

Algorithmic Analysis of Programs with Well Quasi-ordered Domains¹

Parosh Aziz Abdulla

Department of Computer Systems, Uppsala University, P.O. Box 325, 751 05 Uppsala, Sweden
E-mail: parosh@docs.uu.se

Kārlis Čerāns

Institute of Mathematics and Computer Science, University of Latvia, Riga, Latvia
E-mail: karlis@cclu.lv

Bengt Jonsson

Department of Computer Systems, Uppsala University, P.O. Box 325, 751 05 Uppsala, Sweden
E-mail: bengt@docs.uu.se

and

Yih-Kuen Tsay

Department of Information Management, National Taiwan University, Taipei, Taiwan
E-mail: tsay@im.ntu.edu.tw

Over the past few years increasing research effort has been directed towards the automatic verification of infinite-state systems. This paper is concerned with identifying general mathematical structures which can serve as sufficient conditions for achieving decidability. We present decidability results for a class of systems (called *well-structured systems*) which consist of a finite control part operating on an infinite data domain. The results assume that the data domain is equipped with a preorder which is a well quasi-ordering, such that the transition relation is “monotonic” (a simulation) with respect to the preorder. We show that the following properties are decidable for well-structured systems:

- *Reachability*: whether a certain set of control states is reachable. Other safety properties can be reduced to the reachability problem.

¹ Supported in part by the Swedish Board for Industrial and Technical Development (NUTEK) and by the Swedish Research Council for Engineering Sciences (TFR). The work of the second author has been partially supported by Grant 93–596 from the Latvian Council of Science. The work of the fourth author has been partially supported by the National Science Council, Taiwan (Republic of China).

- *Eventuality*: whether all executions eventually reach a given set of control states (represented as AFp in CTL).
- *Simulation*: whether there exists a simulation between a finite automaton and a well-structured system. The simulation problem will be shown to be decidable in both directions.

We also describe how these general principles subsume several decidability results from the literature about timed automata, relational automata, Petri nets, and lossy channel systems. © 2000 Academic Press

1. INTRODUCTION

Over the past few years increasing research effort has been directed towards the automatic verification of infinite-state systems. This has resulted in numerous highly nontrivial algorithms for the verification of different classes of such systems. Examples include timed automata [ACD90, AH89, Čer92a], hybrid automata [Hen95], relational automata [BBK77, Čer92b, Čer94], Petri nets [Jan90, JM95], systems with many identical processes [CG87, PP92], and lossy channel systems [AJ93, AK95]. As interest in this area increases, it will be important to extract common principles that underlie these and related results.

Our goal is to develop general mathematical structures which could serve as sufficient conditions for achieving decidability. Our objective is twofold. We aim on the one hand to give a unified explanation of existing decidability results including those mentioned above and on the other hand to provide guidelines for discovering similar decidability results for other classes of systems.

Existing work on general principles for deciding properties of infinite-state systems is fairly limited. Many existing methods are based on *finite partitioning*, where the state space of the original system is partitioned into a finite number of equivalence classes under bisimulation [ACD90, Hen95, Čer94]. Two states belonging to the same partition are equivalent in the sense that transitions from them lead to equivalent states. The requirement of having an appropriate finite partitioning of the state space is rather restrictive since it implies that the system under consideration is “essentially finite-state.”

In this paper we present substantially more general conditions for decidability of several verification problems. We work with a *preorder* on states instead of an *equivalence*. We consider systems which consist of a finite control part operating on an infinite data domain. The main requirement is that the data domain is equipped with a preorder such that the following properties (which are generalizations of those required by finite partitioning methods) hold: (i) the transition system is “monotonic” with respect to the preorder; i.e., transitions from larger states lead to larger states (this means that smaller states are *simulated* by larger states); and (ii) the preorder on the data domain is a well quasi-ordering, which means that each infinite sequence contains an element which is larger than or equivalent to an earlier element in the sequence. We call the class of systems satisfying these properties *well-structured systems*.

Our method generalizes finite partitioning in the following sense:

- We employ a preorder instead of an equivalence relation. It is clear that having an equivalence relation is a special case of having a preorder (an equivalence relation is a preorder which is symmetric).
- We work with states which are related through simulation, instead of bisimulation in the case of finite partitioning. Observe that by taking the preorder to be an equivalence, the definitions of simulation and bisimulation coincide.
- We require the preorder to be a well quasi-ordering, instead of requiring the number of equivalence classes to be finite. In case the preorder is taken to be an equivalence relation, our requirement implies that the number of equivalence classes is finite.

This means that, apart from systems whose state spaces can be finitely partitioned, e.g., timed automata [ACD90, Čer92a], various classes of hybrid automata [Hen95], and rational relational automata [Čer94], our methods can be used to analyze systems which do not allow for finite partitioning, such as Petri nets [JM95], lossy channel systems [AJ93], and integral relational automata [Čer94].

In this paper, we show that the following properties are decidable for well-structured systems:

- *Reachability*: whether a certain set of control states is reachable. Several properties can be reduced to the reachability problem, notably invariant properties and safety properties represented by the prefix-closed set of traces of a finite automaton.
- *Eventuality*: whether all executions eventually reach a given set of control states (represented as AFp in CTL).
- *Simulation*: whether there is a simulation between a finite automaton and a well-structured system. The simulation problem is shown to be decidable in both directions.

The reachability problem is solved by a backwards reachability analysis. Starting from a set I of states, the reachability of which is to be decided, we generate the set of states from which I can be reached by a sequence of at most j transitions, for successively larger j . The sets that are successively generated in this way are upwards closed with respect to the preorder and form an ascending chain (under set inclusion). Since the preorder is a well quasi-ordering, each set can be represented by a finite set of minimal states, and the chain converges after a finite number of iterations. The problem of whether a well-structured system is simulated by a finite automaton is solved using similar principles. Eventuality properties and the problem of whether a finite automaton is simulated by a well-structured system are checked by a standard tableau method. Again the tableau construction terminates by the well quasi-ordering property.

The iteration method in this paper can also be viewed as an abstract interpretation of the infinite state space. Instead of working with sets of states of the transition system, we work in an abstract domain consisting of finite sets of minimal states. One contribution is that we show, for well-structured systems, that we can

work in this abstract domain without losing precision in our analysis of reachability, and with the additional benefit that fixpoint iterations of this kind always converge.

Related Work. The idea of verifying a system by analyzing a property for an abstraction or a simpler approximation of the system has been considered by several authors [CGL92, LGS⁺95, DGG94]. These papers present conditions such that if the property is satisfied by the abstract program, then it will be satisfied by the original program. Sufficient conditions are given for an abstraction to preserve, e.g., the branching time logic CTL* or fragments thereof. However, these works are not concerned primarily with constructing decision procedures for verification.

Finkel [Fin90] shows that, for well-structured systems, it is decidable whether a system has a finite reachability tree. In this paper we use essentially a variant of his algorithm for checking eventuality properties. In addition, Finkel considers a restricted class of well-structured systems, namely those with *strict monotonicity*. This means that transitions from *strictly* larger states lead to *strictly* larger states. For this class it is shown that the coverability problem and the problem of whether the set of reachable states is finite are both decidable. The coverability problem is equivalent to the control state reachability problem and is solved in [Fin90] using a generalization of the Karp–Miller algorithm [KM69]. The Karp–Miller algorithm can be used to solve the coverability problem for, e.g., Petri nets. However, the algorithm depends on strict monotonicity, which does not hold in general for well-structured systems (e.g., for lossy channel systems), and hence the Karp–Miller algorithm cannot be applied to our class of systems.

Outline. The remainder of the paper is structured as follows. In the following section, we define infinite-state systems as systems with a finite-state control part operating on a possibly infinite domain of data values. In Section 3 we define well-structured systems. Section 4 presents the method for deciding reachability, Section 5 treats eventuality properties, and Section 6 shows how to check simulations. In Section 7 we give examples of several classes of well-structured systems. In Section 8 we give some conclusions and directions for future research.

2. INFINITE-STATE SYSTEMS

In this section we give the basic definitions for infinite-state systems. As a general model of such systems, we adopt labeled transition systems. We assume a finite set A of *labels*. Each label $\lambda \in A$ represent an observable interaction with the environment.

DEFINITION 2.1. A (labeled) transition system \mathcal{L} is a pair $\langle S, \delta \rangle$, where

- S is a set of *states*, formed as the cartesian product $Q \times D$ of a finite set Q of *control states* and a possibly infinite set D of *data values*, and
- $\delta \subseteq S \times A \times S$ is a set of *transitions*.

We use $\langle q, d \rangle$ to denote the state whose control state is q and whose data value is d and $s \xrightarrow{\lambda} s'$ to denote that $\langle s, \lambda, s' \rangle \in \delta$. Intuitively, $s \xrightarrow{\lambda} s'$ means that the system can move from state s to state s' while performing the observable action λ .

We let $s \rightarrow s'$ denote that there is a λ such that $s \xrightarrow{\lambda} s'$ and let $\xrightarrow{*}$ denote the reflexive transitive closure of \rightarrow .

For $q \in Q$ and $A \subseteq D$, we use $\langle q, A \rangle$ to denote the set $\{\langle q, d \rangle \mid d \in A\}$. For $s \in S$ and $T \subseteq S$ we say that T is *reachable from* s (written $s \xrightarrow{*} T$) if there exists a state $s' \in T$ such that $s \xrightarrow{*} s'$.

For $T \subseteq S$ and $\lambda \in A$, we define $pre_\lambda(T)$ to be the set $\{s' \mid \exists s \in T. s' \xrightarrow{\lambda} s\}$. Analogously, we define $post_\lambda(T)$ as $\{s' \mid \exists s \in T. s \xrightarrow{\lambda} s'\}$. By $pre(T)$ ($post(T)$) we mean $\bigcup_{\lambda \in A} pre_\lambda(T)$ ($\bigcup_{\lambda \in A} post_\lambda(T)$). Sometimes we write $pre(s)$ ($post(s)$) instead of $pre(\{s\})$ ($post(\{s\})$).

A *computation from a state* s is a sequence of the form $s_0 s_1 \cdots s_n$, where $s_0 = s$, $s_i \rightarrow s_{i+1}$, and either $n = \infty$ (i.e., the sequence is infinite) or there is no state s' such that $s_n \rightarrow s'$.

3. WELL-STRUCTURED SYSTEMS

In this section, we define a class of transition systems which we call *well-structured systems*, for which we will present our decidability results. First, we recall the notion of *preorders*.

3.1. Preliminaries

A preorder \preceq is a reflexive and transitive (binary) relation on a set D . We say that \preceq is *decidable* if there is a procedure which, given $a, b \in D$, decides whether $a \preceq b$. The relation \preceq is a *well quasi-ordering* if there is no infinite sequence a_0, a_1, a_2, \dots , such that $a_i \not\preceq a_j$ for all $i < j$. A set M is said to be *canonical* if $a, b \in M$ implies $a \preceq b$. We say that $M \subseteq A$ is a *minor set* of A , if (i) for all $a \in A$ there exists $b \in M$ such that $b \preceq a$ and (ii) M is canonical.

A set $I \subseteq D$ is an *ideal* (in D) if $a \in I, b \in D$, and $a \preceq b$ imply $b \in I$; i.e., the set I is upward-closed with respect to the relation \preceq . We define the (*upward*) *closure* of a set $A \subseteq D$, denoted $\mathcal{C}(A)$, as the ideal $\{b \in D \mid \exists a \in A. a \preceq b\}$, which is generated by A .

For sets A and B , we say that $A \equiv B$ if $\mathcal{C}(A) = \mathcal{C}(B)$. Observe that $A \equiv B$ if and only if for all $a \in A$ there is a $b \in B$ such that $b \preceq a$, and vice versa.

LEMMA 3.1. *If a preorder \preceq is a well quasi-ordering, then for each set A there exists at least one finite minor set of A .*

Proof. Suppose that no finite minor set of A exists. We show that \preceq is not a well quasi-ordering. We construct an infinite sequence a_0, a_1, a_2, \dots of elements in A as follows. Let a_0 be any arbitrary element in A . We choose a_{i+1} such that $a_j \not\preceq a_{i+1}$ for each $j: 0 \leq j \leq i$. The element a_{i+1} exists, since otherwise we could easily construct a minor set of the finite set $\{a_0, a_1, \dots, a_i\}$ which would also be a minor set of A , contradicting the assumption that no such sets exist. It is clear that the sequence a_0, a_1, a_2, \dots violates the well quasi-orderedness property. ■

Notice that although Lemma 3.1 implies that each minor set is finite, there may still be infinitely many such minor sets. Also, we observe that if \preceq is a partial

order, then there exists a unique minor set of A . We use min to denote a function which, given a set A , returns a minor set of A .

From Lemma 3.1, and the fact that $\mathcal{C}(\text{min}(I)) = I$ for each ideal I , it follows that we can use $\text{min}(I)$ as a finite representation of I .

LEMMA 3.2. *For a preorder \preceq on a set A , \preceq is a well quasi-ordering iff for each infinite sequence $I_0 \subseteq I_1 \subseteq I_2 \subseteq \dots$ of ideals in A there is a k such that $I_k = I_{k+1}$.*

Proof. (only if) Suppose that we have an infinite sequence $I_0 \subset I_1 \subset I_2 \subset \dots$. It follows that there is a sequence a_0, a_1, a_2, \dots of elements in A such that for all $k \geq 0$ we have $a_k \in I_k$ and $a_k \notin I_j$ for each $j < k$. This means $a_j \not\preceq a_k$ for $j < k$, otherwise $a_k \in I_j$, since I_j is an ideal. This is a contradiction since the sequence a_0, a_1, a_2, \dots will then violate the well quasi-ordering assumption.

(if) Suppose that we have an infinite sequence of elements a_0, a_1, a_2, \dots in A , where $a_j \preceq a_k$ if $j < k$. We define an infinite sequence I_0, I_1, I_2, \dots of ideals, where $I_k = \mathcal{C}(\{a_0, a_1, \dots, a_k\})$. It is clear that $I_0 \subset I_1 \subset I_2 \subset \dots$. ■

In fact, for any sequence I_0, I_1, I_2, \dots we can show that there are j and k such that $j < k$ and $I_j = I_k$. We use the *only if*-direction of Lemma 3.2 to prove termination of some of our verification algorithms.

3.2. Well-Structured Systems

In our framework we require that the set D of data values be equipped with a decidable preorder \preceq and assume that we are given a minor set of D which we henceforth call D_{min} . We extend the preorder \preceq on D to a decidable preorder \preceq on the set S of states defined by $\langle q, d \rangle \preceq \langle q', d' \rangle$ if and only if $q = q'$ and $d \preceq d'$.

A transition system $\langle S, \delta \rangle$ is *monotonic* (with respect to \preceq) if for each $s_1, s_2, s_3 \in S$ and $\lambda \in A$, if $s_1 \preceq s_2$ and $s_1 \xrightarrow{\lambda} s_3$, then there exists s_4 such that $s_3 \preceq s_4$ and $s_2 \xrightarrow{\lambda} s_4$.

LEMMA 3.3. *A transition system $\langle S, \delta \rangle$ is monotonic iff the set of ideals in S is closed under the applications of both pre_λ and pre .*

Proof. We show the claim for pre_λ . The claim for pre follows from the fact that $\text{pre} = \bigcup_{\lambda \in A} \text{pre}_\lambda$.

(only if) Suppose that $\langle S, \delta \rangle$ is monotonic. Take any ideal I in S . Suppose that $s_1 \in \text{pre}_\lambda(I)$ and $s_1 \preceq s_2$. We show that $s_2 \in \text{pre}_\lambda(I)$. We know that there is $s_3 \in I$ such that $s_1 \xrightarrow{\lambda} s_3$. By monotonicity it follows that there is s_4 such that $s_3 \preceq s_4$ and $s_2 \xrightarrow{\lambda} s_4$. Since I is an ideal, we have $s_4 \in I$, and hence $s_2 \in \text{pre}_\lambda(I)$.

(if) Suppose that $\langle S, \delta \rangle$ is not monotonic. It follows that there are states s_1, s_2 , and s_3 , and $\lambda \in A$ such that $s_1 \preceq s_2$, $s_1 \xrightarrow{\lambda} s_3$, but there is no s_4 where $s_3 \preceq s_4$ and $s_2 \xrightarrow{\lambda} s_4$. Define the ideal $I = \mathcal{C}(\{s_3\})$. It is clear that $s_1 \in \text{pre}_\lambda(I)$ but $s_2 \notin \text{pre}_\lambda(I)$. This means that $\text{pre}_\lambda(I)$ is not an ideal. ■

DEFINITION 3.4. A transition system $\mathcal{L} = \langle S, \delta \rangle$, assuming a decidable preorder \preceq on the set D of data values, is said to be *well-structured* if

1. it is monotonic;
2. \preceq is a well quasi-ordering; and
3. for each state $s \in S$ and $\lambda \in A$, the set $\min(\text{pre}_\lambda(\mathcal{C}(\{s\})))$ is computable.

Note that $\min(\text{pre}_\lambda(\mathcal{C}(\{s\})))$ is finite if \preceq is a well quasi-ordering. We define $\text{minpre}_\lambda(s)$ as notation for $\min(\text{pre}_\lambda(\mathcal{C}(\{s\})))$. For a set T of states we use $\text{minpre}_\lambda(T)$ to denote $\bigcup_{s \in T} \text{minpre}_\lambda(s)$. On the concrete models where we shall apply our theory (Section 7) the computability of $\text{minpre}_\lambda(s)$ will be rather obvious given the explicit syntactic representations of the transition relations.

Comment on the Representation of Ideals. In this paper, we will represent ideals by minor sets. An alternative (and more general) representation is in terms of *constraints*. We then assume a set Φ of constraints over the domain D of data values. Each constraint ϕ denotes a subset $\llbracket \phi \rrbracket$ of D . Given a set Φ of constraints, we can define a preorder \preceq_Φ on D by $d \preceq_\Phi d'$ iff for all constraints ϕ in Φ we have that $d \in \llbracket \phi \rrbracket$ implies $d' \in \llbracket \phi \rrbracket$. We observe that the constraint representation is at least as general as the approach we use here, where we start by a preorder: an arbitrary preorder \preceq can be obtained as the preorder \preceq_Φ by letting Φ be the set which for each $d_0 \in D$ contains a constraint that denotes the set $\{d \in D \mid d_0 \preceq d\}$.

Instead of using finite minor sets to represent ideals, we can use finite sets of constraints. A set of constraints denotes the union of the denotations of its elements (recall that each constraint denotes a set). In some cases (e.g., for real-time automata) such a representation is more convenient, since a constraint sometime represents a large minor set.

3.3. Composition

For labeled transition systems $\mathcal{L}_1 = \langle S_1, \delta_1 \rangle$ and $\mathcal{L}_2 = \langle S_2, \delta_2 \rangle$, we define the *composition* $\mathcal{L}_1 \parallel \mathcal{L}_2$ of \mathcal{L}_1 and \mathcal{L}_2 to be the transition system $\langle S, \delta \rangle$, where

- $S = S_1 \times S_2$.
- $\langle s_1, s_2 \rangle \xrightarrow{\lambda} \langle s'_1, s'_2 \rangle$ iff $s_1 \xrightarrow{\lambda} s'_1$ and $s_2 \xrightarrow{\lambda} s'_2$.

Our definition of composition is the standard one taken from process algebras such as CSP or LOTOS.

THEOREM 3.5. *For well structured systems \mathcal{L}_1 and \mathcal{L}_2 , the composition $\mathcal{L}_1 \parallel \mathcal{L}_2$ is well structured.*

Proof. Let $\mathcal{L}_1 = \langle S_1, \delta_1 \rangle$ and $\mathcal{L}_2 = \langle S_2, \delta_2 \rangle$. Let \preceq_1 and \preceq_2 be the respective preorders defined on S_1 and S_2 . Define the preorder \preceq for L by $\langle s_1, s_2 \rangle \preceq \langle s'_1, s'_2 \rangle$ whenever both $s_1 \preceq_1 s'_1$ and $s_2 \preceq_2 s'_2$. Monotonicity and computability of minpre follow directly from their definitions. To prove the well quasi-ordering property of \preceq , consider an infinite sequence of state pairs $\langle s_0, s'_0 \rangle, \langle s_1, s'_1 \rangle, \langle s_2, s'_2 \rangle, \dots$, where $s_i \in S_1$ and $s'_i \in S_2$. It follows from the well quasi-ordering of \preceq_1 that there is an infinite increasing sequence i_0, i_1, i_2, \dots , such that $s_{i_j} \preceq_1 s_{i_k}$ whenever $i_j \leq i_k$. Since \preceq_2 is also a well quasi-ordering we conclude that there are j and k , where $j < k$ and $s'_{i_j} \preceq_2 s'_{i_k}$. This means that $\langle s_{i_j}, s'_{i_j} \rangle \preceq \langle s_{i_k}, s'_{i_k} \rangle$. ■

4. CONTROL STATE REACHABILITY

In this section we describe an algorithm to solve the control state reachability problem for well structured transition systems. More precisely, given a state s and a control state q , we want to check whether $\langle q, D \rangle$ is reachable from s . Our algorithm actually solves the more general problem of deciding whether an ideal I is reachable from a given state s . Since $\langle q, D \rangle$ is an ideal, the control state reachability problem is a special case of the reachability problem for ideals.

To check the reachability of an ideal I , we perform a reachability analysis backwards. Starting from I we define the sequence I_0, I_1, I_2, \dots of sets by $I_0 = I$ and $I_{j+1} = I \cup \text{pre}(I_j)$. Intuitively, I_j denotes the set of states from which I is reachable in at most j steps. Thus, if we define $\text{pre}^*(I)$ to be $\bigcup_{j \geq 0} I_j$, then I is reachable from s if and only if $s \in \text{pre}^*(I)$. Notice that $\text{pre}^*(I)$ is the least fixpoint $\mu X. I \cup \text{pre}(X)$. By Lemma 3.3 each I_j is an ideal in S . We know that $I_0 \subseteq I_1 \subseteq I_2 \subseteq \dots$, and hence from Lemma 3.2 it follows that there is a k such that $I_k = I_{k+1}$. It can easily be seen that $I_\ell = I_k$ for all $\ell \geq k$, implying that $\text{pre}^*(I) = I_k$.

Our method for deciding whether I is reachable is based on generating the above sequence I_0, I_1, I_2, \dots of ideals and checking for convergence. This cannot be carried out directly since I_j is an infinite set. Instead, we represent each I_j by a canonical set $M_j = \text{min}(I_j)$. By Lemma 3.1 each minor set M_j is finite. It is straightforward to show that $M_{j+1} \equiv \text{min}(\text{min}(I) \cup \text{minpre}(M_j))$, which is computable as

$$M_{j+1} = \text{min} \left(\text{min}(I) \cup \bigcup_{s \in M_j} \text{min}(\text{pre}(\mathcal{C}(\{s\}))) \right)$$

since, by the definition of well-structured transition systems, each set $\text{min}(\text{pre}(\mathcal{C}(\{s\})))$ is computable, and the union is taken over a finite set of sets.

From the above discussion we conclude that if we define $\text{minpre}^*(M_0)$ to be $\bigcup_{j \geq 0} M_j$, then there is a k such that $M_{k+1} \equiv M_k$, and $\text{minpre}^*(M_0) \equiv M_k$. This implies that $\text{minpre}^*(M)$ is computable for any minor set M of I and in fact $\mathcal{C}(\text{minpre}^*(M)) = \text{pre}^*(I)$.

THEOREM 4.1. *The control state reachability problem is decidable for well-structured systems.*

Proof. Given a state s and a control state q we compute $\text{minpre}^*(\langle q, D_{\text{min}} \rangle)$. We then check whether there is an $s' \in \text{minpre}^*(\langle q, D_{\text{min}} \rangle)$ such that $s' \preceq s$. ■

Abstract Interpretation. The above analysis algorithm can also be phrased in terms of abstract interpretation [CC77, JN94]. We intend to compute the fixpoint $\mu X. I \cup \text{pre}(X)$ for a set $I \subseteq S$ by iteration. Instead of computing this fixpoint in the lattice $\langle 2^S, \subseteq \rangle$ of sets of states, we move to the abstract lattice $\langle \mathcal{M}, \sqsubseteq \rangle$, where \mathcal{M} is the set of canonical subsets of S , and where $M \sqsubseteq M'$ if $\mathcal{C}(M) \subseteq \mathcal{C}(M')$. The correspondence between the concrete lattice $\langle 2^S, \subseteq \rangle$ and the abstract lattice $\langle \mathcal{M}, \sqsubseteq \rangle$ is expressed by a pair $\langle \alpha, \gamma \rangle$ of functions as follows.

- $\alpha: 2^S \mapsto \mathcal{M}$, defined by $\alpha(T) = \text{min}(T)$, maps each set of states in the concrete lattice to its abstract representation.

- $\gamma: \mathcal{M} \mapsto 2^S$, defined by $\gamma(M) = \mathcal{C}(M)$, recovers the concrete meaning of an element in the abstract lattice.

The pair $\langle \alpha, \gamma \rangle$ forms a *Galois insertion*² of $\langle \mathcal{M}, \sqsubseteq \rangle$ into $\langle 2^S, \subseteq \rangle$.

Our algorithm for deciding reachability can be seen as computing the fixpoint $\mu X. \min(I) \sqcup \min \text{pre}(X)$ in the lattice $\langle \mathcal{M}, \sqsubseteq \rangle$, where $M_1 \sqcup M_2 = \min(M_1 \cup M_2)$. The monotonicity of the transition relation ensures that this computation corresponds exactly to the computation $\mu X. I \cup \text{pre}(X)$ in $\langle 2^S, \subseteq \rangle$ if I is an ideal in S . Exactness follows from the identity

$$\text{pre}(\gamma(M)) = \gamma(\min \text{pre}(M))$$

for all $M \in \mathcal{M}$ and ensures that if the fixpoint computation converges to M_k , then $\gamma(M_k)$ is the least fixpoint of $\mu X. I \cup \text{pre}(X)$ in $\langle 2^S, \subseteq \rangle$. Finally, well quasi-orderedness of \preceq implies that all ascending chains in $\langle \mathcal{M}, \sqsubseteq \rangle$ are finite, thus guaranteeing convergence of any least fixpoint computation.

5. EVENTUALITY PROPERTIES

In this section we describe an algorithm for deciding whether each computation starting from an initial state eventually reaches a certain control state satisfying a predicate p over control states. In CTL, these properties are of the form $\text{AF}p$. We present an algorithm for the dual property $\text{EG}p$ from which an algorithm for $\text{AF}p$ can easily be derived using the correspondence $\text{AF}p \equiv \neg \text{EG} \neg p$. The property $\text{EG}p$ is true in a state s_0 iff there is a computation from s_0 in which all states have a control part that satisfies p . Our algorithm will actually solve the more general problem of whether s_0 satisfies a property of the form EGI for an ideal I . We write this property as $s_0 \models \text{EGI}$.

The algorithm essentially builds a tree of reachable states, starting from the initial state and successively exploring the successors of each state in the tree. We must then consider the possibility that $\text{post}(s)$ is infinite for some states s (i.e., the transition relation is not finite branching). To overcome this difficulty, we say that a transition system is *essentially finite branching* if for each state s we can effectively compute a finite subset of $\text{post}(s)$, denoted $\text{maxpost}(s)$, such that for each state $s' \in \text{post}(s)$ there is a state $s'' \in \text{maxpost}(s)$ with $s' \preceq s''$. If $\text{post}(s)$ is finite, then $\text{maxpost}(s)$ can be taken as $\text{post}(s)$. In the cases where $\text{post}(s)$ is infinite (as can be the case, e.g., for real-time automata), the subset $\text{maxpost}(s)$ can fully represent the set $\text{post}(s)$ for the purposes of this algorithm.

In the algorithm, we build a tree labeled by properties of the form $s \models \text{EGI}$. The root node is labeled by $s_0 \models \text{EGI}$. A node labeled by $s \models \text{EGI}$ is a leaf if either

² In the abstract interpretation literature, a *Galois insertion* is defined as follows. Let $\langle \text{Concr}, \sqsubseteq_{\text{Concr}} \rangle$ be an ordered *concrete* domain, and let $\langle \text{Abs}, \sqsubseteq_{\text{Abs}} \rangle$ be an ordered *abstract* domain. Consider mappings $\alpha: \text{Concr} \rightarrow \text{Abs}$ and $\gamma: \text{Abs} \rightarrow \text{Concr}$. We say that the pair $\langle \alpha, \gamma \rangle$ forms a *Galois insertion* if (i) α and γ are monotonic, (ii) $\forall a \in \text{Abs}: a = \alpha(\gamma(a))$, and (iii) $\forall c \in \text{Concr}: c \sqsubseteq_{\text{Concr}} \gamma(\alpha(c))$.

1. $s \notin I$, in which case, the node is considered *unsuccessful*, or
2. the node has an ancestor labeled $s' \models \text{EGI}$ for some s' with $s' \preceq s$, in which case, the node is considered *successful*, or
3. $s \in I$ and $\text{post}(s)$ is empty, in which case, the node is considered *successful*.

From a non-leaf node labeled $s \models \text{EGI}$ we create a child labeled $s' \models \text{EGI}$ for each state $s' \in \text{maxpost}(s)$. The algorithm answers “yes” if a successful node is encountered; otherwise it answers “no.”

The correctness of the algorithm follows from the fact that when a successful node is encountered according to criterion 2, we can, by monotonicity, construct an infinite path where all states are in I by continuing from the ancestor node. Completeness follows by the observation that the possibly unexplored successors of a state (i.e., those in $\text{maxpost}(s)$ but not in $\text{post}(s)$ for some s) can be satisfactorily represented by “larger” states (with respect to \preceq) in $\text{maxpost}(s)$. The construction of the tree terminates by König’s lemma, since the tree is finite branching and all branches are finite (this follows from well quasi-orderedness). We have thus proved the following theorem:

THEOREM 5.1. *The eventuality problem for control states is decidable for well-structured and essentially finite branching systems.*

In [Fin90] an algorithm is presented to check whether the reachability tree of a well-structured system is finite. The algorithm can be seen as a variant of our algorithm to check eventuality properties as follows. We take I to be the set S of all states. A node labeled by $s \models \text{EGS}$ is a leaf if either

- the node has an ancestor labeled $s' \models \text{EGS}$ for some s' with $s' \preceq s$. In this case, the node is considered *successful*, or
- $\text{post}(s)$ is empty. In this case, the node is considered *unsuccessful*.

The reachability tree is finite iff no successful nodes are encountered.

6. SIMULATIONS BETWEEN INFINITE SYSTEMS AND FINITE SYSTEMS

In this section we consider the problem of whether a well-structured system is simulated by a finite transition system. A transition system is said to be *finite* if it has a finite set of states. In our algorithms we assume that a finite transition system is described by finite sets representing states and transitions.

DEFINITION 6.1. Given two transition systems $\mathcal{L}_1 = \langle S_1, \delta_1 \rangle$ and $\mathcal{L}_2 = \langle S_2, \delta_2 \rangle$, we say that a relation $\mathcal{R} \subseteq S_1 \times S_2$ is a *simulation* (of \mathcal{L}_1 by \mathcal{L}_2) if for each $\langle s_1, s_2 \rangle \in \mathcal{R}$, $s'_1 \in S_1$, and $\lambda \in A$, if $s_1 \xrightarrow{\lambda} s'_1$, then there exists $s'_2 \in S_2$ such that $s_2 \xrightarrow{\lambda} s'_2$ and $\langle s'_1, s'_2 \rangle \in \mathcal{R}$.

Simulating an Infinite System by a Finite System. For $s_1 \in S_1$ and $s_2 \in S_2$, we say that s_1 is *simulated* by s_2 , denoted $s_1 \sqsubseteq s_2$, if there is a simulation \mathcal{R} of \mathcal{L}_1 by \mathcal{L}_2 such that $\langle s_1, s_2 \rangle \in \mathcal{R}$.

A transition system is said to be *intersection effective* if $\min(\mathcal{C}(s_1) \cap \mathcal{C}(s_2))$ is computable for any states s_1 and s_2 .

THEOREM 6.2. *For a state s in an intersection effective well-structured transition system and a state q in a finite transition system, it is decidable whether $s \sqsubseteq q$.*

Proof. The idea is to calculate the set of pairs $\langle s, q \rangle$ of states such that $s \not\sqsubseteq q$. We observe that for each q , the set $\{s \mid s \not\sqsubseteq q\}$ is an ideal. This allows us to compute the set by a fixpoint iteration analogous to that used for the reachability problem. For each state q of the finite transition system, we define a sequence $I_0^q, I_1^q, I_2^q, \dots$, where $I_0^q = \emptyset$, and $s \in I_{j+1}^q$ if and only if either

- $s \in I_j^q$ or
- there are λ and s' such that $s \xrightarrow{\lambda} s'$ and for all q' if $q \xrightarrow{\lambda} q'$, then $s' \in I_j^{q'}$.

Intuitively, I_j^q denotes the set of states (in the infinite transition system), which s can simulate at most $j-1$ steps. It is clear that I_j^q is an ideal and that $I_0^q \subseteq I_1^q \subseteq I_2^q \subseteq \dots$. By Lemma 3.2 it follows that there is a k such that $I_{k+1}^q = I_k^q$ for all q , and $s \not\sqsubseteq q$ iff $s \in I_k^q$.

We represent I_j^q by the canonical set $M_j^q = \min(I_j^q)$, where $M_0^q = \emptyset$, and

$$M_{j+1}^q = \bigcup_{\lambda} \text{minpre}_{\lambda} \left(\bigcap_{q' \in \text{post}_{\lambda}(q)} M_j^{q'} \right).$$

Note that M_{j+1}^q can be computed from M_j^q for intersection effective well-structured transition systems. We iterate until we reach a k such that $M_{k+1}^q \equiv M_k^q$. To decide whether $s \sqsubseteq q$ we check if $\exists s' \preceq s$ such that $s' \in M_k^q$. ■

Weak Simulation. The result of Theorem 6.2 can be generalized to the case of weak simulation as follows. We assume that the set of labels is extended by the silent event τ . Let $\langle S, \delta \rangle$ be a transition system. For $s_1, s_2 \in S$ and $\lambda \neq \tau$, we let $s_1 \xrightarrow{\lambda} s_2$ denote that s_2 is reachable from s_1 through a finite number of τ -transitions, followed by a λ -transition, followed by a finite number of τ -transitions. For $T \subseteq S$ and $\lambda \neq \tau$, we define $\overline{\text{pre}}_{\lambda}(T)$ to be the set $\{s' \mid \exists s \in T. s' \xrightarrow{\lambda} s\}$. Analogously, we define $\overline{\text{post}}_{\lambda}(T)$ as $\{s' \mid \exists s \in T. s \xrightarrow{\lambda} s'\}$. We let $\overline{\text{minpre}}_{\lambda}(T)$ denote $\min(\overline{\text{pre}}_{\lambda}(\mathcal{C}(T)))$. From the discussion in Section 4, we conclude that $\overline{\text{minpre}}_{\lambda}(\{s\})$ is computable for each $s \in S$.

The definition of simulation can be generalized to *weak simulation* by replacing the relation \longrightarrow by \Longrightarrow .

THEOREM 6.3. *For a state s in an intersection effective well-structured transition system and a state q in a finite transition system, it is decidable whether s is weakly simulated by q .*

Proof. We modify the algorithm in the proof of Theorem 6.2 and define

$$M_{j+1}^q = \bigcup_{\lambda} \overline{\text{minpre}}_{\lambda} \left(\bigcup_{q' \in \overline{\text{post}}_{\lambda}(q)} M_j^{q'} \right). \quad \blacksquare$$

Simulating a Finite System by an Infinite System. We consider the problem of whether a finite transition system is simulated by a well-structured system. We present an algorithm which assumes that the well-structured system is essentially

finite branching. For a state s and label λ , define the set $\text{maxpost}_\lambda(s)$ analogously to the definition of $\text{maxpost}(s)$. To determine whether $q_0 \sqsubseteq s_0$, we construct an and-or tree as follows. The root is labeled by $q_0 \sqsubseteq s_0$. A node labeled by $q \sqsubseteq s$ is an and-node. Such a node is a leaf if either

- $\text{post}(q)$ is empty, or
- it has an ancestor labeled $q \sqsubseteq s'$ for some s' with $s \preceq s'$.

Every non-leaf and-node, with a label $q \sqsubseteq s$, has a descendant labeled $q' \ll_\lambda s$, for each $q' \in \text{post}_\lambda(q)$. A node labeled by $q' \ll_\lambda s$ is an or-node. Such a node is a leaf if $\text{post}_\lambda(s)$ is empty. Every non-leaf or-node, with a label $q' \ll_\lambda s$, has a descendant labeled $q' \sqsubseteq s'$, for each $s' \in \text{maxpost}_\lambda(s)$.

By arguments similar to those used in the eventuality algorithm, the and-or tree is always finite, and $q_0 \sqsubseteq s_0$ if and only if the root in the and-or tree generated from $q_0 \sqsubseteq s_0$ evaluates to *true* (where or-node are leaves that evaluate to *false*, and-nodes are leaves that evaluate to *true*). We have thus proved the following theorem:

THEOREM 6.4. *The problem of whether $q_0 \sqsubseteq s_0$, for a state q_0 of a finite-state system and a state s_0 of a well-structured and essentially finite branching system, is decidable.*

7. EXAMPLE MODELS

In this section we give four examples of computation models, namely, lossy channel systems, vector addition systems with states (a model equivalent to Petri nets), relational automata, and timed automata. For each model, we describe how it can be viewed as an infinite-state transition system and show that it is equipped with a preorder such that it satisfies the conditions in Section 3 for being well structured. The set S of states of each model is formed by the Cartesian product of a finite set Q of control states and an infinite set D of *data values*. The set δ of transitions is derived from a finite set Cmd of *commands*, where a command corresponds to an atomic event involving a change of state while performing an observable interaction (represented by an element in the set A of labels) with the environment.

7.1. Lossy Channel Systems

Lossy channel systems (LCSs) [AJ93] are systems of finite-state processes that communicate messages from a finite alphabet M over a finite set C of unreliable, unbounded FIFO channels. The channels are unreliable in the sense that they may lose messages at any time. The set Q of control states of an LCS is typically the Cartesian product of the control states of the finite-state processes in the system. LCSs have been used to model and verify data transfer protocols (e.g., sliding window protocols) that are designed to tolerate message losses in channels [BZ83, Boc78].

A command in Cmd is a quadruple of the form $\langle q, op, \lambda, q' \rangle$, where $q, q' \in Q$, $\lambda \in A$, and op is an operation of one of the forms

- $c!m$: sends $m \in M$ to $c \in C$ (appends m to the end of c).
- $c?m$: receives $m \in M$ from $c \in C$ (removes c from the head of c). This operation may be performed only if m is in the head of c .
- $skip$: is the empty operation (which does not change the channel contents).

The data domain D of an LCS is the set of mappings $C \mapsto M^*$. For strings $x, y \in M^*$, we use $x \bullet y$ to denote the concatenation of x and y . For $d \in D$, we use $d[c := x]$ to denote d' , where $d'(c) = x$, and $d'(c') = d(c')$, for $c' \neq c$.

The preorder \leq on D (which is also a partial order) is given by $d_1 \leq d_2$ iff for each channel c , the string $d_1(c)$ is a (not necessarily contiguous) substring of $d_2(c)$. Now we show that LCSs are well-structured systems.

Let d be any data value. A command of form $\langle q, c!m, \lambda, q' \rangle$ gives rise to the set of transitions of form $\langle q, d \rangle \xrightarrow{\lambda} \langle q', d[c := d(c) \bullet m] \rangle$. A command of form $\langle q, c?m, \lambda, q' \rangle$ gives rise to the set of transitions of form $\langle q, d[c := m \bullet d(c)] \rangle \xrightarrow{\lambda} \langle q', d \rangle$. A command of form $\langle q, skip, \lambda, q' \rangle$ gives rise to the set of transitions of form $\langle q, d \rangle \xrightarrow{\lambda} \langle q', d \rangle$. Also, if $\langle q, d \rangle \xrightarrow{\lambda} \langle q', d' \rangle$ is a (non-lossy) transition obtained from a command in one of the ways just described, then a *lossy transition* of the form $\langle q, d_1 \rangle \xrightarrow{\lambda} \langle q', d'_1 \rangle$ may be performed for any $d_1 \supseteq d$ and any $d'_1 \leq d'$. Intuitively, the messages in d_1 but not in d are unobservably lost immediately before the non-lossy transition and those in d' but not in d'_1 are lost after the non-lossy transition.

- *Monotonicity*: follows from the existence of lossy transitions, implying that a state may lose any number of messages transforming to a smaller state and then may perform all the transitions of the smaller state.

- *Well quasi-orderedness*: holds by a beautiful result which states that the substring relation among strings over a finite alphabet is a well quasi-ordering [Hig52].

- *Computability of \minpre_λ* : the set $\minpre_\lambda(\langle q, d \rangle)$ is defined to be $\min(T)$, where T is the smallest set containing the following states:

- $\langle q', d' \rangle$, if $\langle q', c!m, \lambda, q \rangle \in Cmd$ and $d = d'[c := d'(c) \bullet m]$.
- $\langle q', d' \rangle$, if $\langle q', c!m, \lambda, q \rangle \in Cmd$, $d = d'$ and the last element of $d(c)$ is not equal to m .
- $\langle q', d' \rangle$, if $\langle q', c!m, \lambda, q \rangle \in Cmd$, $d = d'$ and $d(c)$ is equal to the empty string.
- $\langle q', d' \rangle$, if $\langle q', c?m, \lambda, q \rangle \in Cmd$ and $d' = d[c := m \bullet d(c)]$.
- $\langle q', d' \rangle$, if $\langle q', skip, \lambda, q \rangle \in Cmd$ and $d = d'$.

Intersection effectiveness and essentially finite branching are obvious.

The decidability of control state reachability and eventuality properties is shown in [AJ93], while the decidability of simulation with finite transition systems in both directions is shown in [AK95].

7.2. Vector Addition Systems with States (Petri Nets)

A *vector addition system with states* (VASS) models a finite-state machine operating on a finite number of variables which range over the natural numbers.

In a VASS, the data domain D is the set of k -dimensional vectors $N = \langle n_1, n_2, \dots, n_k \rangle$, where each n_i is a natural number. We use $N[i]$ to denote n_i . We apply addition, subtraction, and relational operators pointwise on k -dimensional vectors. We let $\bar{0}$ denote the vector which is constantly 0. The preorder (in fact partial order) \leq on D is defined as $N_1 \leq N_2$ iff $N_1[i] \leq N_2[i]$, for each $i: 1 \leq i \leq k$.

A command in Cmd is of the form $\langle q, N_1, N_2, \lambda, q' \rangle$, where $q, q' \in Q$, $\lambda \in A$, and N_1, N_2 are k -dimensional vectors. The set δ then contains all transitions of form $\langle q, N_3 \rangle \xrightarrow{\lambda} \langle q', N_3 - N_1 + N_2 \rangle$ such that $\langle q, N_1, N_2, \lambda, q' \rangle$ is a command and $\bar{0} \leq N_3 - N_1$.

For a VASS, monotonicity, intersection effectiveness, and essentially finite branching are obvious. Well quasi-orderedness is a special case of that for lossy channel systems. The set $minpre_\lambda(\langle q, d \rangle)$ is defined to be $min(T)$, where T is the set of elements of the form $\langle q', N' \rangle$ such that $\langle q', N_1, N_2, \lambda, q \rangle \in Cmd$ and $N'[i] = \max(0, N[i] - N_2[i] + N_1[i])$ for each $1 \leq i \leq k$.

Control state reachability (often called coverability in the Petri net literature) and eventuality properties for VASS can also be decided by the Karp–Miller algorithm [Km69]. The control state reachability algorithm we present in this paper performs backwards reachability analysis and can be considered as an alternative to the Karp–Miller algorithm which uses forward reachability analysis. The decidability of simulation with finite transition systems in both directions is shown by Jančar and Moller [JM95].

7.3. Real-Time Automata

The data domain of a *real-time automaton* consists of the set of mappings from a finite set X of *clocks* to the set of nonnegative real numbers. Real-time automata have in recent years become important for modeling and analysis of time-dependent systems. The exact definitions of real-time automata vary slightly. Here we give a typical presentation. A command in Cmd is of the form $\langle q, \alpha, \lambda, q' \rangle$, where $q, q' \in Q$, $\lambda \in A$, and α is a guarded command of form $g \rightarrow stmt$ in which

- the guard g is a Boolean combination of inequalities of the form $x \sim n$, where $x \in X$, n is an integer, and \sim is one of the relations $\geq, \leq, <$ or $>$.
- the body $stmt$ for each $x \in X$ contains an assignment of one of the forms $x := x$ or $x := 0$.

Each command $\langle q, (g \rightarrow stmt), \lambda, q' \rangle$ gives rise to the set of transitions of the form $\langle q, d \rangle \xrightarrow{\lambda} \langle q, d' \rangle$, where the values of clocks in d satisfy g , and where the value of x in d' is obtained by performing the assignments in $stmt$. Also, if $\langle q, d \rangle \xrightarrow{\lambda} \langle q', d' \rangle$ is a transition obtained from a command in one of the ways just described, then a transition of the form $\langle q, d_1 \rangle \xrightarrow{\lambda} \langle q', d'_1 \rangle$ may be performed for any d_1 and d'_1 such that d_1 is obtained from d by letting all clocks advance the same (positive real-valued) amount, and d'_1 is analogously obtained from d' .

As the preorder \preceq we take the equivalence relation on clock states, introduced in [ACD90]. This is the largest equivalence induced by predicates of the form $x \sim n$, or of the form $x_1 \leq n \wedge x_2 \leq n \wedge x_1 - x_2 \sim n$ for clocks x , x_1 , and x_2 , a relation $\sim \in \{ \geq, \leq, <, > \}$, and a nonnegative integer n which is at most equal to the largest constant that occurs syntactically in commands.

It can be shown that \preceq is an equivalence, which is also a bisimulation, and hence the transition system is monotonic. The system is well quasi-ordered since there are finitely many equivalence classes. The computation rules for minpre_λ , intersection effectiveness, and the essentially finite branching property are also rather straightforward.

7.4. Relational Automata

A relational automaton (RA) is a computing device which besides having a finite control structure possesses a finite number of data variables, each one taking its value from some (possibly infinite) ordered domain. The operations that the automaton can perform on the data variables are comparison and assignment. The ordering of the data values stored into the variables during runtime may influence the control flow of the automaton.

A *rational relation automaton* (QRA) has a finite set X of *data variables* that assume values in the set \mathcal{Q} of rational numbers. A command in Cmd is a quadruple of the form $\langle q, \alpha, \lambda, q' \rangle$, where $q, q' \in \mathcal{Q}$, $\lambda \in \mathcal{A}$, and α is a guarded command of form $g \rightarrow \text{stmt}$ in which

- the guard g is a Boolean combination of inequalities of form $x < y$, $x < c$, and $c < y$ for $x, y \in X$ and $c \in \mathcal{Q}$, and where
- the body stmt contains, for each $x \in X$, an assignment of one of the forms $x := y$, $x := c$, or $x := \{?\}$ for $y \in X$ and $c \in \mathcal{Q}$.

Given a QRA P , we let $\text{Cons}(P)$ be the set of all constants that occur syntactically in the command of P (clearly, there are only a finite number of those). The set D of data values is the set $X \mapsto \mathcal{Q}$ of possible combinations of values of the data variables. For $c \in \mathcal{Q}$, we use the convention that $d(c) = c$ for any data value d . Each command $\langle q, (g \rightarrow \text{stmt}), \lambda, q' \rangle$ gives rise to the set of transitions of form $\langle q, d \rangle \xrightarrow{\lambda} \langle q', d' \rangle$, where the values of variables in d satisfy g , and where the value of x in d' is equal to

- $d(y)$ if stmt contains $x := y$, for $y \in X \cup \text{Cons}(P)$.
- any element of \mathcal{Q} if stmt contains $x := \{?\}$.

We say that two data values $d_1, d_2 \in (\mathcal{Q} \mapsto \mathcal{Q})$ are *equivalent* if for all x, y in $(X \cup \text{Cons}(P))$ we have $d_1(x) \leq d_1(y)$ iff $d_2(x) \leq d_2(y)$. In other words, two data values are equivalent if the relative ordering between the data variables and constants of the program is the same. It turns out that when \preceq is taken to be this equivalence on D , each QRA becomes a well-structured system. The equivalence \preceq has a finite number of equivalence classes.

We also consider the variation of relational automata, called *integral relational automata* (IRA), obtained by replacing the data domain by the set of integers. In

the case of IRAs, the situation is slightly more complicated, due to the fact that the ordering on the set of integers is not dense. For IRAs there is no equivalence relation which can be taken as \leq (see [Čer94]).

We say that a data value $d_2 \in (X \mapsto I)$ is *sparser than* a data value $d_1 \in (X \mapsto I)$, denoted $d_1 \leq d_2$, if for all x, y in $(X \cup \text{Cons}(P))$ we have that $d_1(x) \leq d_1(y)$ implies $d_1(y) - d_1(x) \leq d_2(y) - d_2(x)$. For instance, if $\text{Cons}(P) = \{0, 1, 2\}$, then the data vector $\langle 2, 10, 12, 1997 \rangle$ is sparser than $\langle 2, 4, 6, 1000 \rangle$, but not sparser than $\langle 1, 10, 12, 1997 \rangle$ since the value of the first variable is no longer equal to the constant 2, and not sparser than $\langle 2, 4, 7, 17 \rangle$ since $7 - 4 > 12 - 10$. If $\text{Cons}(P) \neq \emptyset$, then the *sparser than* relation is a partial order. If $\text{Cons}(P) = \emptyset$, then the relation is a preorder, but neither an equivalence nor a partial order. We note that if $d_1 \leq d_2$, we can define a strictly monotone mapping $\rho: I \mapsto I$ such that

- $\rho(d_1(x)) = d_2(x)$ for all $x \in (X \cup \text{Cons}(P))$, and
- $a \leq b$ implies $b - a \leq \rho(b) - \rho(a)$ for all $a, b \in I$.

We call such a mapping ρ a *sparsifier*.

We now show that \leq meets the requirements for well-structured systems.

- *Monotonicity*: Suppose that $\langle q, d_1 \rangle \xrightarrow{\lambda} \langle q', d'_1 \rangle$ and $d_1 \leq d_2$. We show that there exists $d'_2 = d'_1 \leq d'_2$ such that $\langle q, d_2 \rangle \xrightarrow{\lambda} \langle q', d'_2 \rangle$. We know that there is a tuple $\langle q, g \rightarrow \text{stmt}, \lambda, q' \rangle \in \text{Cmd}$, where d_1 satisfies g and d'_1 is derived from d_1 according to stmt . Since $d_1 \leq d_2$ there is a sparsifier ρ which satisfies the conditions above. For any $x \in X$ we let $d'_2(x) = \rho(d'_1(x))$. It follows that $\langle q, d_2 \rangle \xrightarrow{\lambda} \langle q', d'_2 \rangle$ and $d'_1 \leq d'_2$.

- *Well quasi-orderedness*: Consider an infinite sequence d_0, d_1, d_2, \dots of data values. We observe that there is an infinite increasing sequence i_0, i_1, i_2, \dots such that the variables and constants have the same relative ordering in each d_{i_j} . In other words, $d_{i_j}(x) \leq d_{i_j}(y)$ iff $d_{i_k}(x) \leq d_{i_k}(y)$ for each $j, k \geq 0$ and $x, y \in X \cup \text{Cons}(P)$. This implies that $d_{i_j} \leq d_{i_k}$ iff $|d_{i_j}(x) - d_{i_j}(y)| \leq |d_{i_k}(x) - d_{i_k}(y)|$ for all $x, y \in X \cup \text{Cons}(P)$. Since $|d_{i_j}(x) - d_{i_j}(y)|$ is a natural number, the result follows as a special case of Dickson's lemma [Dic13], which states that for any n , the \leq ordering on n -tuples of natural numbers is a well quasi-ordering.

- *Computability of minpre_λ* : We define $\text{minpre}_\lambda(\{\langle q, d \rangle\})$ to be $\text{min}(T)$, where T is the smallest set with elements satisfying the following properties. Let $\langle q', g \rightarrow \text{stmt}, \lambda, q \rangle \in \text{Cmd}$. Let $\{X_1, X_2\}$ be a partitioning of X such that $x \in X_1$ iff x appears in the right hand side of an assignment operation in stmt . Let k be the number of elements in X_2 . We say that a data element d_0 is k -close to d iff $d \leq d_0$ and $|d_0(x) - d_0(y)| - |d(x) - d(y)| \leq k$ and $|d_0(x) - d(x)| \leq k$ for all $x, y \in X \cup \text{Cons}(P)$. It is clear that for any d the set of k -close d_0 is finite.

The set T contains each element $\langle q', d' \rangle$ such that d' satisfies g and for some d_0 which is k -close to d :

- if $y := x$ occurs in stmt , then $d'(x) = d_0(y)$, and
- if $x \in X_2$, then $d'(x) \in \{d_{\min} - k, \dots, d_{\max} + k\}$, where d_{\min} and d_{\max} are the minimal and maximal values of $d'(y)$ for $y \in X_1 \cup \text{Cons}(P)$.

Intersection effectiveness of IRA is straightforward; however, IRAs are not necessarily essentially finite branching due to the presence of random assignment.

The decidability of control state reachability for IRA is shown in [BBK74, BBK77]. A detailed study of checking properties of IRA, which includes the decidability of eventuality properties, can be found in [Čer94]. The decidability of simulation of an IRA by a finite transition system has not been published before.

8. CONCLUSIONS AND COMMENTS

We provide a unified approach to algorithmic verification of several classes of infinite-state systems. Our method generalizes verification methods based on finite partitioning in the sense that (i) we consider a well quasi-ordered state space instead of a finitely partitioned one, and (ii) we consider states which are related through simulation instead of bisimulation. We are able to explain and derive several seemingly diverse algorithms for verification of different classes of infinite-state systems in a uniform manner. We aim to extend the applicability of our results to derive novel algorithms for verification of new classes of infinite-state states such as parametrized networks of processes and programs with multi-sorted domains.

Although we show decidability of reachability, and hence several classes of safety properties, model checking of, e.g., CTL or PTL formulas is in general undecidable for our class of systems. For example, we show in [AJ96] the undecidability of both these logics for lossy channel systems. Also, there are several computation models, such as basic process algebras [BK85, CHS92] and push-down automata, which have been considered in the literature of infinite-state systems, and which cannot be described within our framework. In the paper we do not offer any complexity analysis of our algorithms. However, the application of the reachability algorithm is feasible for, e.g., lossy channel systems [Kin93]. For example, the verification of a sliding window protocol with more than 10,000 control states and two (unbounded) lossy channels is carried out in a few seconds.

ACKNOWLEDGMENTS

We are grateful to Mats Kindahl, S. Purushothaman-Iyer, and Wang Yi for many interesting comments and discussions.

Final manuscript received June 25, 1998

REFERENCES

- [ACD90] Alur, R., Courcoubetis, C., and Dill, D. (1990), Model-checking for real-time systems, in "Proc. 5th IEEE Int. Symp. on Logic in Computer Science," pp. 414–425, Philadelphia.
- [AH89] Alur, R., and Henzinger, T. (1989), A really temporal logic, in "Proc. 30th Annual Symp. Foundations of Computer Science," pp. 164–169.
- [AJ93] Abdulla, P. A., and Jonsson, B. (1993), Verifying programs with unreliable channels, in "Proc. 8th IEEE Int. Symp. on Logic in Computer Science," pp. 160–170.
- [AJ96] Abdulla, P. A., and Jonsson, B. (1996), Undecidable verification problems for programs with unreliable channels, *Inform. and Comput.* **130**, 71–90.

- [AK95] Abdulla, P. A., and Kindahl, M. (1995), Decidability of simulation and bisimulation between lossy channel systems and finite state systems, in "Proc. CONCUR '95, 6th Int. Conf. on Concurrency Theory" (Lee and Smolka, Eds.), Lecture Notes in Computer Science, Vol. 962, pp. 333–347, Springer-Verlag, Berlin/New York.
- [BBK74] Barzdin, J. M., Bicevskis, J. J., and Kalnins, A. A. (1974), Construction of complete sample systems for program testing, *Latv. Gosudarst. Univ. Uch. Zapiski* **210**. [In Russian]
- [BBK77] Barzdin, J. M., Bicevskis, J. J., and Kalnins, A. A. (1977), Automatic construction of complete systems for program testing, in "IFIP Congress 1977."
- [BK85] Bergstra, J., and Klop, J. (1985), Algebra of communicating process with abstraction, *Theoret. Comput. Sci.* **37**, 77–121.
- [Boc78] Bochman, G. V. (1978), Finite state description of communicating protocols, *Comput. Networks* **2**, 361–371.
- [BZ83] Brand, D., and Zafiropulo, P. (1983), On communicating finite-state machines, *J. Assoc. Comput. Mech.* **2**(5), 323–342.
- [CC77] Cousot, P., and Cousot, R. (1977), Abstract interpretation: A unified model for static analysis of programs by construction or approximation of fixpoints, in "Proc. 4th ACM Symp. on Principles of Programming Languages," pp. 238–252.
- [Čer92a] Čerāns, K. (1992), Decidability of bisimulation equivalence for parallel timer processes, in "Proc. Workshop on Computer Aided Verification," Lecture Notes in Computer Science, Vol. 663, pp. 302–315.
- [Čer92b] Čerāns, K. (1992), Feasibility of finite and infinite paths in data dependent programs, in "LFCS'92," Lecture Notes in Computer Science, Vol. 620, pp. 69–80.
- [Čer94] Čerāns, K. (1994), Deciding properties of integral relational automata, in "Proc. ICALP '94" (Abiteboul and Shamir, Eds.), Lecture Notes in Computer Science, Vol. 820, pp. 35–46, Springer-Verlag, Berlin/New York.
- [CG87] Clarke, E. M., and Grumberg, O. (1987), Avoiding the state explosion problem in temporal logic model checking algorithms, in "Proc. 6th ACM Symp. on Principles of Distributed Computing, Vancouver, Canada," pp. 294–303.
- [CGL92] Clarke, E. M., Grumberg, O., and Long, D. E. (1992), Model checking and abstraction, in "Proc. 19th ACM Symp. on Principles of Programming Languages."
- [CHS92] Christensen, S., Hüttel, H., and Stirling, C. (1992), Bisimulation equivalence is decidable for all context-free processes, in "Proc. CONCUR '92, Theories of Concurrency: Unification and Extension" (Q. R. Cleaveland, Ed.), pp. 138–147.
- [DGG94] Dams, D., Grumberg, O., and Gerth, R. (1994), Abstract interpretation of reactive systems: Abstractions preserving \forall CTL*, \exists CTL* and CTL*, in "Proc. IFIP Working Conference on Programming Concepts, Methods and Calculi (PROCOMET '94)."
- [Dic13] Dickson, L. E. (1913), Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors, *Amer. J. Math.* **35**, 413–422.
- [Fin90] Finkel, A. (1990), Reduction and covering of infinite reachability trees, *Inform. and Comput.* **89**, 144–179.
- [Hen95] Henzinger, T. A. (1995), Hybrid automata with finite bisimulations, in "Proc. ICALP '95," Lecture Notes in Computer Science, Vol. 944, pp. 324–335, Springer-Verlag, Berlin/New York.
- [Hig52] Higman, G. (1952), Ordering by divisibility in abstract algebras, *Proc. London Math. Soc.* **2**, 326–336.
- [Jan90] Jančar, P. (1990), Decidability of a temporal logic problem for Petri nets, *Theoret. Comput. Sci.* **74**, 71–93.
- [JM95] Jančar, P., and Moller, F. (1995), Checking regular properties of Petri nets, in "Proc. CONCUR '95, 6th Int. Conf. on Concurrency Theory," pp. 348–362.
- [JN94] Jones, N. D., and Flemming, N. (1994), Abstract interpretation: A semantics-based tool for program analysis, in "Handbook of Logic in Computer Science," Oxford Univ. Press, London.

- [Kin93] Kindahl, M. (1993), “Implementation of a Reachability Algorithm for Systems with Unreliable Channels,” Master’s thesis, Department of Computer Systems, Uppsala University, Sweden. [Available as Report DoCS 93/44]
- [KM69] Karp, R. M., and Miller, R. E. (1969), Parallel program schemata, *J. Comput. Systems Sci.* **3**, 147–195.
- [LGS⁺95] Loiseaux, C., Graf, S., Sifakis, J., Boujjani, A., and Bensalem, S. (1995), Property preserving abstractions for the verification of concurrent systems, *Formal Methods in System Design* **6**, 11–44.
- [PP92] Peng, W., and Purushothaman, S. (1992), Analysis of a class of communicating finite state machines, *Acta Inform.* **29**(6/7), 499–522.