Abstraction Refinement and Antichains for Trace Inclusion of Infinite State Systems

Radu Iosif, Adam Rogalewicz, Tomáš Vojnar

► To cite this version:

Radu Iosif, Adam Rogalewicz, Tomáš Vojnar. Abstraction Refinement and Antichains for Trace Inclusion of Infinite State Systems. 22nd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), Apr 2016, Eindhoven, Netherlands. pp.71-89, 10.1007/978-3-662-49674-9_5 . hal-01418885

HAL Id: hal-01418885

<https://hal.science/hal-01418885>

Submitted on 17 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Domain

Abstraction Refinement and Antichains for Trace Inclusion of Infinite State Systems

Radu Iosif, Adam Rogalewicz, and Tomáš Vojnar

CNRS/Verimag, France and FIT BUT, Czech Republic

Abstract. A *data automaton* is a finite automaton equipped with variables (counters or registers) ranging over infinite data domains. A trace of a data automaton is an alternating sequence of alphabet symbols and values taken by the counters during an execution of the automaton. The problem addressed in this paper is the inclusion between the sets of traces (data languages) recognized by such automata. Since the problem is undecidable in general, we give a semi-algorithm based on abstraction refinement, which is proved to be sound and complete, but whose termination is not guaranteed. We have implemented our technique in a prototype tool and show promising results on several non-trivial examples.

1 Introduction

In this paper, we address a *trace inclusion* problem for infinite-state systems. Given (i) a network of *data automata* $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ that communicate via a set of shared variables $\mathbf{x}_{\mathcal{A}}$, ranging over an infinite data domain, and a set of input events $\Sigma_{\mathcal{A}}$, and (ii) a data automaton B whose set of variables $\mathbf{x}_B$ is a subset of $\mathbf{x}_{\mathcal{A}}$, does the set of (finite) traces of B contain the traces of A ? Here, by a *trace*, we understand an alternating sequence of valuations of the variables from the set $\mathbf{x}_B$ and input events from the set $\Sigma_{\mathcal{A}} \cap \Sigma_B$, starting and ending with a valuation. Typically, the network of automata $\mathcal{A}$ is an implementation of a concurrent system and B is a specification of the set of good behaviors of the system.

Consider, for instance, the network $\langle A_1, \dots, A_N \rangle$ of data automata equipped with the integer-valued variables x and v shown in Fig. 1–left. The automata synchronize on the **init** symbol and interleave their $\mathbf{a}_{1,\dots,N}$ actions. Each automaton A_i increases the shared variable x and writes its identifier i into the shared variable v as long as the value of x is in the interval $[(i-1)\Delta, i\Delta-1]$, and it is inactive outside this interval, where $\Delta \geq 1$ is an unbounded parameter of the network. A possible specification for this network might require that each firing sequence is of the form **init** $\mathbf{a}_{1,\dots,N}^* \mathbf{a}_2 \mathbf{a}_{2,\dots,N}^* \dots \mathbf{a}_i \mathbf{a}_i^*$ for some $1 \leq i \leq N$, and that v is increased only on the first occurrence of the events $\mathbf{a}_2, \dots, \mathbf{a}_i$, in this order. This condition is encoded by the automaton B (Fig. 1–right). Observe that only the v variable is shared between the network $\langle A_1, \dots, A_N \rangle$ and the specification automaton B —we say that v is *observable* in this case. An example of a trace, for $\Delta = 2$ and $N \geq 3$, is: $(v=0)$ **init** $(v=1)$ $\mathbf{a}_1$ $(v=1)$ $\mathbf{a}_1$ $(v=1)$ $\mathbf{a}_2$ $(v=2)$ $\mathbf{a}_2$ $(v=2)$ $\mathbf{a}_3$ $(v=3)$. Our problem is to check that this, and all other traces of the network, are included in the language of the specification automaton, called the *observer*. The trace inclusion problem has multiple applications, e.g.:

Fig. 1. An instance of the trace inclusion problem.

- Decision procedures for logics describing array structures within imperative programs [16, 15] that use a translation of array formulae to integer counter automata which encode the set of array models of a formula. The expressiveness of such logics is currently limited by the decidability of the emptiness (reachability) problem for counter automata. If we give up on decidability, we can reduce an entailment between two array formulae to the trace inclusion of two integer counter automata, and use the method presented in this paper as a semi-decision procedure. To assess this claim, we have applied our trace inclusion method to several verification conditions for programs with unbounded arrays of integers [6].
- Timed automata and regular specifications of timed languages [2] can be both represented by finite automata extended with real-valued variables [13]. The verification problem boils down to the trace inclusion of two real-valued data automata. In this context, our method has been tested on several timed verification problems, including communication protocols and boolean circuits [27].

When developing a method for checking the inclusion between trace languages of automata extended with variables ranging over infinite data domains, the first problem is the lack of determinisation and/or complementation results. In fact, certain classes of infinite state systems, such as timed automata [2], cannot be determinized and are provably not closed under complement. This is the case due to the fact that the clock variables of a timed automaton are not observable in its timed language, which records only the time lapses between successive events. However, if we require that the values of all variables of a data automaton be part of its trace language, we obtain a determinisation result, which generalizes the classical subset construction by taking into account the data valuations. Building on this first result, we define the complement of a data language and reduce the trace inclusion problem to the emptiness of a product data automaton $\mathcal{L}(A \times \overline{B}) = \emptyset$. It is crucial, for this reduction, that the variables $\mathbf{x}_B$ of the right-hand side data automaton B (the one being determinized) are also controlled by the left-hand side automaton A , in other words, that B has no hidden variables.

The language emptiness problem for data automata is, in general, undecidable [23]. Nevertheless, several semi-algorithms and tools for this problem (better known as the *reachability* problem) have been developed [3, 18, 21, 14]. Among those, the technique of *lazy predicate abstraction* [18] combined with *counterexample-driven refinement* using *interpolants* [21] has been shown to be particularly successful in proving emptiness of rather large infinite-state systems. Moreover, this technique shares similar aspects with the antichain-based algorithm for language inclusion in the case of a finite alpha-

bet [1]. An important similarity is that both techniques use a partial order on states, to prune the state space during the search.

The main result of this paper is a semi-algorithm that combines the principle of the antichain-based language inclusion algorithm [1] with the interpolant-based abstraction refinement semi-algorithm [21], via a general notion of language-based subsumption relation. We have implemented our semi-algorithm in a prototype tool and carried out a number of experiments, involving hardware, real-time systems, and array logic problems. Since our procedure tests inclusion within a set of good traces, instead of empty intersection with a set of error traces, we can encode rather complex verification conditions concisely, using automata of relatively small size.

1.1 Overview

Fig. 2. Sample run of our semi-algorithm.

We introduce the reader to our trace inclusion method by means of an example. Let us consider the network of data automata $\langle A_1, A_2 \rangle$ and the data automaton B from Fig. 1. We prove that, for any value of Δ , any trace of the network $\langle A_1, A_2 \rangle$, obtained as an interleaving of the actions of A_1 and A_2 , is also a trace of the observer B . To this end, our procedure will fire increasingly longer sequences of input events, in search for a counterexample trace. We keep a set of predicates associated with each state $\langle \langle q_1, q_2 \rangle, P \rangle$ of the product automaton where q_i is a state of A_i and P is a set of states of B . These predicates are formulae that define over-approximations of the data values reached simultaneously by the network, when A_i is the state q_i , and by the observer B , in every state from P .

The first input event is **init**, on which A_1 and A_2 synchronize, moving together from the initial state $\langle q_0^1, q_0^2 \rangle$ to $\langle q_1^1, q_1^2 \rangle$. In response, B can choose to either (i) move from

$\{p_0\}$ to $\{p_1\}$, matching the only transition rule from p_0 , or (ii) ignore the transition rule and move to the empty set. In the first case, the values of v match the relation of the rule $p_0 \xrightarrow{\text{init}, v'=1} p_1$, while in the second case, these values match the negated relation $\neg(v' = 1)$. The second case is impossible because the action of the network requires $x' = 0 \wedge v' = 1$. The only successor state is thus $(\langle q_1^1, q_1^2 \rangle, \{p_1\})$ in Fig. 2 (a). Since no predicates are initially available at this state, the best over-approximation of the set of reachable data valuations is the universal set ($\top$).

The second input event is $\mathbf{a}_1$, on which A_1 moves from q_1^1 back to itself, while A_2 makes an idle step because no transition with $\mathbf{a}_1$ is enabled from q_1^2 . Again, B has the choice between moving from $\{p_1\}$ either to $\emptyset$ or $\{p_1\}$. Let us consider the first case, in which the successor state is $(\langle q_1^1, q_1^2 \rangle, \emptyset, \top)$. Since q_1^1 and q_1^2 are final states of A_1 and A_2 , respectively, and no final state of B is present in $\emptyset$, we say that the state is accepting. If the accepting state (in dashed boxes in Fig. 2) is reachable according to the transition constraints along the input sequence $\text{init}.\mathbf{a}_1$, we have found a counterexample trace that is in the language of $\langle A_1, A_2 \rangle$ but not in the language of B .

To verify the reachability of the accepting state, we check the satisfiability of the path formula corresponding to the composition of the transition constraints $x' = 0 \wedge v' = 1$ (init) and $0 \leq x < \Delta \wedge x' = x + 1 \wedge v' = 1 \wedge \neg(v' = v)$ ($\mathbf{a}_1$) in Fig. 2 (a). This formula is unsatisfiable, and the proof of infeasibility provides the interpolant $\langle v = 1 \rangle$. This formula is an explanation for the infeasibility of the path because it is implied by the first constraint and it is unsatisfiable in conjunction with the second constraint. By associating the new predicate $v = 1$ with the state $(\langle q_1^1, q_1^2 \rangle, \{p_1\})$, we ensure that the same spurious path will never be explored again.

We delete the spurious counterexample and recompute the states along the input sequence $\text{init}.\mathbf{a}_1$ with the new predicate. In this case, $(\langle q_1^1, q_1^2 \rangle, \emptyset)$ is unreachable, and the outcome is $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$. However, this state was first encountered after the sequence init , so there is no need to store a second occurrence of this state in the tree. We say that the node $\text{init}.\mathbf{a}_1$ is subsumed by init , and indicate this by a dashed arrow in Fig. 2 (b).

We continue with $\mathbf{a}_2$ from the state $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$. In this case, A_1 makes an idle step and A_2 moves from q_1^2 to itself. In response, B has the choice between moving from $\{p_1\}$ to either (i) $\{p_1\}$ with the constraint $v' = v$, (ii) $\{p_2\}$ with the constraint $v' = v + 1$, (iii) $\{p_1, p_2\}$ with the constraint $v' = v \wedge v' = v + 1 \rightarrow \perp$ (this possibility is discarded), (iv) $\emptyset$ for data values that satisfy $\neg(v' = v) \wedge \neg(v' = v + 1)$. The last case is also discarded because the value of v after init constrained to 1 and the A_2 imposes further the constraint $v' = 2$ and $v = 1 \wedge v' = 2 \wedge \neg(v' = v) \wedge \neg(v' = v + 1) \rightarrow \perp$. Hence, the only $\mathbf{a}_2$ -successor of $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$ is $(\langle q_1^1, q_1^2 \rangle, \{p_2\}, \top)$, in Fig. 2 (b).

By firing the event $\mathbf{a}_1$ from this state, we reach $(\langle q_1^1, q_1^2 \rangle, \emptyset, v = 1)$, which is, again, accepting. We check whether the path $\text{init}.\mathbf{a}_2.\mathbf{a}_1$ is feasible, which turns out not to be the case. For efficiency reasons, we find the shortest suffix of this path that can be proved infeasible. It turns out that the sequence $\mathbf{a}_2.\mathbf{a}_1$ is infeasible starting from the state $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$, which is called the *pivot*. This proof of infeasibility yields the interpolant $\langle v = 1, \Delta < x \rangle$, and a new predicate $\Delta < x$ is associated with $(\langle q_1^1, q_1^2 \rangle, \{p_2\})$. The refinement phase rebuilds only the subtree rooted at the pivot state, in Fig. 2 (b).

The procedure then builds the tree on Fig. 2 (c) starting from the pivot node and finds the accepting state $(\langle q_1^1, q_1^2 \rangle, \emptyset, \Delta < x)$ as the result of firing the sequence $\text{init}.\mathbf{a}_2.\mathbf{a}_2$.

This path is spurious, and the new predicate $v = 2$ is associated with the location $(\langle q_1^1, q_1^2 \rangle, \{p_2\})$. The pivot node is the same as in Fig. 2 (b), and, by recomputing the subtree rooted at this node with the new predicates, we obtain the tree in Fig. 2 (d), in which all frontier nodes are subsumed by their predecessors. Thus, no new event needs to be fired, and the procedure can stop reporting that the trace inclusion holds.

1.2 Related Work

The trace inclusion problem has been previously addressed in the context of timed automata [25]. Although the problem is undecidable in general, decidability is recovered when the left-hand side automaton has at most one clock, or the only constant appearing in the clock constraints is zero. These are essentially the only known decidable cases of language inclusion for timed automata.

The study of *data automata* [5, 11] usually deals with the complexity of decision problems in logics describing data languages for simple theories, typically infinite data domains with equality. Here we provide a semi-decision procedure for the language inclusion between data automata controlled by generic first-order theories, whose language-theoretic problems are undecidable.

Data words are also studied in the context of *symbolic visibly pushdown automata* (SVPA) [10]. Language inclusion is decidable for SVPAs with transition guards from a decidable theory because SVPAs are closed under complement and the emptiness can be reduced to a finite number of queries expressible in the underlying theory of guards. Decidability comes here at the cost of reducing the expressivity and forbidding comparisons between adjacent positions in the input (only comparisons between matching call/return positions of the input nested words are allowed).

Finally, several works on model checking infinite-state systems against CTL [4] and CTL* [8] specifications are related to our problem as they check inclusion between the set of computation trees of an infinite-state system and the set of trees defined by a branching temporal logic specification. The verification of existential CTL formulae [4] is reduced to solving forall-exists quantified Horn clauses by applying counterexample guided refinement to discover witnesses for existentially quantified variables. The work [8] on CTL* verification of infinite systems is based on partial symbolic determination, using prophecy variables to summarize the future program execution. For finite-state systems, automata are a strictly more expressive formalism than temporal logics¹. Such a comparison is, however, non-trivial for infinite-state systems. Nevertheless, we found the data automata considered in this paper to be a natural tool for specifying verification conditions of array programs [16, 15, 6] and regular properties of timed languages [2].

2 Preliminary Definitions

Let $\mathbb{N}$ denote the set of non-negative integers including zero. For any $k, \ell \in \mathbb{N}$, $k \leq \ell$, we write $[k, \ell]$ for the set $\{k, k+1, \dots, \ell\}$. We write $\perp$ and $\top$ for the boolean constants

¹ For (in)finite words, the class of LTL-definable languages coincides with the star-free languages, which are a strict subclass of (ω) -regular languages.

false and true, respectively. Given a possibly infinite data domain $\mathcal{D}$, we denote by $\text{Th}(\mathcal{D}) = \langle \mathcal{D}, f_1, \dots, f_m \rangle$ the set of syntactically correct first-order formulae with function symbols $f_1, \dots, f_m$. A variable x is said to be *free* in a formula ϕ , denoted as $\phi(x)$, iff it does not occur under the scope of a quantifier.

Let $\mathbf{x} = \{x_1, \dots, x_n\}$ be a finite set of variables. A *valuation* $\mathbf{v} : \mathbf{x} \rightarrow \mathcal{D}$ is an assignment of the variables in $\mathbf{x}$ with values from $\mathcal{D}$. We denote by $\mathcal{D}^{\mathbf{x}}$ the set of such valuations. For a formula $\phi(\mathbf{x})$, we denote by $\mathbf{v} \models_{\text{Th}(\mathcal{D})} \phi$ the fact that substituting each variable $x \in \mathbf{x}$ by $\mathbf{v}(x)$ yields a valid formula in the theory $\text{Th}(\mathcal{D})$. In this case, $\mathbf{v}$ is said to be a *model* of ϕ . A formula is said to be *satisfiable* iff it has a model. For a formula $\phi(\mathbf{x}, \mathbf{x}')$ where $\mathbf{x}' = \{x' \mid x \in \mathbf{x}\}$ and two valuations $\mathbf{v}, \mathbf{v}' \in \mathcal{D}^{\mathbf{x}}$, we denote by $(\mathbf{v}, \mathbf{v}') \models_{\text{Th}(\mathcal{D})} \phi$ the fact that the formula obtained from ϕ by substituting each x with $\mathbf{v}(x)$ and each x' with $\mathbf{v}'(x')$ is valid in $\text{Th}(\mathcal{D})$.

Data Automata. *Data Automata* (DA) are extensions of non-deterministic finite automata with variables ranging over an infinite data domain $\mathcal{D}$, equipped with a first order theory $\text{Th}(\mathcal{D})$. Formally, a DA is a tuple $A = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q, \mathbf{1}, F, \Delta \rangle$, where:

- Σ is a finite alphabet of input events and $\diamond \in \Sigma$ is a special padding symbol,
- $\mathbf{x} = \{x_1, \dots, x_n\}$ is a set of variables,
- Q is a finite set of *states*, $\mathbf{1} \in Q$ is an *initial state*, $F \subseteq Q$ are *final states*, and
- Δ is a set of *rules* of the form $q \xrightarrow{\sigma, \phi(\mathbf{x}, \mathbf{x}')} q'$ where $\sigma \in \Sigma$ is an alphabet symbol and $\phi(\mathbf{x}, \mathbf{x}')$ is a formula in $\text{Th}(\mathcal{D})$.

A *configuration* of A is a pair $(q, \mathbf{v}) \in Q \times \mathcal{D}^{\mathbf{x}}$. We say that a configuration $(q', \mathbf{v}')$ is a *successor* of $(q, \mathbf{v})$ if and only if there exists a rule $q \xrightarrow{\sigma, \phi} q' \in \Delta$ and $(\mathbf{v}, \mathbf{v}') \models_{\text{Th}(\mathcal{D})} \phi$. We denote the successor relation by $(q, \mathbf{v}) \xrightarrow{\sigma, \phi}_A (q', \mathbf{v}')$, and we omit writing ϕ and A when no confusion may arise. We denote by $\text{succ}_A(q, \mathbf{v}) = \{(q', \mathbf{v}') \mid (q, \mathbf{v}) \rightarrow_A (q', \mathbf{v}')\}$ the set of successors of a configuration $(q, \mathbf{v})$.

A *trace* is a finite sequence $w = (\mathbf{v}_0, \sigma_0), \dots, (\mathbf{v}_{n-1}, \sigma_{n-1}), (\mathbf{v}_n, \diamond)$ of pairs $(\mathbf{v}_i, \sigma_i)$ taken from the infinite alphabet $\mathcal{D}^{\mathbf{x}} \times \Sigma$. A *run* of A over the *trace* w is a sequence of configurations $\pi : (q_0, \mathbf{v}_0) \xrightarrow{\sigma_0} (q_1, \mathbf{v}_1) \xrightarrow{\sigma_1} \dots \xrightarrow{\sigma_{n-1}} (q_n, \mathbf{v}_n)$. We say that the run π is *accepting* if and only if $q_n \in F$, in which case A *accepts* w . The *language* of A , denoted $\mathcal{L}(A)$, is the set of traces accepted by A .

Data Automata Networks. A *data automata network* (DAN) is a non-empty tuple $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ of data automata $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{1}_i, F_i, \Delta_i \rangle$, $i \in [1, N]$ whose sets of states are pairwise disjoint. A DAN is a succinct representation of an exponentially larger DA $\mathcal{A}^e = \langle \mathcal{D}, \Sigma_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}}, Q_{\mathcal{A}}, \mathbf{1}_{\mathcal{A}}, F_{\mathcal{A}}, \Delta_{\mathcal{A}} \rangle$, called the *expansion* of $\mathcal{A}$, where:

- $\Sigma_{\mathcal{A}} = \Sigma_1 \cup \dots \cup \Sigma_N$ and $\mathbf{x}_{\mathcal{A}} = \mathbf{x}_1 \cup \dots \cup \mathbf{x}_N$,
- $Q_{\mathcal{A}} = Q_1 \times \dots \times Q_N$, $\mathbf{1}_{\mathcal{A}} = \langle \mathbf{1}_1, \dots, \mathbf{1}_N \rangle$ and $F_{\mathcal{A}} = F_1 \times \dots \times F_N$,
- $\langle q_1, \dots, q_N \rangle \xrightarrow{\sigma, \phi} \langle q'_1, \dots, q'_N \rangle$ if and only if (i) for all $i \in I$, $q_i \xrightarrow{\sigma, \phi_i} q'_i$, (ii) for all $i \notin I$, $q_i = q'_i$, and (iii) $\phi \equiv \bigwedge_{i \in I} \phi_i \wedge \bigwedge_{j \notin I} \tau_j$, where $I = \{i \in [1, N] \mid q_i \xrightarrow{\sigma, \phi_i} q'_i \in \Delta_i\}$ is the set of DA that can move from q_i to q'_i while reading the input symbol σ , and $\tau_j \equiv \bigwedge_{x \in \mathbf{x}_j \setminus (\bigcup_{i \in I} \mathbf{x}_i)} x' = x$ propagates the values of the local variables in A_j that are not updated by $\{A_i\}_{i \in I}$.

Intuitively, all automata that can read an input symbol synchronize their actions on that symbol whereas the rest of the automata make an idle step and copy the values of their local variables which are not updated by the active automata. The language of the DAN $\mathcal{A}$ is defined as the language of its expansion $\mathcal{DA}$, i.e. $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}^e)$.

Trace Inclusion. Let $\mathcal{A}$ be a DAN and $\mathcal{A}^e = \langle \mathcal{D}, \Sigma, \mathbf{x}_{\mathcal{A}}, Q_{\mathcal{A}}, \mathbf{t}_{\mathcal{A}}, F_{\mathcal{A}}, \Delta_{\mathcal{A}} \rangle$ be its expansion. For a set of variables $\mathbf{y} \subseteq \mathbf{x}_{\mathcal{A}}$, we denote by $v \downarrow_{\mathbf{y}}$ the restriction of a valuation $v \in \mathcal{D}^{\mathbf{x}_{\mathcal{A}}}$ to the variables in $\mathbf{y}$. For a trace $w = (v_0, \sigma_0), \dots, (v_n, \diamond) \in (\mathcal{D}^{\mathbf{x}_{\mathcal{A}}} \times \Sigma_{\mathcal{A}})^*$, we denote by $w \downarrow_{\mathbf{y}}$ the trace $(v_0 \downarrow_{\mathbf{y}}, \sigma_0), \dots, (v_{n-1} \downarrow_{\mathbf{y}}, \sigma_{n-1}), (v_n \downarrow_{\mathbf{y}}, \diamond) \in (\mathcal{D}^{\mathbf{y}} \times \Sigma)^*$. We lift this notion to sets of words in the natural way, by defining $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{y}} = \{w \downarrow_{\mathbf{y}} \mid w \in \mathcal{L}(\mathcal{A})\}$.

We are now ready to define the trace inclusion problem on which we focus in this paper. Given a DAN $\mathcal{A}$ as before and a DA $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{t}_B, F_B, \Delta_B \rangle$ such that $\mathbf{x}_B \subseteq \mathbf{x}_{\mathcal{A}}$, the *trace inclusion problem* asks whether $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B)$? The right-hand side DA B is called *observer*, and the variables in $\mathbf{x}_B$ are called *observable* variables.

3 Boolean Closure Properties of Data Automata

We show first that data automata are closed under the boolean operations of union, intersection and complement and that they are amenable to determinisation. Clearly, the emptiness problem is, in general, undecidable, due to the result of Minsky on 2-counter machines with integer variables, increment, decrement and zero test [23].

Let $A = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q, \mathbf{t}, F, \Delta \rangle$ be a DA for the rest of this section. A is said to be *deterministic* if and only if, for each trace $w \in \mathcal{L}(A)$, A has at most one run over w . The first result of this section is that, interestingly, any DA can be determinized while preserving its language. The determinisation procedure is a generalization of the classical subset construction for Rabin-Scott word automata on finite alphabets. The reason why determinisation is possible for automata over an infinite data alphabet $\mathcal{D}^{\mathbf{x}} \times \Sigma$ is that the successive values taken by *each variable* $x \in \mathbf{x}$ are tracked by the language $\mathcal{L}(A) \subseteq (\mathcal{D}^{\mathbf{x}} \times \Sigma)^*$. This assumption is crucial: a typical example of automata over an infinite alphabet, that cannot be determinized, are timed automata [2], where only the elapsed time is reflected in the language, and not the values of the variables (clocks).

Formally, the *deterministic* DA accepting the language $\mathcal{L}(A)$ is defined as $A^d = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q^d, \mathbf{t}^d, F^d, \Delta^d \rangle$, where $Q^d = 2^Q$, $\mathbf{t}^d = \{\mathbf{t}\}$, $F^d = \{P \subseteq Q \mid P \cap F \neq \emptyset\}$ and Δ^d is the set of rules $P \xrightarrow{\sigma, \theta} P'$ such that:

- for all $p' \in P'$ there exists $p \in P$ and a rule $p \xrightarrow{\sigma, \Psi} p' \in \Delta$,
- $\theta(\mathbf{x}, \mathbf{x}') \equiv \bigwedge_{p' \in P'} \bigvee_{\substack{p \xrightarrow{\sigma, \Psi} p' \in \Delta \\ p \in P}} \Psi \wedge \bigwedge_{p' \in Q \setminus P'} \bigwedge_{\substack{p \xrightarrow{\sigma, \Phi} p' \in \Delta \\ p \in P}} \neg \Phi$.

The main difference with the classical subset construction for Rabin-Scott automata is that here we consider *all sets* P' of states that have a predecessor in P , not just the maximal such set. This refined subset construction takes into account not just the alphabet symbols in Σ , but also the valuations of the variables in $\mathbf{x}$. Observe, moreover, that A^d can be built for any theory $\text{Th}(\mathcal{D})$ that is closed under conjunction and negation. The following lemma states the main properties of A^d .

Lemma 1. Given a DA $A = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q, \mathbf{1}, F, \Delta \rangle$, (1) for any $w \in (\mathcal{D}^{\mathbf{x}} \times \Sigma)^*$ and $P \in Q^d$, A^d has exactly one run on w that starts in P , and (2) $\mathcal{L}(A) = \mathcal{L}(A^d)$.

Proof. (1) Let $w = (v_0, \sigma_0), \dots, (v_{n-1}, \sigma_{n-1}), (v_n, \diamond)$ be an arbitrary trace and $P \subseteq Q$ be a state of A^d . We first build a run $\pi = (P_0, v_0) \xrightarrow{\sigma_0, \theta_0} (P_1, v_1) \dots \xrightarrow{\sigma_{n-1}, \theta_{n-1}} (P_n, v_n)$ of A^d such that $P_0 = P$, by induction on $n \geq 0$. If $n = 0$, then $w = (v_0, \diamond)$ and $\pi = (P_0, v_0)$ is trivially a run of A^d over w . For the induction step, let $n > 0$ and suppose that A^d has a run $(P_0, v_0) \xrightarrow{\sigma_0, \theta_0} \dots (P_{n-1}, v_{n-1})$ such that $P_0 = P$. We extend this run to a run over w , by considering:

$$P_n = \left\{ p \in Q \mid \exists q \in P_{n-1} . q \xrightarrow{\sigma_{n-1}, \phi} p \in \Delta \text{ and } (v_{n-1}, v_n) \models_{\text{Th}(\mathcal{D})} \phi \right\}$$

$$\theta_n \equiv \bigwedge_{\substack{p' \in P_n \\ p \in P_{n-1}}} \bigvee_{\substack{p \xrightarrow{\sigma, \psi} p' \in \Delta}} \psi \wedge \bigwedge_{\substack{p' \in Q \setminus P_n \\ p \xrightarrow{\sigma, \phi} p' \in \Delta}} \neg \phi .$$

It is not hard to see that $(v_{n-1}, v_n) \models \theta_n$, thus $(P_0, v_0) \xrightarrow{\sigma_0, \theta_0} \dots \xrightarrow{\sigma_n, \theta_n} (P_n, v_n)$ is indeed a run of A^d over w . To show that π is unique, suppose, by contradiction, that there exists a different run $\pi' = (R_0, v_0) \xrightarrow{\sigma_0, \omega_0} (R_1, v_1) \dots \xrightarrow{\sigma_{n-1}, \omega_{n-1}} (R_n, v_n)$ such that $P_0 = R_0 = P$. Notice that the relation labeling any transition rule $P_i \xrightarrow{\sigma_i, \theta_i} P_{i+1}$ is entirely determined by the sets P_i and P_{i+1} , so two runs are different iff they differ in at least one state, i.e. and $P_j \neq R_j$, for some $j \in [1, n]$. Let i denote the smallest such j and suppose that there exists $p \in P_i$ such that $p \notin R_i$ (the symmetrical case $p \in R_i$ and $p \notin P_i$ is left to the reader). By the definition of Δ^d , there exists $q \in P_{i-1} = R_{i-1}$ such that $q \xrightarrow{\sigma_{i-1}, \psi} p \in \Delta$. Since $(v_{i-1}, v_i) \models \theta_{i-1} \wedge \omega_{i-1}$, we obtain that $(v_{i-1}, v_i) \models \bigvee \{ \psi \mid q \xrightarrow{\sigma_{i-1}, \psi} p \in \Delta, q \in P_{i-1} \}$ and $(v_{i-1}, v_i) \models \bigwedge \{ \neg \psi \mid q \xrightarrow{\sigma_{i-1}, \psi} p \in \Delta, q \in P_{i-1} \}$, contradiction. Thus π is the only run of A^d over w , starting in P .

(2) Let $w = (v_0, \sigma_0), \dots, (v_{n-1}, \sigma_{n-1}), (v_n, \diamond)$ be a trace. “ $\subseteq$ ” If $w \in \mathcal{L}(A)$, then A has a run $(q_0, v_0) \xrightarrow{\sigma_0, \phi_0} \dots \xrightarrow{\sigma_{n-1}, \phi_{n-1}} (q_n, v_n)$ such that $q_0 = \mathbf{1}$ and $q_n \in F$. By point (1), A^d has a unique run $(P_0, v_0) \xrightarrow{\sigma_0, \theta_0} \dots \xrightarrow{\sigma_{n-1}, \theta_{n-1}} (P_n, v_n)$ over w . We prove that $q_i \in P_i$, by induction on $i \in [0, n]$. For $i = 0$, we have $P_0 = \{\mathbf{1}\}$, by the definition of A^d . For the induction step, suppose that $i \in [1, n]$ and $q_{i-1} \in P_{i-1}$. By contradiction, assume that $q_i \notin P_i$. Since $(v_{i-1}, v_i) \models_{\text{Th}(\mathcal{D})} \theta_{i-1}$, we obtain $(v_{i-1}, v_i) \models_{\text{Th}(\mathcal{D})} \neg \phi_{i-1}$, contradiction. Thus $q_i \in P_i$ for all $i \in [0, n]$, and $q_n \in P_n$, hence $P_n \cap F \neq \emptyset$. Then $P_n \in F^d$, and $w \in \mathcal{L}(A^d)$. “ $\supseteq$ ” If $w \in \mathcal{L}(A^d)$, then A^d has a (unique) run $(P_0, v_0) \xrightarrow{\sigma_0, \theta_0} (P_1, v_1) \dots \xrightarrow{\sigma_{n-1}, \theta_{n-1}} (P_n, v_n)$ over w , such that $P_0 = \{\mathbf{1}\}$ and $P_n \cap F \neq \emptyset$. Then there exists $p_n \in P_n \cap F$ and, by the definition of A^d , there exists $p_{n-1} \in P_{n-1}$ such that $p_{n-1} \xrightarrow{\sigma_{n-1}, \psi_{n-1}} p_n \in \Delta$ and $(v_{n-1}, v_n) \models_{\text{Th}(\mathcal{D})} \psi_{n-1}$. Continuing this argument backwards, we can find a run $(q_0, v_0) \xrightarrow{\sigma_0, \psi_0} \dots \xrightarrow{\sigma_{n-1}, \psi_{n-1}} (q_n, v_n)$ of A over w , such that $q_i \in P_i$, for all $i \in [0, n]$. Since $P_0 = \{\mathbf{1}\}$ and $q_n \in F$, we obtain that $w \in \mathcal{L}(A)$. $\square$

The construction of a deterministic DA recognizing the language of A is key to defining a DA that recognizes the complement of A . Let $\bar{A} = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q^d, \mathbf{1}^d, Q^d \setminus F^d, \Delta^d \rangle$. In other words, $\bar{A}$ has the same structure as A^d , and the set of final states consists of those subsets that contain no final state, i.e. $\{P \subseteq Q \mid P \cap F = \emptyset\}$. Using Lemma 1, it is not difficult to show that $\mathcal{L}(\bar{A}) = (\mathcal{D}^{\mathbf{x}} \times \Sigma)^* \setminus \mathcal{L}(A)$.

Next, we show closure of DA under intersection. Let $B = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q', \iota', F', \Delta' \rangle$ be a DA and define $A \times B = \langle \mathcal{D}, \Sigma, \mathbf{x}, Q \times Q', (\iota, \iota'), F \times F', \Delta^\times \rangle$, where $(q, q') \xrightarrow{\sigma, \phi} (p, p') \in \Delta^\times$ if and only if $q \xrightarrow{\sigma, \phi} p \in \Delta$, $q' \xrightarrow{\sigma, \psi} p' \in \Delta'$ and $\phi \equiv \phi \wedge \psi$. It is easy to show that $\mathcal{L}(A \times B) = \mathcal{L}(A) \cap \mathcal{L}(B)$. DA are also closed under union, since $\mathcal{L}(A) \cup \mathcal{L}(B) = \mathcal{L}(\overline{A \times B})$.

Let us turn now to the trace inclusion problem. The following lemma shows that this problem can be effectively reduced to an equivalent language emptiness problem. However, note that this reduction does not work when the trace inclusion problem is generalized by removing the condition $\mathbf{x}_B \subseteq \mathbf{x}_A$. In other words, if the observer uses local variables not shared with the network², i.e. $\mathbf{x}_B \setminus \mathbf{x}_A \neq \emptyset$, the generalized trace inclusion problem $\mathcal{L}(A) \downarrow_{\mathbf{x}_A \cap \mathbf{x}_B} \subseteq \mathcal{L}(B) \downarrow_{\mathbf{x}_A \cap \mathbf{x}_B}$ has a negative answer iff *there exists a trace* $w = (v_0, \sigma_0), \dots, (v_n, \diamond) \in \mathcal{L}(A)$ such that, *for all valuations* $\mu_0, \dots, \mu_n \in \mathcal{D}^{\mathbf{x}_B \setminus \mathbf{x}_A}$, we have $w' = (v_0 \downarrow_{\mathbf{x}_A \cap \mathbf{x}_B} \cup \mu_0, \sigma_0), \dots, (v_n \downarrow_{\mathbf{x}_A \cap \mathbf{x}_B} \cup \mu_n, \diamond) \notin \mathcal{L}(B)$. This kind of quantifier alternation cannot be easily accommodated within the framework of language emptiness, in which only one type of (existential) quantifier occurs.

Lemma 2. *Given DA $A = \langle \mathcal{D}, \Sigma, \mathbf{x}_A, Q_A, \iota_A, F_A, \Delta_A \rangle$ and $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \iota_B, F_B, \Delta_B \rangle$ such that $\mathbf{x}_B \subseteq \mathbf{x}_A$. Then $\mathcal{L}(A) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B) \downarrow_{\mathbf{x}_B}$ if and only if $\mathcal{L}(A \times \overline{B}) = \emptyset$.*

Proof. We have $\mathcal{L}(A) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B) \downarrow_{\mathbf{x}_B}$ iff $\mathcal{L}(A) \downarrow_{\mathbf{x}_B} \cap \mathcal{L}(\overline{B}) = \mathcal{L}(A \times \overline{B}) \downarrow_{\mathbf{x}_B} = \emptyset$ iff $\mathcal{L}(A \times \overline{B}) = \emptyset$. $\square$

The trace inclusion problem is undecidable, which can be shown by reduction from the language emptiness problem for DA (take B such that $\mathcal{L}(B) = \emptyset$). However the above lemma shows that any semi-decision procedure for the language emptiness problem can also be used to deal with the trace inclusion problem.

4 Abstract, Check, and Refine for Trace Inclusion

This section describes our semi-algorithm for checking the trace inclusion between a given network $\mathcal{A}$ and an observer B . Let $\mathcal{A}^e$ denote the expansion of $\mathcal{A}$, defined in the previous. In the light of Lemma 2, the trace inclusion problem $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B)$, where the set of observable variables $\mathbf{x}_B$ is included in the set of network variables, can be reduced to the language emptiness problem $\mathcal{L}(\mathcal{A}^e \times \overline{B}) = \emptyset$.

Although language emptiness is undecidable for data automata [23], several cost-effective semi-algorithms and tools [17, 21, 14, 3] have been developed, showing that it is possible, in many practical cases, to provide a yes/no answer to this problem. However, to apply one of the existing off-the-shelf tools to our problem, one needs to build the product automaton $\mathcal{A}^e \times \overline{B}$ prior to the analysis. Due to the inherent state explosion caused by the interleaving semantics of the network as well as by the complementation of the observer, such a solution would not be efficient in practice.

To avoid building the product automaton, our procedure builds *on-the-fly* an over-approximation of the (possibly infinite) set of reachable configurations of $\mathcal{A}^e \times \overline{B}$. This over-approximation is defined using the approach of *lazy predicate abstraction* [17],

² For timed automata, this is the case since the only shared variable is the time, and the observer may have local clocks.

combined with *counterexample-driven abstraction refinement* using *interpolants* [21]. We store the explored abstract states in a structure called an *antichain tree*. In general, antichain-based algorithms [28] store only states which are incomparable w.r.t. a partial order called *subsumption*. Our method can be thus seen as an extension of the antichain-based language inclusion algorithm [1] to infinite state systems by means of predicate abstraction and interpolation-based refinement. Since the trace inclusion problem is undecidable in general, termination of our procedure is not guaranteed; in the following, we shall, however, call our procedure an algorithm for the sake of brevity.

4.1 Antichain Trees

In this section, we define antichain trees, which are the main data structure of the trace inclusion (semi-)algorithm. Let $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ be a network of automata where $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{t}_i, F_i, \Delta_i \rangle$, for all $i \in [1, N]$, and let $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{t}_B, F_B, \Delta_B \rangle$ be an observer such that $\mathbf{x}_B \subseteq \bigcup_{i=1}^N \mathbf{x}_i$. We also denote by $\mathcal{A}^e = \langle \mathcal{D}, \Sigma_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}}, Q_{\mathcal{A}}, \mathbf{t}_{\mathcal{A}}, F_{\mathcal{A}}, \Delta_{\mathcal{A}} \rangle$ the expansion of the network $\mathcal{A}$ and by $\mathcal{A}^e \times B = \langle \mathcal{D}, \Sigma_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}}, Q^p, \mathbf{t}^p, F^p, \Delta^p \rangle$ the product automaton used for checking language inclusion.

An *antichain tree* for the network $\mathcal{A}$ and the observer B is a tree whose nodes are labeled by *product states* (see Fig. 2 for examples). Intuitively, a product state is an over-approximation of the set of configurations of the product automaton $\mathcal{A}^e \times B$ that share the same control state. Formally, a *product state for $\mathcal{A}$ and B* is a tuple $s = (\mathbf{q}, P, \Phi)$ where (i) $(\mathbf{q}, P)$ is a state of $\mathcal{A}^e \times B$ with $\mathbf{q} = \langle q_1, \dots, q_N \rangle$ being a state of the network expansion $\mathcal{A}^e$ and P being a set of states of the observer B , and (ii) $\Phi(\mathbf{x}_{\mathcal{A}}) \in \text{Th}(\mathcal{D})$ is a formula which defines an over-approximation of the set of valuations of the variables $\mathbf{x}_{\mathcal{A}}$ that reach the state $(\mathbf{q}, P)$ in $\mathcal{A}^e \times B$. A product state $s = (\mathbf{q}, P, \Phi)$ is a finite representation of a possibly infinite set of configurations of $\mathcal{A}^e \times B$, denoted as $\llbracket s \rrbracket = \{(\mathbf{q}, P, \mathbf{v}) \mid \mathbf{v} \models_{\text{Th}(\mathcal{D})} \Phi\}$.

To build an over-approximation of the set of reachable states of the product automaton, we need to compute, for a product state s , an over-approximation of the set of configurations that can be reached in one step from s . To this end, we define first a finite abstract domain of product states, based on the notion of *predicate map*. A predicate map is a partial function that associates sets of facts about the values of the variables used in the product automaton, called *predicates*, with components of a product state, called *substates*. The reason behind the distribution of predicates over substates is two-fold. First, we would like the abstraction to be *local*, i.e. the predicates needed to define a certain subtree in the antichain must be associated with the labels of that subtree only. Second, once a predicate appears in the context of a substate, it should be subsequently reused whenever that same substate occurs as part of another product state.

Formally, a *substate* of a state $(\langle q_1, \dots, q_N \rangle, P) \in Q^p$ of the product automaton $\mathcal{A}^e \times B$ is a pair $(\langle q_{i_1}, \dots, q_{i_k} \rangle, S)$ such that (i) $\langle q_{i_1}, \dots, q_{i_k} \rangle$ is a subsequence of $\langle q_1, \dots, q_N \rangle$, and (ii) $S \neq \emptyset$ only if $S \cap P \neq \emptyset$. We denote the substate relation by $(\langle q_{i_1}, \dots, q_{i_k} \rangle, S) \triangleleft (\langle q_1, \dots, q_N \rangle, P)$. The substate relation requires the automata $A_{i_1}, \dots, A_{i_k}$ of the network $\mathcal{A}$ to be in the control states $q_{i_1}, \dots, q_{i_k}$ simultaneously, and the observer B to be in at least some state of S provided $S \neq \emptyset$ (if $S = \emptyset$, the state of B is considered to be irrelevant). Let $\mathcal{S}_{\langle \mathcal{A}, B \rangle} = \{r \mid \exists q \in Q^p. r \triangleleft q\}$ be the set of substates of a state of $\mathcal{A}^e \times B$.

A predicate map $\Pi : \mathcal{S}_{\langle \mathcal{A}, B \rangle} \rightarrow 2^{\text{Th}(\mathcal{D})}$ associates each substate $(\mathbf{r}, S) \in \mathcal{Q}_{i_1} \times \dots \times \mathcal{Q}_{i_k} \times 2^{Q_B}$ with a set of formulae $\pi(\mathbf{x})$ where (i) $\mathbf{x} = \mathbf{x}_{i_1} \cup \dots \cup \mathbf{x}_{i_k} \cup \mathbf{x}_B$ if $S \neq \emptyset$, and (ii) $\mathbf{x} = \mathbf{x}_{i_1} \cup \dots \cup \mathbf{x}_{i_k}$ if $S = \emptyset$. Notice that a predicate associated with a substate refers only to the local variables of those network components $A_{i_1}, \dots, A_{i_k}$ and of the observer B that occur in the particular substate.

Example 1. The antichain in Fig. 2 (d) uses the predicate map $(\langle q_1^1, q_1^2 \rangle, \{p_1\}) \mapsto \{v = 1\}$, $(\langle q_1^1, q_1^2 \rangle, \{p_2\}) \mapsto \{\Delta < x, v = 2\}$. ■

We are now ready to define the abstract semantics of the product automaton $\mathcal{A}^e \times \bar{B}$, induced by a given predicate map. For convenience, we define first a set $\text{Post}(s)$ of concrete successors of a product state $s = (\mathbf{q}, P, \Phi)$ such that $(\mathbf{r}, S, \Psi) \in \text{Post}(s)$ if and only if (i) the product automaton $\mathcal{A}^e \times \bar{B}$ has a rule $(\mathbf{q}, P) \xrightarrow{\sigma, \theta} (\mathbf{r}, S) \in \Delta^p$ and $\Psi(\mathbf{x}_{\mathcal{A}}) \equiv \exists \mathbf{x}'_{\mathcal{A}} . \Phi(\mathbf{x}'_{\mathcal{A}}) \wedge \theta(\mathbf{x}'_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}}) \rightarrow \perp$. The set of concrete successors does not contain states with empty set of valuations because these states are unreachable in $\mathcal{A}^e \times \bar{B}$.

Given a predicate map Π , the set $\text{Post}_{\Pi}(s)$ of abstract successors of a product state s is defined as follows: $(\mathbf{r}, S, \Psi^{\#}) \in \text{Post}_{\Pi}(s)$ if and only if (i) there exists a product state $(\mathbf{r}, S, \Psi) \in \text{Post}(s)$ and (ii) $\Psi^{\#}(\mathbf{x}_{\mathcal{A}}) \equiv \bigwedge_{r \triangleleft (\mathbf{r}, S)} \bigwedge \{\pi \in \Pi(r) \mid \Psi \rightarrow \pi\}$. In other words, the set of data valuations that are reachable by an abstract successor is the tightest over-approximation of the concrete set of reachable valuations, obtained as the conjunction of the available predicates from the predicate map that over-approximate this set.

Example 2. (contd. from Ex. 1) Consider the antichain from Fig. 2 (d). The concrete successors of $s = (\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$ are $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, \Psi_1)$ and $(\langle q_1^1, q_1^2 \rangle, \{p_2\}, \Psi_2)$:

$$\begin{aligned} \Psi_1 &\equiv \exists v', x', \Delta' . v' = 1 \wedge x = x' + 1 \wedge v = 1 \wedge \Delta = \Delta' \wedge 0 \leq x' < \Delta \wedge v = v' \\ \Psi_2 &\equiv \exists v', x', \Delta' . v' = 1 \wedge x = x' + 1 \wedge v = 2 \wedge \Delta = \Delta' \wedge \Delta \leq x' < 2\Delta \wedge v = v' + 1 \end{aligned}$$

With predicate map Π from Ex. 1, $\text{Post}_{\Pi}(s) = \{(\langle q_1^1, q_1^2 \rangle, \{p_1\}, \Psi_1^{\#}), (\langle q_1^1, q_1^2 \rangle, \{p_2\}, \Psi_2^{\#})\}$:

$$\begin{aligned} \Psi_1 \rightarrow v = 1 &\Rightarrow \Psi_1^{\#} \equiv v = 1 \\ \Psi_2 \rightarrow v = 2 \text{ and } \Psi_2 \rightarrow \Delta < x &\Rightarrow \Psi_2^{\#} \equiv v = 2 \wedge \Delta < x . \quad \blacksquare \end{aligned}$$

Finally, an *antichain tree* (or, simply antichain) $\mathcal{T}$ for $\mathcal{A}$ and B is a tree whose nodes are labeled with product states and whose edges are labeled by input symbols and concrete transition relations. Let $\mathbb{N}^*$ be the set of finite sequences of natural numbers that denote the positions in the tree. For a tree position $p \in \mathbb{N}^*$ and $i \in \mathbb{N}$, the position $p.i$ is a *child* of p . A set $S \subseteq \mathbb{N}^*$ is said to be *prefix-closed* if and only if, for each $p \in S$ and each prefix q of p , we have $q \in S$ as well. The root of the tree is denoted by the empty sequence ϵ .

Formally, an antichain $\mathcal{T}$ is a set of pairs $\langle s, p \rangle$, where s is a product state and $p \in \mathbb{N}^*$ is a tree position, such that (1) for each position $p \in \mathbb{N}^*$ there exists at most one product state s such that $\langle s, p \rangle \in \mathcal{T}$, (2) the set $\{p \mid \langle s, p \rangle \in \mathcal{T}\}$ is prefix-closed, (3) $(\text{root}_{\langle \mathcal{A}, B \rangle}, \epsilon) \in \mathcal{T}$ where $\text{root}_{\langle \mathcal{A}, B \rangle} = (\langle \mathbf{u}_1, \dots, \mathbf{u}_N \rangle, \{\mathbf{u}_B\}, \top)$ is the label of the root, and (4) for each edge $(\langle s, p \rangle, \langle t, p.i \rangle)$ in $\mathcal{T}$, there exists a predicate map Π such that $t \in \text{Post}_{\Pi}(s)$. For the latter condition, if $s = (\mathbf{q}, P, \Phi)$ and $t = (\mathbf{r}, S, \Psi)$, there exists a unique rule $(\mathbf{q}, P) \xrightarrow{\sigma, \theta} (\mathbf{r}, S) \in \Delta^p$, and we shall sometimes denote the edge as $s \xrightarrow{\sigma, \theta} t$ or simply $s \xrightarrow{\theta} t$ when the tree positions are not important.

Each antichain node $n = (s, d_1 \dots d_k) \in \mathcal{T}$ is naturally associated with a path from the root to itself $\rho: n_0 \xrightarrow{\sigma_1, \theta_1} n_1 \xrightarrow{\sigma_2, \theta_2} \dots \xrightarrow{\sigma_k, \theta_k} n_k$. We denote by ρ_i the node n_i for each $i \in [0, k]$, and by $|\rho| = k$ the length of the path. The *path formula* associated with ρ is $\Theta(\rho) \equiv \bigwedge_{i=1}^k \theta(\mathbf{x}_{\mathcal{A}}^{i-1}, \mathbf{x}_{\mathcal{A}}^i)$ where $\mathbf{x}_{\mathcal{A}}^i = \{x^i \mid x \in \mathbf{x}_{\mathcal{A}}\}$ is a set of indexed variables for each $i \in [0, k]$.

Example 3. Consider the path $\rho : (\langle q_0^1, q_0^2 \rangle, \{p_0\}, \top) \xrightarrow{\text{init}} (\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1) \xrightarrow{a_2} (\langle q_1^1, q_1^2 \rangle, \{p_2\}, \Delta < x) \xrightarrow{a_2} (\langle q_1^1, q_1^2 \rangle, \emptyset, \Delta < x)$ in the antichain from Fig. 2 (c). The path formula of ρ is $\Theta(\rho) \equiv \theta_1 \wedge \theta_2 \wedge \theta_3$ where:

$$\begin{aligned} \theta_1 &\equiv v_1 = 1 \wedge x_1 = 0 \wedge 0 < \Delta_1 \\ \theta_2 &\equiv v_2 = v_1 + 1 \wedge \Delta_2 = \Delta_1 \wedge v_2 = 2 \wedge x_2 = x_1 + 1 \wedge \Delta_1 \leq x_1 < 2\Delta_1 \wedge \neg(v_2 = v_1) \\ \theta_3 &\equiv v_3 = 2 \wedge \Delta_3 = \Delta_2 \wedge x_3 = x_2 + 1 \wedge \Delta_2 \leq x_2 < 2\Delta_2 \wedge \neg(v_3 = v_2) \quad \blacksquare \end{aligned}$$

4.2 Counterexample-driven Abstraction Refinement

A *counterexample* is a path from the root of the antichain to a node which is labeled by an *accepting* product state. A product state $(\mathbf{q}, P, \Phi)$ is said to be *accepting* iff $(\mathbf{q}, P)$ is an accepting state of the product automaton $\mathcal{A}^e \times \bar{B}$, i.e. $\mathbf{q} \in F_{\mathcal{A}}$ and $P \cap F_B = \emptyset$. A counterexample is said to be *spurious* if its path formula is unsatisfiable, i.e. the path does not correspond to a concrete execution of $\mathcal{A}^e \times \bar{B}$. In this case, we need to (i) remove the path ρ from the current antichain and (ii) refine the abstract domain in order to exclude the occurrence of ρ from future state space exploration.

Let $\rho : \text{root}_{\langle \mathcal{A}, B \rangle} = (\mathbf{q}_0, P_0, \Phi_0) \xrightarrow{\theta_1} (\mathbf{q}_1, P_1, \Phi_1) \xrightarrow{\theta_2} \dots \xrightarrow{\theta_k} (\mathbf{q}_k, P_k, \Phi_k)$ be a spurious counterexample in the following. For efficiency reasons, we would like to save as much work as possible and remove only the smallest suffix of ρ which caused the spuriousness. For some $j \in [0, k]$, let $\Theta^j(\rho) \equiv \Phi_j(\mathbf{x}_{\mathcal{A}}^0) \wedge \bigwedge_{i=j}^k \theta_i(\mathbf{x}_{\mathcal{A}}^{i-j}, \mathbf{x}_{\mathcal{A}}^{i-j+1})$ be the formula defining all sequences of data valuations that start in the set Φ_j and proceed along the suffix $(\mathbf{q}_j, P_j, \Phi_j) \rightarrow \dots \rightarrow (\mathbf{q}_k, P_k, \Phi_k)$ of ρ . The *pivot* of a path ρ is the maximal position $j \in [0, k]$ such that $\Theta^j(\rho) = \perp$, and -1 if ρ is not spurious.

Example 4. (contd. from Ex. 3) The path formula $\Theta(\rho) \equiv \theta_1 \wedge \theta_2 \wedge \theta_3$ from Ex. 3 is unsatisfiable, thus ρ is a spurious counterexample. Moreover, we have $\Theta^1(\rho) \equiv \top \wedge \theta_2 \wedge \theta_3 \rightarrow \perp$ because of the unsatisfiable subformula $v_2 = 2 \wedge v_3 = 2 \wedge \neg(v_3 = v_2)$. Since $\Theta^2(\rho)$ is satisfiable, the pivot of ρ is 1. $\blacksquare$

Finally, we describe the refinement of the predicate map, which ensures that a given spurious counterexample will never be found in a future iteration of the abstract state space exploration. The refinement is based on the notion of *interpolant* [21].

Definition 1. Given a formula $\Phi(\mathbf{x})$ and a sequence $\langle \theta_1(\mathbf{x}, \mathbf{x}'), \dots, \theta_k(\mathbf{x}, \mathbf{x}') \rangle$ of formulae, an *interpolant* is a sequence of formulae $\mathbf{I} = \langle I_0(\mathbf{x}), \dots, I_k(\mathbf{x}) \rangle$ where: (1) $\Phi \rightarrow I_0$, (2) $I_k \rightarrow \perp$, and (3) $I_{i-1}(\mathbf{x}) \wedge \theta_i(\mathbf{x}, \mathbf{x}') \rightarrow I_i(\mathbf{x}')$ for all $i \in [1, k]$.

Any given interpolant is a witness for the unsatisfiability of a (suffix) path formula $\Theta^j(\rho)$. Dually, if *Craig's Interpolation Lemma* [9] holds for the considered first-order data theory $\text{Th}(\mathcal{D})$, any infeasible path formula is guaranteed to have an interpolant.

Example 5. (contd. from Ex. 4) The path formula $\Theta^1(\rho)$ in Ex. 4 has the interpolant $I = \langle \top, v = 2, \perp \rangle$. ■

Given a spurious counterexample ρ with pivot $j \geq 0$, an interpolant $\mathbf{I} = \langle I_0, \dots, I_{k-j} \rangle$ for the infeasible path formula $\Theta^j(\rho)$ can be used to refine the abstract domain by augmenting the predicate map Π . As an effect of this refinement, the antichain construction algorithm will avoid every path with the suffix $(\mathbf{q}_j, P_j, \Phi_j) \rightarrow \dots \rightarrow (\mathbf{q}_k, P_k, \Phi_k)$ in a future iteration. If $I_i \Leftrightarrow C_i^1(\mathbf{y}_1) \wedge \dots \wedge C_i^{m_i}(\mathbf{y}_{m_i})$ is a conjunctive normal form (CNF) of the i -th component of the interpolant, we consider the substate $(\mathbf{r}_i^\ell, S_i^\ell)$ for each $C_i^\ell(\mathbf{y}_\ell)$ where $\ell \in [1, m_i]$:

- $\mathbf{r}_i^\ell = \langle q_{i_1}, \dots, q_{i_h} \rangle$ where $1 \leq i_1 < \dots < i_h \leq N$ is the largest sequence of indices such that $\mathbf{x}_{i_g} \cap \mathbf{y}_\ell \neq \emptyset$ for each $g \in [1, h]$ and the set $\mathbf{x}_{i_g}$ of variables of the network component $\text{DA } A_{i_g}$,
- $S_i^\ell = P_j$ if $\mathbf{x}_B \cap \mathbf{y}_\ell \neq \emptyset$, and $S_i^\ell = \emptyset$, otherwise.

A predicate map Π is said to be *compatible* with a spurious path $\rho : s_0 \xrightarrow{\theta_1} \dots \xrightarrow{\theta_k} s_k$ with pivot $j \geq 0$ if $s_j = (\mathbf{q}_j, P_j, \Phi_j)$ and there is an interpolant $\mathbf{I} = \langle I_0, \dots, I_{k-j} \rangle$ of the suffix $\langle \theta_1, \dots, \theta_k \rangle$ wrt. Φ_j such that, for each clause C of some equivalent CNF of I_i , $i \in [0, k-j]$, it holds that $C \in \Pi(r)$ for some substate $r \triangleleft s_{i+j}$. The following lemma proves that, under a predicate map compatible with a spurious path ρ , the antichain construction will exclude further paths that share the suffix of ρ starting with its pivot.

Lemma 3. *Let $\rho : (\mathbf{q}_0, P_0, \Phi_0) \xrightarrow{\theta_0} (\mathbf{q}_1, P_1, \Phi_1) \xrightarrow{\theta_1} \dots \xrightarrow{\theta_{k-1}} (\mathbf{q}_k, P_k, \Phi_k)$ be a spurious counterexample and Π be a predicate map compatible with ρ . Then, there is no sequence of product states $(\mathbf{q}_j, P_j, \Psi_0), \dots, (\mathbf{q}_k, P_k, \Psi_{k-j})$ such that: (1) $\Psi_0 \rightarrow \Phi_j$ and (2) $(\mathbf{q}_{i+1}, P_{i+1}, \Psi_{i-j+1}) \in \text{Post}_\Pi((\mathbf{q}_i, P_i, \Psi_{i-j}))$ for all $i \in [j, k-1]$.*

Proof. Let $j \in [0, k]$ be the pivot of ρ . Since ρ is spurious, there exists an interpolant $\mathbf{I} = \langle I_0, \dots, I_{k-j} \rangle$ for $\Phi_j \rightarrow \langle \theta_j, \dots, \theta_k \rangle$. It is sufficient to prove that $\Psi_i \rightarrow I_i$ for all $i \in [0, k-j]$. Since $I_{k-j} = \perp$, we obtain $\Psi_{k-j} = \perp$, and consequently $(\mathbf{q}_{k-j}, P_{k-j}, \perp) \in \text{Post}_\Pi((\mathbf{q}_{k-j-1}, P_{k-j-1}, \Psi_{k-j-1}))$. By the definition of Post_Π , we have $(\mathbf{q}_{k-j}, P_{k-j}, \perp) \in \text{Post}((\mathbf{q}_{k-j-1}, P_{k-j-1}, \Psi_{k-j-1}))$, which contradicts with the definition of Post . We show that $\Psi_i \rightarrow I_i$ for all $i \in [0, k-j]$, by induction on $k-j$. For the base case $k-j = 0$, we have $\Psi_0 \rightarrow \Phi_j \rightarrow I_0$. For the induction step, we assume $\Psi_i \rightarrow I_i$ for all $i \in [0, k-j-1]$ and prove $\Psi_{k-j} \rightarrow I_{k-j}$. By the induction hypothesis, we have:

$$\begin{aligned} \Psi_{k-j-1}(\mathbf{x}_{\mathcal{A}}) &\rightarrow I_{k-j-1}(\mathbf{x}_{\mathcal{A}}) \\ \Psi_{k-j-1}(\mathbf{x}_{\mathcal{A}}) \wedge \theta_{k-j-1}(\mathbf{x}_{\mathcal{A}}, \mathbf{x}'_{\mathcal{A}}) &\rightarrow I_{k-j-1}(\mathbf{x}_{\mathcal{A}}) \wedge \theta_{k-j-1}(\mathbf{x}_{\mathcal{A}}, \mathbf{x}'_{\mathcal{A}}) \rightarrow I_{k-j}(\mathbf{x}'_{\mathcal{A}}) . \end{aligned}$$

Let $C_1 \wedge \dots \wedge C_\ell$ be the CNF of I_{k-j} . Since Π is compatible with ρ , for each clause C_i , there exists a substate $r \triangleleft (\mathbf{q}_k, P_k)$ such that $C_i \in \Pi(r)$. By the definition of Post_Π , we obtain that $\Psi_{k-j} \rightarrow C_i$ for each $i \in [1, \ell]$, hence $\Psi_{k-j} \rightarrow I_{k-j}$. □

Observe that the refinement induced by interpolation is *local* since Π associates sets of predicates with substates of the states in $\mathcal{A}^e \times \bar{B}$, and the update impacts only the states occurring within the suffix of that particular spurious counterexample.

4.3 Subsumption

The main optimization of antichain-based algorithms [1] for checking language inclusion of automata over finite alphabets is that product states that are *subsets* of already visited states are never stored in the antichain. On the other hand, language emptiness semi-algorithms, based on *predicate abstraction* [21] use a similar notion to cover newly generated abstract successor states by those that were visited sooner and that represent larger sets of configurations. In this case, state coverage does not only increase efficiency but also ensures termination of the semi-algorithm in many practical cases.

In this section, we generalize the subset relation used in classical antichain algorithms with the notion of coverage from predicate abstraction, and we define a more general notion of *subsumption* for data automata. Given a state $(\mathbf{q}, P)$ of the product automaton $\mathcal{A}^e \times \bar{B}$ and a valuation $\mathbf{v} \in \mathcal{D}^{\mathbf{x}_A}$, the *residual language* $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \bar{B})$ is the set of traces w accepted by $\mathcal{A}^e \times \bar{B}$ from the state $(\mathbf{q}, P)$ such that $\mathbf{v}$ is the first valuation which occurs on w . This notion is then lifted to product states as follows: $\mathcal{L}_s(\mathcal{A}^e \times \bar{B}) = \bigcup_{(\mathbf{q}, P, \mathbf{v}) \in \llbracket s \rrbracket} \mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \bar{B})$ where $\llbracket s \rrbracket$ is the set of configurations of the product automaton $\mathcal{A}^e \times \bar{B}$ represented by the given product state s .

Definition 2. Given a DAN $\mathcal{A}$ and a DA B , a partial order $\sqsubseteq$ is a subsumption provided that, for any two product states s and t , we have $s \sqsubseteq t$ only if $\mathcal{L}_s(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \bar{B})$.

A procedure for checking the emptiness of $\mathcal{A}^e \times \bar{B}$ needs not continue the search from a product state s if it has already visited a product state t that subsumes s . The intuition is that any counterexample discovered from s can also be discovered from t . The trace inclusion semi-algorithm described below in Section 4.4 works, in principle, with any given subsumption relation. In practice, our implementation uses the subsumption relation defined by the lemma below:

Lemma 4. The relation defined s.t. $(\mathbf{q}, P, \Phi) \sqsubseteq_{img} (\mathbf{r}, S, \Psi) \Leftrightarrow \mathbf{q} = \mathbf{r}, P \supseteq S$, and $\Phi \rightarrow \Psi$ is a subsumption.

Proof. For any valuation $\mathbf{v} \in \mathcal{D}^{\mathbf{x}_A}$, we have $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \bar{B}) = \mathcal{L}_{(\mathbf{q}, \mathbf{v})}(\mathcal{A}^e) \cap \mathcal{L}_{(P, \mathbf{v})}(\bar{B})$. Since $P \supseteq S$, we have $\mathcal{L}_{(P, \mathbf{v})}(\bar{B}) \supseteq \mathcal{L}_{(S, \mathbf{v})}(\bar{B})$, thus $\mathcal{L}_{(P, \mathbf{v})}(\bar{B}) \subseteq \mathcal{L}_{(S, \mathbf{v})}(\bar{B})$. We obtain that $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_{(\mathbf{r}, \mathbf{v})}(\mathcal{A}^e) \cap \mathcal{L}_{(S, \mathbf{v})}(\bar{B}) = \mathcal{L}_{(\mathbf{r}, S, \mathbf{v})}(\mathcal{A}^e \times \bar{B})$. Since moreover $\Phi \rightarrow \Psi$, we have that $\mathcal{L}_{(\mathbf{q}, P, \Phi)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_{(\mathbf{r}, S, \Phi)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_{(\mathbf{r}, S, \Psi)}(\mathcal{A}^e \times \bar{B})$. $\square$

Example 6. In the antichain from Fig. 2 (d), $(\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1) \sqsubseteq_{img} (\langle q_1^1, q_1^2 \rangle, \{p_1\}, v = 1)$ because $\langle q_1^1, q_1^2 \rangle = \langle q_1^1, q_1^2 \rangle$, $\{p_1\} \supseteq \{p_1\}$, and $v = 1 \rightarrow v = 1$. $\blacksquare$

As a remark, the language inclusion algorithm for non-deterministic automata on finite alphabets [1] uses also a more sophisticated subsumption relation based on a pre-computed simulation [22] between the states of the automata. We have defined a similar notion of simulation for data automata and an algorithm for computing such simulations [20]. The integration of data simulations within the framework of antichain-based abstraction refinement and its practical assessment are considered as future work.

Algorithm 1 Trace Inclusion Semi-algorithm

```

input:
  1. a DAN  $\mathcal{A} = \langle A_1, \dots, A_N \rangle$  such that  $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{l}_i, F_i, \Delta_i \rangle$  for all  $i \in [1, N]$ ,
  2. a DA  $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{l}_B, F_B, \Delta_B \rangle$  such that  $\mathbf{x}_B \subseteq \bigcup_{i=1}^N \mathbf{x}_i$ .
output: true if  $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B)$ , otherwise a trace  $\tau \in \mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \setminus \mathcal{L}(B)$ .

1:  $\Pi \leftarrow \emptyset, \text{Visited} \leftarrow \emptyset, \text{Next} \leftarrow \langle \text{root}_{\langle \mathcal{A}, B \rangle}, \epsilon \rangle, \text{Subsume} \leftarrow \emptyset$ 
2: while  $\text{Next} \neq \emptyset$  do
3:   chose  $\text{curr} \in \text{Next}$  and move  $\text{curr}$  from  $\text{Next}$  to  $\text{Visited}$ 
4:   match  $\text{curr}$  with  $\langle s, p \rangle$ 
5:   if  $s$  is an accepting product state then
6:     let  $\rho$  be the path from the root to  $\text{curr}$  and  $k$  be the pivot of  $\rho$ 
7:     if  $k \geq 0$  then
8:        $\Pi \leftarrow \text{REFINEPREDICATEMAPBYINTERPOLATION}(\Pi, \rho, k)$ 
9:        $\text{rem} \leftarrow \text{SUBTREE}(\rho_k)$ 
10:      for  $(n, m) \in \text{Subsume}$  such that  $m \in \text{rem}$  do
11:        move  $n$  from  $\text{Visited}$  to  $\text{Next}$ 
12:      remove  $\text{rem}$  from  $(\text{Visited}, \text{Next}, \text{Subsume})$ 
13:      add  $\rho_k$  to  $\text{Next}$ 
14:    else
15:      return  $\text{EXTRACTCOUNTEREXAMPLE}(\rho)$ 
16:   else
17:      $i \leftarrow 0$ 
18:     for  $t \in \text{Post}_{\Pi}(s)$  do
19:       if there exists  $m = \langle t', p' \rangle \in \text{Visited}$  such that  $t \sqsubseteq t'$  then
20:         add  $(\text{curr}, m)$  to  $\text{Subsume}$ 
21:       else
22:          $\text{rem} \leftarrow \{n \in \text{Next} \mid n = \langle t', p' \rangle \text{ and } t' \sqsubset t\}$ 
23:          $\text{succ} \leftarrow \langle t, p.i \rangle$ 
24:          $i \leftarrow i + 1$ 
25:         for  $n \in \text{Visited}$  such that  $n$  has a successor  $m \in \text{rem}$  do
26:           add  $(n, \text{succ})$  to  $\text{Subsume}$ 
27:         for  $(n, m) \in \text{Subsume}$  such that  $m \in \text{rem}$  do
28:           add  $(n, \text{succ})$  to  $\text{Subsume}$ 
29:         remove  $\text{rem}$  from  $(\text{Visited}, \text{Next}, \text{Subsume})$ 
30:         add  $\text{succ}$  to  $\text{Next}$ 

```

4.4 The Trace Inclusion Semi-algorithm

With the previous definitions, Algorithm 1 describes the procedure for checking trace inclusion. It uses a classical worklist iteration loop (lines 2-30) that builds an antichain tree by simultaneously unfolding the expansion $\mathcal{A}^e$ of the network $\mathcal{A}$ and the complement $\bar{B}$ of the the observer B , while searching for a counterexample trace $w \in \mathcal{L}(\mathcal{A}^e \times \bar{B})$. Both $\mathcal{A}^e$ and $\bar{B}$ are built on-the-fly, during the abstract state space exploration.

The processed antichain nodes are kept in the set Visited , and their abstract successors, not yet processed, are kept in the set Next . Initially, $\text{Visited} = \emptyset$ and $\text{Next} = \{\text{root}_{\langle \mathcal{A}, B \rangle}\}$. The algorithm uses a predicate map Π , which is initially empty (line 1).

We keep a set of subsumption edges $\text{Subsume} \subseteq \text{Visited} \times (\text{Visited} \cup \text{Next})$ with the following meaning: $(\langle s, p \rangle, \langle t, q \rangle) \in \text{Subsume}$ for two antichain nodes, where s, t are product states and $p, q \in \mathbb{N}^*$ are tree positions, if and only if there exists an abstract successor $s' \in \text{Post}_{\Pi}(s)$ such that $s' \sqsubseteq t$ (Definition 2). Observe that we do not explicitly store a subsumed successor of a product state s from the antichain; instead, we add a subsumption edge between the node labeled with s and the node that subsumes that particular successor. The algorithm terminates when each abstract successors of a node from Next is subsumed by some node from Visited .

An iteration of Algorithm 1 starts by choosing a current antichain node $\text{curr} = \langle s, p \rangle$ from Next and moving it to Visited (line 3). If the product state s is accepting (line 5) we check the counterexample path ρ , from the root of the antichain to curr , for spuriousness, by computing its pivot k . If $k \geq 0$, then ρ is a spurious counterexample (line 7), and the path formula of the suffix of ρ , which starts with position k , is infeasible. In this case, we compute an interpolant for the suffix and refine the current predicate map Π by adding the predicates from the interpolant to the corresponding substates of the product states from the suffix (line 8).

The computation of the interpolant and the update of the predicate map are done by the `REFINEPREDICATEMAPBYINTERPOLATION` function using the approach described in Section 4.2. Subsequently, we remove (line 12) from the current antichain the subtree rooted at the pivot node ρ_k , i.e. the k -th node on the path ρ (line 9), and add ρ_k to Next in order to trigger a recomputation of this subtree with the new predicate map. Moreover, all nodes with a successor previously subsumed by a node in the removed subtree are moved from Visited back to Next in order to reprocess them (line 11).

On the other hand, if the counterexample ρ is found to be real ($k = -1$), any valuation $v \in \bigcup_{i=0}^{|\rho|} \mathcal{D}^{\mathbf{x}_i}$ that satisfies the path formula $\Theta(\rho)$ yields a counterexample trace $w \in \mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \setminus \mathcal{L}(B)$, obtained by ignoring all variables from $\mathbf{x}_A \setminus \mathbf{x}_B$ (line 15).

If the current node is not accepting, we generate its abstract successors (line 18). In order to keep in the antichain only nodes that are incomparable w.r.t. the subsumption relation $\sqsubseteq$, we add a successor t of s to Next (lines 23 and 30) only if it is not subsumed by another product state from a node $m \in \text{Visited}$. Otherwise, we add a subsumption edge (curr, m) to the set Subsume (line 20). Furthermore, if t is not subsumed by another state in Visited , we remove from Next all nodes $\langle t', p' \rangle$ such that t strictly subsumes t' (lines 22 and 29) and add subsumption edges to the node storing t from all nodes with a removed successor (line 26) or a removed subsumption edge (line 28).

The following theorem states the soundness of our trace inclusion semi-algorithm. The theorem is proved in the appendix together with the other results presented above.

Theorem 1. *Let $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ be a DAN such that $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{l}_i, F_i, \Delta_i \rangle$ for all $i \in [1, N]$, and let $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{l}_B, F_B, \Delta_B \rangle$ be a DA such that $\mathbf{x}_B \subseteq \bigcup_{i=1}^N \mathbf{x}_i$. If Algorithm 1 terminates and returns true on input $\mathcal{A}$ and B , then $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B)$.*

The dual question “if there exists a counterexample trace $w \in \mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \setminus \mathcal{L}(B)$, will Algorithm 1 discover it?” can also be answered positively, using an implementation that enumerates the abstract paths in a systematic way, e.g. by using a breadth-first path exploration. This can be done using a queue to implement the Next set in Algorithm 1.

4.5 Proof of Theorem 1

Given a network $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ where $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{l}_i, F_i, \Delta_i \rangle$ for all $i \in [1, N]$ and an observer $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{l}_B, F_B, \Delta_B \rangle$, we recall that a configuration of the product automaton $\mathcal{A}^e \times \bar{B}$ is a tuple $(\langle q_1, \dots, q_N \rangle, P, v) \in Q_1 \times \dots \times Q_N \times 2^{Q_B} \times \mathcal{D}^{\mathbf{x}_A}$, and a node of the antichain $\mathcal{T}$ is a pair $\langle s, p \rangle$ where s is a product state for $\mathcal{A}$ and B and $p \in \mathbb{N}^*$ is a tree position. Moreover, $\text{root}_{\langle \mathcal{A}, B \rangle} = (\langle \mathbf{l}_1, \dots, \mathbf{l}_N \rangle, \{\mathbf{l}_B\}, \top)$ is the product state that labels the root of $\mathcal{T}$. In the following, let $\Gamma = (\Pi, \text{Visited}, \text{Next}, \text{Subsume})$

be an *antichain state* where Π is the predicate map, and **Visited**, **Next**, and **Subsume** are the sets of antichain nodes handled by Algorithm 1.

We say that Γ is a *closed antichain state* if and only if, for all nodes $\langle s, p \rangle \in \mathbf{Visited}$ and every successor $(\mathbf{q}, P, \mathbf{v}) \in \text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket)$ of a configuration of the product automaton $\mathcal{A}^e \times \bar{B}$ represented by the product state s , there exists a node $\langle t, r \rangle \in \mathbf{Visited} \cup \mathbf{Next}$ such that $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \bar{B})$ and one of the following holds:

- $r = p.i$ for some $i \in \mathbb{N}$, i.e. $\langle t, r \rangle$ is a child of $\langle s, p \rangle$ in the antichain $\mathcal{T} = \mathbf{Visited} \cup \mathbf{Next}$, or
- $(\langle s, p \rangle, \langle t, r \rangle) \in \mathbf{Subsume}$.

In other words, the current antichain $\mathcal{T}$, defined as the union of the sets **Visited** and **Next**, is in a closed state, if the residual language of every successor of a configuration of the product automaton $\mathcal{A}^e \times \bar{B}$ that is covered by a visited product state must be included in the residual language of a product state stored in the antichain, either as a direct successor in the tree or via a subsumption edge.

For a product state s , we define $\text{Dist}(s) = \min \{|w| \mid w \in \mathcal{L}_s(\mathcal{A}^e \times \bar{B})\}$, and $\text{Dist}(s) = \infty$ if and only if $\mathcal{L}_s(\mathcal{A}^e \times \bar{B}) = \emptyset$. For a finite non-empty set of antichain nodes S , we define $\text{Dist}(S) = \min \{\text{Dist}(s) \mid \langle s, p \rangle \in S\}$ with $\text{Dist}(\emptyset) = \infty$.

Lemma 5. *Given a network $\mathcal{A}$ and an observer B , for any product state s of $\mathcal{A}$ and B , we have $\text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket) = \bigcup_{t \in \text{Post}(s)} \llbracket t \rrbracket$.*

Proof. Let $s = (\mathbf{q}, P, \Phi)$. “ $\subseteq$ ” Let $(\mathbf{r}, S, \mu) \in \text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket)$ be a configuration of $\mathcal{A}^e \times \bar{B}$ for which there exists $(\mathbf{q}, P, \mathbf{v}) \in \llbracket s \rrbracket$ such that $(\mathbf{q}, P, \mathbf{v}) \xrightarrow{\sigma, \theta} (\mathbf{r}, S, \mu)$. Then there exists a unique rule $(\mathbf{q}, P) \xrightarrow{\sigma, \theta} (\mathbf{r}, S) \in \Delta^p$ such that $(\mathbf{v}, \mu) \models_{\text{Th}(\mathcal{D})} \theta$. Moreover, if $(\mathbf{q}, P, \mathbf{v}) \in \llbracket s \rrbracket$, we have $\mathbf{v} \models_{\text{Th}(\mathcal{D})} \Phi$. Let $t = (\mathbf{r}, S, \Psi) \in \text{Post}(s)$ where $\Psi(\mathbf{x}_{\mathcal{A}}) \equiv \exists \mathbf{x}'_{\mathcal{A}}. \Phi(\mathbf{x}'_{\mathcal{A}}) \wedge \theta(\mathbf{x}'_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}})$. We have $\mu \models_{\text{Th}(\mathcal{D})} \Psi$, hence $(\mathbf{r}, S, \mu) \in \llbracket t \rrbracket$. “ $\supseteq$ ” Let $(\mathbf{r}, S, \mu) \in \llbracket t \rrbracket$ for some $t \in \text{Post}(s)$. Then we have $t = (\mathbf{r}, S, \Psi)$ where $\Psi(\mathbf{x}_{\mathcal{A}}) \equiv \exists \mathbf{x}'_{\mathcal{A}}. \Phi(\mathbf{x}'_{\mathcal{A}}) \wedge \theta(\mathbf{x}'_{\mathcal{A}}, \mathbf{x}_{\mathcal{A}})$. Since $\mu \models_{\text{Th}(\mathcal{D})} \Psi$, there exists $\mathbf{v} \models_{\text{Th}(\mathcal{D})} \Phi$ such that $(\mathbf{q}, P, \mathbf{v}) \xrightarrow{\sigma, \theta} (\mathbf{r}, S, \mu)$. Hence $(\mathbf{q}, P, \mathbf{v}) \in \llbracket s \rrbracket$, thus $(\mathbf{r}, S, \mu) \in \text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket)$. $\square$

Lemma 6. *Given a network $\mathcal{A}$, an observer B , and a predicate map Π , for any product state s of $\mathcal{A}^e \times \bar{B}$ and any product state $t \in \text{Post}(s)$ there exists $t' \in \text{Post}_{\Pi}(s)$ such that $\llbracket t \rrbracket \subseteq \llbracket t' \rrbracket$.*

Proof. Let $t = (\mathbf{r}, S, \Psi) \in \text{Post}(s)$. By the definition of Post_{Π} , we have $t' = (\mathbf{r}, S, \Psi^{\sharp}) \in \text{Post}_{\Pi}(s)$, where $\Psi \rightarrow \Psi^{\sharp}$, thus $\llbracket t \rrbracket \subseteq \llbracket t' \rrbracket$. $\square$

Lemma 7. *Given a network $\mathcal{A}$, an observer B , and a predicate map Π , for each product state s and each configuration $(\mathbf{q}, P, \mathbf{v}) \in \text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket)$ there exists a product state $t \in \text{Post}_{\Pi}(s)$ such that $(\mathbf{q}, P, \mathbf{v}) \in \llbracket t \rrbracket$.*

Proof. We use the fact that $\text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket) = \bigcup_{t \in \text{Post}(s)} \llbracket t \rrbracket$ (Lemma 5) and that for each $t \in \text{Post}(s)$ there exists $t' \in \text{Post}_{\Pi}(s)$ such that $\llbracket t \rrbracket \subseteq \llbracket t' \rrbracket$ (Lemma 6). $\square$

The proof of soundness of Algorithm 1 relies on the inductive invariants (Inv_1) and (Inv_2) from the following lemma.

Lemma 8. *The following invariants hold each time line 2 is reached in Algorithm 1:*

- (Inv₁) $\Gamma = (\Pi, \text{Visited}, \text{Next}, \text{Subsume})$ is closed,
- (Inv₂) $\text{Dist}(\text{root}_{\langle \mathcal{A}, \mathcal{B} \rangle}) < \infty \Rightarrow \text{Dist}(\text{Visited}) > \text{Dist}(\text{Next})$.

Proof. Initially, when coming to line 2 for the first time, we have $\text{Visited} = \emptyset$, thus $\text{Dist}(\text{Visited}) = \infty$, and both invariants hold trivially. For the case when coming to line 2 after executing the loop body, we denote by:

$$\begin{aligned}\Gamma_{old} &= (\Pi_{old}, \text{Visited}_{old}, \text{Next}_{old}, \text{Subsume}_{old}) \\ \Gamma_{new} &= (\Pi_{new}, \text{Visited}_{new}, \text{Next}_{new}, \text{Subsume}_{new})\end{aligned}$$

the antichain states before and after the execution of the main loop. We assume that both invariants hold for Γ_{old} .

(Inv₁) Let $\langle s, p \rangle \in \text{Visited}_{new}$ and $(\mathbf{q}, P, \mathbf{v}) \in \text{succ}_{\mathcal{A}^e \times \overline{\mathcal{B}}}(\llbracket s \rrbracket)$. We distinguish two cases according to the control path taken inside the main loop:

- (1) If the test on line 5 is positive, the predicate map is augmented, i.e. $\Pi_{new} \supseteq \Pi_{old}$ (line 8). Let $\Gamma' = (\Pi_{new}, \text{Visited}_{old}, \text{Next}_{old}, \text{Subsume}_{old})$ be the next antichain state. Clearly Γ' is closed provided that Γ_{old} is. Next, let $n_{pivot} \in \text{Visited}_{old}$ be the pivot of the path to the current node (line 6) and define the following sets of nodes:

$$\begin{aligned}T &= \text{SUBTREE}(n_{pivot}) \\ S &= \{n \in \text{Visited}_{old} \mid \exists m \in T. (n, m) \in \text{Subsume}_{old}\}\end{aligned}$$

Then we obtain (lines 10–13):

$$\begin{aligned}\text{Visited}_{new} &= \text{Visited}_{old} \setminus (S \cup T) \\ \text{Next}_{new} &= ((\text{Next}_{old} \cup S) \setminus T) \cup \{n_{pivot}\} \\ \text{Visited}_{new} \cup \text{Next}_{new} &= ((\text{Visited}_{old} \cup \text{Next}_{old}) \setminus T) \cup \{n_{pivot}\}\end{aligned}$$

Since Γ' is closed, there exists a node $\langle t, r \rangle \in \text{Visited}_{old} \cup \text{Next}_{old}$ such that $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \overline{\mathcal{B}}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \overline{\mathcal{B}})$ and either $r = p.i$ for some $i \in \mathbb{N}$ or $\langle s, p \rangle, \langle t, r \rangle \in \text{Subsume}_{old}$. We distinguish two cases:

- (a) $\langle t, r \rangle \notin T$. Then $\langle t, r \rangle \in \text{Visited}_{new} \cup \text{Next}_{new}$ and, because $\text{Subsume}_{new} = \text{Subsume}_{old} \cap (\text{Visited}_{new} \times (\text{Visited}_{new} \cup \text{Next}_{new}))$, we obtain that Γ_{new} is closed as well.
- (b) $\langle t, r \rangle \in T$. Then we distinguish two further cases:
 - (i) If $r = p.i$ for some $i \in \mathbb{N}$, since we have assumed that $\langle s, p \rangle \in \text{Visited}_{new}$, we have $\langle s, p \rangle \notin T$. The only possibility is then $\langle t, r \rangle = n_{pivot}$ and $\langle s, p \rangle$ is the parent of n_{pivot} . In this case, we have $\langle t, r \rangle \in \text{Next}_{new}$.
 - (ii) If $(\langle s, p \rangle, \langle t, r \rangle) \in \text{Subsume}_{old}$, then $\langle s, p \rangle \in S$, which contradicts the assumption $\langle s, p \rangle \in \text{Visited}_{new}$.
- (2) Otherwise, the test on line 5 is negative, in which case we have $\Pi_{new} = \Pi_{old}$ and $\text{Visited}_{new} = \text{Visited}_{old} \cup \{\text{curr}\}$. For each $(\mathbf{q}, P, \mathbf{v}) \in \text{succ}_{\mathcal{A}^e \times \overline{\mathcal{B}}}(\llbracket s \rrbracket)$ there exists $t \in \text{Post}_{\Pi}(s)$ such that $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \overline{\mathcal{B}}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \overline{\mathcal{B}})$ (by Lemma 7). We distinguish two cases:
 - (a) $\langle s, p \rangle = \text{curr}$. In this case, either (i) there is $\langle t', p' \rangle \in \text{Visited}_{old}$ such that $t \sqsubseteq t'$, and then we also have $\mathcal{L}_{(\mathbf{q}, P, \mathbf{v})}(\mathcal{A}^e \times \overline{\mathcal{B}}) \subseteq \mathcal{L}_{t'}(\mathcal{A}^e \times \overline{\mathcal{B}})$ (Definition 2) and

- $\langle\langle s, p \rangle, \langle t', p' \rangle\rangle \in \text{Subsume}_{new}$ (added on line 20), or (ii) $(t, p.i) \in \text{Next}_{new}$ for some $i \in \mathbb{N}$ (added on lines 23 and 30).
- (b) Otherwise $\langle s, p \rangle \in \text{Visited}_{old}$. As Γ' is closed, there is $\langle u, r \rangle \in \text{Visited}_{old} \cup \text{Next}_{old}$ such that $\mathcal{L}_{(q,p,v)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_u(\mathcal{A}^e \times \bar{B})$ and either $r = p.i$ for some $i \in \mathbb{N}$ or $\langle\langle s, p \rangle, \langle u, r \rangle\rangle \in \text{Subsume}_{old}$. We distinguish two sub-cases:
- (i) $\langle u, r \rangle \in \text{rem}$ (line 22). Then $\mathcal{L}_u(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \bar{B})$ (Definition 2), hence $\mathcal{L}_{(q,p,v)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \bar{B})$. If $r = p.i$, then $\langle\langle s, p \rangle, \langle t, r' \rangle\rangle \in \text{Subsume}_{new}$ for some $r' \in \mathbb{N}^*$ (added on line 26). Else, if $\langle\langle s, p \rangle, \langle u, r \rangle\rangle \in \text{Subsume}_{old}$, we have $\langle\langle s, p \rangle, \langle t, r' \rangle\rangle \in \text{Subsume}_{new}$ for some $r' \in \mathbb{N}^*$ (added on line 28). In both cases, we obtain that Γ_{new} is closed.
 - (ii) $\langle u, r \rangle \notin \text{rem}$. Then $\langle u, r \rangle \in \text{Visited}_{new} \cup \text{Next}_{new}$. Since $\text{Subsume}_{new} = \text{Subsume}_{old} \cap (\text{Visited}_{new} \times (\text{Visited}_{new} \cup \text{Next}_{new}))$, we obtain that Γ_{new} is closed.

(Inv₂) We distinguish two cases:

1. If $\text{Dist}(\text{Visited}_{new}) = \infty$, it is sufficient to show that $\text{Dist}(\text{Next}_{new}) < \infty$. Suppose, by contradiction, that $\text{Dist}(\text{Next}_{new}) = \infty$, hence $\text{Dist}(\text{Visited}_{new} \cup \text{Next}_{new}) = \infty$, and since $\text{root}_{\langle \mathcal{A}, B \rangle} \in \text{Visited}_{new} \cup \text{Next}_{new}$, we obtain $\text{Dist}(\text{root}_{\langle \mathcal{A}, B \rangle}) = \infty$, contradiction.
2. Otherwise, $\text{Dist}(\text{Visited}_{new}) < \infty$ and there exists a node $\langle s, p \rangle \in \text{Visited}_{new}$ such that $\text{Dist}(\text{Visited}_{new}) = \text{Dist}(s) < \infty$. Let $w = (v_0, \sigma_0), (v_1, \sigma_1), \dots, (v_n, \diamond) \in \mathcal{L}_s(\mathcal{A}^e \times \bar{B})$ be a trace such that $\text{Dist}(\text{Visited}_{new}) = n$. Then there exists a run $(q_0, P_0, v_0) \xrightarrow{\sigma_0} (q_1, P_1, v_1) \xrightarrow{\sigma_1} \dots \xrightarrow{\sigma_{n-1}} (q_n, P_n, v_n)$ of $\mathcal{A}^e \times \bar{B}$ over w such that $(q_0, P_0, v_0) \in \llbracket s \rrbracket$ and (q_n, P_n) a final state of $\mathcal{A}^e \times \bar{B}$. Since Γ_{new} is closed due to (Inv₁) and $(q_1, P_1, v_1) \in \text{succ}_{\mathcal{A}^e \times \bar{B}}(\llbracket s \rrbracket)$, there exists a node $\langle s_1, p_1 \rangle \in \text{Visited}_{new} \cup \text{Next}_{new}$ such that $\mathcal{L}_{(q_1, P_1, v_1)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_{s_1}(\mathcal{A}^e \times \bar{B})$. If $\langle s_1, p_1 \rangle \in \text{Next}_{new}$, we obtain that $\text{Dist}(\text{Next}_{new}) \leq n-1$, and we are done. Otherwise, $\langle s_1, p_1 \rangle \in \text{Visited}_{new}$, and we can repeat the same argument inductively, to discover a sequence of nodes $\langle s_1, p_1 \rangle, \dots, \langle s_n, p_n \rangle \in \text{Visited}_{new}$ such that $\mathcal{L}_{(q_i, P_i, v_i)}(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_{s_n}(\mathcal{A}^e \times \bar{B})$ for all $i \in [1, n]$. Since (q_n, P_n) is a final state of $\mathcal{A}^e \times \bar{B}$, we have $(v_n, \diamond) \in \mathcal{L}_{(q_i, P_i, v_i)}(\mathcal{A}^e \times \bar{B})$, thus $(v_n, \diamond) \in \mathcal{L}_{s_n}(\mathcal{A}^e \times \bar{B})$, and s_n is an accepting product state. But this contradicts with the fact that accepting product states are never stored in the antichain. $\square$

Back to the proof of Theorem 1:

Proof. If Algorithm 1 terminates and reports true, this is because $\text{Next} = \emptyset$, hence $\text{Dist}(\text{Next}) = \infty$. By Lemma 8 (Inv₂), we obtain that $\text{Dist}(\text{root}_{\langle \mathcal{A}, B \rangle}) = \infty$. Suppose, by contradiction, that $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \not\subseteq \mathcal{L}(B)$. By Lemma 2, there exists a trace

$$w = (v_0, \sigma_0)(v_1, \sigma_1) \dots (v_n, \diamond) \in \mathcal{L}(\mathcal{A}^e \times \bar{B}) .$$

Thus we have a run of $\mathcal{A}^e \times \bar{B}$ over w :

$$(q_0, P_0, v_0) \xrightarrow{\sigma_0} (q_1, P_1, v_1) \xrightarrow{\sigma_1} \dots \xrightarrow{\sigma_{n-1}} (q_n, P_n, v_n)$$

where $\mathbf{q}_0 = \langle \mathbf{1}_1, \dots, \mathbf{1}_N \rangle$, $P_0 = \{\mathbf{1}_B\}$, $\mathbf{q}_n$ is final in $\mathcal{A}^e$, $P_n \cap F_B = \emptyset$. But, since $(\mathbf{q}_0, P_0, \mathbf{v}_0) \in \llbracket \text{root}_{\langle \mathcal{A}, B \rangle} \rrbracket$, we have $w \in \mathcal{L}_{\text{root}_{\langle \mathcal{A}, B \rangle}}(\mathcal{A}^e \times \overline{B})$. Hence, $\text{Dist}(\text{root}_{\langle \mathcal{A}, B \rangle}) \leq n$, which is in contradiction with the fact that $\text{Dist}(\text{root}_{\langle \mathcal{A}, B \rangle}) = \infty$. Consequently, it must be the case that $\mathcal{L}(\mathcal{A}) \downarrow_{\mathbf{x}_B} \subseteq \mathcal{L}(B)$. $\square$

5 Computing Simulations of Data Automata

In the case of classical Rabin-Scott finite automata over finite alphabets, a *simulation* [22] is a relation on the states of an automaton, which is invariant with respect to the transition relation of the automaton. The simulation-based approach for checking language inclusion between two automata A and B first computes a simulation relation on the union of the states of A and B , then checks whether the pair of initial states is a member of the simulation relation. This method is not complete because there exist automata, such that $\mathcal{L}(A) \subseteq \mathcal{L}(B)$, but the initial state of A is not simulated by the initial state of B . A pre-computed simulation relation can be used however to speed up the convergence of the antichain-based method, by weakening the subsumption relation used by the antichain construction algorithm [1].

In this section, we define a notion of simulation between data automata and give an algorithm that computes useful under-approximations of the weakest simulation on a data automaton. The simulation relation can be used to enhance the convergence of Algorithm 1, similar to the way in which classical simulations are integrated with the antichain-based language inclusion algorithm for automata over finite alphabets [1].

Definition 3. A relation $R \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$ is a data simulation for a DA $A = \langle \Sigma, \mathcal{D}, \mathbf{x}, Q, \mathbf{1}, F, \Delta \rangle$ if and only if, for all $(q, \mathbf{v}, q') \in R$ the following hold:

1. $q \in F$ only if $q' \in F$, and
2. for all $\sigma \in \Sigma$ and all $(r, \mathbf{v}') \in Q \times \mathcal{D}^{\mathbf{x}}$ such that $(q, \mathbf{v}) \xrightarrow{\sigma}_A (r, \mathbf{v}')$ there exists $r' \in Q$ such that $(q', \mathbf{v}) \xrightarrow{\sigma}_A (r', \mathbf{v}')$ and $(r, \mathbf{v}', r') \in R$.

Observe that, while a classical simulation is a binary relation on states, a data simulation is a ternary relation involving also a valuation of the variables. The following lemma shows that a data simulation preserves residual languages.

Lemma 9. Given a DA $A = \langle \Sigma, \mathcal{D}, \mathbf{x}, Q, \mathbf{1}, F, \Delta \rangle$ and $R \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$ a data simulation for A , for any tuple $(q, \mathbf{v}, q') \in R$ we have $\mathcal{L}_{(q, \mathbf{v})}(A) \subseteq \mathcal{L}_{(q', \mathbf{v})}(A)$.

Proof. Let $(q, \mathbf{v}) = (q_0, \mathbf{v}_0) \xrightarrow{\sigma_0} \dots \xrightarrow{\sigma_{n-1}} (q_n, \mathbf{v}_n)$ be a run of A such that $q_n \in F$. By induction on $n \geq 0$ we find a run $(q', \mathbf{v}) = (q'_0, \mathbf{v}_0) \xrightarrow{\sigma_0} \dots \xrightarrow{\sigma_{n-1}} (q'_n, \mathbf{v}_n)$ of A , such that $(q_i, \mathbf{v}_i, q'_i) \in R$, for all $i \in [0, n]$. Moreover, since $q_n \in F$, we also obtain q'_n . $\square$

Let $A = \langle \Sigma, \mathcal{D}, \mathbf{x}, Q, \mathbf{1}, F, \Delta \rangle$, where $Q = \{q_1, \dots, q_k\}$, for some $k > 0$, be a DA for the rest of this section. The data simulation algorithm (Algorithm 2) given in this section manipulates sets of valuations from $\mathcal{D}^{\mathbf{x}}$ that are definable by first-order formulae in $\text{Th}(\mathcal{D})$. A relation $R \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$ is said to be *definable* if and only if there exists a matrix $\Phi = [\phi_{ij}]_{i,j=1}^k$ of formulae $\phi_{ij}(\mathbf{x}) \in \text{Th}(\mathcal{D})$ such that $(q_i, \mathbf{v}, q_j) \in R \Leftrightarrow \mathbf{v} \models \phi_{ij}$. For $\ell \in [1, k]$, we denote by Φ_ℓ the ℓ -th row of the matrix Φ .

Algorithm 2 is a refinement algorithm which handles two matrices of formulae that define the relations $Sim, PrevSim \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$. In the following we shall use the same names to denote the relations and their matrix representations. Intuitively, $PrevSim$ is the previous candidate for simulation, whereas Sim is an entry-wise stronger relation, that refines $PrevSim$. The refinement step is performed backwards, with respect to each transition rule $q_i \xrightarrow{\sigma, \phi} q_\ell$ of the automaton. Namely, for each pair of valuations such that $(v, v') \models_{Th(\mathcal{D})} \phi$ and $(q_\ell, v', q_m) \in PrevSim$ for some state q_m , we add the tuple $(q_i, v, q_j) \in Sim$ for all predecessors q_j of q_m , such that $q_j \xrightarrow{\sigma, \psi} q_m$ and $(v, v') \models_{Th(\mathcal{D})} \psi$. This update guarantees that, for every transition $(q_i, v) \xrightarrow{\sigma} A(q_\ell, v')$, where $(q_i, v, q_j) \in Sim$ there exists a state q_m such that $(q_j, v) \xrightarrow{\sigma} A(q_m, v')$ and $(q_\ell, v', q_m) \in PrevSim$. The algorithm stops when Sim and $PrevSim$ define the same relation, which is, moreover, a simulation.

To define the update, we use the following function:

$$PrevSim_\sigma(i, j, \ell, R) \equiv \forall \mathbf{x}' . \phi(\mathbf{x}, \mathbf{x}') \rightarrow \bigvee_{q_j \xrightarrow{\sigma, \psi} q_m} \psi(\mathbf{x}, \mathbf{x}') \wedge R_{\ell m}(\mathbf{x}'), \text{ where } q \xrightarrow{\sigma, \phi} q' \in \Delta .$$

We define also the sets $post_\sigma(q) = \{q' \mid q \xrightarrow{\sigma, \phi} q' \in \Delta\}$ and $pre_\sigma(q) = \{q' \mid q' \xrightarrow{\sigma, \phi} q \in \Delta\}$. With this notation, Algorithm 2 describes the procedure that computes a data simulation for a given data automaton.

Initially, the matrix $PrevSim$ is true everywhere (line 3). The current simulation candidate Sim is initialized to false for all $i, j \in [1, k]$ such that $q_i \in F$ and $q_j \notin F$ (line 7). Observe that, in this case q_j cannot simulate q_i , by Definition 3 (1). Otherwise, we initialize Sim_{ij} to the strongest pre-simulation with respect to $PrevSim$ (line 9). In the iterative loop (lines 10–20) the algorithm choses a state q_ℓ for which the current simulation candidate Sim_ℓ is not equivalent to the previous one $PrevSim_\ell$ (line 10) and sharpens the set Sim_{ij} , with respect to the transition rule $q_i \xrightarrow{\sigma, \phi} q_\ell$, for all input symbols $\sigma \in \Sigma$ and all peer states $q_j, j \in [1, k]$ (line 14). The following invariants are key to proving the correctness of Algorithm 2.

Lemma 10. *The following invariants hold each time Algorithm 2 reaches line 10:*

- ($SimInv_1$) $\forall i, j \in [1, k]: Sim_{ij} \rightarrow PrevSim_{ij}$.
- ($SimInv_2$) $\forall \sigma \in \Sigma \forall i, j, \ell \in [1, k] \forall v, v' \in \mathcal{D}^{\mathbf{x}}: v \models_{Th(\mathcal{D})} Sim_{ij} \text{ and } (q_i, v) \xrightarrow{\sigma} (q_\ell, v') \Rightarrow \exists m \in [1, k]: (q_j, v) \xrightarrow{\sigma} (q_m, v') \text{ and } v' \models_{Th(\mathcal{D})} PrevSim_{\ell m}$.

Proof. Let Sim' and $PrevSim'$ denote the global matrices after one iteration of the loop on lines 10–20.

($SimInv_1$) When line 10 is reached for the first time, $PrevSim_{ij} = \top$, for all $i, j \in [1, k]$, thus $SimInv_1$ holds initially. Since Sim is modified only at lines 14 or 17, we have $Sim_{ij} \rightarrow Sim'_{ij}$, for all $i, j \in [1, k]$. Moreover, either $PrevSim'_{ij} = Sim_{ij}$, or $PrevSim'_{ij} = PrevSim_{ij}$, for all $i, j \in [1, k]$ (line 20). Thus $PrevSim'_{ij} \rightarrow Sim_{ij} \rightarrow Sim'_{ij}$, for all $i, j \in [1, k]$, by an application of the induction hypothesis.

($SimInv_2$) We show that this invariant holds the first time the control reaches line 10. Let $\sigma \in \Sigma, i, j, \ell \in [1, k]$ and $v, v' \in \mathcal{D}^{\mathbf{x}}$ such that $v \models_{Th(\mathcal{D})} Sim_{ij}$ and $(q_i, v) \xrightarrow{\sigma} (q_\ell, v')$. Since

Algorithm 2 Data Simulation Algorithm

input: a data automaton $A = \langle \Sigma, \mathcal{D}, \mathbf{x}, Q, \mathfrak{t}, F, \Delta \rangle$, where $Q = \{q_1, \dots, q_k\}$ and a constant $K > 0$
output: a data simulation $R \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$ for A
global vars $[Sim_{ij}]_{i,j=1}^k, [PrevSim_{ij}]_{i,j=1}^k, [Cnt_{ij}]_{i,j=1}^k$

- 1: **for** $i = 1, \dots, k$ **do**
- 2: **for** $j = 1, \dots, k$ **do**
- 3: $PrevSim_{ij} \leftarrow \top$
- 4: $Cnt_{ij} \leftarrow K$
- 5: **for** $j = 1, \dots, k$ **do**
- 6: **if** $q_i \in F$ and $q_j \notin F$ **then**
- 7: $Sim_{ij} \leftarrow \perp$
- 8: **else**
- 9: $Sim_{ij} \leftarrow \bigwedge_{\sigma \in \Sigma} \bigwedge_{q_\ell \in post_\sigma(q_i)} PreSim_\sigma(i, j, \ell, PrevSim)$
- 10: **for all** $\ell \in [1, k]$ such that $Sim_\ell \not\Leftarrow PrevSim_\ell$ **do**
- 11: **for** $\sigma \in \Sigma$ **do**
- 12: **for** $q_i \in pre_\sigma(q_\ell)$ **do**
- 13: **for** $j = 1, \dots, k$ **do**
- 14: $Sim_{ij} \leftarrow Sim_{ij} \wedge PreSim_\sigma(i, j, \ell, Sim)$
- 15: **for all** $j = 1, \dots, k$ such that $Sim_{\ell j} \not\Leftarrow PrevSim_{\ell j}$ **do**
- 16: **if** $Cnt_{\ell j} = 0$ **then**
- 17: $Sim_{\ell j} \leftarrow \perp$
- 18: **else**
- 19: $Cnt_{\ell j} \leftarrow Cnt_{\ell j} - 1$
- 20: $PrevSim_\ell \leftarrow Sim_\ell$
- 21: **return** Sim

$v \models_{Th(\mathcal{D})} Sim_{ij}$ (thus $Sim_{ij} \neq \perp$) and $q_\ell \in post_\sigma(q_i)$ we have that $v \models_{Th(\mathcal{D})} PreSim_\sigma(i, j, \ell, PrevSim)$, where $q_i \xrightarrow{\sigma, \phi} q_\ell \in \Delta$. Since $(q_i, v) \xrightarrow{\sigma} (q_\ell, v')$ we obtain that $(v, v') \models_{Th(\mathcal{D})} \phi(\mathbf{x}, \mathbf{x}')$, and consequently $(v, v') \models_{Th(\mathcal{D})} \Psi(\mathbf{x}, \mathbf{x}') \wedge PrevSim_{\ell m}(\mathbf{x}')$, for some $m \in [1, k]$, such that $q_j \xrightarrow{\sigma, \Psi} q_m \in \Delta$. Hence $SimInv_2$ holds when the control first reaches line 10.

For the induction step, let us assume that $SimInv_2$ holds at line 10 and we prove that it holds also after executing line 20. Let $\sigma \in \Sigma, i, j, \ell \in [1, k]$ and $v, v' \in \mathcal{D}^{\mathbf{x}}$ such that $v \models_{Th(\mathcal{D})} Sim'_{ij}$ and $(q_i, v) \xrightarrow{\sigma} (q_\ell, v')$. We distinguish two cases:

1. if $Sim_\ell \not\Leftarrow PrevSim_\ell$ on line 10, since $q_i \in pre_\sigma(q_\ell)$, then Sim'_{ij} was updated at line 14. Since $v \models_{Th(\mathcal{D})} Sim'_{ij}$, we obtain $v \models_{Th(\mathcal{D})} PreSim_\sigma(q_i, q_j, q_\ell, Sim)$. Moreover, $PrevSim'_\ell$ is updated to Sim'_ℓ at line 20, hence $v \models_{Th(\mathcal{D})} PreSim_\sigma(q_i, q_j, q_\ell, PrevSim')$ as well. Since $(q_i, v) \xrightarrow{\sigma} (q_\ell, v')$, we obtain that $(v, v') \models_{Th(\mathcal{D})} \Psi(\mathbf{x}, \mathbf{x}') \wedge PrevSim'_{\ell m}(\mathbf{x}')$, for some $m \in [1, k]$ such that $q_j \xrightarrow{\sigma, \Psi} q_m \in \Delta$, thus $(v, v') \models_{Th(\mathcal{D})} \Psi(\mathbf{x}, \mathbf{x}')$ and $v' \models_{Th(\mathcal{D})} PrevSim'_{\ell m}$. Thus $SimInv_2$ holds for Sim' and $PrevSim'$.
2. else $Sim_\ell \Leftarrow PrevSim_\ell$ on line 10, $PrevSim'_\ell \Leftarrow PrevSim_\ell$ because the update on line 20 is skipped, and for all $q_i \in pre_\sigma(q_\ell)$ and all $j \in [1, k]$, we have $Sim'_{ij} \Leftarrow Sim_{ij}$. Then $SimInv_2$ holds for Sim' and $PrevSim'$ because it holds for Sim and $PrevSim$, by the induction hypothesis.

□

The algorithm iterates the loop on lines (10–20) until Sim and $PrevSim$ define the same relation. Since, in general the data constraints Sim_{ij} , at each iteration step, might form an infinitely decreasing sequence, we use the matrix Cnt of integer counters, initially set to some input value $K > 0$ (line 4). Observe that each entry Cnt_{ij} decreases every time $Sim_{ij} \not\Leftarrow PrevSim_{ij}$ (line 19). When the counter Cnt_{ij} reaches zero, we set Sim_{ij} to false (line 17), which guarantees that $Sim_{ij} \Leftarrow PrevSim_{ij}$ always in the future. Since the number of entries in the counter matrix is finite, the algorithm is guaranteed to terminate. The following theorem summarizes the main result of this section.

Theorem 2. *Let $A = \langle \Sigma, \mathcal{D}, \mathbf{x}, Q, \mathbf{v}, F, \Delta \rangle$ be a DA. Then Algorithm 2 terminates on input A and the output is a data simulation $R \subseteq Q \times \mathcal{D}^{\mathbf{x}} \times Q$ for A .*

Proof. Let Sim^n and $PrevSim^n$ denote the matrices Sim and $PrevSim$ at the n -th iteration of the loop on lines 10–20, for $n \geq 0$. Algorithm 2 terminates whenever $Sim_{ij}^n \Leftarrow PrevSim_{ij}^n$, for all $i, j \in [1, k]$ (line 10). Suppose, by contradiction, that this never happens, thus there exist $i, j \in [1, k]$ such that $Sim_{ij}^n \not\Leftarrow PrevSim_{ij}^n$, for all $n \geq 0$. Then $Cnt_{ij}^K = 0$ (line 19) and $Sim_{ij}^{K+1} = PrevSim_{ij}^{K+2} = \perp$ (lines 17 and 20). Since $Sim_{ij}^n \rightarrow PrevSim_{ij}^n$, by Lemma 10 ($SimInv_1$), we obtain that $Sim_{ij}^{K+2} = PrevSim_{ij}^{K+2}$, contradiction.

To prove that the output of Algorithm 2 is a data simulation for A , we use Lemma 10 ($SimInv_2$) and the fact that, upon termination, we have $Sim_{ij} \Leftarrow PrevSim_{ij}$, for all $i, j \in [1, k]$. $\square$

5.1 Simulation and Subsumption

Finally, we explain how a data simulation relation computed by Algorithm 2 can be used to optimize the trace inclusion semi-algorithm. Let $\mathcal{A} = \langle A_1, \dots, A_N \rangle$ be DAN, where $A_i = \langle \mathcal{D}, \Sigma_i, \mathbf{x}_i, Q_i, \mathbf{v}_i, F_i, \Delta_i \rangle$, for all $i \in [1, N]$, and $B = \langle \mathcal{D}, \Sigma, \mathbf{x}_B, Q_B, \mathbf{v}_B, F_B, \Delta_B \rangle$ be an observer DA such that $\mathbf{x}_B \subseteq \bigcup_{i=1}^N \mathbf{x}_i$.

The main problem in using data simulation to enhance the convergence of our trace inclusion semi-algorithm is related to the fact that simulation relations are, in general, not compositional w.r.t. the interleaving semantics of the network. In other words, if we have N data simulations $R_i \subseteq Q_i \times \mathcal{D}^{\mathbf{x}_i} \times Q_i$, then their cross-product $R \subseteq Q_{\mathcal{A}} \times \mathcal{D}^{\mathbf{x}_{\mathcal{A}}} \times Q_{\mathcal{A}}$ defined as:

$$\forall q_1, r_1 \in Q_1 \dots \forall q_N, r_N \in Q_N \forall \mathbf{v} \in \mathcal{D}^{\mathbf{x}_{\mathcal{A}}}: (\langle q_1, \dots, q_N \rangle, \mathbf{v}, \langle r_1, \dots, r_N \rangle) \in R \Leftrightarrow (q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i$$

is not necessarily a simulation on the network expansion $\mathcal{A}^e$. The reason for this can be seen for $N = 2$. Let $\sigma_1, \sigma_2 \in \Sigma_{\mathcal{A}}$, such that $\sigma_1 \notin \Sigma_2$ and $\sigma_2 \notin \Sigma_1$. The execution of $\mathcal{A}^e$ on the sequence of input symbols $\sigma_1 \sigma_2$ is $(\langle q_1, q_2 \rangle, \mathbf{v}) \xrightarrow{\sigma_1} (\langle q'_1, q_2 \rangle, \mathbf{v}') \xrightarrow{\sigma_2} (\langle q'_1, q'_2 \rangle, \mathbf{v}'')$. Suppose that $(q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i, i = 1, 2$. Then there exists $r'_1 \in Q_1$ such that $(\langle r_1, r_2 \rangle, \mathbf{v}) \xrightarrow{\sigma_1} (\langle r'_1, r_2 \rangle, \mathbf{v}')$ and $(q'_1, \mathbf{v}' \downarrow_{\mathbf{x}_1}, r'_1) \in R_1$. In order to use the simulation and build the continuation $(\langle r'_1, r_2 \rangle, \mathbf{v}') \xrightarrow{\sigma_2} (\langle r'_1, r'_2 \rangle, \mathbf{v}')$, we would need that $(q_2, \mathbf{v}' \downarrow_{\mathbf{x}_2}, r_2) \in R_2$, which is not necessarily ensured by the hypothesis $(q_2, \mathbf{v} \downarrow_{\mathbf{x}_2}, r_2) \in R_2$.

We propose a partial solution to this problem, based on a restriction concerning the distribution of the network variables $\mathbf{x}_{\mathcal{A}} = \bigcup_{i=1}^N \mathbf{x}_i$ over the components $A_1, \dots, A_N$: for each $i \in [1, N]$, we have $\mathbf{x}_i = \mathbf{x}^g \cup \mathbf{x}_i^\ell$, where $\mathbf{x}^g$ is a set of *global* variables, and $\mathbf{x}_i^\ell$ are

the *local* variables of A_i . In other words, the only variables shared between more than one component are $\mathbf{x}^g$, which, moreover, are visible to all components.

Then the problem can be bypassed if none of the simulation relations $R_i \subseteq Q_i \times \mathcal{D}^{\mathbf{x}_i} \times Q_i$ may constrain the global variables:

Assumption 3. For each $i \in [1, N]$ and each $(q_i, \mathbf{v}, r_i) \in R_i$ we also have $(q_i, \mathbf{v}', r_i) \in R_i$ for each $\mathbf{v}' \in \mathcal{D}^{\mathbf{x}_i}$ such that $\mathbf{v} \downarrow_{\mathbf{x}_i^\ell} = \mathbf{v}' \downarrow_{\mathbf{x}_i^\ell}$.

Under this assumption, we use pre-computed data simulations $R_i \subseteq Q_i \times \mathcal{D}^{\mathbf{x}^g} \times Q_i$ and $R_B \subseteq Q_B \times \mathcal{D}^{\mathbf{x}^B} \times Q_B$ to generalize the basic subsumption relation between product states (defined by Lemma 4) thus speeding up the convergence of Algorithm 1.

Lemma 11. Under assumption 3, the relation defined as:

$$\begin{aligned} (\langle q_1, \dots, q_N \rangle, P, \Phi) \sqsubseteq_{sim} (\langle r_1, \dots, r_N \rangle, S, \Psi) \\ \Leftrightarrow \\ \forall i \in [1, N] \forall \mathbf{v} \in \mathcal{D}^{\mathbf{x}^A} : \mathbf{v} \models \Phi \Rightarrow \mathbf{v} \models \Psi \text{ and } \left\{ \begin{array}{l} (q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i \\ \forall p \in S \exists q \in P . (p, \mathbf{v} \downarrow_{\mathbf{x}_B}, q) \in R_B \end{array} \right. \end{aligned}$$

is a subsumption relation.

Proof. Let $s = (\langle q_1, \dots, q_N \rangle, P, \Phi)$ and $t = (\langle r_1, \dots, r_N \rangle, S, \Psi)$ be two product states, such that $s \sqsubseteq_{sim} t$. According to Definition 2, we need to prove that $\mathcal{L}_s(\mathcal{A}^e \times \bar{B}) \subseteq \mathcal{L}_t(\mathcal{A}^e \times \bar{B})$. To this end, it is sufficient to prove that for each $\mathbf{v} \in \mathcal{D}^{\mathbf{x}^A}$ such that $\mathbf{v} \models \Phi$:

1. $\mathcal{L}_{(\langle q_1, \dots, q_N \rangle, \mathbf{v})}(\mathcal{A}^e) \subseteq \mathcal{L}_{(\langle r_1, \dots, r_N \rangle, \mathbf{v})}(\mathcal{A}^e)$, and
2. for all $p \in S$ there exists $q \in P$ such that $\mathcal{L}_{(p, \mathbf{v} \downarrow_{\mathbf{x}_B})}(B) \subseteq \mathcal{L}_{(q, \mathbf{v} \downarrow_{\mathbf{x}_B})}(B)$.

Assuming that the above statements hold, we have:

$$\begin{aligned} \mathcal{L}_s(\mathcal{A}^e \times \bar{B}) &= \bigcup_{\mathbf{v} \models \Phi} \left(\mathcal{L}_{(\langle q_1, \dots, q_N \rangle, \mathbf{v})}(\mathcal{A}^e) \cap \bigcap_{q \in P} \mathcal{L}_{(q, \mathbf{v} \downarrow_{\mathbf{x}_B})}(\bar{B}) \right) \\ &\subseteq \bigcup_{\mathbf{v} \models \Phi} \left(\mathcal{L}_{(\langle r_1, \dots, r_N \rangle, \mathbf{v})}(\mathcal{A}^e) \cap \bigcap_{p \in S} \mathcal{L}_{(p, \mathbf{v} \downarrow_{\mathbf{x}_B})}(\bar{B}) \right) \\ &\subseteq \bigcup_{\mathbf{v} \models \Psi} \left(\mathcal{L}_{(\langle r_1, \dots, r_N \rangle, \mathbf{v})}(\mathcal{A}^e) \cap \bigcap_{p \in S} \mathcal{L}_{(p, \mathbf{v} \downarrow_{\mathbf{x}_B})}(\bar{B}) \right) \\ &= \mathcal{L}_t(\mathcal{A}^e \times \bar{B}) \end{aligned}$$

and we are done. Moreover, the second point above is a direct consequence of the second point of the definition of $\sqsubseteq_{sim}$ and Definition 2. We are left with proving the first point. Let $(\langle q_1, \dots, q_N \rangle, \mathbf{v}) \xrightarrow{\sigma, \Phi_i} (\langle q'_1, \dots, q'_N \rangle, \mathbf{v}')$ be a transition of $\mathcal{A}^e$ and let $(\langle r_1, \dots, r_N \rangle, \mathbf{v})$ be a configuration of $\mathcal{A}^e$ such that $(q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i$, for each $i \in [1, N]$. Let $i \in [1, N]$ be an arbitrary component, and distinguish two cases:

- if $q_i \xrightarrow{\sigma, \Phi_i} q'_i \in \Delta_i$ and $(\mathbf{v} \downarrow_{\mathbf{x}_i}, \mathbf{v}' \downarrow_{\mathbf{x}_i}) \models_{\text{Th}(\mathcal{D})} \Phi_i$, i.e. $(q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}) \xrightarrow{\sigma} (q'_i, \mathbf{v}' \downarrow_{\mathbf{x}_i})$, then, since $(q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i$ there exists $r'_i \in Q_i$ s.t. $(r_i, \mathbf{v} \downarrow_{\mathbf{x}_i}) \xrightarrow{\sigma} (r'_i, \mathbf{v}' \downarrow_{\mathbf{x}_i})$ and $(q'_i, \mathbf{v}' \downarrow_{\mathbf{x}_i}, r'_i) \in R_i$.
- otherwise, $q_i = q'_i$ and $\mathbf{v} \downarrow_{\mathbf{x}_i^\ell} = \mathbf{v}' \downarrow_{\mathbf{x}_i^\ell}$. By Assumption 3, we obtain $(q_i, \mathbf{v}' \downarrow_{\mathbf{x}_i}, r_i) \in R_i$.

By choosing $r'_i = r_i$, we obtain $(q'_i, \mathbf{v}' \downarrow_{\mathbf{x}_i}, r'_i) \in R_i$.

Hence $(q'_i, \mathbf{v}' \downarrow_{\mathbf{x}_i}, r'_i) \in R_i$, for all $i \in [1, N]$. Thus, the relation defined as:

$$(\langle q_1, \dots, q_N \rangle, \mathbf{v}, \langle r_1, \dots, r_N \rangle) \in R \Leftrightarrow \forall i \in [1, N] : (q_i, \mathbf{v} \downarrow_{\mathbf{x}_i}, r_i) \in R_i$$

is a data simulation (Definition 3), thus, by Lemma 9, we obtain that $\mathcal{L}_{(\langle q_1, \dots, q_N \rangle, \mathbf{v})}(\mathcal{A}^e) \subseteq \mathcal{L}_{(\langle r_1, \dots, r_N \rangle, \mathbf{v})}(\mathcal{A}^e)$, for all $\mathbf{v} \in \mathcal{D}^{\mathbf{x}^A}$, such that $\mathbf{v} \models_{\text{Th}(\mathcal{D})} \Phi$, and the first point above holds. $\square$

6 Experimental Results

We have implemented Algorithm 1 in a prototype tool³ using the MATHSAT SMT solver [7] for answering the satisfiability queries and computing the interpolants. The results of the experiments are given in Tables 1 and 2. The results were obtained on an Intel i7-4770 CPU @ 3.40GHz machine with 32GB RAM.

Table 1. Experiments with single-component networks.

Example	$A(Q /\Delta)$	$B(Q /\Delta)$	Vars.	Res.	Time
Arrays shift	3/3	3/4	5	ok	< 0.1s
Array rotation 1	4/5	4/5	7	ok	0.1s
Array rotation 2	8/21	6/24	11	ok	34s
Array split	20/103	6/26	14	ok	4m32s
HW counter 1	2/3	1/2	2	ok	0.2s
HW counter 2	6/12	1/2	2	ok	0.4s
Synchr. LIFO	4/34	2/15	4	ok	2.5s
ABP-error	14/20	2/6	14	cex	2s
ABP-correct	14/20	2/6	14	ok	3s

Table 1 contains experiments where the network $\mathcal{A}$ consists of a single component. We applied the tool on several verification conditions generated from imperative programs with arrays [6] (Array shift, Array rotation 1+2, Array split) available online [24]. Then, we applied it on models of hardware circuits (HW Counter 1+2, Synchronous LIFO) [26]. Finally, we checked two versions (correct and faulty) of the timed Alternating Bit Protocol [29].

Table 2 provides a list of experiments where the network $\mathcal{A}$ has $N > 1$ components. First, we have the example of Fig. 1 (Running). Next, we have several examples of real-time verification problems [27]: a controller of a railroad crossing [19] (Train) with T trains, the Fischer Mutual Exclusion protocol with deadlines Δ and Γ (Fischer), and a hardware communication circuit with K stages, composed of timed NOR gates (Stari). Third, we have modelled a Producer-Consumer example [12] with a fixed buffer size B . Fourth, we have experimented with several models of parallel programs that manipulate arrays (Array init, Array copy, Array join) with window size Δ .

For the time being, our implementation is a proof-of-concept prototype that leaves plenty of room for optimization (e.g. caching intermediate computation results) likely to improve the performance on more complicated examples. Despite that, we found the results from Tables 1 and 2 rather encouraging.

7 Conclusions

We have presented an interpolation-based abstraction refinement method for trace inclusion between a network of data automata and an observer where the variables used

³ <http://www.fit.vutbr.cz/research/groups/verifit/tools/includer/>

Table 2. Experiments with multiple-component networks (e.g., $2 \times 2/2 + 2 \times 3/3$ in column $\mathcal{A}$ means that $\mathcal{A}$ is a network with 4 components, of which 2 DA with 2 states and 2 rules, and 2 DA with 3 states and 3 rules).

Example	N	$\mathcal{A}(Q / \Delta)$	$B(Q / \Delta)$	Vars.	Res.	Time
Running	2	$2 \times 2/2$	3/4	3	ok	0.2s
Running	10	$10 \times 2/2$	11/20	3	ok	25s
Train ($T = 5$)	7	$5 \times 3/3 + 4/4 + 4/4$	2/38	1	ok	4s
Train ($T = 20$)	22	$20 \times 3/3 + 4/4 + 4/4$	2/128	1	ok	6m26s
Fischer ($\Delta = 1, \Gamma = 2$)	2	$2 \times 5/6$	1/10	4	ok	8s
Fischer ($\Delta = 1, \Gamma = 2$)	3	$3 \times 5/6$	1/15	4	ok	2m48s
Fischer ($\Delta = 2, \Gamma = 1$)	2	$2 \times 5/6$	1/10	4	cex	3s
Fischer ($\Delta = 2, \Gamma = 1$)	3	$3 \times 5/6$	1/15	4	cex	32s
Stari ($K = 1$)	5	$4/5 + 2/4 + 5/7 + 5/7 + 5/7$	3/6	3	ok	0.5s
Stari ($K = 2$)	8	$4/5 + 2/4 + 2 \times 5/7 + 2 \times 5/7 + 2 \times 5/7$	3/6	3	ok	0.5s
Prod-Cons ($B = 3$)	2	$4/4 + 4/4$	2/7	2	ok	10s
Prod-Cons ($B = 6$)	2	$4/4 + 4/4$	2/7	2	ok	2m32s
Array init ($\Delta = 2$)	5	$5 \times 2/2$	2/6	2	ok	3s
Array init ($\Delta = 2$)	15	$15 \times 2/2$	2/16	2	ok	3m15s
Array copy ($\Delta = 20$)	20	$20 \times 2/2$	2/21	3	ok	0.3s
Array copy ($\Delta = 20$)	150	$150 \times 2/2$	2/151	3	ok	43s
Array join ($\Delta = 10$)	4	$2 \times 2/2 + 2 \times 3/3$	2/3	2	ok	6s
Array join ($\Delta = 20$)	6	$3 \times 2/2 + 3 \times 3/3$	2/4	2	ok	1m9s

by the observer are a subset of those used by the network. The procedure builds on a new determinization result for DAs and combines in a novel way predicate abstraction and interpolation with antichain-based inclusion checking. The procedure has been successfully applied to several examples, including verification problems for array programs, real-time systems, and hardware designs. Future work includes an extension of the method to data tree automata and its application to logics for heaps with data. Also, we foresee an extension of the method to handle infinite traces.

Acknowledgement. This work was supported by the Czech Science Foundation project 14-11384S, the BUT FIT project FIT-S-14-2486, and the French ANR-14-CE28-0018 project Vecolib.

References

1. Abdulla, P., Chen, Y.F., Holik, L., Mayr, R., Vojnar, T.: When simulation meets antichains. In: Proc. of TACAS’10, LNCS, vol. 6015, pp. 158–174. Springer (2010)
2. Alur, R., Dill, D.L.: A theory of timed automata. Theor. Comput. Sci. 126(2), 183–235 (1994)
3. Bardin, S., Finkel, A., Leroux, J., Petrucci, L.: Fast: Fast acceleration of symbolic transition systems. In: Proc. of CAV’03. LNCS, vol. 2725. Springer (2003)
4. Beyene, T.A., Popeea, C., Rybalchenko, A.: Solving existentially quantified horn clauses. In: Proc. of CAV’13. LNCS, vol. 8044. Springer (2013)
5. Bojańczyk, M., David, C., Muscholl, A., Schwentick, T., Segoufin, L.: Two-variable logic on data words. ACM Trans. Comput. Logic 12(4), 27:1–27:26 (2011)
6. Bozga, M., Habermehl, P., Iosif, R., Konečný, F., Vojnar, T.: Automatic verification of integer array programs. In: Proc. of CAV’09. LNCS, vol. 5643, pp. 157–172 (2009)
7. Cimatti, A., Griggio, A., Schaafsma, B., Sebastiani, R.: The MathSAT5 SMT Solver. In: Proc. of TACAS. LNCS, vol. 7795 (2013)

8. Cook, B., Khlaaf, H., Piterman, N.: On automation of ctl^* verification for infinite-state systems. In: Proc. of CAV'15. LNCS, vol. 9206. Springer (2015)
9. Craig, W.: Three uses of the herbrand-gentzen theorem in relating model theory and proof theory. *J. Symb. Log.* 22(3), 269–285 (1957)
10. D'Antoni, L., Alur, R.: Symbolic visibly pushdown automata. In: Proc. of CAV'14. LNCS, vol. 8559. Springer (2014)
11. Decker, N., Habermehl, P., Leucker, M., Thoma, D.: Ordered navigation on multi-attributed data words. In: Proc. of CONCUR'14. LNCS, vol. 8704, pp. 497–511 (2014)
12. Dhar, A.: Algorithms For Model-Checking Flat Counter Systems. Ph.D. thesis, Univ. Paris 7 (2014)
13. Fribourg, L.: A closed-form evaluation for extended timed automata. Tech. rep., CNRS et Ecole Normale Supérieure de Cachan (1998)
14. Grebenshchikov, S., Lopes, N.P., Popeea, C., Rybalchenko, A.: Synthesizing software verifiers from proof rules. In: ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012. pp. 405–416 (2012)
15. Habermehl, P., Iosif, R., Vojnar, T.: A logic of singly indexed arrays. In: Proc. of LPAR'08. LNCS, vol. 5330, pp. 558–573 (2008)
16. Habermehl, P., Iosif, R., Vojnar, T.: What else is decidable about integer arrays? In: Proc. of FOSSACS'08. LNCS, vol. 4962, pp. 474–489 (2008)
17. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: Proc. of POPL'02. ACM (2002)
18. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Software Verification with Blast. In: Proc. of 10th SPIN Workshop. LNCS, vol. 2648 (2003)
19. Henzinger, T.A., Nicollin, X., Sifakis, J., Yovine, S.: Symbolic model checking for real-time systems. *Information and Computation* 111, 394–406 (1992)
20. Iosif, R., Rogalewicz, A., Vojnar, T.: Abstraction refinement for trace inclusion of data automata. CoRR abs/1410.5056 (2014), <http://arxiv.org/abs/1410.5056>
21. McMillan, K.L.: Lazy abstraction with interpolants. In: Proc. of CAV'06. LNCS, vol. 4144. Springer (2006)
22. Milner, R.: An algebraic definition of simulation between programs. In: Proc. of IJCAI'71. Morgan Kaufmann Publishers Inc. (1971)
23. Minsky, M.: *Computation: Finite and Infinite Machines*. Prentice-Hall (1967)
24. Numerical Transition Systems Repository. <http://nts.imag.fr/index.php/Flata> (2012)
25. Ouaknine, J., Worrell, J.: On the language inclusion problem for timed automata: Closing a decidability gap. In: Proc of LICS'04. IEEE Computer Society (2004)
26. Smrcka, A., Vojnar, T.: Verifying parametrised hardware designs via counter automata. In: HVC'07. pp. 51–68 (2007)
27. Tripakis, S.: The analysis of timed systems in practice. Ph.D. thesis, Univ. Joseph Fourier, Grenoble (December 1998)
28. Wulf, M.D., Doyen, L., Henzinger, T.A., Raskin, J.: Antichains: A new algorithm for checking universality of finite automata. In: Proc. of CAV'06. LNCS, vol. 4144. Springer (2006)
29. Zbrzezny, A., Polrola, A.: Sat-based reachability checking for timed automata with discrete data. *Fundamenta Informaticae* 79, 1–15 (2007)