

A Theory of Normed Simulations

W.O.D. Griffioen, F.W. Vaandrager

Computing Science Institute/

CSI-R0013 July 2000

Computing Science Institute Nijmegen
Faculty of Mathematics and Informatics
Catholic University of Nijmegen
Toernooiveld 1
6525 ED Nijmegen
The Netherlands

A Theory of Normed Simulations*

David Griffioen[†]

Frits Vaandrager[‡]

Computing Science Institute, University of Nijmegen
P.O. Box 9010, 9500 GL Nijmegen, The Netherlands

Abstract

In existing simulation proof techniques, a single step in a lower-level specification may be simulated by an extended execution fragment in a higher-level one. As a result, it is cumbersome to mechanize these techniques using general purpose theorem provers. Moreover, it is undecidable whether a given relation is a simulation, even if tautology checking is decidable for the underlying specification logic. This paper introduces various types of *normed simulations*. In a normed simulation, each step in a lower-level specification can be simulated by at most one step in the higher-level one, for any related pair of states. We show that, under some reasonable assumptions, it is decidable whether a given relation is a normed forward simulation, provided tautology checking is decidable for the underlying logic. We also prove that, at the semantic level, normed forward and backward simulations together form a complete proof method for establishing behavior inclusion, provided that the higher-level specification has finite invisible nondeterminism. Apart from their pleasant theoretical properties, normed simulations also have turned out to be quite useful as a vehicle for the formalization of refinement proofs via theorem provers.

AMS Subject Classification (1991): 68Q60, 68Q68.

CR Subject Classification (1991): F.1.1, F.3.1.

Keywords & Phrases: Computer aided verification, normed simulations, automata, refinement mappings, forward simulations, backward simulations, history variables, prophecy variables.

*A preliminary version of this paper appeared as Sections 1 and 2 in [GV98].

[†]Supported by the Netherlands Organization for Scientific Research (NWO) under contract SION 612-316-125. Current affiliation: Baan, Dept. BIS, Baron van Nagellstraat 89, 3771 LK Barneveld, The Netherlands, E-mail: dgriffio@baan.nl.

[‡]E-mail: fvaan@cs.kun.nl.

Contents

1	Introduction	3
2	Preliminaries	5
3	Step Refinements and Execution Correspondence	6
3.1	Step Refinements	6
3.2	Execution Correspondence	8
3.3	Soundness and Partial Completeness	9
4	Normed Forward Simulations	10
5	Normed Backward Simulations	15
6	Normed History Relations	18
7	Normed Prophecy Relations	21
8	Decidability	21
9	Reachability	25

1 Introduction

Simulation relations and refinement functions are widely used to prove that a lower-level specification of a reactive system correctly implements a higher-level one; see [AL91, LV95, Lyn96, RE98] for overviews and numerous references to applications. Technically, a *simulation* (or *refinement*) is a relation (or function) R between the states of a lower-level specification A and a higher-level specification B , that satisfies a condition like

$$(s, u) \in R \wedge s \xrightarrow{a}_A t \Rightarrow \exists v : u \xrightarrow{a}_B v \wedge (t, v) \in R \quad (1)$$

(If lower-level state s and higher-level state u are related, and in A there is a transition from s to t , then there is a matching transition in B from u to a state v that relates to t ; see also Figure 1.) The existence of a simulation implies that any behavior of A can also be exhibited by B .




Figure 1: Transfer condition (1).

The main reason why simulations are useful is that they reduce *global* reasoning about behaviors and executions to *local* reasoning about states and transitions. However, to the best of our knowledge, all complete simulation proof methods that appear in the literature fall back on some form of global reasoning in the case of specifications containing internal (or stuttering) transitions. The usual transfer condition for *forward simulations* [LV95], for instance, says

$$(s, u) \in R \wedge s \xrightarrow{a}_A t \Rightarrow \exists \text{ execution fragment } \alpha : \text{first}(\alpha) = u \quad (2) \\ \wedge \text{trace}(\alpha) = \text{trace}(a) \wedge (t, \text{last}(\alpha)) \in R$$

(Each lower-level transition can be simulated by a sequence of higher-level transitions which, apart from the action that has to be matched, may also contain an arbitrary number of internal “ τ ” transitions; see also Figure 2.) Thus the research program to reduce global reasoning to local reasoning has not been carried out to its completion. In manual proofs of simulation relations,




Figure 2: Transfer condition (2).

this is usually not a problem: in practice lower-level transitions are typically matched by at most

one higher-level transition; moreover humans tend to be quite good in reasoning about sequences, and move effortlessly from transitions to executions and back. In contrast, it turns out to be rather cumbersome to formalize arguments involving sequences using existing theorem provers (see [DGM97] for a comparative study). In fact, in several papers in which formalizations of simulation proofs are described, the authors only consider a restricted type of simulation in which each lower-level transition is matched by at most one higher-level transition [HSV94, NS95, DGRV00]. However, there are many examples of situations where these restricted types of simulations cannot be applied. In approaches, such as [SAGG⁺93], where the full transfer condition (2) is formalized, the user has to supply the simulating execution fragments α to the prover explicitly, which makes the verification process highly interactive.

In this paper, we introduce a simulation proof method which remedies the above problems. The idea is to define a function n that assigns a norm $n(s \xrightarrow{a} t, u)$, in some well-founded domain, to each pair of a transition in A and a state of B . If u has to simulate transition $s \xrightarrow{a} t$ then it may either do nothing (if a is internal and t is related to u), or it may do a matching a -transition, or it may perform an internal transition $u \xrightarrow{b} v$ such that the norm decreases, i.e.,

$$n(s \xrightarrow{a} t, v) < n(s \xrightarrow{a} t, u).$$

We establish that *normed forward simulations* and *normed backward simulations* together constitute a complete proof method for establishing trace inclusion. In addition we show how *history* and *prophecy relations* (which are closely related to the history and prophecy variables of [AL91]) can be enriched with a norm function, to obtain another complete proof method in combination with a simple notion of refinement mapping.

The preorders generated by normed forward simulations are strictly finer than the preorders induced by the forward simulations of [LV95]. In fact, we will characterize normed forward simulations in terms of *branching forward simulations* [GW96], and present a similar characterization for the backward case. It is possible to come up with a variant of normed forward simulation that induces the same preorder as forward simulations, but technically this is somewhat more involved (see [Gri00][Section 6.5.10]).

When proving invariance properties of programs, one is faced with two problems. The first problem is related to the necessity of proving tautologies of the assertion logic, whereas the second manifests in the need of finding sufficiently strong invariants. In order to address the first problem, powerful decision procedures have been incorporated in theorem provers such as PVS [ORSH95]. If tautology checking is decidable then it is decidable whether a given state predicate is valid for the initial states and preserved by all transitions. The task of finding such a predicate, i.e. solving the second problem, is the responsibility of the user, even though some very powerful heuristics have been devised to automate this search [BLS96, MBSU98]. Analogously, if specifications A and B , a conjectured forward simulation relation R and norm function n can all be expressed within a decidable assertion logic, and if the specification of B only contains a finite number of deterministic transition predicates, then it is decidable whether the pair (R, n) is a normed forward simulation. This result, which does not hold for earlier approaches such as [LV95], is a distinct advantage of normed forward simulations.

The idea of using norm functions to prove simulation relations also occurs in [GS95], where it is used to prove branching bisimilarity in the context of the process algebra μCRL . However, in [GS95] the norm function is defined on the states of B only and does not involve the transitions of A . As a consequence, the method of [GS95] does not always apply to diverging processes. Norm functions very similar to ours were also studied by Namjoshi [Nam97]. He uses them to obtain a characterization of the stuttering bisimulation of [BCG88], which is the equivalent of branching bisimulation in a setting where states rather than actions are labeled (see [DNV95]). Both [GS95] and [Nam97] do not address effectiveness issues. Although we present normed simulations in a setting of labeled transition systems, it should not be difficult to transfer our results to a process algebraic setting such as [GS95] or a state based setting such as [Nam97]. Inspired by our approach, norm functions have been used in [BS00] to define a new bisimulation equivalence for probabilistic systems.

In this paper, we only present maximally simple examples to illustrate the various definitions and results. In [Gri00][Chapter 6], normed simulations have been applied in a substantial case study, namely the verification of the leader election protocol that is part of the IEEE 1394 “Firewire” standard. This verification has been mechanically checked using PVS.¹

In the presentation of our results, we will closely follow [LV95] and stick to the notations introduced in that paper. In fact, our aim will be (amongst others) to derive analogous results as in [LV95], only for different types of simulations. However, we decided not to present normed versions of the forward-backward and backward-forward simulations of [LV95], since these simulations have thus far not been used in practice and technically this would bring nothing new. Apart from the notion of a norm function, a major technical innovation in the present paper is a new, simple definition of execution correspondence [GSSL93, SALL93], and the systematic use of this concept in the technical development.

2 Preliminaries

In this section, we briefly recall some basic concurrency theory definitions, taken from [LV95]. An *automaton* (or *labeled transition system*) A consists of

- a (possibly infinite) set $states(A)$ of states,
- a nonempty set $start(A) \subseteq states(A)$ of start states,
- a set $acts(A)$ of actions that includes the internal (or stuttering) action τ , and
- a set $steps(A) \subseteq states(A) \times acts(A) \times states(A)$ of steps.

Write $s \xrightarrow{a}_A t$ as a shorthand for $(s, a, t) \in steps(A)$. We let $ext(A)$, the *external actions*, denote $acts(A) - \{\tau\}$. An *execution fragment* of A is a finite or infinite alternating sequence, $s_0 a_1 s_1 a_2 s_2 \dots$, of states and actions of A , beginning with a state, and if it is finite also ending with a state, such that for all $i > 0$, $s_{i-1} \xrightarrow{a_i} s_i$. An *execution* of A is an execution fragment that begins with a start state. We denote by $execs^*(A)$ and $execs(A)$ the sets of finite and all executions of A , respectively. A state s of A is *reachable* if s occurs as the last state in some finite execution α of A . In this case we write $reachable(A, s)$. Also, we write $reachable(A)$ for the set of reachable states of A .

The *trace* of an execution fragment α , notation $trace(\alpha)$, is the subsequence of non- τ actions occurring in α . A finite or infinite sequence β of external actions is a *trace* of A if A has an execution α with $\beta = trace(\alpha)$. Write $traces^*(A)$ and $traces(A)$ for the sets of finite and all traces of A , respectively. Write $A \leq_{*T} B$ if $traces^*(A) \subseteq traces^*(B)$, and $A \leq_T B$ if $traces(A) \subseteq traces(B)$.

Suppose A is an automaton, s and t are states of A , and β is a finite sequence over $ext(A)$. We say that (s, β, t) is a *move* of A , and write $s \xrightarrow{\beta}_A t$, or just $s \xrightarrow{\beta} t$ when A is clear, if A has a finite execution fragment α that starts in s , has trace β and ends in t .

Three restricted kinds of automata play an important role in this paper:

1. A is *deterministic* if $|start(A)| = 1$, and for any state s and any finite sequence β over $ext(A)$, there is at most one state t such that $s \xrightarrow{\beta} t$. A deterministic automaton is characterized uniquely by the properties that $|start(A)| = 1$, every τ -step is of the form (s, τ, s) for some s , and for each state s and each action a there is at most one state t such that $s \xrightarrow{a}_A t$.
2. A has *finite invisible nondeterminism (fin)* if $start(A)$ is finite, and for any state s and any finite sequence β over $ext(A)$, there are only finitely many states t such that $s \xrightarrow{\beta} t$.
3. A is a *forest* if, for each state s of A , there is exactly one execution that leads to s . A forest is characterized uniquely by the property that all states of A are reachable, start states have no incoming steps, and each of the other states has exactly one incoming step.

¹ Actually, we discovered the notion of a normed simulation while formalizing the correctness proof of this leader election protocol.

The relation $after(A)$ consists of the pairs (β, s) for which there is a finite execution of A with trace β and last state s :

$$after(A) \triangleq \{(\beta, s) \mid \exists \alpha \in execs^*(A) : trace(\alpha) = \beta \text{ and } last(\alpha) = s\}.$$

(Here $last$ denotes the function that returns the last element of a finite, nonempty sequence.) We also define $past(A)$ to be the inverse of $after(A)$, $past(A) \triangleq after(A)^{-1}$; this relates a state s of A to the traces of finite executions of A that lead to s .

The following elementary lemma from [LV95] states that for the restricted kinds of automata defined above, the relations $after$ and $past$ satisfy certain nice properties.

Lemma 2.1

1. If A is deterministic then $after(A)$ is a function from $traces^*(A)$ to $states(A)$.
2. If A has fin then $after(A)$ is image-finite, i.e., each trace in the domain of $after(A)$ is only related to finitely many states.
3. If A is a forest then $past(A)$ is a function from $states(A)$ to $traces^*(A)$.

3 Step Refinements and Execution Correspondence

In this section, we present *step refinements*, the simplest notion of simulation that we consider in this paper. In order to prove soundness of step refinements, we also introduce the auxiliary notion of *execution correspondence*. This notion plays a key role in this paper; the technical lemmas that we prove in this section will also be used repeatedly in subsequent sections.

3.1 Step Refinements

Let A and B be automata. A *step refinement* from A to B is a partial function r from $states(A)$ to $states(B)$ that satisfies the following two conditions:

1. If $s \in start(A)$ then $s \in domain(r)$ and $r(s) \in start(B)$.
2. If $s \xrightarrow{a}_A t \wedge s \in domain(r)$ then $t \in domain(r)$ and
 - (a) $r(s) = r(t) \wedge a = \tau$, or
 - (b) $r(s) \xrightarrow{a}_B r(t)$.

Note that, by a trivial inductive argument, the set of states for which r is defined contains all the reachable states of A (and is thus an *invariant* of this automaton). We write $A \leq_R B$ if there exists a step refinement from A to B .

As far as we know, the notion of step refinements was first proposed in [NS95]. However, if we insist on the presence of stuttering steps $s \xrightarrow{\tau} s$ for each state s (a common assumption in models of reactive systems) then clause (2a) in the above definition becomes superfluous and the notion of a step refinement reduces to that of a homomorphism between reachable subautomata [Gin68]. Step refinements are slightly more restrictive than the *possibility mappings* from [LT87] (called *weak refinements* in [LV95]). In the case of a possibility mapping each (reachable) step of A may be matched by a sequence of steps in B with the same trace. This means that in the above definition condition (2) is replaced by:

2. If $s \xrightarrow{a}_A t \wedge s \in domain(r)$ then $t \in domain(r)$ and B has an execution fragment α with $first(\alpha) = r(s)$, $trace(\alpha) = trace(a)$ and $last(\alpha) = r(t)$.




Figure 3: A step refinement.

Observe that, unlike step refinements, possibility mappings do not reduce global reasoning to local reasoning.

Example 3.1 Figure 3 illustrates the notion of a step refinement. Note that the τ -steps in A are not matched by any step in B . Also the c -step in A is not matched by any step in B : both source and target states of this step are outside the domain of the step refinement. This is allowed since both states are unreachable. Observe that there is no step refinement from B to A , but that there exists a possibility mapping from B to A .

Figure 4 gives another example. In this case there is a step refinement from A' to B' but not from B' to A' . There is not even a possibility mapping from B' to A' .




Figure 4: Another step refinement.

The following proposition states a basic sanity property of step refinements.

Proposition 3.2 \leq_R is a preorder (i.e., is transitive and reflexive).

Proof: The identity function from $states(A)$ to itself trivially is a step refinement from A to itself. Hence \leq_R is reflexive. Transitivity follows from the observation that if r is a step refinement from A to B and r' is a step refinement from B to C , then the function composition $r' \circ r$ is a step refinement from A to C . ■

3.2 Execution Correspondence

If there exists a step refinement from A to B then we can construct, for each execution fragment of A , a corresponding execution fragment of B with the same trace. The notion of ‘corresponding’ is formalized below.

Suppose A and B are automata, $R \subseteq \text{states}(A) \times \text{states}(B)$, and $\alpha = s_0 a_1 s_1 a_2 s_2 \dots$ and $\alpha' = u_0 b_1 u_1 b_2 u_2 \dots$ are execution fragments of A and B , respectively. Let $\text{index}(\alpha)$ and $\text{index}(\alpha')$ denote the index sets of α and α' . Then α and α' *correspond via R* and are *R -related*, notation $(\alpha, \alpha') \in R$, if there exists an *index relation* over R , i.e., a relation $I \subseteq \text{index}(\alpha) \times \text{index}(\alpha')$ such that (1) if two indices are related by I then the corresponding states are related by R ; (2) I is monotone; (3) each index of α is related to an index of α' and vice versa; (4) sides of “squares” always have the same label and sides of “triangles” are labeled with τ . Formally we require, for $i, i' \in \text{index}(\alpha)$ and $j, j' \in \text{index}(\alpha')$,

1. $(i, j) \in I \Rightarrow (s_i, u_j) \in R$
2. $(i, j) \in I \wedge (i', j') \in I \wedge i < i' \Rightarrow j \leq j'$
3. I and I^{-1} are total
4. $(i, j) \in I \wedge (i + 1, j + 1) \in I \Rightarrow a_{i+1} = b_{j+1}$
 $(i, j) \in I \wedge (i + 1, j) \in I \Rightarrow a_{i+1} = \tau$
 $(i, j) \in I \wedge (i, j + 1) \in I \Rightarrow b_{j+1} = \tau$

We write $(A, B) \in R$ if for every execution α of A there is an execution α' of B such that $(\alpha, \alpha') \in R$, and $[A, B] \in R$ if for every finite execution α of A there is a finite execution α' of B with $(\alpha, \alpha') \in R$. Figure 5 illustrates the correspondence between two executions of automata A and B from Figure 3.




Figure 5: Execution correspondence.

Another notion of correspondence has been presented in [GSSL93, SALL93] and formalized in [Mue98]. Within the theory of I/O automata, execution correspondence plays a crucial role in proofs of preservation of both safety and liveness properties. Our notion is more restrictive than the one of [GSSL93, SALL93], but technically simpler. Moreover it has the advantage that it preserves until properties. In this paper, we only study safety properties and it suffices to know that corresponding executions have the same trace. The latter fact is established in the next lemma.

Lemma 3.3 (*Corresponding execution fragments have the same trace*)

1. Suppose I is an index relation as above and $(i, j) \in I$. Then $\text{trace}(s_0 a_1 s_1 \dots a_i s_i) = \text{trace}(u_0 b_1 u_1 \dots b_j u_j)$.
2. If $(\alpha, \alpha') \in R$ then $\text{trace}(\alpha) = \text{trace}(\alpha')$.

Proof: For (1), suppose $(i, j) \in I$. By induction on $i + j$ we prove

$$\text{trace}(s_0 a_1 s_1 \cdots a_i s_i) = \text{trace}(u_0 b_1 u_1 \cdots b_j u_j).$$

If $i + j = 0$ then both i and j are 0. Clearly, $\text{trace}(s_0) = \text{trace}(u_0) = \lambda$.

For the induction step, suppose $i + j > 0$. For reasons of symmetry we may assume w.l.o.g. that $i > 0$. Let j' be the largest index with $j' \leq j$ and $(i - 1, j') \in I$. (By monotonicity, $i - 1$ can only be related to indices less than or equal to j , and by totality there is at least one such an index.) We distinguish between three cases:

1. $j' = j$. Then by condition (4b), $a_i = \tau$. By induction hypothesis,

$$\text{trace}(s_0 a_1 s_1 \cdots a_{i-1} s_{i-1}) = \text{trace}(u_0 b_1 u_1 \cdots b_j u_j).$$

$$\text{Hence } \text{trace}(s_0 a_1 s_1 \cdots a_i s_i) = \text{trace}(u_0 b_1 u_1 \cdots b_j u_j).$$

2. $j' = j - 1$. Then by condition (4a), $a_i = b_j$. By induction hypothesis,

$$\text{trace}(s_0 a_1 s_1 \cdots a_{i-1} s_{i-1}) = \text{trace}(u_0 b_1 u_1 \cdots b_{j-1} u_{j-1}).$$

$$\text{Hence } \text{trace}(s_0 a_1 s_1 \cdots a_i s_i) = \text{trace}(u_0 b_1 u_1 \cdots b_j u_j).$$

3. $j' < j - 1$. Then by conditions (2) and (3), $(i, j - 1) \in I$. By condition (4c), this implies $b_j = \tau$. By induction hypothesis,

$$\text{trace}(s_0 a_1 s_1 \cdots a_i s_i) = \text{trace}(u_0 b_1 u_1 \cdots b_{j-1} u_{j-1}).$$

$$\text{Hence } \text{trace}(s_0 a_1 s_1 \cdots a_i s_i) = \text{trace}(u_0 b_1 u_1 \cdots b_j u_j).$$

This completes the proof of the induction step.

For (2), suppose that $(\alpha, \alpha') \in R$. Then there exists an index relation I that relates α and α' . Using (1) and the fact that both I and I^{-1} are total, it follows that each finite prefix of $\text{trace}(\alpha)$ is also a finite prefix of $\text{trace}(\alpha')$, and vice versa. This implies $\text{trace}(\alpha) = \text{trace}(\alpha')$. \blacksquare

The next corollary will be used repeatedly in the rest of this paper. It states that in order to prove trace inclusion between automata A and B it suffices to find for each execution of A a corresponding execution of B . Depending on whether one wants to prove inclusion of all traces or of finite traces only, a stronger resp. weaker type of execution correspondence is required.

Corollary 3.4 (*Execution correspondence implies trace inclusion*)

1. If $(A, B) \in R$ then $[A, B] \in R$.
2. If $[A, B] \in R$ then $A \leq_{*T} B$.
3. If $(A, B) \in R$ then $A \leq_T B$.

Proof: Statement (1) follows from the definitions. Statements (2) and (3) follow immediately from Lemma 3.3 and the definitions. \blacksquare

3.3 Soundness and Partial Completeness

The next theorem states that if there is a step refinement from A to B , it is possible to construct, for each execution of A , a corresponding execution of B . Using Corollary 3.4, this implies that step refinements constitute a sound technique for proving trace inclusion. In addition, the next theorem also allows us to use step refinements as a sound technique for proving implementation relations between live automata, as in [GSSL93, SALL93, Mue98].

Theorem 3.5 (*Soundness of step refinements*)

If r is a step refinement from A to B then $(A, B) \in r$.

Proof: Suppose r is a step refinement from A to B . Let $\alpha = s_0 a_1 s_1 \dots$ be an execution of A . Inductively, we define an execution $\alpha' = u_0 b_1 u_1 \dots$ of B and an index relation I such that α and α' are r -related via I .

To start with, define $u_0 = r(s_0)$ and declare $(0, 0)$ to be an element of I .

Now suppose $(i, j) \in I$ and i is a nonfinal index of α . We distinguish between two cases:

1. If $r(s_i) \xrightarrow{a_{i+1}}_B r(s_{i+1})$ then define $b_{j+1} = a_{i+1}$, $u_{j+1} = r(s_{i+1})$, and declare $(i+1, j+1)$ to be an element of I ;
2. otherwise, declare $(i+1, j)$ to be an element of I .

By construction, using the defining properties of a step refinement, it follows that I is an index relation. This implies $(A, B) \in r$. ■

Step refinements alone do not provide a complete method for proving trace inclusion. There is a partial completeness result, however.

Theorem 3.6 (*Partial completeness of step refinements*)

Suppose A is a forest, B is deterministic and $A \leq_{*T} B$. Then $A \leq_R B$.

Proof: The relation $r \triangleq \text{after}(B) \circ \text{past}(A)$ is a step refinement from A to B . ■

Actually, we can even slightly strengthen the above theorem. It suffices to assume that A restricted to its reachable states is a forest, and that B restricted to its reachable states is deterministic. In Figure 3, automaton A restricted to its reachable states is a forest and automaton B is deterministic. As we observed already, there is a step refinement from A to B . Even if we restrict to reachable states, automaton B is not a forest and automaton A is not deterministic. As we observed, there is no step refinement from B to A .

In practice, the preconditions of Theorem 3.6 are seldom met. The higher-level specification often is deterministic, but it rarely occurs that the lower-level specification is a forest. Nevertheless, step refinements have been used in several substantial case studies, see for instance [HSV94, NS95, DGRV00].

4 Normed Forward Simulations

Even though there exists no step refinement from automaton B' to automaton A' in Figure 4, these automata do have the same traces. By moving from functions to relations it becomes possible to prove that each trace of B' is also a trace of A' . This idea is formalized in the following definition.

A *normed forward simulation* from A to B consists of a relation $f \subseteq \text{states}(A) \times \text{states}(B)$ and a function $n : \text{steps}(A) \times \text{states}(B) \rightarrow S$, for some well-founded set S , such that (here $f[s]$ denotes the set $\{u \mid (s, u) \in f\}$):

1. If $s \in \text{start}(A)$ then $f[s] \cap \text{start}(B) \neq \emptyset$.
2. If $s \xrightarrow{a}_A t \wedge u \in f[s]$ then
 - (a) $u \in f[t] \wedge a = \tau$, or
 - (b) $\exists v \in f[t] : u \xrightarrow{a}_B v$, or
 - (c) $\exists v \in f[s] : u \xrightarrow{\tau}_B v \wedge n(s \xrightarrow{a}_B t, v) < n(s \xrightarrow{a}_B t, u)$.

Write $A \leq_F B$ if there exists a normed forward simulation from A to B .

The intuition behind this definition is that if $s \xrightarrow{a}_A t$ and $(s, u) \in f$, then either (a) the transition in A is a stuttering step that does not have to be matched, or (b) there is a matching step in B , or (c) B can do a stuttering step which decreases the norm. Since the norm decreases at

each application of clause (c), this clause can only be applied a finite number of times. In general, the norm function may depend both on the transitions in A and on the states of B . However, if B is *convergent*, i.e., there are no infinite τ -paths, then one can simplify the type of the norm function (though not necessarily the definition of the norm function itself) to $n : \text{states}(B) \rightarrow S$. In fact, in the approach of [GS95], which not always applies to divergent processes, the norm function is required to be of this restricted type.

Example 4.1 In Figure 4, the relation indicated by the dashed lines, together with an arbitrary norm function, is a normed forward simulation from B' to A' .

Consider automata A and B in Figure 3. Let n be the function that assigns norm 1 to state $s0$ and norm 0 to all other states of A . Then n together with the relation indicated by the dashed lines constitutes a normed forward simulation from B to A .




Figure 6: Norm function must take steps of C into account.

Now consider the automata C and D in Figure 6. Let m be a norm function satisfying

$$\begin{aligned} m(s0 \xrightarrow{a} s1, u0) &= 0 & m(s0 \xrightarrow{a} s1, u1) &= 1 \\ m(s0 \xrightarrow{b} s3, u0) &= 1 & m(s0 \xrightarrow{b} s3, u1) &= 0 \end{aligned}$$

Then m together with the relation indicated by the dashed lines constitutes a normed forward simulation from C to D . It is not hard to see that in this example, where D is not convergent, the norm necessarily depends on the selected step in C .

The next proposition asserts that normed forward simulations indeed generalize step refinements.

Proposition 4.2 $A \leq_R B \Rightarrow A \leq_F B$.

Proof: Together with an arbitrary norm function, any step refinement (viewed as a relation) is a normed forward simulation. ■

The soundness of normed forward simulations is trivially implied by the following lemma and Corollary 3.4.

Lemma 4.3 Suppose (f, n) is a normed forward simulation from A to B , A has an execution fragment α with first state s , and u is a state of B with $u \in f[s]$. Then B has an execution fragment α' that starts in u such that $(\alpha, \alpha') \in f$.

Proof: Let $c : \text{steps}(A) \times \text{states}(B) \rightarrow \{L, C, R\} \times \text{states}(B)$ be a function such that $c(s \xrightarrow{a} t, u) = (x, v)$ and $u \in f[s]$ implies

1. If $x = L$ then $u \in f[t] \wedge a = \tau$.

2. If $x = C$ then $v \in f[t] \wedge u \xrightarrow{a}_B v$.
3. If $x = R$ then $v \in f[s] \wedge u \xrightarrow{\tau}_B v \wedge n(s \xrightarrow{a} t, v) < n(s \xrightarrow{a} t, u)$.

The existence of c , which chooses between a left move (L) of A , a common move (C) of A and B , or a right move (R) of B , is guaranteed by the fact that (f, n) is a normed forward simulation.

Let $\alpha = s_0 a_1 s_1 a_2 s_2 \dots$. Then $s = s_0$. Inductively, we define a sequence $\sigma = z_0 z_1 z_2 \dots$ of 4-tuples in $\mathbb{N} \times \mathbb{N} \times \text{acts}(B) \times \text{states}(B)$. The first element in the sequence is $z_0 = (0, 0, \tau, u)$. If $z_k = (i, j, b, u)$ is an element of the sequence, and i is a nonfinal index of α , then we define z_{k+1} as follows

1. If $c(s_i \xrightarrow{a_{i+1}} s_{i+1}, u) = (L, v)$ then $z_{k+1} = (i + 1, j, b, u)$.
2. If $c(s_i \xrightarrow{a_{i+1}} s_{i+1}, u) = (C, v)$ then $z_{k+1} = (i + 1, j + 1, a_{i+1}, v)$.
3. If $c(s_i \xrightarrow{a_{i+1}} s_{i+1}, u) = (R, v)$ then $z_{k+1} = (i, j + 1, \tau, v)$.

Suppose that both (i, j, b, u) and (i', j, b', u') occur in sequence σ . We claim that $b = b'$ and $u = u'$. To see why this is true assume w.l.o.g. that (i, j, b, u) occurs before (i', j, b', u') . Now observe that the values of both the first and second component of elements in σ increase monotonically. This means that each successor of (i, j, b, u) up to and including (i', j, b', u') has been obtained from its predecessor by applying rule (1). This implies that the the second resp. third components of all elements in the sequence from (i, j, b, u) until (i', j, b', u') coincide. Hence $b = b'$ and $u = u'$.

Using this property, we can define for each element (i, j, b, u) in σ , $b_j = b$ and $u_j = u$. Let $\alpha' = u_0 b_1 u_1 b_2 u_2 \dots$ and let $I = \{(i, j) \mid \exists b, u : (i, j, b, u) \text{ occurs in } \sigma\}$. By construction of σ , using the properties of c , it follows that α' is an execution fragment of B that starts in u , and that I is an index relation over f . This implies $(\alpha, \alpha') \in f$. \blacksquare

Theorem 4.4 (*Soundness of normed forward simulations*)

If f is a normed forward simulation from A to B then $(A, B) \in f$.

Proof: Immediate from the definitions and Lemma 4.3. \blacksquare

Example 4.5 Consider automata C and E in Figure 7. There does not exist a normed forward




Figure 7: Difference between forward simulations and normed forward simulations.

simulation from C to E . Such a simulation would have to relate states s_0 and u_0 . But in order for E to simulate the step $s_0 \xrightarrow{b} s_3$, it would also have to relate states s_0 and u_2 . But this is impossible since from state u_2 there is no way to simulate the step $s_0 \xrightarrow{a} s_1$.

It turns out that there does exist a *forward simulation* in the sense of [LV95] from C to E . In the case of a forward simulation, a step of A may be matched by a sequence of steps in B with

the same trace. This means that in the definition of a normed forward simulation condition (2) is replaced by:

2. If $s \xrightarrow{a}_A t \wedge u \in f[s]$ then B has an execution fragment α with $first(\alpha) = u$, $trace(\alpha) = trace(a)$ and $last(\alpha) \in f[t]$.

The dashed lines in Figure 7 indicate a forward simulation from C to E .

The automata A and B in Figure 3 provide us with a similar example: there exists a forward simulation from B to A , but no normed forward simulation.

The difference between forward simulations and normed forward simulations is very similar to the difference between Milner's *observation equivalence* [Mil89] and the *branching bisimulation* of Van Glabbeek and Weijland [GW96]. In fact, we can characterize normed forward simulations in terms of "branching forward simulations", a notion that is inspired by the branching bisimulations of [GW96]. A similar characterization has been obtained by Namjoshi [Nam97] in the setting of stuttering bisimulations.

Formally, a *branching forward simulation* from A to B is a relation $f \subseteq states(A) \times states(B)$ such that

1. If $s \in start(A)$ then $f[s] \cap start(B) \neq \emptyset$.
2. If $s \xrightarrow{a}_A t$ and $u \in f[s]$ then B has an execution fragment that starts in u and that is f -related to $s \xrightarrow{a} t$.

The following theorem implies that there exists a normed forward simulation between two automata if and only if there is a branching forward simulation between them.

Theorem 4.6

1. Suppose (f, n) is a normed forward simulation from A to B . Then f is a branching forward simulation from A to B .
2. Suppose f is a branching forward simulation from A to B . Let $n(s \xrightarrow{a} t, u)$ be 0 if $u \notin f[s]$ and otherwise be equal to the length of the shortest execution fragment that starts in u and that is f -related to $s \xrightarrow{a} t$. Then (f, n) is a normed forward simulation from A to B .

Proof: Part (1) follows by Lemma 4.3. The proof of part (2) is routine. ■

An interesting corollary of Theorem 4.6 is that if there exists a normed forward simulation between two automata, there is in fact a normed forward simulation with a norm that has the natural numbers as its range.

The proof that branching bisimilarity is an equivalence is known to be tricky [Bas96]. Likewise, the proof that branching forward simulations induce a preorder is nontrivial. We first need to define the auxiliary concept of a *reduced* index relation and to prove a lemma about it.

Suppose that α and α' are R -related via index relation I . We say that I is *reduced* if the following two conditions are satisfied:

1. If α is finite then I relates the final index of α only to the final index of α' .
2. I is N -free: $(i, j) \in I \wedge (i + 1, j + 1) \in I \Rightarrow (i + 1, j) \notin I \wedge (i, j + 1) \notin I$.

Observe that if α is finite and I is reduced, then α' is also finite. The following technical lemma states that index relations can always be reduced.

Lemma 4.7 *Suppose that α and α' are R -related via index relation I . Then α' has a prefix α'' that is R -related to α via a reduced index relation $J \subseteq I$.*

Proof: If α is infinite then let $\alpha'' = \alpha'$. If α is finite then let α'' be the finite prefix of α' up to and including the first state whose index is related by I to the final index of α .

Inductively we define a sequence $\sigma = z_0 z_1 z_2 \dots$ of pairs in $\mathbb{N} \times \mathbb{N}$. The first element of the sequence is $z_0 = (0, 0)$. If $z_k = (i, j)$ is an element of the sequence and i is a nonfinal index then we define z_{k+1} as follows:

1. $(i + 1, j + 1) \in I \Rightarrow z_{k+1} = (i + 1, j + 1)$
2. $(i + 1, j) \in I \wedge (i + 1, j + 1) \notin I \Rightarrow z_{k+1} = (i + 1, j)$
3. $(i, j + 1) \in I \wedge (i + 1, j + 1) \notin I \Rightarrow z_{k+1} = (i, j + 1)$

Note that since I is an index relation, z_{k+1} is properly defined. Let $J = \{(i, j) \mid (i, j) \text{ occurs in } \sigma\}$. It is routine to check that $J \subseteq I$, that α and α'' are R -related via J , and that J is reduced. A tricky point is the totality of J and J^{-1} . We prove that J is total by contradiction. Suppose that J is not total. Let i be the smallest index of α with $J[i] = \emptyset$. Let j be the smallest index of α' with $(i, j) \in I$ (j exists since index relation I is total). Let l be the maximal index of α' with $(i - 1, l) \in J$ (there is a maximal index since $(i - 1, l) \in J$ implies $(i - 1, l) \in I$, which implies $l \leq j$ by monotonicity of index relation I). Let $z_k = (i - 1, l)$. Since $J[i] = \emptyset$, $z_{k+1} = (i - 1, l + 1)$. Hence $(i - 1, l + 1) \in J$. But this contradicts the fact that l be the maximal index of α' with $(i - 1, l) \in J$.

In a similar way also the totality of J^{-1} and \mathbb{N} -freeness can be proved by contradiction. ■

We are now prepared to prove that branching forward simulations (and hence also normed forward simulations) induce a preorder.

Proposition 4.8 \leq_F is a preorder.

Proof: For reflexivity, observe that the identity function from $\text{states}(A)$ to itself is a branching forward simulation from A to itself.

For transitivity, suppose f and g are branching forward simulations from A to B and from B to C , respectively. We claim that $g \circ f$ is a branching forward simulation from A to C . It is trivial to check that $g \circ f$ satisfies condition (1) in the definition of a branching forward simulation. For condition (2), suppose that $s \xrightarrow{a}_A t \wedge u \in (g \circ f)[s]$. Then there exists a state w of B such that $w \in f[s]$ and $u \in g[w]$. Hence there is an execution fragment α starting in w such that $s \xrightarrow{a} t$ and α are f -related via some index relation I . By Lemma 4.7, we may assume that I is reduced. Also, there is an execution fragment α' starting in u such that α and α' are g -related via some index relation J . Again by Lemma 4.7, we may assume that J is reduced. Using the fact that both I and J are reduced, it is routine to check that $s \xrightarrow{a} t$ and α' are $g \circ f$ -related via index relation $J \circ I$. Thus $g \circ f$ satisfies condition (2) in the definition of a branching forward simulation. ■

Variants of the partial completeness result below appear in several papers, see for instance [Jon87, LV95]. Since higher-level specifications are often deterministic, this result explains why in practice (normed) forward simulations can so often be used to prove behavior inclusion.

Theorem 4.9 (Partial completeness of normed/branching forward simulations)

If B is deterministic and $A \leq_{*T} B$ then $A \leq_F B$.

Proof: The relation $f \triangleq \text{after}(B) \circ \text{past}(A)$ is a branching forward simulation from A to B . ■

It is interesting to note that there is one result from [LV95] concerning forward simulations that does not carry over to the normed/branching simulations of this paper. This result, Proposition 3.12, states that if A is a forest and $A \leq_F B$ then $A \leq_R B$. The automata C and D of Figure 6 constitute a counterexample. Actually, the same Proposition 3.12 also does not carry over to the setting of timed automata studied in [LV96].

5 Normed Backward Simulations

As we observed, there exists no normed forward simulation from automaton B to automaton A in Figure 3, even though both automata have the same traces. Also, there does not exist a normed forward simulation from automaton C to the trace equivalent automaton E in Figure 7. In both cases a forward simulation in the sense of [LV95] exists. However, the example in Figure 8 below shows that also forward simulations do not yet provide us with a complete method for proving trace inclusion. It is well-known from the literature that completeness can be obtained by adding some form of *backward simulation*.

Example 5.1 There exists no (normed/branching) forward simulation from automaton C to automaton F in Figure 8. The relation indicated by the dashed lines fails since from state $u0$ the b -step from $s0$ can not be simulated, whereas from $u2$ the a -step from $s0$ can not be simulated.




Figure 8: The need for backward simulations.

In many respects, backward simulations are the dual of forward simulations. Whereas a forward simulation requires that *some* state in the image of each start state should be a start state, a backward simulation essentially requires that *all* states in the image of a start state be start states. Also, a forward simulation requires that *forward* steps in the source automaton can be simulated from related states in the target automaton, whereas the corresponding condition for a backward simulation requires that *backward* steps can be simulated. However, the two notions are not completely dual: the definition of a backward simulation contains a nonemptiness condition, and also, in order to obtain soundness for general trace inclusion, backward simulations also require a finite image condition. The mismatch is due to the asymmetry in our automata between the future and the past: from any given state, all the possible histories are finite executions, whereas the possible futures can be infinite.

Formally, we define a *normed backward simulation* from A to B to be a pair of a total relation $b \subseteq \text{states}(A) \times \text{states}(B)$ and a function $n : (\text{steps}(A) \cup \text{start}(A)) \times \text{states}(B) \rightarrow S$, for some well-founded set S , satisfying

1. If $s \in \text{start}(A) \wedge u \in b[s]$ then
 - (a) $u \in \text{start}(B)$, or
 - (b) $\exists v \in b[s] : v \xrightarrow{\tau}_B u \wedge n(s, v) < n(s, u)$.
2. If $t \xrightarrow{a}_A s \wedge u \in b[s]$ then
 - (a) $u \in b[t] \wedge a = \tau$, or
 - (b) $\exists v \in b[t] : v \xrightarrow{a}_B u$, or
 - (c) $\exists v \in b[s] : v \xrightarrow{\tau}_B u \wedge n(t \xrightarrow{a} s, v) < n(t \xrightarrow{a} s, u)$.

Write $A \leq_B B$ if there is a normed backward simulation from A to B , and $A \leq_{iB} B$ if there is a normed backward simulation from A to B that is image-finite.

Example 5.2 In Figure 8, the relation indicated by the dashed lines is a normed backward simulation from C to E , for arbitrary norm functions. It is not difficult to construct normed backward simulations from automaton B to automaton A in Figure 3, and from automaton C to automaton E in Figure 7.




Figure 9: No image-finite normed backward simulation.

Figure 9 illustrates the difference between \leq_B and \leq_{iB} . Relation $states(G) \times states(H)$ together with an arbitrary norm function constitutes a normed backward simulation from G to H . We claim that no image-finite normed backward simulation exist. Because suppose that b is such a relation. Then, for all $i, j \in \mathbb{N}$ with $i > 0$,

$$(s_i, u_j) \in b \Rightarrow (s_{i-1}, u_{j+1}) \in b$$

This implies that

$$(s_i, u_j) \in b \Rightarrow (s_0, u_{i+j}) \in b$$

Since each state s_i is related to at least one state u_j , it follows that state s_0 is related to infinitely many states, which is a contradiction.

The following proposition states some trivial connections between the preorders induced by normed backward simulations and step refinements.

Proposition 5.3

1. If all states of A are reachable and $A \leq_R B$ then $A \leq_{iB} B$.
2. If $A \leq_{iB} B$ then $A \leq_B B$.

Proof: Trivial. ■

The next lemma is required to prove soundness of normed backward simulations.

Lemma 5.4 Suppose (b, n) is a normed backward simulation from A to B , A has a finite execution fragment α with last state s , and u is a state of B with $u \in b[s]$. Then B has a finite execution fragment α' that ends in u such that $(\alpha, \alpha') \in b$. Moreover, if α is an execution then α' can be chosen to be an execution as well.

Proof: Similar to the proof of Lemma 4.3. ■

By Lemma 5.4 and Corollary 3.4, the existence of a normed backward simulation implies inclusion of finite traces. Normed backward simulations, however, are in general not a sound method for proving inclusion of infinite traces. As a counterexample, consider automata G and

H from Figure 9. There exists a normed backward simulation from G to H , but the infinite trace a^ω of G is not a trace of H . As is well-known from the literature, a sound method for proving inclusion of infinite traces can be obtained by requiring image finiteness of the simulation relation.

Theorem 5.5 (*Soundness of normed backward simulations*)

1. If b is a normed backward simulation from A to B then $[A, B] \in b$.
2. If moreover b is image-finite then $(A, B) \in b$.

Proof: Statement (1) follows immediately by Lemma 5.4 and the totality of b . In order to prove (2), suppose that b is image-finite. Let α be an execution of A . We have to establish the existence of an execution α' of B with $(\alpha, \alpha') \in b$. If α is finite then this follows by Lemma 5.4 and the totality of b . So assume that α is infinite. We use a minor variation of König's Lemma [Knu97] that occurs in [LV95]:

Let G be an infinite digraph such that (1) G has finitely many roots, i.e., nodes without incoming edges, (2) each node of G has finite outdegree, and (3) each node of G is reachable from some root. Then there is an infinite path in G starting from some root.

The nodes of the graph G that we consider are pairs (I, γ) where γ is a finite execution of B and I is an index relation that relates γ to some finite prefix of α . There is an edge from a node (I, γ) to a node (I', γ') iff γ is a prefix of γ' and I' extends I with precisely one element. It is straightforward to check that G satisfies the conditions of König's Lemma. Hence G has an infinite path. Let J be the union of all the index relations occurring on nodes in this path, and let α' be the limit of the finite executions of the nodes in this path. Observe that, by image-finiteness of b , each index of α occurs in the domain of J . Hence $(\alpha, \alpha') \in b$. ■

The following Proposition 5.6 is in a sense the converse of Proposition 5.3. The proof is similar to that of the corresponding result in [LV95].

Proposition 5.6

1. If B is deterministic and $A \leq_B B$ then $A \leq_R B$.
2. If all states of A are reachable, B has fin and $A \leq_B B$, then $A \leq_{iB} B$.

Proof: For (1), suppose that B is deterministic and that b is a normed backward simulation from A to B . Suppose that s is a reachable state of A . We will prove that $b[s]$ contains exactly one element. Since any normed backward simulation that is functional on the reachable states trivially induces a step refinement, this gives us $A \leq_R B$.

Because b is a normed backward simulation it is a total relation, so we know $b[s]$ contains at least one element. Suppose that both $u_1 \in b[s]$ and $u_2 \in b[s]$; we prove $u_1 = u_2$. Since s is reachable, A has an execution α that ends in s . By Lemma 5.4, B has executions α_1 and α_2 which end in u_1 and u_2 , respectively, such that $(\alpha, \alpha_1) \in b$ and $(\alpha, \alpha_2) \in b$. By Lemma 3.3, $\text{trace}(\alpha) = \text{trace}(\alpha_1) = \text{trace}(\alpha_2)$. Now $u_1 = u_2$ follows by Lemma 2.1(1), using the fact the B is deterministic.

For (2), suppose that all states of A are reachable, B has fin, and b is a normed backward simulation from A to B . Suppose that s is a state of A . Since s is reachable, there is an execution α that ends in s . Let β be trace of α . By Lemma 5.4 there exists, for each $u \in b[s]$, an execution α_u of B that ends in u such that $(\alpha, \alpha_u) \in b$. By Lemma 3.3, $\text{trace}(\alpha_u) = \beta$. Hence $b[s] \subseteq \text{after}(B)[\beta]$. But since B has fin, $\text{after}(B)[\beta]$ is finite by Lemma 2.1(2). Hence b is image-finite. ■

As in the forward case, we will now characterize normed backward simulations in terms of “branching backward simulations”, and use this characterization to establish that \leq_B and \leq_{iB} are preorders.

A *branching backward simulation* from A to B is a total relation $b \subseteq \text{states}(A) \times \text{states}(B)$ such that

1. If $s \in \text{start}(A)$ and $u \in b[s]$ then B has an execution that ends in u and is b -related to s .
2. If $t \xrightarrow{a}_A s$ and $u \in f[s]$ then B has an execution fragment that ends in u and is b -related to $t \xrightarrow{a} s$.

Theorem 5.7

1. Suppose (b, n) is a normed backward simulation from A to B . Then b is a branching backward simulation from A to B .
2. Suppose b is a branching backward simulation from A to B . Let $n(s, u)$ be 0 if s is not a start state or $u \notin b[s]$ and otherwise be equal to the length of the shortest execution that ends in u and is b -related to s . Furthermore, let $n(t \xrightarrow{a} s, u)$ be 0 if $u \notin f[s]$ and otherwise equal to the length of the shortest execution fragment ending in u that is b -related to $t \xrightarrow{a}_A s$. Then (b, n) is a normed forward simulation from A to B .

Proof: Statement (1) follows by Lemma 5.4. The proof of statement (2) is routine. ■

As in the forward case, we see that if there exists a normed backward simulation between two automata, there is in fact a normed backward simulation with a norm that has the natural numbers as its range.

Proposition 5.8 \leq_B and \leq_{iB} are preorders.

Proof: Similar to the proof of Proposition 4.8. ■

The following partial completeness result is a variation of results of [Jon90, LV95].

Theorem 5.9 (*Partial completeness of normed backward simulations*)

If A is a forest and $A \leq_{*T} B$ then $A \leq_B B$.

Proof: The relation $b \triangleq \text{after}(B) \circ \text{past}(A)$ is a branching backward simulation from A to B . ■

Note that by Proposition 5.6 we can strengthen the conclusion of Theorem 5.9 to $A \leq_{iB} B$ in case B has finite invisible nondeterminism.

Example 5.10 Consider the automata A' and B' in Figure 4. There exists no normed backward simulation from B' to A' . The relation indicated by the dashed lines fails since the backward transition from state $u0$ cannot be simulated from the related state $s0$. By consequence, also normed backward simulations do not provide a complete proof method for establishing trace inclusion. In the next section, we will see that completeness can be obtained by combining normed forward and backward simulations.

6 Normed History Relations

In this section we define *normed history relations*. These provide an abstract view of the *history variables* of Abadi and Lamport [AL91], which in turn are abstractions of the *auxiliary variables* of Owicki and Gries [OG76].

A pair (r, n) is a *normed history relation* from A to B if r is a step refinement from B to A , and (r^{-1}, n) is a normed forward simulation from A to B . Write $A \leq_H B$ if there exists a normed history relation from A to B .

Clearly $A \leq_H B$ implies $A \leq_F B$ and $B \leq_R A$. Through these implications, the preorder and soundness results for normed forward simulations and step refinements carry over to normed history relations. In fact, if (r, n) is a normed history relation from A to B then r is just a functional *branching bisimulation* from B to A in the sense of Van Glabbeek and Weijland [GW96]. Hence, history relations preserve behavior of automata in a very strong sense. Intuitively, there is a

history relation from A to B if B can be obtained from A by adding an extra state variable that records information about the history of an execution.

Example 6.1 Consider again the automata A' and B' in Figure 4. Together with an arbitrary norm function, the dashed lines constitute a normed history relation from B' to A' . Because, as we observed, there is no step refinement from B' to A' , there exists no normed history relation from A' to B' .

An important example of a history relation is provided by the “unfolding” construction. The *unfolding* of an automaton A , notation $unfold(A)$, is the automaton obtained from A by recording the complete history of an execution. Formally, $unfold(A)$ is the automaton B defined by

- $states(B) = execs^*(A)$,
- $start(B) =$ the set of executions of A that consist of a single start state,
- $acts(B) = acts(A)$, and
- for $\alpha', \alpha \in states(B)$ and $a \in acts(B)$, $\alpha' \xrightarrow{a}_B \alpha \Leftrightarrow \alpha = \alpha' a \text{ last}(\alpha)$.

The next proposition relates an automaton to its unfolding.

Proposition 6.2 $unfold(A)$ is a forest and $A \leq_H unfold(A)$.

Proof: Clearly, $unfold(A)$ is a forest. The function $last$ which maps each finite execution of A to its last state is a step refinement from $unfold(A)$ to A , and the relation $last^{-1}$, together with an arbitrary norm function, is a normed forward simulation from A to $unfold(A)$. ■

The following completeness theorem, a variation of a result due to Sistla [Sis91], asserts that normed history relations together with normed backward simulations constitute a complete proof method for establishing trace inclusion. By consequence, also normed forward simulations together with normed backward simulations constitute a complete proof method.

Theorem 6.3 (*Completeness of normed history relations and normed backward simulations*)
If $A \leq_{*T} B$ then there exists an automaton C such that $A \leq_H C \leq_B B$.

Proof: Take $C = unfold(A)$. By Proposition 6.2, C is a forest and $A \leq_H C$. Since $A \leq_{*T} B$, also $C \leq_{*T} B$ by soundness of history relations. Next apply the partial completeness result for backward simulations (Theorem 5.9) to conclude $C \leq_B B$. ■

Observe that if we can assume in addition that B has fin, we may replace \leq_B by \leq_{iB} in the conclusion using Proposition 5.6.

Normed forward simulations are equivalent to normed history variables combined with step refinements: whenever there is a normed forward simulation from A to B , we can find an intermediate automaton C such that there is a normed history relation from A to C and a step refinement from C to B . The converse implication trivially holds since normed history relations and step refinements are special cases of normed forward simulations. In order to prove the existence of automaton C , we need to define a notion of “superposition” of automata and to prove a technical lemma.

Let $R \subseteq states(A) \times states(B)$ be a relation with $R \cap (start(A) \times start(B)) \neq \emptyset$. The *superposition* $sup(A, B, R)$ of A and B via R is the automaton C defined by

- $states(C) = R$,
- $start(C) = R \cap (start(A) \times start(B))$,
- $acts(C) = acts(A) \cap acts(B)$, and

- for $(s, u), (t, v) \in \text{states}(C)$ and $a \in \text{acts}(C)$, $(s, u) \xrightarrow{a}_C (t, v) \Leftrightarrow$

$$\begin{aligned} & a = \tau \wedge s = t \wedge u \xrightarrow{\tau}_B v \\ \vee & a = \tau \wedge u = v \wedge s \xrightarrow{\tau}_A t \\ \vee & s \xrightarrow{a}_A t \wedge u \xrightarrow{a}_B v. \end{aligned}$$

Essentially, the superposition $\text{sup}(A, B, R)$ is just the usual parallel composition of A and B with the set of states restricted to R .

Lemma 6.4 *Suppose (f, n) is a normed forward simulation from A to B . Let $C = \text{sup}(A, B, f)$ and let π_1 and π_2 be the projection functions that map states of C to their first and second components, respectively. Let n' be the norm function given by $n'(\delta, u) = n(\delta, \pi_2(u))$. Then (π_1, n') is a normed history relation from A to C , and π_2 is a step refinement from C to B .*

Proof: Straightforward from the definitions. ■

Theorem 6.5 $A \leq_F B \Leftrightarrow (\exists C : A \leq_H C \leq_R B)$.

Proof: Implication “ \Rightarrow ” follows by Lemma 6.4. For implication “ \Leftarrow ”, suppose $A \leq_H C \leq_R B$. Then $A \leq_F C$ by the definition of history relations, and $C \leq_F B$ because any step refinement is a normed forward simulation. Now $A \leq_F B$ follows by the fact that \leq_F is a preorder. ■

In [AK96][Intermezzo 3.23], Klop and Ariola state a remarkable result: on a domain of finitely branching process graphs (i.e., automata considered modulo isomorphism) the preorder induced by functional bisimulations (i.e., history relations) is in fact a partial order: $A \leq_H B$ and $B \leq_H A$ implies $A = B$. The authors also present a counterexample to show that the finite branching property is needed to prove this result. Below we present a slight generalization of the result of [AK96] in the setting of our paper. It turns out to be sufficient to assume that automata have finite invisible nondeterminism (fin).

Theorem 6.6 *Suppose A and B have fin, $A \leq_H B$ and $B \leq_H A$. Then the reachable subautomata of A and B are isomorphic.*

Proof: Suppose that (f, n) is a normed history relation from A to B , and (g, m) is a normed history relation from B to A . Because A and B have fin, both $\text{start}(A)$ and $\text{start}(B)$ are finite. Since f is a step refinement, it maps start states of B to start states of A . Using that f^{-1} is a forward simulation, we infer that f is surjective on start states. Hence $|\text{start}(B)| \leq |\text{start}(A)|$. By a similar argument, using that (g, m) is a normed history relation from B to A , we obtain $|\text{start}(A)| \leq |\text{start}(B)|$. This means that f is also injective on start states.

Let β be an arbitrary trace of A and B . Using a similar argument as above, we infer

$$\begin{aligned} f(\text{after}(A)[\beta]) &= \text{after}(B)[\beta] \\ g(\text{after}(B)[\beta]) &= \text{after}(A)[\beta] \end{aligned}$$

Since, by Lemma 2.1(2), both $\text{after}(A)[\beta]$ and $\text{after}(B)[\beta]$ are finite, it follows that

$$|\text{after}(A)[\beta]| = |\text{after}(B)[\beta]|$$

This means that f and g are injective on the sets $\text{after}(B)[\beta]$ and $\text{after}(A)[\beta]$, respectively. But since each reachable state is in a set $\text{after}(B)[\beta]$ or $\text{after}(A)[\beta]$, for some β , it follows that f and g are injective on all states. Now the required isomorphism property follows from the fact that f and g are step refinements. ■

Intuitively, one may interpret the above result as follows: if $A \leq_H B$ then B contains more *history information* than A . If B contains more history information than A , and A contains more history information than B , then they are equal.

7 Normed Prophecy Relations

In this section, we will define normed prophecy relations and show that they correspond to normed backward simulations, very similarly to the way in which normed history relations correspond to normed forward simulations.

A pair (r, n) is a *normed prophecy relation* from A to B if r is a step refinement from B to A and (r^{-1}, n) is a normed backward simulation from A to B . We write $A \leq_P B$ if there is a normed prophecy relation from A to B , and $A \leq_{iP} B$ if there is a normed prophecy relation (r, n) with r^{-1} image-finite. Thus $A \leq_{iP} B$ implies $A \leq_{iB} B$ and $A \leq_P B$, and $A \leq_P B$ implies $A \leq_B B$ and $B \leq_R A$. Moreover, if all states of A are reachable, B has finite invisible nondeterminism and $A \leq_P B$, then $A \leq_{iP} B$. It is easy to check that the preorder and soundness results for backward simulations and refinements carry over to prophecy relations.

The following lemma is the analogue of Lemma 6.4 in the backward setting. Using this lemma, we can prove that normed backward simulations are equivalent to normed prophecy variables combined with step refinements.

Lemma 7.1 *Suppose (b, n) is a normed backward simulation from A to B . Let $C = \text{sup}(A, B, b)$ and let π_1 and π_2 be the projection functions that map states of C to their first and second components, respectively. Let n' be the norm function given by $n'(\delta, u) = n(\delta, \pi_2(u))$. Then (π_1, n') is a normed prophecy relation from A to C , and π_2 is a step refinement from C to B . If b is image-finite then so is π_1^{-1} .*

Theorem 7.2

1. $A \leq_B B \Leftrightarrow (\exists C : A \leq_P C \leq_R B)$.
2. $A \leq_{iB} B \Leftrightarrow (\exists C : A \leq_{iP} C \leq_R B)$.

Proof: Analogous to that of Theorem 6.5, using Lemma 7.1. ■

We can now prove variants of the well-known completeness result of Abadi and Lamport [AL91].

Theorem 7.3 *(Completeness of normed history+prophecy relations and step refinements)*
*Suppose $A \leq_{*T} B$. Then*

1. $\exists C, D : A \leq_H C \leq_P D \leq_R B$.
2. *If B has fin then $\exists C, D : A \leq_H C \leq_{iP} D \leq_R B$.*

Proof: By Theorem 6.3, there exists an automaton C with $A \leq_H C \leq_B B$. Next, Theorem 7.2 yields the required automaton D with $C \leq_P D \leq_R B$, which proves (1). The proof of (2) is similar, but uses Proposition 5.6. ■

The following theorem states that \leq_P is a partial order on the class of automata with fin, considered modulo isomorphism of reachable subautomata. The proof is analogous to that of Theorem 6.6, the corresponding result for normed history relations.

Theorem 7.4 *Suppose A and B have fin, $A \leq_P B$ and $B \leq_P A$. Then the reachable subautomata of A and B are isomorphic.*

8 Decidability

Thus far, our exposition has been purely semantic. In the words of [AL91]: “We have considered specifications, but not the languages in which they are expressed. We proved the existence of refinement mappings, but said nothing about whether they are expressible in any language.” In this section, we move to the syntactic world and discuss some decidability issues. To this end we

have to fix a language for defining automata. The language below can be viewed as a simplified version of the IOA language of [GLV97].

We assume an underlying assertion language \mathcal{L} which is a first-order language over interpreted symbols for expressing functions and predicates over some concrete domains such as integers, arrays, and lists of integers. If X is a set of (typed) variables then we write $F(X)$ and $E(X)$ for the collection of formulas and expressions, respectively, in which variables from X may occur free. An automaton can be described syntactically by first specifying a finite set X of variables, referred to as the *state variables*. For each state variable x we assume the presence of a copy x' , called the *primed version* of x . We write X' for the set $\{x' \mid x \in X\}$ and, if ϕ is a formula then we write ϕ' for the formula obtained from ϕ by replacing each occurrence of a state variable by its primed version. The set of states of the automaton is defined as the set of all valuations of the state variables in X . The set of initial states is specified by a predicate in $F(X)$, called the *initial condition*. The actions are specified via a finite number of *action names* with, for each action name a , a finite list \vec{v} of variables called the *parameters* of a . We assume $\{\vec{v}\} \cap X = \emptyset$. The set of actions of the automaton is defined as the union, for each action name a , of all tuples $a(\vec{d})$, where \vec{d} is a valuation of the parameters \vec{v} in their respective domains. The transition relation is specified by providing, for each action name a with parameters \vec{v} , a *transition predicate* in $F(X \cup \{\vec{v}\} \cup X')$, i.e., a predicate that may contain action parameters as well as primed and unprimed state variables.

Example 8.1 Below we specify a FIFO channel in a self-explanatory IOA like syntax [GLV97].

```

automaton Channel
  states
    buffer: Seq[Nat]
  initial condition
    buffer = {}
  actions
    send(v: Nat),
    receive(v: Nat),
    tau
  transitions
    action send(v)
      predicate buffer' = buffer |- v
    action receive(v)
      predicate buffer ~= {} /\ v = head(buffer)
        /\ buffer' = tail(buffer)
    action tau
      predicate false

```

This piece of syntax defines an automaton A with

- $states(A) = \mathbb{N}^*$,
- $start(A) = \{\lambda\}$,
- $acts(A) = \{\text{send}(d), \text{receive}(d) \mid d \in \mathbb{N}\} \cup \{\tau\}$,
- $steps(A)$ is the least set that contains the following elements, for all $\sigma \in \mathbb{N}^*$ and $d \in \mathbb{N}$,

$$\begin{array}{ccc}
 \sigma & \xrightarrow{\text{send}(d)} & \sigma d \\
 d \sigma & \xrightarrow{\text{receive}(d)} & \sigma.
 \end{array}$$

Now assume that we have specified two automata A and B , using state variables \vec{x} and \vec{y} , respectively. Let $X = \{\vec{x}\}$ and $Y = \{\vec{y}\}$. Assume $X \cap Y = \emptyset$.

A step refinement from A to B can be specified by a formula of the form $\theta \wedge \vec{y} = \vec{e}$, with $\theta \in E(X)$ and \vec{e} a list of expressions in $E(X)$ that matches \vec{y} in terms of length and types. In this formula, the first conjunct defines the domain of the step refinement whereas the second conjunct defines a map from states of A to states of B by specifying, for each state variable of B , its value in terms of the values of the state variables of A .

A normed forward simulation can be described by a predicate in $F(X \cup Y)$ together with, for each action type a with parameters \vec{v} , an expression in $E(X \cup \{\vec{v}\} \cup X' \cup Y)$ that specifies the norm function. In practice, norm functions often only depend on the states of B , which means that they can be specified by means of a single expression in $E(Y)$.

Example 8.2 Consider the following specification, essentially just the chaining of two FIFO channels.

```

automaton TwoChannels
  states
    buffer1: Seq[Nat],
    buffer2: Seq[Nat]
  initial condition
    buffer1 = {} /\ buffer2 = {}
  actions
    send(v: Nat),
    receive(v: Nat),
    tau
  transitions
    action send(v)
      predicate buffer1' = buffer1 |- v /\ buffer2' = buffer2
    action receive(v)
      predicate buffer2 ~= {} /\ v = head(buffer2)
        /\ buffer2' = tail(buffer2) /\ buffer1' = buffer1
    action tau
      predicate buffer1 ~= {} /\ buffer1' = tail(buffer1) /\
        /\ buffer2' = buffer2 |- head(buffer1)

```

Let B be the automaton denoted by this specification. It is easy to prove that the formula below defines a step refinement from B to the automaton A of Example 8.1.

$$\text{buffer} = \text{buffer2} \parallel \text{buffer1}$$

It is also routine to check that this formula together with the norm on states of B defined by

$$\text{if } \text{buffer1} \sim = \{\} \wedge \text{buffer2} = \{\} \text{ then } 1 \text{ else } 0$$

defines a normed forward simulation from A to B .

We will now show that, under some reasonable (sufficient but certainly not necessary) assumptions, it is in fact decidable whether a given predicate/expression indeed corresponds to a step refinement or normed forward simulation. Assume that automaton A is described using state variables \vec{x} , initial condition φ_0 and, for each action name a , a transition predicate φ_a . Likewise, assume that automaton B is described using state variables \vec{y} , initial condition ψ_0 and, for each action name a , a transition predicate ψ_a . Assume further that each action name a of A is also an action name of B , and that a has the same parameters in both A and B . Write P_a for the list of parameters of a . We require that $P_\tau = \emptyset$.

Suppose that we want to check whether a formula $\rho \triangleq \theta \wedge \vec{y} = \vec{e}$ denotes a step refinement. This is equivalent to proving validity of the following formula:

$$\begin{aligned} & \varphi_0 \Rightarrow \theta \\ \bigwedge & \varphi_0 \wedge \rho \Rightarrow \psi_0 \end{aligned}$$

$$\begin{aligned}
\bigwedge_a \quad \varphi_a \wedge \theta &\Rightarrow \theta' \\
\bigwedge_{a \neq \tau} \varphi_a \wedge \rho \wedge \rho' &\Rightarrow \psi_a \\
\bigwedge \quad \varphi_\tau \wedge \rho \wedge \rho' &\Rightarrow \psi_\tau \vee \vec{y}' = \vec{y}'
\end{aligned}$$

In this formula, the first conjunct asserts that the function is defined for start states of A ; the second conjunct that start states of A are mapped onto start states of B ; the third conjunct that if the function is defined for the source of a transition then it is also defined for the target state of a transition; and the two final conjuncts encode the transfer condition. Thus checking whether a partial function is a step refinement from A to B is decidable if the partial function as well as A and B can all be expressed within a fragment of \mathcal{L} for which tautology checking is decidable.

Next suppose that we want to check whether a formula ρ together with norm expressions n_a , for each action name a , denotes a normed forward simulation from A to B . In order to turn this into a decidable question, we have to make some additional assumptions about the specification of B . We assume that B has finitely many start states², which are listed explicitly, i.e., we require that the initial condition ψ_0 is of the form

$$\psi_0 = \bigvee_{i \in I_0} \vec{y}' = e_0^i \quad (3)$$

where I_0 is a finite index set and, for each i , e_0^i is a list of closed terms. In addition we assume that in any state and for any given value of the action parameters, only finitely many transitions are possible in B , which are listed explicitly. Formally we require that, for each action type a , transition predicate ψ_a is of the form

$$\psi_a = \bigvee_{i \in I_a} (\chi_a^i \wedge \vec{y}' = e_a^i) \quad (4)$$

where I_a is a finite index set and, for each i , χ_a^i is a formula in $F(Y \cup \{P_a\})$ and e_a^i is a list of expressions in $E(Y \cup \{P_a\})$. Basically, χ_a^i gives the precondition of the i -th instance of transition a and $\vec{y}' = e_a^i$ specifies the effect of taking it. Both assumption (3) and (4) are satisfied by most automaton specifications that one encounters in practice. In particular, the assumptions hold for the channels specified in Examples 8.1 and 8.1. Only specifications that involve a nondeterministic choice that is not a priori bounded fall outside of our format. An example of this, described in [SAGG⁺93], is a FIFO channel in which a crash action may result in the loss of an arbitrary subset of the messages contained in a buffer. Under assumptions (3) and (4), we can eliminate the existential quantifiers that occur in the definition of a normed forward simulation, and checking the conditions in this definition becomes equivalent to proving validity of the following formula:

$$\begin{aligned}
\varphi_0 &\Rightarrow \bigvee_{i \in I_0} \rho[e_0^i/\vec{y}'] \\
\bigwedge_{a \neq \tau} \varphi_a \wedge \rho &\Rightarrow \bigvee_{i \in I_a} (\chi_a^i \wedge \rho'[e_a^i/\vec{y}']) \vee \bigvee_{i \in I_\tau} (\chi_\tau^i \wedge \rho[e_\tau^i/\vec{y}'] \wedge n_a[e_\tau^i/\vec{y}'] < n_a) \\
\bigwedge \quad \varphi_\tau \wedge \rho &\Rightarrow \rho'[\vec{y}'/\vec{y}'] \vee \bigvee_{i \in I_\tau} (\chi_\tau^i \wedge \rho'[e_\tau^i/\vec{y}']) \vee \bigvee_{i \in I_\tau} (\chi_\tau^i \wedge \rho[e_\tau^i/\vec{y}'] \wedge n_\tau[e_\tau^i/\vec{y}'] < n_\tau)
\end{aligned}$$

If this formula can be expressed within a fragment of \mathcal{L} for which tautology checking is decidable then it is decidable whether ρ together with expressions n_a constitutes a normed forward simulation. It is easy to see that a similar result can also be obtained for normed history variables. Thus far, however, we have not been able to come up with plausible syntactic restrictions, applicable in practical cases, that ensure decidability of normed backward simulations and/or normed prophecy relations. It is for instance not clear how one can eliminate the existential quantifier in the formula

²This assumption can be relaxed if we assume that the value of certain state variables of B is fully determined by ρ and the state of A : for those state variables the initial value can be left unspecified.

that asserts that in a normed backward simulation for each state of A there exists a related state of B .

Our decidability results for step refinements and normed forward simulations do not carry over to the refinements and forward simulations as described, for instance, in [LV95]. In order to see this, let A be a system with two states, an initial and a final one, and a single transition labeled *halt* from the initial to the final state. Let B be a system that simulates the n -th Turing machine such that each computation step of the Turing machine corresponds with a τ -move, and that moves via a *halt*-action to a designated final state if and only if the computation of the Turing machine terminates. The function that maps the initial state of A to the initial state of B and the final state of A to the final state of B is a weak refinement iff the n -th Turing machine halts. It is straightforward to specify A , B and the function from states of A to states of B in a decidable logic. Hence it is undecidable whether a given function is a weak refinement, even in a setting where the underlying logic is decidable.

9 Reachability

For the sake of simplicity, all definitions of simulations and refinements so far have been presented without any mention of reachability or invariants. However, in practical verifications it is almost always the case that first some invariants (properties that hold for all reachable states) are established for the lower-level and/or higher-level specification. These invariants are then used in proving the step correspondence. In this section we show how to integrate reachability concerns into the simulation definitions. More specifically, we present adapted versions of step refinements, normed forward simulations and normed backward simulations which include reachability concerns, and discuss their relationship with the original definitions. For examples of the use of these adapted definitions and their formalization in PVS, we refer to [Gri00].

An adapted *step refinement* from A to B consists of a partial function $r : \text{states}(A) \rightarrow \text{states}(B)$ satisfying the following two conditions:

1. If $s \in \text{start}(A)$ then $s \in \text{domain}(r)$ and $r(s) \in \text{start}(B)$.
2. If $s \xrightarrow{a}_A t \wedge s \in \text{domain}(r) \wedge \text{reachable}(A, s) \wedge \text{reachable}(B, r(s))$ then $t \in \text{domain}(r)$ and
 - (a) $r(s) = r(t) \wedge a = \tau$, or
 - (b) $r(s) \xrightarrow{a}_B r(t)$.

Clause $\text{reachable}(A, s)$ in condition (2) allows us to reuse invariants that have previously been established for lower-level specification A , whereas clause $\text{reachable}(B, r(s))$ in condition (2) makes it possible to reuse known invariants of higher-level specification B . The adapted definition can easily be seen as a special case of the original definition in Section 3.1: if r is an adapted step refinement then the restriction r' of r defined by

$$s \in \text{domain}(r') \stackrel{\Delta}{=} s \in \text{domain}(r) \wedge \text{reachable}(A, s) \wedge \text{reachable}(B, r(s)),$$

is a regular step refinement. Conversely, any regular step refinement trivially satisfies the conditions of the adapted version.

An adapted *normed forward simulation* from A to B consists of a relation $f \subseteq \text{states}(A) \times \text{states}(B)$ and a function $n : \text{steps}(A) \times \text{states}(B) \rightarrow S$, for some well-founded set S , such that:

1. If $s \in \text{start}(A)$ then $f[s] \cap \text{start}(B) \neq \emptyset$.
2. If $s \xrightarrow{a}_A t \wedge u \in f[s] \wedge \text{reachable}(A, s) \wedge \text{reachable}(B, u)$ then
 - (a) $u \in f[t] \wedge a = \tau$, or
 - (b) $\exists v \in f[t] : u \xrightarrow{a}_B v$, or
 - (c) $\exists v \in f[s] : u \xrightarrow{\tau}_B v \wedge n(s \xrightarrow{a} t, v) < n(s \xrightarrow{a} t, u)$.

Again, the clause $reachable(A, s)$ in condition (2) allows us to reuse invariants that have previously been established for A , whereas clause $reachable(B, u)$ in condition (2) permits reuse of invariants of B . And again the adapted definition can easily be seen as a special case of the original definition (in Section 4): if (f, n) is an adapted normed forward simulation then the pair (g, n) , where $g = f \cap (reachable(A) \times reachable(B))$, is a regular normed forward simulation. Conversely, any regular normed forward simulation trivially is an adapted normed forward simulation.

An adapted *normed backward simulation* from A to B consists of a relation $b \subseteq states(A) \times states(B)$, a predicate $Q \subseteq states(B)$, and a function $n : (steps(A) \cup start(A)) \times states(B) \rightarrow S$, for some well-founded set S , such that:

1. If $s \in start(A) \wedge u \in b[s] \wedge Q(u)$ then
 - (a) $u \in start(B)$, or
 - (b) $\exists v \in b[s] : v \xrightarrow{\tau}_B u \wedge n(s, v) < n(s, u) \wedge Q(v)$.
2. If $t \xrightarrow{a}_A s \wedge u \in b[s] \wedge reachable(A, t) \wedge Q(u)$ then
 - (a) $u \in b[t] \wedge a = \tau$, or
 - (b) $\exists v \in b[t] : v \xrightarrow{a}_B u \wedge Q(v)$, or
 - (c) $\exists v \in b[s] : v \xrightarrow{\tau}_B u \wedge n(t \xrightarrow{a} s, v) < n(t \xrightarrow{a} s, u) \wedge Q(v)$.
3. If $reachable(A, s)$ then $\exists u \in b[s] : Q(u)$.

Clause $reachable(A, t)$ in condition (2) allows us to reuse invariants that have previously been established for A , and clause $Q(u)$ in condition (2) permits reuse of invariants of B . Note that by a trivial inductive argument a backward simulation can never relate a reachable state of A to a non-reachable state of B . Thus we can safely restrict the range of any backward simulation by all invariants proven for B . To this end predicate Q has been included in the definition of the adapted normed backward simulation, even though strictly speaking (1) Q need not be an invariant, and (2) Q can always be eliminated by restricting the range of b . Once more the adapted definition is a special case of the original definition (in Section 5): if (b, n) is an adapted normed backward simulation then (b, n) is also a regular normed backward simulation from the automaton A' , that restricts A to its reachable states, to the automaton B' , that restricts B to the states in Q . Conversely, any regular normed backward simulation trivially is an adapted normed backward simulation with $Q = states(B)$.

We leave it up to the reader to work out adapted versions of the normed history and prophecy relations.

Acknowledgement

We thank Mariëlle Stoelinga for spotting a mistake in an earlier version of this paper, and Jan Willem Klop for discussions that led us to Theorem 6.6.

References

- [AK96] Z.M. Ariola and J.W. Klop. Equational term graph rewriting. *Fundamenta Informaticae*, 26(3/4):207–240, 1996. Extended version as University of Oregon Technical Report CIS-TR-95-16.
- [AL91] M. Abadi and L. Lamport. The existence of refinement mappings. *Theoretical Computer Science*, 82(2):253–284, 1991.
- [Bas96] T. Basten. Branching bisimilarity is an equivalence indeed! *Information Processing Letters*, 58(3):141–147, 1996.

- [BCG88] M.C. Browne, E.M. Clarke, and O. Grümberg. Characterizing finite Kripke structures in propositional temporal logic. *Theoretical Computer Science*, 59(1,2):115–131, 1988.
- [BLS96] S. Bensalem, Y. Lakhnech, and H. Saidi. Powerful techniques for the automatic generation of invariants. In R. Alur and T.A. Henzinger, editors, *Proceedings of the 8th International Conference on Computer Aided Verification*, New Brunswick, NJ, USA, volume 1102 of *Lecture Notes in Computer Science*, pages 323–335. Springer-Verlag, July/August 1996.
- [BS00] C. Baier and M.I.A. Stoelinga. Norm functions for probabilistic bisimulations with delays. In J. Tiuryn, editor, *Proceedings of 3rd International Conference on Foundations of Science and Computation Structures (FOSSACS)*, Berlin, Germany, March 2000, volume 1784 of *Lecture Notes in Computer Science*, pages 1–16. Springer-Verlag, 2000.
- [DGM97] M.C.A. Devillers, W.O.D. Griffioen, and O. Müller. Possibly infinite sequences: A comparative case study. In E.L. Gunter and A. Felty, editors, *10th International Conference on Theorem Proving in Higher Order Logics (TPHOLs'97)*, volume 1275 of *Lecture Notes in Computer Science*, pages 89–104. Springer-Verlag, 1997.
- [DGRV00] M.C.A. Devillers, W.O.D. Griffioen, J.M.T. Romijn, and F.W. Vaandrager. Verification of a leader election protocol: Formal methods applied to IEEE 1394. *Formal Methods in System Design*, 16(3):307–320, June 2000.
- [DNV95] R. De Nicola and F.W. Vaandrager. Three logics for branching bisimulation. *Journal of the ACM*, 42(2):458–487, March 1995.
- [Gin68] A. Ginzburg. *Algebraic Theory of Automata*. Academic Press, New York – London, 1968.
- [GLV97] S.J. Garland, N.A. Lynch, and M. Vaziri. IOA: A language for specifying, programming, and validating distributed systems, September 1997. Available through URL <http://larch.lcs.mit.edu:8001/~garland/ioaLanguage.html>.
- [Gri00] W.O.D. Griffioen. *Studies in Computer Aided Verification of Protocols*. PhD thesis, University of Nijmegen, May 2000. Postscript and PVS sources available via http://www.cs.kun.nl/ita/former_members/davidg/.
- [GS95] J.F. Groote and J.S. Springintveld. Focus points and convergent process operators — a proof strategy for protocol verification. Report CS-R9566, Department of Software Technology, CWI, Amsterdam, November 1995.
- [GSSL93] R. Gawlick, R. Segala, J.F. Søgaaard-Andersen, and N.A. Lynch. Liveness in timed and untimed systems. Technical Report MIT/LCS/TR-587, Laboratory for Computer Science, MIT, Cambridge, MA, December 1993.
- [GV98] W.O.D. Griffioen and F.W. Vaandrager. Normed simulations. In A.J. Hu and M.Y. Vardi, editors, *Proceedings of the 10th International Conference on Computer Aided Verification*, Vancouver, BC, Canada, volume 1427 of *Lecture Notes in Computer Science*, pages 332–344. Springer-Verlag, June/July 1998.
- [GW96] R.J. van Glabbeek and W.P. Weijland. Branching time and abstraction in bisimulation semantics. *Journal of the ACM*, 43(3):555–600, 1996.
- [HSV94] L. Helmink, M.P.A. Sellink, and F.W. Vaandrager. Proof-checking a data link protocol. In H. Barendregt and T. Nipkow, editors, *Proceedings International Workshop TYPES'93*, Nijmegen, The Netherlands, May 1993, volume 806 of *Lecture Notes in Computer Science*, pages 127–165. Springer-Verlag, 1994.

- [Jon87] B. Jonsson. *Compositional Verification of Distributed Systems*. PhD thesis, Department of Computer Systems, Uppsala University, 1987. DoCS 87/09.
- [Jon90] B. Jonsson. On decomposing and refining specifications of distributed systems. In J.W. de Bakker, W.P. de Roever, and G. Rozenberg, editors, *Proceedings REX Workshop on Stepwise Refinement of Distributed Systems: Models, Formalism, Correctness*, Mook, The Netherlands, May/June 1989, volume 430 of *Lecture Notes in Computer Science*, pages 361–387. Springer-Verlag, 1990.
- [Knu97] D.E. Knuth. *Fundamental Algorithms*, volume 1 of *The Art of Computer Programming*. Addison-Wesley, Reading, Massachusetts, 1997. Third edition.
- [LT87] N.A. Lynch and M.R. Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Proceedings of the 6th Annual ACM Symposium on Principles of Distributed Computing*, pages 137–151, August 1987. A full version is available as MIT Technical Report MIT/LCS/TR-387.
- [LV95] N.A. Lynch and F.W. Vaandrager. Forward and backward simulations, I: Untimed systems. *Information and Computation*, 121(2):214–233, September 1995.
- [LV96] N.A. Lynch and F.W. Vaandrager. Forward and backward simulations, II: Timing-based systems. *Information and Computation*, 128(1):1–25, July 1996.
- [Lyn96] N.A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, Inc., San Francisco, California, 1996.
- [MBSU98] Z. Manna, A. Browne, H.B. Sipma, and T.E. Uribe. Visual abstraction for temporal verification. In A.M. Haeberer, editor, *Proceedings AMAST’98*, volume 1548 of *Lecture Notes in Computer Science*, pages 28–41. Springer-Verlag, 1998.
- [Mil89] R. Milner. *Communication and Concurrency*. Prentice-Hall International, Englewood Cliffs, 1989.
- [Mue98] O. Mueller. *A Verification Environment for I/O Automata Based on Formalized Meta-Theory*. PhD thesis, Technical University of Munich, September 1998.
- [Nam97] K.S. Namjoshi. A simple characterization of stuttering bisimulation. In S. Ramesh and G. Sivakumar, editors, *Proceedings 17th Conference on Foundations of Software Technology and Theoretical Computer Science*, Kharagpur, India, volume 1346 of *Lecture Notes in Computer Science*, pages 284–296. Springer-Verlag, December 1997.
- [NS95] T. Nipkow and K. Slind. I/O automata in Isabelle/HOL. In P. Dybjer, B. Nordström, and J. Smith, editors, *Types for Proofs and Programs*, volume 996 of *Lecture Notes in Computer Science*, pages 101–119. Springer-Verlag, 1995.
- [OG76] S. Owicki and D. Gries. An axiomatic proof technique for parallel programs. *Acta Informatica*, 6(4):319–340, 1976.
- [ORSH95] S. Owre, J. Rushby, N. Shankar, and F. von Henke. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 21(2):107–125, February 1995.
- [RE98] W.P. de Roever and K. Engelhardt. *Data Refinement: Model-Oriented Proof Methods and their Comparison*. Cambridge Tracts in Theoretical Computer Science 47. Cambridge University Press, 1998.
- [SAGG⁺93] J. Søgaard-Andersen, S. Garland, J. Guttag, N.A. Lynch, and A. Pogoyants. Computer-assisted simulation proofs. In C. Courcoubetis, editor, *Proceedings of the 5th International Conference on Computer Aided Verification*, Elounda, Greece, volume 697 of *Lecture Notes in Computer Science*, pages 305–319. Springer-Verlag, 1993.

- [SALL93] J.F. Sogaard-Andersen, N.A. Lynch, and B.W. Lampson. Correctness of communication protocols – a case study. Technical Report MIT/LCS/TR-589, Laboratory for Computer Science, MIT, Cambridge, MA, November 1993.
- [Sis91] A.P. Sistla. Proving correctness with respect to nondeterministic safety specifications. *Information Processing Letters*, 39(1):45–49, July 1991.