  
Telephone: +45 8942 3360  
Telefax: +45 8942 3255  
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through WWW and  
anonymous FTP:**

**<http://www.brics.dk/>  
[ftp ftp.brics.dk \(cd pub/BRICS\)](ftp://ftp.brics.dk/cd/pub/BRICS)**

# Automated logical verification based on trace abstractions

Nils Klarlund Mogens Nielsen Kim Sunesen

**BRICS\***

Department of Computer Science

University of Aarhus

Ny Munkegade

DK-8000 Aarhus C.

{klarlund,mnielsen,ksunesen}@daimi.aau.dk

## Abstract

We propose a new and practical framework for integrating the behavioral reasoning about distributed systems with model-checking methods.

Our proof methods are based on *trace abstractions*, which relate the behaviors of the program and the specification. We show that for finite-state systems such symbolic abstractions can be specified conveniently in Monadic Second-Order Logic (M2L). Model-checking is then made possible by the reduction of non-determinism implied by the trace abstraction.

Our method has been applied to a recent verification problem by Broy and Lamport. We have transcribed their behavioral description of a distributed program into temporal logic and verified it against another distributed system without constructing the global program state space. The reasoning is expressed entirely within M2L and is carried out by a decision procedure. Thus M2L is a practical vehicle for handling complex temporal logic specifications, where formulas decided by a push of a button are as long as 10-15 pages.

## 1 Introduction

This paper is concerned with the specification and verification of distributed systems. Often, the relationship between a program and a specification is expressed in terms of a state-based refinement mapping, see [18] for a survey. Thus, when systems are specified by behavioral or temporal constraints, it is necessary first to find state-representations. In this process, important information may be lost or misconstrued.

---

\*Basic Research in Computer Science,  
Centre of the Danish National Research Foundation.

In this paper, we exhibit a logic of *traces* (i.e. finite computation sequences) that allows compositional reasoning directly about behaviors. We formulate *trace abstractions* and their proof rules as an alternative to the use of refinement mappings for the verification of distributed systems.

Our main goal is to show that our method is useful in practice for complicated examples. Thus, use of our logic and proof rules must be supported by a decision procedure that will give answers to logical questions about the systems, such as “Does trace abstraction  $R$  show that program  $P$  implements  $S$ ?”

To this end, we formulate a sound and complete verification method based on trace abstractions. We show that our method for finite-state systems can be formulated in a very succinct formalism: the *Monadic Second-order Logic* (M2L).

We address the important problem of relating a distributed program to a non-deterministic specification that also is a distributed system. Non-determinism arises when systems have alphabets that are partitioned into observable and internal actions. Abstracting away internal actions generally introduces non-determinism.

Our contribution is to show an alternative to usual techniques, which tend to involve rather involved concepts such as *prophecy variables* or mappings to sets of sets of states. These can be replaced by behavioral predicates that need only to partially link the program and the specification. The remaining information is then calculated automata-theoretically by means of the subset construction.

We formulate a compositional rule to avoid the explicit construction of the global program space.

Using the Mona implementation of M2L, we have verified a recent verification problem by Broy and Lamport by transcribing several pages of informally stated temporal properties. The formulas resulting are decided in minutes despite their size ( $10^5$  characters). A detailed treatment of this problem can be found in [11].

## 1.1 Relations to previous work

The systems we define are closely related to those described by Hoare in [9], where an alphabet  $\Sigma$  and a set of traces over  $\Sigma$  is associated with every process. We use a composition operator, similar to Hoare’s parallel operator ( $\parallel$  [9]) forcing systems to synchronize on events or actions shared by both alphabets.

We do not know of any earlier work using relations directly on traces. In fact, the use of state mappings—one of the most advocated methods for proving refinement, see e.g. [17, 15, 16, 10] and for a survey [18]—were introduced as a way to avoid behavioral reasoning, often regarded as being too complex. The theory of state mappings is by now well-understood, but not simple, with the completeness results in [1, 12, 20]. In the finite state case, an important difference between the state mapping approach and ours is that in the traditional approaches, the mapping is to be exactly specified state by state whereas in our approach the relation between behaviors may be specified partially leaving the rest to our verification tool. In [2], Lamport and Abadi gave a proof rule

for proving correctness of implementation of compound systems based on an assumption/guarantee method. A closed compound system is split into a number of open systems by factoring out dependencies as assumptions. Our rule is very different in that dependencies are reflected in a requirement about the relationship between trace abstractions for components.

The TLA formalism by Lamport [16] and the temporal logic of Manna and Pnueli [10] offer unified frameworks for specifying systems and state mappings, and for proving the correctness of implementation. Both logics are undecidable, but work has been done on establishing mechanical support, see [6, 7].

Clarke, Browne, and Kurshan [4] have applied model checking techniques to the language containment problem ( $\mathcal{L}(M_1) \subseteq \mathcal{L}(M_2)$ ), where  $M_1$  and  $M_2$  are  $\omega$ -automata. They reduce the containment problem to a model-checking problem by forming a product of the automata and checking whether the product is a model for a certain *CTL\** formula. The method is applicable to any common kind of  $\omega$ -automata. Thus it deals with liveness properties unlike our method, which only deals with logic over finite prefixes. However, the method in [4] suffers from the restriction that  $M_2$  be deterministic.

Kurshan, see [13], has devised an automata-theoretic framework for modeling and verifying synchronous transition systems. His use of homomorphisms allow complex properties to be reduced to ones that can be verified by means of model-checking.

Kurshan's methods were extended in [14] to the asynchronous *input/output automata* of [17]. There, Kurshan et al. give an account of interleaving composition in terms of conventional, synchronous automata. Our treatment of concurrency is similar in its use of stuttering for modeling asynchrony except that we do not consider fairness (which is a property of infinite sequences). A principal difference is that our proposal is based on comparing sequences of events, whereas the method of [14] is essentially state-based or event-based.

Binary Decision Diagrams (BDDs) are usually used in verification to compactify representations of state-spaces, see e.g. [5]. The *Mona* implementation [8] of a decision procedure for M2L uses BDDs to handle large alphabets.

## 1.2 Overview

In Section 2, we discuss our formal framework, which is based on an interleaving semantics of processes that work in a global space of events. M2L is explained in Section 3. We show in Section 4 that with some additional concepts, it is possible to formulate the verification method of Section 2 in M2L. In Section 5, we explain the role of trace abstractions in our solution of the Broy and Lamport verification problem.

## 2 Traces and abstractions

We regard systems in a fairly standard way: they are devices that produce sequences of events that are either observable or internal. Systems exist in

a universe. They can be composed and compared. Trace abstractions relate a program to a specification. These abstractions form a sound and complete verification method, and a simple decomposition rule is easy to formulate.

## 2.1 Systems, universes and normalization

A system  $A$  determines an alphabet  $\Sigma_A$  of *events*, which is partitioned into *observable events*  $\Sigma_A^{Obs}$  and *internal events*  $\Sigma_A^{Int}$ . A *behavior* of  $A$  is a finite sequence over  $\Sigma_A$ . The system  $A$  also determines a prefix-closed language  $L_A$  of behaviors called *traces* of  $A$ . We write  $A = (L_A, \Sigma_A^{Obs}, \Sigma_A^{Int})$ . The *projection*  $\pi$  from a set  $\Sigma^*$  to a set  $\Sigma'^*$  ( $\Sigma' \subseteq \Sigma$ ) is the unique string homomorphism from  $\Sigma^*$  to  $\Sigma'^*$  given by  $\pi(a) = a$ , if  $a$  is in  $\Sigma'$ , and  $\pi(a) = \epsilon$  otherwise, where  $\epsilon$  is the empty string. The *observable behaviors* of a system  $A$ ,  $Obs(A)$ , are the projections on  $\Sigma_A^{Obs}$  of the traces of  $A$ , that is  $Obs(A) = \{\pi(\alpha) \mid \alpha \in L_A\}$ , where  $\pi$  is the projection from  $\Sigma_A^*$  onto  $(\Sigma_A^{Obs})^*$ .

A system  $A$  is thought of as existing in a *universe* which contains the systems with which it is composed and compared. The events possible in this universe constitute a global alphabet  $\mathcal{U}$ , which contains  $\Sigma_A$  and all other alphabets of interest. Moreover,  $\mathcal{U}$  is assumed to contain the distinguished event  $\tau$ , which is not in the alphabet of any system. The set  $N_\Sigma(A)$  of *normalized traces* over an alphabet  $\Sigma \supseteq \Sigma_A$  is the set  $h^{-1}(L_A)$ , where  $h$  is the projection from  $\Sigma^*$  onto  $\Sigma_A^*$ . Normalization plays an essential rôle when composing systems and when proving correctness of implementation of systems with internal events.

## 2.2 Composition

We say that systems  $A$  and  $B$  are *composable* if they do not disagree on the partition of events, that is, if no internal event of  $A$  is an observable event of  $B$  and vice versa, or symbolically, if  $\Sigma_A^{Int} \cap \Sigma_B^{Obs} = \emptyset$  and  $\Sigma_B^{Int} \cap \Sigma_A^{Obs} = \emptyset$ . Given composable systems  $A$  and  $B$ , we define their *composition*  $A \parallel B = (L_{A \parallel B}, \Sigma_{A \parallel B}^{Obs}, \Sigma_{A \parallel B}^{Int})$ , where

- the set of observable events is the union of the sets of observable events of the components, that is,  $\Sigma_{A \parallel B}^{Obs} = \Sigma_A^{Obs} \cup \Sigma_B^{Obs}$ ,
- the set of internal events is the union of the sets of internal events of the components, that is,  $\Sigma_{A \parallel B}^{Int} = \Sigma_A^{Int} \cup \Sigma_B^{Int}$ , and
- the set of traces is the intersection of the sets of normalized traces with respect to the alphabet  $\Sigma_{A \parallel B}$ , i.e.  $L_{A \parallel B} = N_{\Sigma_{A \parallel B}}(A) \cap N_{\Sigma_{A \parallel B}}(B)$ .

(Note that the restriction above for composability ensures that  $A \parallel B$  has also disjoint observable and internal events.)

A trace of  $A \parallel B$  is the interleaving of a trace of  $A$  with a trace of  $B$  in which common events are synchronized. The projection of a trace of  $A \parallel B$  onto the alphabet of any of the components is a trace of the component. Composition is commutative, idempotent, and associative, and extends straightforwardly to any number  $n$  of composable systems  $A_i$ . We write  $A_1 \parallel \dots \parallel A_n$  or just  $\parallel A_i$ .

**Example 2.1** To make the concepts clearer, we show how to present the well-known scheduler [19] of Milner in terms of our systems. The distributed scheduler is based on passing a token consecutively between a number of computing agents. We consider a three-agent version of the scheduler. The  $i$ th agent  $S_i$  performs observable events  $a_i$  and  $b_i$  to indicate the beginning and the end of computing, respectively, and it synchronizes with its neighbor agents by interacting on the internal events  $c_i$  and  $c_{i\ominus 1}$ , where  $i$  is 0, 1, or 2, and  $\ominus$  is subtraction modulo 3.

For a regular expression  $r$ , we denote by  $\mathcal{L}^{Pre}(r)$  the regular language obtained by taking the prefix-closure of the language associated with  $r$ . Thus the agents may be described by:

$$\begin{aligned} S_0 &= (\mathcal{L}^{Pre}((a_0c_0(b_0c_2 + c_2b_0))^*), \{a_0, b_0\}, \{c_0, c_2\}), \\ S_1 &= (\mathcal{L}^{Pre}(c_0(a_1c_1(b_1c_0 + c_0b_1))^*), \{a_1, b_1\}, \{c_0, c_1\}), \\ S_2 &= (\mathcal{L}^{Pre}(c_1(a_2c_2(b_2c_1 + c_1b_2))^*), \{a_2, b_2\}, \{c_1, c_2\}) \end{aligned}$$

The scheduler is defined in terms of the compound system:

$$S = S_0 \parallel S_1 \parallel S_2$$

where the set of observable events then consists of the  $a_i$ 's and  $b_i$ 's.

### 2.3 Implementation

We say that systems  $A$  and  $B$  are *comparable* if they have the same set of observable events  $\Sigma^{Obs}$ , that is,  $\Sigma^{Obs} = \Sigma_A^{Obs} = \Sigma_B^{Obs}$ . In the following,  $A$  and  $B$  denote comparable systems and  $\pi$  denotes the projection from  $\mathcal{U}^*$  onto  $(\Sigma^{Obs})^*$ .

**Definition 2.1**  $A$  implements  $B$  if and only if  $Obs(A) \subseteq Obs(B)$

**Example 2.2** Another way of defining a scheduler is to use a central agent  $C$ . The  $i$ th agent still performs observable events  $a_i$  and  $b_i$  but now synchronizes with the agent  $C$  by interacting on the internal event  $d_i$ . The agents are given by the systems

$$\begin{aligned} C &= (\mathcal{L}^{Pre}((d_0d_0d_1d_1d_2d_2)^*), \emptyset, \{d_0, d_1, d_2\}), \\ P_i &= (\mathcal{L}^{Pre}((d_i a_i d_i b_i)^*), \{a_i, b_i\}, \{d_i\}), \quad i = 0, 1, 2 \end{aligned}$$

and the scheduler is defined by the compound system:

$$P = P_0 \parallel P_1 \parallel P_2 \parallel C$$

where the observable events are the  $a_i$ 's and  $b_i$ 's and internal events are the  $d_i$ 's.

The systems  $S$  and  $P$  may be seen as existing in the universe  $\mathcal{U} = \{a_i, b_i, c_i, d_i, \tau \mid i = 0, 1, 2\}$  and are clearly comparable. The reader may convince himself that  $P$  implements  $S$ , but in general this not an easy task.

## 2.4 Relational trace abstractions

A *trace abstraction* is a relation on traces preserving observable behaviors.

**Definition 2.2** A trace abstraction  $\mathcal{R}$  from  $A$  to  $B$  is a relation on  $\mathcal{U}^* \times \mathcal{U}^*$  such that

1. If  $\alpha \mathcal{R} \beta$  then  $\pi(\alpha) = \pi(\beta)$
2.  $N_{\mathcal{U}}(A) \subseteq \text{dom } \mathcal{R}$
3.  $\text{rng } \mathcal{R} \subseteq N_{\mathcal{U}}(B)$

The first condition states that any pair of related traces must agree on observable events. The second and third condition require that any normalized trace of  $A$  should be related to some normalized trace of  $B$ , and only to normalized traces of  $B$ . The use of trace abstractions forms a sound and complete method in the sense that there exists a trace abstraction from  $A$  to  $B$  if and only if  $A$  implements  $B$ .

**Theorem 2.1** There exists a trace abstraction from  $A$  to  $B$  if and only if  $A$  implements  $B$ .

We would like to prove that a compound system  $\|A_i$  implements another compound system  $\|B_i$  by exhibiting trace abstractions  $\mathcal{R}_i$  from  $A_i$  to  $B_i$ . A simple extra condition is needed for this to work:

**Theorem 2.2** Let  $A_i$  and  $B_i$  be pairwise comparable systems forming the compound systems  $\|A_i$  and  $\|B_i$ . If

$$\mathcal{R}_i \text{ is a trace abstraction from } A_i \text{ to } B_i \quad (1)$$

$$\bigcap_i \text{dom } \mathcal{R}_i \subseteq \text{dom } \bigcap_i \mathcal{R}_i \quad (2)$$

then

$$\|A_i \text{ implements } \|B_i$$

Intuitively, the extra condition 2., which we call the *compatibility requirement*, ensures that the choices defined by the trace abstractions can be made to agree on internal events.

Due to the possibility of non-trivial interference on internal events among the component systems, the first premise alone of the composition rule is not sufficient to ensure the conclusion. Consider e.g. the following systems

$$\begin{aligned} A_1 &= (\{a\}^*, \{a\}, \emptyset), & B_1 &= (\{ac\}^* \{\epsilon, a\}, \{a\}, \{c\}) \\ A_2 &= (\{b\}^*, \{b\}, \emptyset), & B_2 &= (\{bc\}^* \{\epsilon, b\}, \{b\}, \{c\}) \end{aligned}$$

where  $\epsilon$  is the empty string.  $\text{Obs}(A_i) \subseteq \text{Obs}(B_i)$ , but  $\text{Obs}(A_1 \parallel A_2) \not\subseteq \text{Obs}(B_1 \parallel B_2)$ , since  $aa \in \text{Obs}(A_1 \parallel A_2)$ , but  $aa \notin \text{Obs}(B_1 \parallel B_2)$ .

The next example illustrates that even when significant internal interaction exists among the components, the decomposition theorem may be applied.

**Example 2.3** Consider the schedulers from before. For each  $i = 0, 1, 2$  let  $\chi_i$  be the string homomorphism from  $\mathcal{U}^*$  to  $\mathcal{U}^*$  inductively defined by  $\chi_i(\epsilon) = \epsilon$ , and for  $\alpha \in \mathcal{U}^*$  and  $u \in \mathcal{U}$ ,

$$\chi_i(\alpha u) = \begin{cases} \chi_i(\alpha)c_i & \text{if } u = d_i \text{ \& } \\ & \text{the number of } d_i\text{s in } \alpha \text{ is odd} \\ \chi_i(\alpha) & \text{if } u = c_i \\ \chi_i(\alpha)u & \text{otherwise} \end{cases}$$

where  $\epsilon$  is the empty string.  $\chi_i$  maps every string  $\alpha$  into a string identical to  $\alpha$  except for every occurrence of  $c_i$  being erased and every even occurrence of  $d_i$  being replaced by  $c_i$ . Let  $\chi = \chi_0 \circ \chi_1 \circ \chi_2$ . We use  $\chi$  to relate the internal behaviors of  $P$  and  $S$ . It is not hard to check that the relations  $\mathcal{R}_i = \{(\alpha, \chi(\alpha)) \mid \chi(\alpha) \in N_{\mathcal{U}}(S_i)\}$  are trace abstractions from  $P_i \parallel C$  to  $S_i$ , respectively. (Requirements 1. and 3. are satisfied by definition. To see that 2. holds, we consider some  $\alpha \in N_{\mathcal{U}}(P_i \parallel C)$  and argue that  $\chi(\alpha) \in N_{\mathcal{U}}(S_i)$ .) Also, it is not hard to see that  $\bigcap_i \text{dom } \mathcal{R}_i \subseteq \text{dom } \bigcap_i \mathcal{R}_i$ . (For each  $\alpha$  there is exactly one  $\chi(\alpha)$ .) Hence by Theorem 2.2, it follows that  $(P_0 \parallel C) \parallel (P_1 \parallel C) \parallel (P_2 \parallel C)$  implements  $S$  and therefore that  $P$  implements  $S$ .

An almost trivial observation is:

**Corollary 2.1** If additionally the components of the specification are non-interfering on internal events, that is,  $\Sigma_{B_i}^{Int} \cap \Sigma_{B_j}^{Int} = \emptyset$ , for every  $i \neq j$ , then  $A_i$  implements  $B_i$  implies  $\|A_i$  implements  $\|B_i$ .

### 3 Monadic second-order logic on strings

The logical language we use is the monadic second-order logic (M2L) on strings, where a closed formula is interpreted relative to a natural number  $n$  (the *length*). First-order variables  $p, q, \dots$  range over the set  $\{0, \dots, n-1\}$  (the set of *positions*), and second-order variables  $P, Q, \dots, P_1, P_2, \dots$  range over subsets of  $\{0, \dots, n-1\}$ . Atomic formulas are of the form  $p = q$ ,  $p = q + 1$ ,  $p < q$  and  $q \in P$ . Formulas are constructed in the standard way from atomic formulas by means of the Boolean connectives  $\neg, \wedge, \vee, \Rightarrow$  and  $\Leftrightarrow$ , and first and second-order quantifiers  $\forall$  and  $\exists$ . We adopt the standard notation of writing  $\phi(P_1, \dots, P_k, p_1, \dots, p_l)$  to denote an open formula  $\phi$  whose free variables are among  $P_1, \dots, P_k, p_1, \dots, p_l$ . Let 0 and \$ be the M2L definable constants denoting the positions 0 and  $n-1$ , respectively. The expressive power of M2L is illustrated by the formula

$$\exists P. 0 \in P \wedge (\forall p. p < \$ \Rightarrow (p \in P \Leftrightarrow p+1 \notin P))$$

which defines the set of even numbers. A second-order variable  $P$  can be seen as denoting a string of bits  $b_0 \dots b_{n-1}$  such that  $b_i = 1$  if and only if  $i \in P$ . This leads to a natural way of associating a language  $L(\phi)$  over  $\Sigma = \mathbb{B}^m$  of satisfying interpretations to an open formula  $\phi(P_1, \dots, P_m)$  having only second-order

variables occurring free ( $\mathcal{B}$  denotes the set  $\{0, 1\}$ ). As an example, consider the formula  $\phi \equiv \forall p.p \in P_1 \Leftrightarrow p \notin P_2$ . Then  $L(\phi)$  is a language over the alphabet  $\Sigma = \mathcal{B}^2$ , where each  $(b_1, b_2) \in \mathcal{B}^2$  denotes the membership status of the current position relative to  $P_1$  and  $P_2$ . For example, writing the tuples as columns, we have

$$\begin{array}{l} P_1: 11010 \\ P_2: 00101 \end{array} \in L(\phi) \text{ and } \begin{array}{l} P_1: 11010 \\ P_2: 01000 \end{array} \notin L(\phi)$$

Any language defined by a M2L formula is regular and conversely any regular language can be defined by a M2L formula. Given a formula  $\phi$ , a minimal finite automaton accepting  $L(\phi)$  can effectively be constructed using the standard operations of complementation, product, subset construction, and projection. In particular, the existential quantifier becomes associated with a subset construction—and a potential exponential blow-up in the number of states. The construction of automata constitutes a decision procedure for M2L, since  $\phi$  is a tautology if and only if  $L(\phi)$  is the set of all strings. In case  $\phi$  is not a tautology, a witness in terms of a minimal interpretation falsifying  $\phi$  can be derived from the minimum deterministic automaton recognizing  $L(\phi)$ . We use the tool *Mona* [8], which implements the decision procedure and the counter-example facility.

## 4 The finite state case

We now restrict attention to systems with regular trace languages. We show for a large class of finite-state systems that trace abstractions definable by regular languages constitute a complete method for proving the implementation property.

Given strings  $\alpha = \alpha_0 \dots \alpha_n \in \Sigma_1^*$  and  $\beta = \beta_0 \dots \beta_n \in \Sigma_2^*$ , we write  $\alpha \wedge \beta$  for the string  $(\alpha_0, \beta_0) \dots (\alpha_n, \beta_n) \in (\Sigma_1 \times \Sigma_2)^*$ . Every language  $L_{\mathcal{R}}$  over a product alphabet  $\Sigma_1 \times \Sigma_2$  has a canonical embedding as a relation  $\mathcal{R}_L \subseteq \Sigma_1^* \times \Sigma_2^*$  on strings of equal length given by  $\alpha \wedge \beta \in L_{\mathcal{R}} \stackrel{\text{def}}{\Leftrightarrow} \alpha \mathcal{R}_L \beta$ . Hence in the following we shall use the two representations interchangeably. Accordingly, we say that a trace abstraction is *regular* if it is the embedding of a regular language over  $\mathcal{U} \times \mathcal{U}$ .

Not all trace abstractions between finite-state systems are regular, since there may be an unbounded number of internal events between pairs of corresponding observable events. The next definition is an essential step towards the identification of regular trace abstractions.

**Definition 4.1** Given a subset  $\Sigma'$  of  $\Sigma$  we say that strings  $\alpha, \beta \in \Sigma^*$  are  $\Sigma'$ -*synchronized* if they are of equal length and if whenever the  $i$ th position in  $\alpha$  contains a letter in  $\Sigma'$  then the  $i$ th position in  $\beta$  contains the same letter, and vice versa.

**Definition 4.2** Let  $\hat{\mathcal{R}}$  be the language over  $\mathcal{U} \times \mathcal{U}$  given by  $\alpha \wedge \beta \in \hat{\mathcal{R}}$  if and only if

$$\beta \in N_{\mathcal{U}}(B) \text{ and } \alpha, \beta \text{ are } \Sigma^{Obs}\text{-synchronized}$$

Since  $N_{\mathcal{U}}(B)$  is a regular language (by assumption of this Section), so is  $\hat{\mathcal{R}}$ . The

next proposition gives a sufficient condition for  $\hat{\mathcal{R}}$  and any regular subset of  $\hat{\mathcal{R}}$  to be a trace abstraction. We return to the significance of the last part when dealing with automated proofs.

**Proposition 4.1** If  $N_{\mathcal{U}}(A) \subseteq \text{dom } \hat{\mathcal{R}}$  then  $\hat{\mathcal{R}}$  is a regular trace abstraction from  $A$  to  $B$ . Moreover, in general for any regular language  $\mathcal{C} \subseteq (\mathcal{U} \times \mathcal{U})^*$ , if  $N_{\mathcal{U}}(A) \subseteq \text{dom } \hat{\mathcal{R}} \cap \mathcal{C}$ , then  $\hat{\mathcal{R}} \cap \mathcal{C}$  is a regular trace abstraction from  $A$  to  $B$ .

It is not hard to see that if  $\hat{\mathcal{R}}$  is a regular trace abstraction, then it is the largest such relating  $\Sigma^{Obs}$ -synchronized traces. In this case we denote  $\hat{\mathcal{R}}$  the *canonical* trace abstraction.

Non-regularity of trace abstractions occurs if for example there are arbitrarily many non-observable events between any two observable events. However, it may also happen that a behavior of the program may have too few internal events between two observable events in the sense that any behavior of the specification with the same observable behavior may require more internal events. We next give a precise definition of this phenomenon. Let  $\pi_A$  and  $\pi_B$  be the projections from  $\Sigma_A^*$  and  $\Sigma_B^*$ , respectively, onto  $(\Sigma^{Obs})^*$ .

**Definition 4.3** A trace  $\alpha \in L_A$  is *internally finer* than a trace  $\beta \in L_B$  if  $\pi_A(\alpha) = \pi_B(\beta)$ , and for all  $e, e' \in \Sigma^{Obs}$ ,  $u \in (\Sigma_A^{Int})^*$ ,  $v \in (\Sigma_B^{Int})^*$ ,  $\alpha_1, \alpha_2 \in \Sigma_A^*$  and  $\beta_1, \beta_2 \in \Sigma_B^*$ , such that  $\pi_A(\alpha_1) = \pi_B(\beta_1)$

$$\vee \left. \begin{array}{l} \alpha = \alpha_1 e u e' \alpha_2 \quad \wedge \quad \beta = \beta_1 e v e' \beta_2 \\ \alpha = u e' \alpha_2 \quad \wedge \quad \beta = v e' \beta_2 \end{array} \right\} \Rightarrow |u| \geq |v|$$

A system  $A$  is internally finer than a system  $B$  if for any trace  $\alpha$  of  $A$  such that  $\pi_A(\alpha) \in \text{Obs}(B)$ , there exists a trace  $\beta$  of  $B$  such that  $\alpha$  is internally finer than  $\beta$ .

Consider the scheduler example. System  $P$  is internally finer than  $S$ , whereas the converse is not true. We restate the soundness and completeness result from the general case for a constrained class of systems and regular trace abstractions.

**Theorem 4.1** Assume that  $A$  is internally finer than  $B$ . There exists a canonical trace abstraction from  $A$  to  $B$  if and only if  $A$  implements  $B$ .

The restriction on programs to be internally finer than their specifications can be overcome simply by adding more internal behavior to the program. More precisely, given systems  $A$  and  $B$  there always exists a system  $A'$  such that  $A$  and  $A'$  have the same observable behaviors, that is,  $\text{Obs}(A) = \text{Obs}(A')$ , and such that  $A'$  is internally finer than  $B$ . E.g. using  $S'_0 = (\mathcal{L}^{Pre}((d_0 a_0 d_0 c_0 (b_0 c_2 + c_2 b_0))^*), \{a_0, b_0\}, \{c_0, c_2, d_0\})$  instead of  $S_0$  and with similar changes using  $S'_1$  and  $S'_2$  for  $S_1$  and  $S_2$ , respectively, we have that  $S' = S'_0 \parallel S'_1 \parallel S'_2$  is internally finer than  $P$  and that  $\text{Obs}(S) = \text{Obs}(S')$ .

## 4.1 A uniform logical framework

In the finite setting, reasoning about systems can conveniently be expressed in M2L. Let  $\mathcal{U} = \mathcal{B}^m$  be the universe, where  $m$  is a natural number. Any behavior  $\alpha$  over  $\mathcal{U}$  can be viewed as an interpretation of a sequence of second-order variables  $U_1^\alpha, \dots, U_m^\alpha$ . So behaviors over, say, 1024 different events may be coded using just 10 variables.

We use for each event  $\sigma = (b_1, \dots, b_m) \in \mathcal{U}$  and  $\alpha$  the notation  $\alpha(t) = \sigma$  for the M2L predicate

$$\left( \bigwedge_{b_i=1} t \in U_i^\alpha \right) \wedge \left( \bigwedge_{b_i=0} t \notin U_i^\alpha \right),$$

which states that the behavior denoted by  $\alpha$  has a  $\sigma$  event in the position denoted by  $t$ . A system  $A = (L_A, \Sigma_A^{Obs}, \Sigma_A^{Int})$  is represented by a triple

$$A = (\phi_A, \phi_A^{Obs}, \phi_A^{Int})$$

of formulas defining the normalized traces of the system,  $\phi_A(\alpha)$ , the observable events,  $\phi_A^{Obs}(\alpha, t)$ , and the internal events,  $\phi_A^{Int}(\alpha, t)$ . That is,  $N_{\mathcal{U}}(A) = L(\phi_A)$  and  $\phi_A^{Obs}(\alpha, t)$  and  $\phi_A^{Int}(\alpha, t)$  are predicates that are true if and only if the position denoted by  $t$  in the behavior denoted by  $\alpha$  is an element of  $\Sigma_A^{Obs}$  and  $\Sigma_A^{Int}$ , respectively. Given composable systems  $A$  and  $B$ , composition is represented by

$$A \parallel B = (\phi_A \wedge \phi_B, \phi_A^{Obs} \vee \phi_B^{Obs}, \phi_A^{Int} \vee \phi_B^{Int})$$

We have that  $L(\phi_A \wedge \phi_B) = L(\phi_A) \cap L(\phi_B) = N_{\mathcal{U}}(A \parallel B)$  and that  $\phi_A^{Obs} \vee \phi_B^{Obs}$  and  $\phi_A^{Int} \vee \phi_B^{Int}$  defines the union of the observable and the internal events, respectively. Let now behavior  $\beta$  be represented by  $U_1^\beta, \dots, U_m^\beta$ . The property that behaviors  $\alpha$  and  $\beta$  in  $\mathcal{U}^*$  are  $\Sigma^{Obs}$ -synchronized is expressed by predicate  $\phi_{A,B}^{Obs}(\alpha, \beta)$  defined by

$$\forall t : (\phi_A^{Obs}(\alpha, t) \vee \phi_B^{Obs}(\alpha, t)) \Rightarrow \alpha(t) = \beta(t)$$

The canonical trace abstraction  $\hat{\mathcal{R}}$  of Definition 4.2 is defined by

$$\hat{\mathcal{R}}_{A,B}(\alpha, \beta) \stackrel{\text{def}}{=} \phi_B(\beta) \wedge \phi_{A,B}^{Obs}(\alpha, \beta)$$

By Proposition 4.1 and Theorem 4.1, the implementation property is implied by  $N_{\mathcal{U}}(A) \subseteq \text{dom } \hat{\mathcal{R}}$  and hence by the validity of

$$\phi_A(\alpha) \Rightarrow \exists \beta : \hat{\mathcal{R}}_{A,B}(\alpha, \beta) \tag{3}$$

where  $\exists \beta$  is defined as  $\exists U_1^\beta \dots \exists U_m^\beta$ . Let  $\mathcal{R}_i(\alpha, \beta) \stackrel{\text{def}}{=} \hat{\mathcal{R}}_{A_i, B_i}(\alpha, \beta) \wedge \phi_i(\alpha, \beta)$ . The premises of the decomposition rule 2.2 are expressed by

$$\bigwedge_i (\phi_{A_i}(\alpha) \Rightarrow \exists \beta : \mathcal{R}_i(\alpha, \beta)) \tag{4}$$

$$\bigwedge_i \exists \beta_i : \mathcal{R}_i(\alpha, \beta_i) \Rightarrow \exists \beta : \bigwedge_i \mathcal{R}_i(\alpha, \beta) \tag{5}$$

To express the premise of Corollary 2.1 simply replace equation (5) above by

$$\bigwedge_{i \neq j} \forall t : \phi_{B_i}^{Int}(\alpha, t) \Rightarrow \neg \phi_{B_j}^{Obs}(\alpha, t)$$

Also, properties like composability and comparability can be expressed. The former by

$$\forall t : \quad (\phi_A^{Int}(\alpha, t) \Rightarrow \neg \phi_B^{Obs}(\alpha, t)) \wedge \\ (\phi_B^{Int}(\alpha, t) \Rightarrow \neg \phi_A^{Obs}(\alpha, t))$$

and the latter by

$$\forall t : \phi_A^{Obs}(\alpha, t) \Leftrightarrow \phi_B^{Obs}(\alpha, t)$$

In general M2L is a very flexible logical language making it easy to write tense time and interval temporal logic operators in a straightforward manner. As examples, consider the past operator  $\phi_{\sigma, \mu}^{Before}(\alpha)$  defined by

$$\forall t_1. \alpha(t_1) = \mu \Rightarrow \exists t_0. t_0 < t_1 \wedge \alpha(t_0) = \sigma,$$

and the interval operator  $\phi_{\sigma}^{Between}(\alpha, t_1, t_2)$

$$\exists t. t_1 < t < t_2 \wedge \alpha(t) = \sigma$$

## 4.2 Automated proofs

Formulas (3), (4), and (5) are potentially very difficult, since they involve quantification over behaviors, that is, over  $m$  second-order variables. Each quantification can lead to an exponential blow-up. But if  $A$  has much internal behavior, then it seems reasonable to use a more clever trace abstraction guided by  $A$ 's internal events. In fact, it must be suspected that it is inappropriate that the definition of  $\hat{\mathcal{R}}$  does not involve  $A$  at all.

The canonical trace abstraction can be constrained by adding more precise information about the connection between the internal behavior of system  $A$  and  $B$ . This may reduce the blow-up—or even avoid it in the case a functional regular trace abstraction is formulated.

We next turn to a substantial verification problem to illustrate our technique.

## 5 A specification problem

In this section, we consider the problem proposed by Broy and Lamport [3]. The first part of [3] calls for a specification of a reactive system consisting of a number of sequential processes issuing blocking read and write calls to a memory server. The memory server maintains its memory by performing special atomic reads and writes whenever requested to do so by read and write calls. Depending on the success of atomic reads and writes, return events contain the answers to

read and write calls. The memory must be able to handle several calls (from different processes) concurrently.

The second part of [3] calls for an implementation based on a remote procedure call protocol. The protocol involves a local and a remote party. Calls received locally are forwarded to the remote site, where they are executed. The resulting return events are propagated back to the local site. Altogether, we deal here with four levels of calls and returns.

The goal of [3] is now to verify that every observable trace of the implementation (where atomic read and writes and the remote events are abstracted away) is an observable trace of the specification.

The full informal description [3] includes many technical complications concerning the parameters passed and different kinds of erroneous behaviors. A detailed presentation of our solution can be found in [11].

In performing the verifications, we have limited ourselves to finite domains. We have chosen to have two locations, two kinds of values, two kinds of flags, and two process identities (in addition to the memory process). The resulting program has approximately a hundred thousand states and the specification approximately a thousand states. The systems allow thousands of different events. The systems are modelled as deterministic automata. The full specification amounts to more than 13 pages of M2L formulas (written in a macro language).

The aspect that we are interested in here is the use of trace abstractions. Without going into any further details, we assume that the M2L formulas  $\phi_{P_1} \wedge \phi_{P_2}$  and  $\phi_{S_1} \wedge \phi_{S_2}$  define the implementation and the specification, respectively, of our solution. The universe  $\mathcal{U}$  consists of  $\tau$  and a number of parameterized events: *rd*, *wrt*, *rtn*, *atmrd*, *atmwrt*, *rpcCall*, *rpcRtn* denoting reads, writes, returns, atomic reads, atomic writes, rpc calls and rpc returns respectively. For example, *rd*: [*?*, *obs*, 1] is a read event, where the first parameter is unspecified, the second is *obs*, which stands for an observable event, and the last parameter 1 denotes the process id. A similar notation is used for other events.

The *Mona* tool is currently not able to handle automata of the size corresponding to the distributed program just discussed. Hence we prove the correctness of the implementation by using our composition rule. The obvious idea is to try whether

$$\phi_{P_i}(\alpha) \Rightarrow \phi_{S_i}(\alpha)$$

holds (for  $i = 1$  or  $i = 2$ ; the formulas are symmetric). The *Mona* tool, however, quickly determines that this formula is not valid. There is a counter-example of length 12:

$$\begin{aligned} &rd:[obs], rpcCall, rd, atmrd, rtn, rpcRtn, \\ &rpcCall, rd, atmrd, rtn, rpcRtn, rtn:[obs] \end{aligned}$$

where we have left out most of the parameters. The counter-example arises because the specification system requires exactly one atomic read in every successful read call, whereas the implementation is allowed to retry on failure.

Fortunately, we can let **Mona** establish

$$\phi_{P_i}(\alpha) \Rightarrow \exists \beta. \hat{\mathcal{R}}_i(\alpha, \beta), \quad (6)$$

where  $\hat{\mathcal{R}}_i(\alpha, \beta) \stackrel{\text{def}}{=} \phi_{P_i, S_i}^{Obs}(\alpha, \beta) \wedge \phi_{S_i}(\alpha, \beta)$  is the canonical trace abstraction. Thus,  $\phi_{P_i}$  implements  $\phi_{S_i}$ .

To avoid explicitly modeling the whole system at the implementation level, we use the proof rule for compound systems. The compatibility premise of Theorem 2.2 becomes:

$$\bigwedge_i \exists \beta_i : \hat{\mathcal{R}}_i(\alpha, \beta_i) \Rightarrow \exists \beta : \bigwedge_i \hat{\mathcal{R}}_i(\alpha, \beta) \quad (7)$$

However, the existential quantification on the right hand side of the implication leads to a state explosion that cannot be handled by the **Mona** tool.

Instead, we can exploit the information that the counter-example provided to formulate a more precise trace abstraction. So we have defined predicates that in more detail describe how internal events at one level relate to internal events at the other level. For example, we may add our intuition that between any successful read and its corresponding return at the implementation level only the last atomic read is mapped to an atomic read on the specification level. This formula, which we denote by  $\psi_i$ , looks like:

$$\begin{aligned} & \forall t_1, t_2. (t_1 < t_2 \wedge \\ & \alpha(t_1) = rd : [?, obs, i] \wedge \\ & \alpha(t_2) = rtn : [?, ?, normal, obs, i] \wedge \\ & \neg \phi_{rd:[?, obs, i]}^{Between}(\alpha, t_1, t_2) \wedge \\ & \neg \phi_{wrt:[?, ?, obs, i]}^{Between}(\alpha, t_1, t_2)) \\ \Rightarrow & \\ & (\exists t. t_1 < t < t_2 \wedge \\ & \alpha(t) = \beta(t) = atmrd : [?, ?, ?, i] \wedge \\ & \neg \phi_{atmrd:[?, ?, ?, i]}^{Between}(\alpha, t, t_2) \wedge \\ & \neg \phi_{atmrd:[?, ?, ?, i]}^{Between}(\beta, t_1, t) \wedge \\ & \neg \phi_{atmrd:[?, ?, ?, i]}^{Between}(\beta, t, t_2)) \end{aligned}$$

We define the new trace abstractions  $\mathcal{R}_i(\alpha, \beta)$  to be equal to  $\hat{\mathcal{R}}_i$  conjoined with the  $\psi_i$  and two other similar predicates (one stating that any event on the program level—but an atomic read—is matched by the same event on the specification level; the other stating that an atomic read event on the program level is matched by either an atomic read event or a  $\tau$  event on the specification level). With  $\mathcal{R}_i$ , the **Mona** tool proves formulas (6) and (7) within minutes.

The compatibility property (7) is stated in a single M2L formula of size  $10^5$  with approximately 32 visible variables at the level of deepest nesting (corresponding to an alphabet size of  $2^{32}$ ). During its processing automata with

millions of BDD nodes are created. The proof required user intervention in the form of an explicit (but natural) ordering of BDD variables. Also, we have supplied a little information about evaluation order in the form of parentheses.

## 6 Conclusion

We have offered a practical alternative to the use of refinement mappings. We have indicated how the user contribution of information about behavioral similarities directly can be used to reduce the computational work involved in guessing internal events when two distributed systems are compared.

Our method is entirely formulated within M2L: state machines, temporal properties, finite domains, and verification rules all take on the syntax of the Mona system.

Our experiments show that very complex temporal logic formulas on finite segments of time can be decided in practice—quite in contrast to the situation for temporal logic on the natural numbers.

## References

- [1] M. Abadi and L. Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.
- [2] M. Abadi and L. Lamport. Conjoining specifications. Technical Report Report 118, Digital Equipment Corporation, Systems Research Center, 1993.
- [3] M. Broy and L. Lamport. Specification problem, 1994. A case study for the Dagstuhl Seminar 9439.
- [4] E. M. Clarke, I.A. Browne, and R.P Kurshan. A unified approach for showing language containment and equivalence between various types of  $\omega$ -automata. In A. Arnold, editor, *CAAP, LNCS 431*, pages 103–116, 1990.
- [5] E. M. Clarke, O. Grumberg, and D. E. Long. Verification tools for finite-state concurrent systems. In *A Decade of Concurrency*, pages 124–175. Springer-Verlag, 1993. Lecture Notes in Computer Science, Vol. 803, Proceedings of the REX School/Symposium, Noordwijkerhout, The Netherlands, June 1993.
- [6] U. Engberg, Grønning P., and Lamport L. Mechanical verification of concurrent systems with tla. In *Computer Aided Verification, CAV '92*. Springer-Verlag, 1993. Lecture Notes in Computer Science, Vol. 663.
- [7] Manna Z. et al. Step: The stanford temporal prover. In *Theory and Practice of Software Development (TAPSOFT)*. Springer-Verlag, 1995. Lecture Notes in Computer Science, Vol. 915.

- [8] J.G. Henriksen, O.J.L. Jensen, M.E. Jørgensen, N. Klarlund, R. Paige, T. Rauhe, and A.B. Sandholm. Mona: Monadic second-order logic in practice. In U.H. Engberg, K.G. Larsen, and A. Skou, editors, *Proceedings of the Workshop on Tools and Algorithms for the Construction and Analysis of Systems*, pages 58–73, 1995. BRICS Notes Series NS-95-2.
- [9] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
- [10] Y. Kesten, Z. Manna, and A. Pnueli. Temporal verification and simulation and refinement. In *A Decade of Concurrency*, pages 273–346. ACM, Springer-Verlag, 1993. Lecture Notes in Computer Science, Vol. 803, Proceedings of the REX School/Symposium, Noordwijkerhout, The Netherlands, June 1993.
- [11] N. Klarlund, M. Nielsen, and K. Sunesen. Using monadic second-order logic with finite domains for automated verification. Technical Report RS-95-54, BRICS, Aarhus University, 1995.
- [12] N. Klarlund and F.B. Schneider. Proving nondeterministically specified safety properties using progress measures. *Information and Computation*, 107(1):151–170, 1993.
- [13] R. Kurshan. *Computer-Aided Verification of Coordinating Processes*. Princeton Univ. Press, 1994.
- [14] R. P. Kurshan, M. Merritt, A. Orda, and S. R. Sachs. Modelling asynchrony with a synchronous model. In *Computer Aided Verification, CAV '95, LNCS*, 1995. Lecture Notes in Computer Science.
- [15] L. Lamport. Specifying concurrent program modules. *ACM Transactions on Programming Languages and Systems*, 5(2):190–222, 1983.
- [16] L. Lamport. The temporal logic of actions. Technical Report Report 70, Digital Equipment Corporation, Systems Research Center, 1994. To appear in *Transactions on Programming Language and Systems*.
- [17] N. Lynch and M. Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Proc. Sixth Symp. on the Principles of Distributed Computing*, pages 137–151. ACM, 1987.
- [18] N. Lynch and F. W. Vaandrager. Forward and backward simulations – part i: untimed systems. Technical Report CS-R9313, Centrum voor Wiskunde en Informatica, CWI, Computer Science/Department of Software Technology, 1993.
- [19] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [20] A.P. Sistla. On verifying that a concurrent program satisfies a nondeterministic specification. *Information Processing Letters*, 32(1):17–24, July 1989.

## A Proof of Theorem 2.1

**Proof:** Only if:

Assume that  $\mathcal{R}$  is a trace abstraction from  $A$  to  $B$ .

Let  $\eta \in Obs(A)$ . Then there exists a trace  $\alpha \in L_A \subseteq N_{\mathcal{U}}(A)$  such that  $\eta = \pi(\alpha)$ . By 2.) and 3.) of Definition 2.2 there exists a  $\beta \in N_{\mathcal{U}}(B)$  such that  $\alpha \mathcal{R} \beta$  and hence by 1.) of Definition 2.2 such that  $\pi(\alpha) = \pi(\beta)$ . Hence  $\eta = \pi(\alpha) = \pi(\beta) \in Obs(B)$ .

If:

Assume  $Obs(A) \subseteq Obs(B)$ . Define  $\mathcal{R} = \{(\alpha, \beta) \mid \beta \in N_{\mathcal{U}}(B) \wedge \pi(\alpha) = \pi(\beta)\} \subseteq \mathcal{U}^* \times \mathcal{U}^*$ . We prove that  $\mathcal{R}$  is a trace abstraction from  $A$  to  $B$ . Clearly, 1) and 3) of Definition 2.2 are satisfied. To see that 2) is satisfied let  $\alpha \in N_{\mathcal{U}}(A)$ . Then  $\pi(\alpha) \in Obs(A) \subseteq Obs(B)$  by assumption. Hence there exists a  $\beta \in L_B \subseteq N_{\mathcal{U}}(B)$  such that  $\pi(\alpha) = \pi(\beta)$ . Thus  $\alpha \mathcal{R} \beta$  and therefore  $\alpha \in dom \mathcal{R}$ .  $\square$

## B Proof of Theorem 2.2

**Proof:** The key is that the  $\mathcal{R}_i$ 's are trace abstractions on  $\mathcal{U}^* \times \mathcal{U}^*$ . Assume (1) and (2).

By Theorem 2.1, it is enough to prove that  $\bigcap_i \mathcal{R}_i$  is a trace abstraction from  $\|A_i$  to  $\|B_i$  on  $\mathcal{U}^* \times \mathcal{U}^*$ .

We show that  $\bigcap_i \mathcal{R}_i$  satisfies 1.)-3.) of Definition 2.2.

1.) Obvious.

2.)  $N_{\mathcal{U}}(\|A_i) = \bigcap_i N_{\mathcal{U}}(A_i) \subseteq \bigcap_i dom \mathcal{R}_i \subseteq dom \bigcap_i \mathcal{R}_i$ .

The first inclusion follows from assumption (1) and Definition 2.2(2). The second inclusion from assumption (2).

3.)  $rng \bigcap_i \mathcal{R}_i \subseteq \bigcap_i rng \mathcal{R}_i \subseteq \bigcap_i N_{\mathcal{U}}(B_i) = N_{\mathcal{U}}(\|B_i)$ .

The second inclusion follows from assumption (1) and Definition 2.2(3).  $\square$

## C Proof of Corollary 2.1

**Proof:** Assume that  $\Sigma_{B_i}^{Int} \cap \Sigma_{B_j}^{Int} = \emptyset$  for every  $i \neq j$  and that  $A_i$  implements  $B_i$ . Let  $\Sigma_i^{Obs} = \Sigma_{A_i}^{Obs} = \Sigma_{B_i}^{Obs}$ ,  $\Sigma_i^{Obs} = \Sigma_{\|A_i}^{Obs} = \Sigma_{\|B_i}^{Obs}$ , and let  $\pi$ ,  $\pi_i$  and  $h_i$  be the projections from  $\mathcal{U}^*$  to  $(\Sigma^{Obs})^*$ ,  $(\Sigma_i^{Obs})^*$  and  $(\Sigma_{B_i})^*$ , respectively. By Theorem 2.1, there exist trace abstractions from  $A_i$  to  $B_i$ . Let  $\mathcal{R}_i$  be the largest such, that is  $\mathcal{R}_i = \{(\alpha, \beta) \mid \beta \in N_{\mathcal{U}}(B_i) \wedge \pi_i(\alpha) = \pi_i(\beta)\}$ .

We first prove that if  $\mathcal{R}_i$  is a trace abstraction then also  $\mathcal{R}'_i$  given by  $\mathcal{R}_i \cap \{(\alpha, \beta) \mid \pi(\alpha) = \pi(\beta)\}$  is a trace abstraction. Clearly, it is sufficient to show that for all  $\alpha \in \mathcal{U}^*$  if there exists a  $\beta \in N_{\mathcal{U}}(B_i)$  such that  $\pi_i(\alpha) = \pi_i(\beta)$  then there exists a  $\beta' \in N_{\mathcal{U}}(B_i)$  such that  $\pi(\alpha) = \pi(\beta')$ . To show this assume that  $\alpha \in \mathcal{U}^*$  and  $\beta \in N_{\mathcal{U}}(B_i)$  such that  $\pi_i(\alpha) = \pi_i(\beta)$ . Then  $h_i(\beta) \in L_{B_i} \subseteq N_{\mathcal{U}}(B_i)$  and  $\pi_i(h_i(\beta)) = \pi_i(\alpha) = o_1 \dots o_m$ , where  $o_j \in \Sigma_i^{Obs}$ . Let  $\alpha_1, \dots, \alpha_{m+1} \in (\mathcal{U} - \Sigma_i^{Obs})^*$  and  $\beta_1, \dots, \beta_{m+1} \in (\Sigma_{B_i}^{Int})^*$  be such that

$$\begin{aligned}\alpha &= \alpha_1 o_1 \alpha_2 \dots \alpha_m o_m \alpha_{m+1}, \\ h_i(\beta) &= \beta_1 o_1 \beta_2 \dots \beta_m o_m \beta_{m+1}\end{aligned}$$

then

$$\beta' = \pi(\alpha_1) \beta_1 o_1 \dots \pi(\alpha_m) \beta_m o_m \pi(\alpha_{m+1}) \beta_{m+1}$$

is in  $N_{\mathcal{U}}(B_i)$ , since  $h_i(\beta') = h_i(\beta)$  as  $\pi(\alpha_i) \in (\Sigma^{Obs} - \Sigma_i^{Obs})^*$  and clearly  $\pi(\alpha) = \pi(\beta')$ .

We next prove that  $\bigcap_i \text{dom } \mathcal{R}'_i \subseteq \text{dom } \bigcap_i \mathcal{R}'_i$  and hence the result follows from Theorem 2.2. Given  $\alpha \in \bigcap_i \text{dom } \mathcal{R}'_i$ . Let

$$\alpha = \alpha_1 o_1 \alpha_2 \dots \alpha_m o_m \alpha_{m+1}$$

where  $o_j \in \Sigma^{Obs}$  and  $\alpha_j \in (\mathcal{U} - \Sigma^{Obs})^*$ . Then there exists

$$\begin{aligned}\beta^1 &= \beta_1^1 o_1 \dots \beta_m^1 o_m \beta_{m+1}^1 \in N_{\mathcal{U}}(B_1) \\ &\vdots \\ \beta^n &= \beta_1^n o_1 \dots \beta_m^n o_m \beta_{m+1}^n \in N_{\mathcal{U}}(B_n)\end{aligned}$$

where  $\beta_1^i, \dots, \beta_{m+1}^i \in (\Sigma_{B_i}^{Int})^*$  such that  $\alpha \mathcal{R}'_i \beta^i$ . Hence

$$\eta = \beta_1^1 \dots \beta_1^n o_1 \dots \beta_m^1 \dots \beta_m^n o_m \beta_{m+1}^1 \dots \beta_{m+1}^n$$

is in  $N_{\mathcal{U}}(\|B_i\|)$ , since  $h_i(\eta) = \beta^i$  as  $\Sigma_{B_i}^{Int} \cap \Sigma_{B_j}^{Int} = \emptyset$  for  $i \neq j$ , and thus  $\alpha(\bigcap_i \mathcal{R}_i)\eta$ .  $\square$

## D Proof of Proposition 4.1

**Proof:** Assume that  $N_{\mathcal{U}}(A) \subseteq \text{dom } \hat{\mathcal{R}}$  then by Definition 4.2,  $\hat{\mathcal{R}}$  is a trace abstraction, c.f. Definition 2.2.

To see that  $\hat{\mathcal{R}}$  is regular consider the following definition.

**Definition D.1** Let  $D_{\mathcal{U}}$  and  $D_B$  be the deterministic finite automata associated with the sets of normalized traces of the finite system  $B$  and the universe  $\mathcal{U}$ , respectively. Define from these a non-deterministic automaton  $D_{\mathcal{U}} \otimes D_B$  on the alphabet  $\mathcal{U} \times \mathcal{U}$  as follows. The initial state is the pair of initial states. A state is accepting iff it is a pair of accepting states. The transition relation is given by the set  $\{(s, s') \xrightarrow{(a,b)} (t, t') \mid s \xrightarrow{a} t \wedge s' \xrightarrow{b} t' \wedge ((a \notin \Sigma^{Obs} \wedge b \notin \Sigma^{Obs}) \vee a = b)\}$

The product automaton lock steps the two automata forcing observable events to be synchronized and guessing the right pair otherwise.

Let  $\alpha = \alpha_1 \dots \alpha_n \in \mathcal{U}$  and  $\beta = \beta_1 \dots \beta_n \in \mathcal{U}$ . Assume  $\alpha \wedge \beta \in L(D_{\mathcal{U}} \otimes D_B)$ . Then there exists an accepting run

$$(s_0, t_0) \xrightarrow{(\alpha_1, \beta_1)} (s_1, t_1) \xrightarrow{(\alpha_2, \beta_2)} \dots \xrightarrow{(\alpha_n, \beta_n)} (s_n, t_n)$$

of  $D_{\mathcal{U}} \otimes D_B$  and it is now trivial to check that  $\alpha$  and  $\beta$  are  $\Sigma^{Obs}$ -synchronized and that  $\beta \in N_{\mathcal{U}}(B)$ . Hence  $\alpha \wedge \beta \in \hat{\mathcal{R}}$ . Conversely, assume that 1)  $\alpha$  and  $\beta$  are  $\Sigma^{Obs}$ -synchronized, 2)  $\beta \in N_{\mathcal{U}}(B)$ . We prove that  $\alpha \wedge \beta \in L(D_{\mathcal{U}} \otimes D_B)$ . Since  $\alpha \in \mathcal{U}^*$  there exists an accepting run

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} s_n$$

of  $D_{\mathcal{U}}$  and likewise, since  $\beta \in N_{\mathcal{U}}(B)$  there exists an accepting run

$$t_0 \xrightarrow{\beta_1} t_1 \xrightarrow{\beta_2} \dots \xrightarrow{\beta_n} t_n$$

of  $D_B$ . Due to 1.) and Definition D.1 it follows that

$$(s_0, t_0) \xrightarrow{(\alpha_1, \beta_1)} (s_1, t_1) \xrightarrow{(\alpha_2, \beta_2)} \dots \xrightarrow{(\alpha_n, \beta_n)} (s_n, t_n)$$

is an accepting run of  $D_{\mathcal{U}} \otimes D_B$  accepting  $\alpha \wedge \beta$ .

Finally, since  $\mathcal{C}$  is regular the last result follows trivially.  $\square$

## E Proof of Theorem 4.1

**Proof:** The only if direction is as in the proof of Theorem 2.1, c.f. appendix A. We prove the if direction.

Let  $\pi$  be the projection from  $\mathcal{U}^*$  onto  $(\Sigma^{Obs})^*$ .

Assume that  $A$  implements  $B$ . Let  $\hat{\mathcal{R}}$  be the language defined in Definition 4.2. According to Proposition 4.1, we only need to prove that  $N_{\mathcal{U}}(A) \subseteq \text{dom } \hat{\mathcal{R}}$ . Let  $\alpha \in N_{\mathcal{U}}(A)$ . Then, since  $A$  implements  $B$ ,  $\pi(\alpha) \in \text{Obs}(A) \subseteq \text{Obs}(B)$ , and hence by assumption there exists a trace  $\beta \in L_B$  such that  $\alpha' = h(\alpha) \in L_A$  is internally finer than  $\beta$ , where  $h$  is the projection from  $\mathcal{U}^*$  onto  $\Sigma_A^*$ .

Let  $\pi(\alpha') = \pi(\alpha) = \pi(\beta) = o_1 \dots o_n$  and

$$\begin{aligned} \alpha &= \alpha_1 o_1 \alpha_2 \dots \alpha_n o_n \alpha_{n+1}, \\ \alpha' &= \alpha'_1 o_1 \alpha'_2 \dots \alpha'_n o_n, \\ \beta &= \beta_1 o_1 \beta_2 \dots \beta_n o_n \end{aligned}$$

where  $o_j \in \Sigma^{Obs}$ ,  $\alpha_j \in (\mathcal{U} - \Sigma^{Obs})^*$ ,  $\alpha'_j \in (\Sigma_A^{Int})^*$  and  $\beta_j \in (\Sigma_B^{Int})^*$ . Note, that since  $L_A$  and  $L_B$  are prefix-closed we can with out lose of generality assume that  $\alpha'$  and  $\beta$  end in observable events. Since  $\alpha'$  is internally finer than  $\beta$ , it follows that  $|\alpha'_i| \geq |\beta_i|$  for every  $i = 1, \dots, n$ , and since  $\alpha'$  is the projection of  $\alpha$  onto  $\Sigma_A^*$ , it follows that  $|\alpha_i| \geq |\alpha'_i|$  for every  $i = 1, \dots, n$ . Define

$$\eta = \eta_1 o_1 \eta_2 \dots \eta_n o_n \eta_{n+1}$$

where  $\eta_i = \beta_i \tau^{|\alpha_i| - |\beta_i|}$  for  $i = 1, \dots, n$  and  $\eta_{n+1} = \tau^{|\alpha_{n+1}|}$ . Hence we have the required  $\eta$  such that  $\alpha \mathcal{R} \eta$ , since  $\eta$  and  $\alpha$  are  $\Sigma^{Obs}$ -synchronized and  $\eta \in N_{\mathcal{U}}(B)$ .  
 $\square$

## Recent Publications in the BRICS Report Series

- RS-95-53 Nils Klarlund, Mogens Nielsen, and Kim Sunesen. *Automated Logical Verification based on Trace Abstractions*. November 1995. 19 pp.
- RS-95-52 Antonín Kucera. *Deciding Regularity in Process Algebras*. October 1995. 42 pp.
- RS-95-51 Rowan Davies. *A Temporal-Logic Approach to Binding-Time Analysis*. October 1995. 11 pp.
- RS-95-50 Dany Breslauer. *On Competitive On-Line Paging with Lookahead*. September 1995. 12 pp.
- RS-95-49 Mayer Goldberg. *Solving Equations in the  $\lambda$ -Calculus using Syntactic Encapsulation*. September 1995. 13 pp.
- RS-95-48 Devdatt P. Dubhashi. *Simple Proofs of Occupancy Tail Bounds*. September 1995. 7 pp. To appear in *Random Structures and Algorithms*.
- RS-95-47 Dany Breslauer. *The Suffix Tree of a Tree and Minimizing Sequential Transducers*. September 1995. 15 pp.
- RS-95-46 Dany Breslauer, Livio Colussi, and Laura Toniolo. *On the Comparison Complexity of the String Prefix-Matching Problem*. August 1995. 39 pp. Appears in Leeuwen, editor, *Algorithms - ESA '94: Second Annual European Symposium proceedings*, LNCS 855, 1994, pages 483–494.
- RS-95-45 Gudmund Skovbjerg Frandsen and Sven Skyum. *Dynamic Maintenance of Majority Information in Constant Time per Update*. August 1995. 9 pp.
- RS-95-44 Bruno Courcelle and Igor Walukiewicz. *Monadic Second-Order Logic, Graphs and Unfoldings of Transition Systems*. August 1995. 39 pp. To be presented at CSL '95.
- RS-95-43 Noam Nisan and Avi Wigderson. *Lower Bounds on Arithmetic Circuits via Partial Derivatives (Preliminary Version)*. August 1995. 17 pp. To appear in *36th Annual Conference on Foundations of Computer Science, FOCS '95*, IEEE, 1995.