

Efficient Guided Symbolic Reachability using Reachability Expressions

Dina Thomas¹, Supratik Chakraborty¹, Paritosh Pandya²

¹ Indian Institute of Technology, Bombay, India
e-mail: dina@cfavs.iitb.ac.in, supratik@cse.iitb.ac.in

² Tata Institute of Fundamental Research, India
e-mail: pandya@tifr.res.in

Received: date / Revised version: date

Abstract. Asynchronous systems consist of a set of transitions which are non-deterministically chosen and executed. We present a theory of guiding symbolic reachability in such systems by scheduling clusters of transitions. A theory of reachability expressions which specify the schedules is presented. This theory allows proving equivalence of different schedules which may have radically different performance in BDD-based search. We present experimental evidence to show that optimized reachability expressions give rise to significant performance advantages. The profiling is carried out in the NuSMV framework using examples from discrete timed automata and circuits with delays. A variant tool called NuSMV-DP has been developed for interpreting reachability expressions to carry out the experiments.

1 Introduction

Asynchronous systems consist of a set of processes which execute independently of each other and synchronize occasionally. A standard model of their execution consists of non-deterministically interleaving the actions of individual processes. Activities of such processes can be modeled by a *global transition system* consisting of a set of guarded transitions ($G \mapsto A$). The system starts non-deterministically in one of a set of designated initial states. In any state, one of the enabled transitions is non-deterministically chosen and executed atomically. This causes a state change. This process is then repeated until no new state is reached. Safety verification of such systems typically reduces to exploring whether some undesirable state is reachable by some execution.

Symbolic model checking [10] has emerged as an important technique for program verification and for finding deep logical bugs in reactive systems. In this technique, programs are modeled as finite state transition

systems and the state space of the model is searched for undesirable behaviour. BDD-based [3] symbolic search techniques, e.g. those used in NuSMV [9], can explore very large but finite state spaces efficiently. However, there is wide variability in the computational efficiency of BDD-based searches. It is well-known that the performance of these techniques strongly depends on the size of the BDD representation of transition relations and of intermediate sets of states.

Earlier work in this area has addressed this issue by identifying good variable orders for BDDs representing transition relations [1], by partitioning the transition relation conjunctively or disjunctively [11], and by determining good quantification schedules for conjunctive partitioning [6, 7]. Yet other techniques include guided search where hints [2] are used to direct initial parts of the search, iterative squaring [4, 5] where the number of image computation steps is reduced exponentially at the cost of transition relation squaring, and saturation-based methods [8] that guide the search to exploit local transition clusters so that peak memory requirements are reduced.

In this paper, we propose using clusters of individual guarded transitions within the global transition system in ways that generalize conjunctive and disjunctive partitioning. We introduce a notation called *reachability expressions* and investigate its algebraic properties. Each reachability expression denotes a way of computing the set of final states from a set of initial states. Thus, each reachability expression is a predicate transformer. Reachability expressions are rich enough to specify diverse search strategies such as symbolic breadth-first, round-robin etc. They also include familiar operations like sequential composition, union and Kleene closure. Importantly, they allow us to encode more efficient ways of computing the set of reachable states than symbolic breadth-first search.

We have implemented an interpreter for reachability expressions in a tool called NuSMV-DP¹ that works as a wrapper on top of the reachability engine of NuSMV [9]. Using the semantics of reachability expressions, we show that several distinct reachability expressions are equivalent, i.e. they compute the same predicate transformation. However, as our experiments show, the computational effort involved in applying these equivalent predicate transformers to a given predicate using our BDD-based NuSMV-DP tool can vary significantly. We discuss some equivalence transformations which improve the efficiency of evaluating reachability expressions on finite state systems. A Kleene algebra of reachability expressions has also been investigated by Pandya and Raut [12].

We apply our theory to some examples drawn from discrete timed automata. Such automata can be represented as finite state global transition systems, as discussed in [14]. We experimentally evaluate the performance of BDD-based symbolic reachability analysis on these examples using multiple equivalent reachability expressions. A previous technical report [14] gives details of experiments carried out using NuSMV-DP with diverse problems such as analyzing the Fischer protocol, the job-shop scheduling problem and some asynchronous circuits with delays. These experiments show a significant improvement in the efficiency of computing reachable states using our technique. For example, we have been able to model check the Fischer protocol with 100 processes, whereas classical techniques such as symbolic search using polyhedra or difference bound matrices can handle instances of this protocol with only up to 20 processes. On other examples, the relative gains are significant but more modest.

The remainder of this paper is organized as follows. In Section 2, we present the basic theory of reachability expressions. Section 3 gives experimental results obtained by applying this theory to improve the efficiency of reachability analysis. The experiments are carried out using the NuSMV-DP tool. Finally, we conclude the paper in Section 4.

2 A Theory of Reachability Expressions

We consider a (possibly infinite) state transition system as a 4-tuple (V, Q, Q_0, \mathcal{Y}) , where V is a finite set of *state variables*, Q is the set of states, $Q_0 \subseteq Q$ is the set of initial states, and \mathcal{Y} is a finite set of *guarded actions*. Each variable $v_i \in V$ has an associated domain \mathcal{D}_i . A state $q \in Q$ is an assignment of a value from \mathcal{D}_i to each variable v_i in V . The set of all such assignments constitutes the set of states Q . A *guarded action* is a pair $(G \mapsto A)$, where the guard G denotes a Boolean combination of predicates on the variables in V or a constant

in $\{\text{True}, \text{False}\}$. The action A is either a multiple assignment statement denoting simultaneous assignments to a finite set of variables in V , or the special action *skip*. We leave the concrete syntax of guarded actions unspecified. Note that each state variable $v_i \in V$ may be assigned a value from \mathcal{D}_i at most once in A . The execution semantics of guarded actions is as usual: If A consists of the special action *skip*, the values of all state variables remain unchanged. If the system is in state s_1 , and if the corresponding assignment of values to variables in V satisfies the guard G , we say that the guarded action $(G \mapsto A)$ is *enabled* in state s_1 . The new state reached after executing the action A from state s_1 is obtained by simultaneously assigning to all variables that have been assigned in A , their corresponding values. All state variables that have not been assigned in A retain their values from state s_1 in the next state.

Let $B = (V, Q, Q_0, \mathcal{Y})$ be a state transition system. We define a *cluster* to be a non-empty set of guarded actions of B . We also define a special singleton set δ consisting of the guarded action $(\text{True} \mapsto \text{skip})$. Thus, the action in δ can be executed from every state, and its execution takes every state to itself. The empty set of guarded actions is denoted by \emptyset . An *extended cluster* of the state transition system B is either a subset of \mathcal{Y} or δ or \emptyset . Every extended cluster T defines a relation, R_T , on the set Q of states. We say that $(s_1, s_2) \in R_T$ iff there exists a guarded action $(G \mapsto A) \in T$ such that G evaluates to *True* in s_1 , and action A takes the system from s_1 to s_2 . Given an extended cluster T and a set of states $S (\subseteq Q)$, the *image* of S under T is $Im_T(S) = \{s \mid \exists s' \in S, R_T(s', s)\}$. It is easy to see that $Im_T : 2^Q \rightarrow 2^Q$ is a monotone function with $Im_T(\emptyset) = \emptyset$.

2.1 Syntax and Semantics

Let $\mathcal{T} = \{\tau_1, \dots, \tau_k\}$ be a set of extended clusters of B . Syntactically, a *reachability expression* over \mathcal{T} is a terminal string obtained from the following grammar:

$$E \rightarrow E + E \mid E \mid E \circ E \mid (E) \mid *E \mid \mathbf{T}_1 \mid \dots \mid \mathbf{T}_k$$

In the above syntax, we have used $\mathbf{T}_1, \dots, \mathbf{T}_k$ to denote reachability expressions corresponding to extended clusters τ_1, \dots, τ_k in B . We will use this notation throughout this paper.

The notion of evaluating reachability expressions can be formalized by defining their semantics. The semantics of a reachability expression is defined with respect to an underlying state transition system, and is naturally described as a mapping from sets of states to sets of states. Let $B = (V, Q, q_0, \mathcal{Y})$ be a state transition system, and \mathcal{T} be a set of extended clusters of B . Let σ be a reachability expression over \mathcal{T} and let $S \subseteq Q$. The semantics of σ with respect to B , denoted by $\llbracket \sigma \rrbracket_B$, is a mapping from 2^Q to 2^Q defined inductively as follows. We shall henceforth omit the subscript B when it is clear from the context.

¹ DP stands for “Disjunctive Partitioning” for historical reasons, although the tool can interpret arbitrary reachability expressions.

- $\llbracket \mathbf{T}_i \rrbracket (S) = \text{Im}_{\tau_i}(S)$, for all $\tau_i \in \mathcal{T}$
- $\llbracket \sigma_1 + \sigma_2 \rrbracket (S) = \llbracket \sigma_1 \rrbracket (S) \cup \llbracket \sigma_2 \rrbracket (S)$
- $\llbracket \sigma_1 \circ \sigma_2 \rrbracket (S) = \llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (S))$
- $\llbracket \sigma_1 ; \sigma_2 \rrbracket (S) = \llbracket (\sigma_1 + \delta) \circ (\sigma_2 + \delta) \rrbracket (S)$
- $\llbracket (\sigma) \rrbracket (S) = \llbracket \sigma \rrbracket (S)$
- $\llbracket \sigma^0 \rrbracket (S) = (S)$ and $\llbracket \sigma^{i+1} \rrbracket (S) = \llbracket \sigma \circ \sigma^i \rrbracket (S)$ for all $i \geq 0$.
- $\llbracket * \sigma \rrbracket (S) = \bigcup_{i=0}^{\infty} \llbracket \sigma^i \rrbracket (S)$

Although “;” is seen to be a derived operator, we retain it for notational convenience.

2.2 Properties of Reachability Expressions

Lemma 1. *Let σ be an arbitrary reachability expression and $B = (V, Q, Q_0, \Upsilon)$ be a state transition system. Then (i) $\llbracket \sigma \rrbracket (\emptyset) = \emptyset$, and (ii) $\llbracket \sigma \rrbracket (S_1 \cup S_2) = \llbracket \sigma \rrbracket (S_1) \cup \llbracket \sigma \rrbracket (S_2)$ for all $S_1, S_2 \subseteq Q$.*

Proof. We use induction on the structure of σ .

Basis: If $\sigma = \mathbf{T}_i$, both parts of the lemma hold from the definition of Im_{τ_i} . If $\sigma = \delta$, $\llbracket \delta \rrbracket (\emptyset) = \emptyset$ by definition of δ , and $\llbracket \delta \rrbracket (S_1 \cup S_2) = S_1 \cup S_2 = \llbracket \delta \rrbracket (S_1) \cup \llbracket \delta \rrbracket (S_2)$. If $\sigma = \emptyset$, $\llbracket \sigma \rrbracket (S) = \emptyset$ for all $S \subseteq Q$; hence both parts of the lemma hold.

Hypothesis: Let σ_1 and σ_2 be reachability expressions satisfying both parts of the lemma. *Induction Step:* We consider the following cases:

1. $\sigma = \sigma_1 + \sigma_2$:
 - (i) Emptiness of $\llbracket \sigma_1 + \sigma_2 \rrbracket (\emptyset)$ follows from the definition of the semantics and from the induction hypothesis.
 - (ii) By definition, $\llbracket \sigma_1 + \sigma_2 \rrbracket (S_1 \cup S_2) = \llbracket \sigma_1 \rrbracket (S_1 \cup S_2) \cup \llbracket \sigma_2 \rrbracket (S_1 \cup S_2)$. By induction hypothesis, the latter is equal to $\llbracket \sigma_1 \rrbracket (S_1) \cup \llbracket \sigma_1 \rrbracket (S_2) \cup \llbracket \sigma_2 \rrbracket (S_1) \cup \llbracket \sigma_2 \rrbracket (S_2)$. By associativity of set union, the above is equal to $(\llbracket \sigma_1 \rrbracket (S_1) \cup \llbracket \sigma_2 \rrbracket (S_1)) \cup (\llbracket \sigma_1 \rrbracket (S_2) \cup \llbracket \sigma_2 \rrbracket (S_2))$. By the semantics defined above, this is equal to $\llbracket \sigma_1 + \sigma_2 \rrbracket (S_1) \cup \llbracket \sigma_1 + \sigma_2 \rrbracket (S_2)$.
2. $\sigma = \sigma_1 \circ \sigma_2$:
 - (i) By definition, $\llbracket \sigma \rrbracket (\emptyset) = \llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (\emptyset))$. It therefore follows from the induction hypothesis that this set is empty.
 - (ii) To evaluate $\llbracket \sigma_1 \circ \sigma_2 \rrbracket (S_1 \cup S_2)$, we notice from the definition of the semantics that this is $\llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (S_1 \cup S_2))$. By the induction hypothesis, this is equal to $\llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (S_1) \cup \llbracket \sigma_1 \rrbracket (S_2)) = \llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (S_1) \cup \llbracket \sigma_2 \rrbracket (\llbracket \sigma_1 \rrbracket (S_2)))$. By the definition of the semantics, this equals $\llbracket \sigma_1 \circ \sigma_2 \rrbracket (S_1) \cup \llbracket \sigma_1 \circ \sigma_2 \rrbracket (S_2)$.
3. $\sigma = \sigma_1 ; \sigma_2$: Since $\llbracket \sigma \rrbracket = \llbracket (\sigma_1 + \delta) \circ (\sigma_2 + \delta) \rrbracket$, both parts of the lemma are easily proven from the induction hypothesis and from the two cases proved above.
4. $\sigma = (\sigma_1)$: The proof follows from the induction hypothesis and from the definition of $\llbracket (\sigma_1) \rrbracket$.

5. $\sigma = * \sigma_1$: We first show by induction on i that $\llbracket \sigma_1^i \rrbracket (S_1 \cup S_2) = \llbracket \sigma_1^i \rrbracket (S_1) \cup \llbracket \sigma_1^i \rrbracket (S_2)$ and $\llbracket \sigma_1^i \rrbracket (\emptyset) = \emptyset$ for all $i \geq 0$.
 - *Basis:* If $i = 0$, $\sigma_1^i = \delta$ by definition, and we have already argued above that δ satisfies both parts of the lemma.
 - *Hypothesis:* Assume that $\llbracket \sigma_1^i \rrbracket$ satisfies both parts of the lemma for all i such that $0 \leq i \leq k$.
 - *Induction Step:* By definition, $\sigma_1^{k+1} = \sigma_1 \circ \sigma_1^k$. By hypotheses, both $\llbracket \sigma_1 \rrbracket$ and $\llbracket \sigma_1^k \rrbracket$ satisfy both parts of the lemma. Therefore, by case 2 proved above, $\llbracket \sigma_1^{k+1} \rrbracket$ also satisfies both parts of the lemma. Hence, $\llbracket \sigma_1^i \rrbracket (S_1 \cup S_2) = \llbracket \sigma_1^i \rrbracket (S_1) \cup \llbracket \sigma_1^i \rrbracket (S_2)$, and $\llbracket \sigma_1^i \rrbracket (\emptyset) = \emptyset$ for all $i \geq 0$.

We now prove the two parts of the lemma.

(i) To show that $\llbracket * \sigma_1 \rrbracket (\emptyset) = \emptyset$, we use proof by contradiction. Let $q \in \llbracket * \sigma_1 \rrbracket (\emptyset)$. Then there exists $r \geq 0$ such that $q \in \llbracket \sigma_1^r \rrbracket (\emptyset)$. This means $\llbracket \sigma_1^r \rrbracket (\emptyset)$ is non-empty, which contradicts $\llbracket \sigma_1^r \rrbracket (\emptyset) = \emptyset$ proved above.

(ii) Let us now assume that $q \in \llbracket * \sigma_1 \rrbracket (S_1 \cup S_2)$. By definition of the semantics, there exists an integer $r (\geq 0)$ such that $q \in \llbracket \sigma_1^r \rrbracket (S_1 \cup S_2)$. We have already seen above that $\llbracket \sigma_1^r \rrbracket (S_1 \cup S_2) = \llbracket \sigma_1^r \rrbracket (S_1) \cup \llbracket \sigma_1^r \rrbracket (S_2)$ for all $i \geq 0$. Therefore, $q \in \llbracket \sigma_1^r \rrbracket (S_1) \cup \llbracket \sigma_1^r \rrbracket (S_2)$, and hence $q \in \llbracket * \sigma_1 \rrbracket (S_1) \cup \llbracket * \sigma_1 \rrbracket (S_2)$. This shows that $\llbracket * \sigma_1 \rrbracket (S_1 \cup S_2) \subseteq \llbracket * \sigma_1 \rrbracket (S_1) \cup \llbracket * \sigma_1 \rrbracket (S_2)$. Similarly, if $q \in \llbracket * \sigma_1 \rrbracket (S_1) \cup \llbracket * \sigma_1 \rrbracket (S_2)$. Then there exists an integer $r (\geq 0)$ such that either $q \in \llbracket \sigma_1^r \rrbracket (S_1 \cup S_2)$ or $q \in \llbracket \sigma_2^r \rrbracket (S_1 \cup S_2)$. By the induction hypothesis, this implies that $q \in \llbracket \sigma_1^r \rrbracket (S_1 \cup S_2)$ and hence $q \in \llbracket * \sigma_1 \rrbracket (S_1 \cup S_2)$. It then follows that $\llbracket * \sigma_1 \rrbracket (S_1) \cup \llbracket * \sigma_1 \rrbracket (S_2) \subseteq \llbracket * \sigma_1 \rrbracket (S_1 \cup S_2)$, which along with our earlier result implies $\llbracket * \sigma_1 \rrbracket (S_1 \cup S_2) = \llbracket * \sigma_1 \rrbracket (S_1) \cup \llbracket * \sigma_1 \rrbracket (S_2)$.

Corollary 1. *For all reachability expressions σ , $\llbracket \sigma \rrbracket : 2^Q \rightarrow 2^Q$ is a monotone function.*

Proof. Let $S_1 \subseteq S_2 \subseteq Q$. Then there exists S_3 such that $S_2 = S_1 \cup S_3$. By Lemma 1, $\llbracket \sigma \rrbracket (S_2) = \llbracket \sigma \rrbracket (S_1) \cup \llbracket \sigma \rrbracket (S_3)$. Hence, $\llbracket \sigma \rrbracket (S_1) \subseteq \llbracket \sigma \rrbracket (S_2)$.

Let σ_1 and σ_2 be reachability expressions over a set \mathcal{T} of symbolic extended clusters. We say that σ_1 is *covered by* σ_2 iff $\llbracket \sigma_1 \rrbracket_B(S) \subseteq \llbracket \sigma_2 \rrbracket_B(S)$ for every state transition system B , for every subset S of states of B , and for every instantiation of symbolic clusters in \mathcal{T} with extended clusters of B . We denote this by $\sigma_1 \sqsubseteq \sigma_2$. We say that $\sigma_1 = \sigma_2$, iff $\sigma_1 \sqsubseteq \sigma_2$ and $\sigma_2 \sqsubseteq \sigma_1$.

The following properties of reachability expressions follow from the definition of their semantics and from basic set theory. In the following, $\sigma, \sigma_1, \sigma_2, \sigma_3$ refer to arbitrary reachability expressions over a set of extended clusters. S is some set of states ($S \subseteq Q$).

P1: $(\sigma + \sigma) = \sigma$.

By definition, $\llbracket (\sigma + \sigma) \rrbracket (S) = \llbracket \sigma \rrbracket (S) \cup \llbracket \sigma \rrbracket (S) = \llbracket \sigma \rrbracket (S)$.

P2: $(\sigma_1 + \sigma_2) = (\sigma_2 + \sigma_1)$.

This follows from the definition of the semantics and commutativity of set union.

P3: $(\sigma_1 + (\sigma_2 + \sigma_3)) = ((\sigma_1 + \sigma_2) + \sigma_3)$.

This follows from the definition of the semantics and associativity of set union.

P4: $(\sigma + \Theta) = \sigma$.

By definition, $\llbracket (\sigma + \Theta) \rrbracket(S) = \llbracket \sigma \rrbracket(S) \cup \llbracket \Theta \rrbracket(S) = \llbracket \sigma \rrbracket(S) \cup \phi = \llbracket \sigma \rrbracket(S)$. Hence $(\sigma + \Theta) = \sigma$.

P5: $(\sigma_1 \circ (\sigma_2 \circ \sigma_3)) = ((\sigma_1 \circ \sigma_2) \circ \sigma_3)$.

By definition, $\llbracket (\sigma_1 \circ (\sigma_2 \circ \sigma_3)) \rrbracket(S) = \llbracket (\sigma_2 \circ \sigma_3) \rrbracket(\llbracket \sigma_1 \rrbracket(S)) = \llbracket \sigma_3 \rrbracket(\llbracket \sigma_2 \rrbracket(\llbracket \sigma_1 \rrbracket(S)))$. Similarly, $\llbracket ((\sigma_1 \circ \sigma_2) \circ \sigma_3) \rrbracket(S) = \llbracket \sigma_3 \rrbracket(\llbracket (\sigma_1 + \sigma_2) \rrbracket(S)) = \llbracket \sigma_3 \rrbracket(\llbracket \sigma_2 \rrbracket(\llbracket \sigma_1 \rrbracket(S)))$.

P6: $(\sigma \circ \delta) = (\delta \circ \sigma) = \sigma$.

By definition, $\llbracket (\sigma \circ \delta) \rrbracket(S) = \llbracket \delta \rrbracket(\llbracket \sigma \rrbracket(S))$. From the definition of δ , $\llbracket \delta \rrbracket(\llbracket \sigma \rrbracket(S)) = \llbracket \sigma \rrbracket(S)$. Similarly, $\llbracket (\delta \circ \sigma) \rrbracket(S) = \llbracket \sigma \rrbracket(\llbracket \delta \rrbracket(S)) = \llbracket \sigma \rrbracket(S)$.

P7: $(\sigma \circ \Theta) = (\Theta \circ \sigma) = \Theta$.

By definition of Θ , $\llbracket \Theta \rrbracket(S) = \phi$. Therefore $\llbracket (\sigma \circ \Theta) \rrbracket(S) = \llbracket \Theta \rrbracket(\llbracket \sigma \rrbracket(S)) = \phi$. Similarly, $\llbracket (\Theta \circ \sigma) \rrbracket(S) = \llbracket \sigma \rrbracket(\llbracket \Theta \rrbracket(S)) = \llbracket \sigma \rrbracket(\phi)$. By Lemma 1, $\llbracket \sigma \rrbracket(\phi) = \phi$. Hence $(\sigma \circ \Theta) = (\Theta \circ \sigma) = \Theta$.

P8: $(\sigma_1 \circ (\sigma_2 + \sigma_3)) = ((\sigma_1 \circ \sigma_2) + (\sigma_1 \circ \sigma_3))$.

By definition of the semantics, $\llbracket (\sigma_1 \circ (\sigma_2 + \sigma_3)) \rrbracket(S) = \llbracket (\sigma_2 + \sigma_3) \rrbracket(\llbracket \sigma_1 \rrbracket(S)) = \llbracket \sigma_2 \rrbracket(\llbracket \sigma_1 \rrbracket(S)) \cup \llbracket \sigma_3 \rrbracket(\llbracket \sigma_1 \rrbracket(S)) = \llbracket (\sigma_1 \circ \sigma_2) \rrbracket(S) \cup \llbracket (\sigma_1 \circ \sigma_3) \rrbracket(S)$. By definition of the semantics, the last term is $\llbracket (\sigma_1 \circ \sigma_2) + (\sigma_1 \circ \sigma_3) \rrbracket(S)$. Thus “ \circ ” distributes over “ $+$ ” on the left.

P9: $((\sigma_1 + \sigma_2) \circ \sigma_3) = ((\sigma_1 \circ \sigma_3) + (\sigma_2 \circ \sigma_3))$.

By definition of the semantics, $\llbracket ((\sigma_1 + \sigma_2) \circ \sigma_3) \rrbracket(S) = \llbracket \sigma_3 \rrbracket(\llbracket (\sigma_1 + \sigma_2) \rrbracket(S)) = \llbracket \sigma_3 \rrbracket(\llbracket \sigma_1 \rrbracket(S) \cup \llbracket \sigma_2 \rrbracket(S))$. By Lemma 1, the latter is equal to $\llbracket \sigma_3 \rrbracket(\llbracket \sigma_1 \rrbracket(S)) \cup \llbracket \sigma_3 \rrbracket(\llbracket \sigma_2 \rrbracket(S))$. By definition of the semantics, this is $\llbracket (\sigma_1 \circ \sigma_3) \rrbracket(S) \cup \llbracket (\sigma_2 \circ \sigma_3) \rrbracket(S) = \llbracket ((\sigma_1 \circ \sigma_3) + (\sigma_2 \circ \sigma_3)) \rrbracket(S)$. Thus, “ \circ ” distributes over “ $+$ ” on the right as well.

P10: $(\sigma_1 ; \sigma_2) = \delta + \sigma_1 + \sigma_2 + (\sigma_1 \circ \sigma_2)$.

Follows from the definition of the semantics and the distributivity properties P8 and P9 above.

P11: $(\sigma_1 ; (\sigma_2 ; \sigma_3)) = ((\sigma_1 ; \sigma_2) ; \sigma_3)$.

By definition, $(\sigma_1 ; (\sigma_2 ; \sigma_3)) = (\sigma_1 + \delta) \circ ((\sigma_2 ; \sigma_3) + \delta) = (\sigma_1 + \delta) \circ ((\sigma_2 + \delta) \circ (\sigma_3 + \delta) + \delta)$. Since $\delta \sqsubseteq ((\sigma_2 + \delta) \circ (\sigma_3 + \delta))$, the above simplifies to $(\sigma_1 + \delta) \circ ((\sigma_2 + \delta) \circ (\sigma_3 + \delta))$. From property P5 we get, $(\sigma_1 + \delta) \circ ((\sigma_2 + \delta) \circ (\sigma_3 + \delta)) = ((\sigma_1 + \delta) \circ (\sigma_2 + \delta)) \circ (\sigma_3 + \delta)$. Once again, since $\delta \sqsubseteq ((\sigma_1 + \delta) \circ (\sigma_2 + \delta))$, the above expression is equal to $((\sigma_1 + \delta) \circ (\sigma_2 + \delta) + \delta) \circ (\sigma_3 + \delta) = ((\sigma_1 ; \sigma_2) ; \sigma_3)$. This proves that “ $;$ ” is associative.

P11: $(\sigma ; \Theta) = (\Theta ; \sigma) = (\sigma + \delta)$.

Using the definition of “ $;$ ” and applying property P4 we get $(\sigma ; \Theta) = (\sigma + \delta) \circ (\Theta + \delta) = (\sigma + \delta) \circ \delta$. Applying property P6 we get $(\sigma + \delta) \circ \delta = (\sigma + \delta)$.

Similarly $(\Theta ; \sigma) = (\Theta + \delta) \circ (\sigma + \delta) = \delta \circ (\sigma + \delta) = (\sigma + \delta)$. Hence $(\sigma ; \Theta) = (\Theta ; \sigma) = (\sigma + \delta)$.

Since the “ $+$ ”, “ $;$ ” and “ \circ ” operators are associative (properties P3, P5, and P11), we omit parentheses for grouping arguments of these operators.

Given a set \mathcal{T} of extended clusters, let $\Pi(\mathcal{T})$ denote the set of all reachability expressions over \mathcal{T} . From properties P1, P2, P3 and P4 above, $(\Pi(\mathcal{T}), +)$ forms an idempotent, commutative monoid with Θ as the identity element. Similarly, from properties P5 and P6 $(\Pi(\mathcal{T}), \circ)$ forms a monoid with δ as the identity element. It follows that $(\Pi(\mathcal{T}), +, \circ)$ forms an idempotent semiring. Note that $(\Pi(\mathcal{T}), ;)$ does not have an identity element, and therefore forms a semigroup.

Lemma 2. *Let $\{T_i \mid 1 \leq i \leq k\}$ be a set of extended clusters. If $T = \bigcup_{i=1}^k T_i$, then $\mathbf{T} = \mathbf{T}_1 + \dots + \mathbf{T}_k$.*

Proof. Follows directly from the definition of T and from the definition of $\llbracket \mathbf{T}_1 + \dots + \mathbf{T}_k \rrbracket$.

Lemma 3. *Let $B = (V, Q, Q_0, \Upsilon)$ be a state transition system, σ be a reachability expression over a set \mathcal{T} of extended clusters of B , and S_0 be a subset of Q . Let S_j denote $\llbracket \bigoplus_{i=0}^j (\sigma)^i \rrbracket_B(S_0)$ for all $j \geq 0$. If $S_t = S_{t+1}$ for some $t (\geq 0)$, then $\llbracket *\sigma \rrbracket_B(S_0) = S_t$.*

Proof. Since $S_j = \llbracket \bigoplus_{i=0}^j (\sigma)^i \rrbracket_B(S_0)$, it follows from the definition of the semantics that $S_j \subseteq S_{j+1}$ for all $j \geq 0$. We now show using induction on r that if $S_t = S_{t+1}$, then $S_r = S_t$ for all $r \geq t + 1$.

Basis: Since $S_{t+1} = S_t$, the result holds for $r = t + 1$.

Hypothesis: Let $S_r = S_t$ for $t + 1 \leq r \leq n$.

Induction Step: Using the definition of the semantics we get $\llbracket (\sigma)^{n+1} \rrbracket_B(S_0) = \llbracket \sigma \circ (\sigma)^n \rrbracket_B(S_0)$. Using the associativity of “ \circ ”, $\llbracket \sigma \circ (\sigma)^n \rrbracket_B(S_0) = \llbracket (\sigma)^n \circ \sigma \rrbracket_B(S_0)$. By the induction hypothesis, $\llbracket \bigoplus_{i=0}^n (\sigma)^i \rrbracket_B(S_0) =$

$\llbracket \bigoplus_{i=0}^t (\sigma)^i \rrbracket_B(S_0)$. This implies that $\llbracket (\sigma)^n \rrbracket_B(S_0) \subseteq$

$\llbracket \bigoplus_{i=0}^t (\sigma)^i \rrbracket_B(S_0)$. Using Corollary 1 and the definition

of the semantics, $\llbracket (\sigma)^n \circ \sigma \rrbracket_B(S_0) \subseteq \llbracket (\bigoplus_{i=0}^t (\sigma)^i) \circ \sigma \rrbracket_B(S_0)$. Using the distributivity property P8,

$\llbracket (\sigma)^{n+1} \rrbracket_B(S_0) \subseteq \llbracket \bigoplus_{i=1}^{t+1} (\sigma)^i \rrbracket_B(S_0)$. By the induction

hypothesis, $\llbracket \bigoplus_{i=1}^{t+1} (\sigma)^i \rrbracket_B(S_0) \subseteq \llbracket \bigoplus_{i=0}^{t+1} (\sigma)^i \rrbracket_B(S_0) = S_t$.

Therefore, by transitivity of set inclusion, $\llbracket \sigma^{n+1} \rrbracket_B(S_0) \subseteq S_t$. From the induction hypothesis, $S_n = S_t$. Hence, $S_{n+1} = S_n \cup \llbracket (\sigma)^{n+1} \rrbracket_B(S_0) = S_t$. This proves that $S_r = S_t$ for all $r \geq t + 1$.

To show that $\llbracket *\sigma \rrbracket_B(S_0) = S_t$, consider a state $q \in \llbracket *\sigma \rrbracket_B(S_0)$. By the definition of the semantics, there exists an $r (\geq 0)$ such that $q \in \llbracket (\sigma)^r \rrbracket_B(S_0)$. In other words, $q \in S_r$. If $r < t$, we have $q \in S_t$ since $S_j \subseteq S_{j+1}$ for all $j \geq 0$. If $r \geq t$, we again have $q \in S_t$ since

$S_j = S_t$ for all $j \geq t$. Therefore, $q \in S_t$ in this case too. Hence, $\llbracket * \sigma \rrbracket_B(S_0) \sqsubseteq S_t$. Also, from the definition of $\llbracket * \sigma \rrbracket_B(S_0)$, we have $S_t \sqsubseteq \llbracket * \sigma \rrbracket_B(S_0)$. Therefore, $\llbracket * \sigma \rrbracket_B(S_0) = S_t$.

Lemma 4. *Let $B = (V, Q, Q_0, \mathcal{Y})$ be a finite state transition system, and let $\{T_1, \dots, T_k\}$ be a set of extended clusters of B such that $\bigcup_{i=1}^k T_i = \mathcal{Y}$. The set of reachable states of B , denoted as $\text{reach}(B)$, is given by $\llbracket *(\mathbf{T}_1 + \dots + \mathbf{T}_k) \rrbracket(Q_0)$.*

Proof. Let \mathbf{T} be the reachability expression corresponding to the cluster \mathcal{Y} . By the definition of the semantics, $\llbracket (\mathbf{T})^i \rrbracket(Q_0)$ gives the set of states reachable from some state in Q_0 in exactly i steps. Therefore $\text{reach}(B) = \bigcup_{i=0}^{\infty} \llbracket (\mathbf{T})^i \rrbracket(Q_0)$. By definition, the right-hand side of the above equality is $\llbracket * \mathbf{T} \rrbracket(Q_0)$. The lemma now follows from noting that $\mathbf{T} = (\mathbf{T}_1 + \dots + \mathbf{T}_k)$ by lemma 2.

Lemma 5. *Let $\sigma_1, \sigma_2, \sigma_3, \sigma_4$ be reachability expressions.*

- (a) *If $\sigma_1 \sqsubseteq \sigma_2$ and $\sigma_3 \sqsubseteq \sigma_4$, then $(\sigma_1 \text{ op } \sigma_3) \sqsubseteq (\sigma_2 \text{ op } \sigma_4)$, where $\text{op} \in \{+, \circ, ;\}$.*
- (b) *$(\sigma_1 ; \sigma_2)^i \sqsubseteq (\sigma_1 ; \sigma_2)^{i+1}$ for all $i \geq 0$.*
- (c) *If $\sigma_1 \sqsubseteq \sigma_2$, then $(\sigma_1)^i \sqsubseteq (\sigma_2)^i$ for all $i \geq 0$, and $(*\sigma_1) \sqsubseteq (*\sigma_2)$.*
- (d) *$(*\sigma_1)^i = (*\sigma_1)$ for all $i \geq 1$, and $*(*\sigma_1) = (*\sigma_1)$.*

Proof. (a) Since $\sigma_1 \sqsubseteq \sigma_2$ and $\sigma_3 \sqsubseteq \sigma_4$, for every set S of states, $\llbracket \sigma_1 \rrbracket(S) \subseteq \llbracket \sigma_2 \rrbracket(S)$ and $\llbracket \sigma_3 \rrbracket(S) \subseteq \llbracket \sigma_4 \rrbracket(S)$. Therefore, $\llbracket \sigma_1 \rrbracket(S) \cup \llbracket \sigma_3 \rrbracket(S) \subseteq \llbracket \sigma_2 \rrbracket(S) \cup \llbracket \sigma_4 \rrbracket(S)$. In other words, $(\sigma_1 + \sigma_3) \sqsubseteq (\sigma_2 + \sigma_4)$. Since $\sigma_1 \sqsubseteq \sigma_2$, therefore $\llbracket \sigma_1 \rrbracket(S) \subseteq \llbracket \sigma_2 \rrbracket(S)$. Applying Corollary 1, we get $\llbracket \sigma_3 \rrbracket(\llbracket \sigma_1 \rrbracket(S)) \subseteq \llbracket \sigma_3 \rrbracket(\llbracket \sigma_2 \rrbracket(S))$. However, $\sigma_3 \sqsubseteq \sigma_4$ implies that $\llbracket \sigma_3 \rrbracket(\llbracket \sigma_2 \rrbracket(S)) \subseteq \llbracket \sigma_4 \rrbracket(\llbracket \sigma_2 \rrbracket(S))$. From the transitivity of set inclusion, we get $\llbracket \sigma_3 \rrbracket(\llbracket \sigma_1 \rrbracket(S)) \subseteq \llbracket \sigma_4 \rrbracket(\llbracket \sigma_2 \rrbracket(S))$. Hence, $(\sigma_1 \circ \sigma_3) \sqsubseteq (\sigma_2 \circ \sigma_4)$. Since $(\sigma_i ; \sigma_j) = (\sigma_i + \delta) \circ (\sigma_j + \delta)$, it follows from the results above (for $\text{op} \in \{+, \circ\}$) that $(\sigma_1 ; \sigma_3) \sqsubseteq (\sigma_2 ; \sigma_4)$.

- (b) By definition, $(\sigma_1 ; \sigma_2)^{i+1} = (\sigma_1 ; \sigma_2) \circ (\sigma_1 ; \sigma_2)^i$ for all $i \geq 0$. From property P10 listed above, $\delta \sqsubseteq (\sigma_1 ; \sigma_2)$. Applying Lemma 5(a), we get $\delta \circ (\sigma_1 ; \sigma_2)^i \sqsubseteq (\sigma_1 ; \sigma_2)^{i+1}$. Using property P6, $\delta \circ (\sigma_1 ; \sigma_2)^i = (\sigma_1 ; \sigma_2)^i \sqsubseteq (\sigma_1 ; \sigma_2)^{i+1}$.

- (c) We show by induction on i that if $\sigma_1 \sqsubseteq \sigma_2$, then $(\sigma_1)^i \sqsubseteq (\sigma_2)^i$ for all $i \geq 0$.

Basis: If $i = 0$, $(\sigma_1)^i = (\sigma_2)^i = \delta$ by definition. Therefore, $(\sigma_1)^i \sqsubseteq (\sigma_2)^i$.

Hypothesis: Let $(\sigma_1)^i \sqsubseteq (\sigma_2)^i$ for $0 \leq i \leq k$.

Induction step: Since $\sigma_1 \sqsubseteq \sigma_2$ and $(\sigma_1)^k \sqsubseteq (\sigma_2)^k$ (from induction hypothesis), by applying Lemma 5(a), we get $(\sigma_1 \circ (\sigma_1)^k) \sqsubseteq (\sigma_2 \circ (\sigma_2)^k)$. Combining this with the definition of $(\sigma_j)^{k+1}$, we get $(\sigma_1)^{k+1} \sqsubseteq (\sigma_2)^{k+1}$. Therefore, $(\sigma_1)^i \sqsubseteq (\sigma_2)^i$ for all $i \geq 0$.

Using the above result, we get $\bigcup_{i=0}^{\infty} \llbracket (\sigma_1)^i \rrbracket(S) \subseteq \bigcup_{i=0}^{\infty} \llbracket (\sigma_2)^i \rrbracket(S)$. Hence, $\llbracket * \sigma_1 \rrbracket(S) \subseteq \llbracket * \sigma_2 \rrbracket(S)$ for all sets S of states. In other words, $(*\sigma_1) \sqsubseteq (*\sigma_2)$.

- (d) We first show that $(*\sigma_1)^2 = (*\sigma_1)$. Let S be an arbitrary set of states, and let $S_1 = \llbracket * \sigma_1 \rrbracket(S)$ and $S_2 = \llbracket (*\sigma_1)^2 \rrbracket(S)$. By definition, $S_2 = \llbracket * \sigma_1 \rrbracket(S_1)$. Thus, $S_2 = \bigcup_{i=0}^{\infty} \llbracket (\sigma_1)^i \rrbracket(S_1)$ and hence, $S_1 \subseteq S_2$. Recalling the definitions of S_1 and S_2 , we get $(*\sigma_1) \sqsubseteq (*\sigma_1)^2$.

To show that $(*\sigma_1)^2 \sqsubseteq (*\sigma_1)$, we consider an arbitrary state $q \in \llbracket (*\sigma_1)^2 \rrbracket(S)$. Since $\llbracket (*\sigma_1)^2 \rrbracket(S) = \llbracket * \sigma_1 \rrbracket(S_1)$, there exists a state $q' \in S_1$ and an integer $r \geq 0$ such that $q \in \llbracket (\sigma_1)^r \rrbracket(\{q'\})$. Since $q' \in S_1$ and $S_1 = \llbracket * \sigma_1 \rrbracket(S)$, there exists an integer $t \geq 0$ such that $q' \in \llbracket (\sigma_1)^t \rrbracket(S)$. Thus, $q \in \llbracket (\sigma_1)^r \rrbracket(\llbracket (\sigma_1)^t \rrbracket(S))$. In other words, $q \in \llbracket (\sigma_1)^{t+r} \rrbracket(S)$. Since “ \circ ” is associative, we can express the above as $q \in \llbracket (\sigma_1)^{t+r} \rrbracket(S)$. It follows that $q \in \llbracket * \sigma_1 \rrbracket(S)$. Therefore, $\llbracket (*\sigma_1)^2 \rrbracket(S) \subseteq \llbracket * \sigma_1 \rrbracket(S)$ and hence, $(*\sigma_1)^2 \sqsubseteq (*\sigma_1)$. Since we have also shown above that $(*\sigma_1) \sqsubseteq (*\sigma_1)^2$, we get $(*\sigma_1)^2 = (*\sigma_1)$.

We now show by induction on i that $(*\sigma_1)^i = (*\sigma_1)$ for all $i \geq 1$.

Basis: For $i = 1$, the result holds trivially.

Hypothesis: Let us assume that $(*\sigma_1)^i = (*\sigma_1)$ for $1 \leq i \leq k$.

Induction step: By definition, $(*\sigma_1)^{k+1} = (*\sigma_1) \circ (*\sigma_1)^k$. By the induction hypothesis, $(*\sigma_1)^k = (*\sigma_1)$. Therefore, $(*\sigma_1)^{k+1} = (*\sigma_1) \circ (*\sigma_1) = (*\sigma_1)^2$. However, we have shown above that $(*\sigma_1)^2 = (*\sigma_1)$. Therefore, $(*\sigma_1)^{k+1} = (*\sigma_1)$. We have thus shown by induction that $(*\sigma_1)^i = (*\sigma_1)$ for all $i \geq 1$.

To show that $*(*\sigma_1) = (*\sigma_1)$, we note that $\llbracket *(*\sigma_1) \rrbracket(S) = \bigcup_{j=0}^{\infty} \llbracket (*\sigma_1)^j \rrbracket(S)$. However, we have seen above that $(*\sigma_1)^j = (*\sigma_1)$ for all $j \geq 1$. Hence, $\bigcup_{j=0}^{\infty} \llbracket (*\sigma_1)^j \rrbracket(S) = S \cup \bigcup_{j=1}^{\infty} \llbracket * \sigma_1 \rrbracket(S)$. Since $S \subseteq \llbracket * \sigma_1 \rrbracket(S)$, the right hand side of the above equality reduces to $\llbracket * \sigma_1 \rrbracket(S)$. Therefore, $*(*\sigma_1) = (*\sigma_1)$.

Lemma 6. *Let σ be a reachability expression over a set, $\{T_1, \dots, T_k\}$, of extended clusters.*

*Then $\sigma \sqsubseteq *(\mathbf{T}_1 + \dots + \mathbf{T}_k)$.*

Proof. We use induction on the structure of σ to prove the lemma.

Basis: If $\sigma = \mathbf{T}_i$ for $i \in \{1, \dots, k\}$, or $\sigma = \delta$ or $\sigma = \Theta$, the \sqsubseteq relation holds trivially.

Hypothesis: Let σ_1 and σ_2 be reachability expressions over $\{T_1, \dots, T_k\}$, such that $\sigma_1 \sqsubseteq *(\mathbf{T}_1 + \dots + \mathbf{T}_k)$ and $\sigma_2 \sqsubseteq *(\mathbf{T}_1 + \dots + \mathbf{T}_k)$.

Induction Step: We consider the following cases:

1. $\sigma = \sigma_1 + \sigma_2$: The result follows from the induction hypothesis, Lemma 5(a) and the idempotency of the “+” operator.
2. $\sigma = \sigma_1 \circ \sigma_2$: The result follows from the induction hypothesis, Lemma 5(a) and Lemma 5(d).
3. $\sigma = \sigma_1 ; \sigma_2$: The result follows from the induction hypothesis, Lemma 5(a), the definition of the semantics of the “;” operator, and Lemma 5(d).
4. $\sigma = (\sigma_1)$: The result follows trivially from the induction hypothesis and the definition of $\llbracket (\sigma_1) \rrbracket$.

5. $\sigma = *\sigma_1$: The result follows from the induction hypothesis, Lemma 5(c) and Lemma 5(d).

This completes the induction, proving that all reachability expressions over $\{T_1, \dots, T_k\}$ are covered by $*(\mathbf{T}_1 + \dots + \mathbf{T}_k)$

Lemma 7. *For all reachability expressions σ_1 and σ_2 , $*(\sigma_1; \sigma_2) = *(\sigma_1; (*\sigma_2)) = *((*\sigma_1); \sigma_2) = *((*\sigma_1); (*\sigma_2))$.*

Proof. By definition of the semantics, $\sigma_2 \sqsubseteq (*\sigma_2)$. Therefore, by Lemma 5(a), $(\sigma_1; \sigma_2) \sqsubseteq (\sigma_1; (*\sigma_2))$, and by Lemma 5(c), $*(\sigma_1; \sigma_2) \sqsubseteq *((\sigma_1; (*\sigma_2)))$.

To show that $*(\sigma_1; (*\sigma_2)) \sqsubseteq *((\sigma_1; \sigma_2))$, consider a set S of states, and let $q \in \llbracket \sigma_1; (*\sigma_2) \rrbracket(S)$. Since $(\sigma_1; (*\sigma_2)) = (\delta + \sigma_1 + (*\sigma_2) + (\sigma_1 \circ (*\sigma_2)))$, we have the following cases:

- $q \in \llbracket \delta + \sigma_1 \rrbracket(S)$: Since $(\delta + \sigma_1) \sqsubseteq (\sigma_1; \sigma_2)$, $q \in \llbracket *(\sigma_1; \sigma_2) \rrbracket(S)$.
- $q \in \llbracket (*\sigma_2) \rrbracket(S)$: There exists an integer $r \geq 0$ such that $q \in \llbracket (\sigma_2)^r \rrbracket(S)$. Since $\sigma_2 \sqsubseteq (\sigma_1; \sigma_2)$, by Lemma 5(c), $(\sigma_2)^r \sqsubseteq (\sigma_1; \sigma_2)^r$. Hence, $q \in \llbracket *(\sigma_1; \sigma_2) \rrbracket(S)$.
- $q \in \llbracket \sigma_1 \circ (*\sigma_2) \rrbracket(S)$: There exists an integer $r \geq 0$ such that $q \in \llbracket \sigma_1 \circ (\sigma_2)^r \rrbracket(S)$. Since $(\sigma_2)^r \sqsubseteq (\sigma_1; \sigma_2)^r$ and $\sigma_1 \sqsubseteq (\sigma_1; \sigma_2)$, by Lemma 5(a), $(\sigma_1 \circ (\sigma_2)^r) \sqsubseteq (\sigma_1; \sigma_2)^{r+1}$. Therefore we get, $q \in \llbracket (\sigma_1; \sigma_2)^{r+1} \rrbracket(S)$ and hence, $q \in \llbracket *(\sigma_1; \sigma_2) \rrbracket(S)$.

We have thus shown above that $(\sigma_1; (*\sigma_2)) \sqsubseteq *((\sigma_1; \sigma_2))$. Applying Lemmas 5(c) and (d), we get $*(\sigma_1; (*\sigma_2)) \sqsubseteq *(\sigma_1; \sigma_2)$. Since we also have $*(\sigma_1; \sigma_2) \sqsubseteq *((\sigma_1; (*\sigma_2)))$, it follows that $*(\sigma_1; \sigma_2) = *((\sigma_1; (*\sigma_2)))$. The proof of $*(\sigma_1; \sigma_2) = *((*\sigma_1); \sigma_2)$ is similar with the roles of σ_1 and σ_2 interchanged.

If we substitute $*\sigma_1$ for σ_1 in the result proved above, we get $*((*\sigma_1); \sigma_2) = *((*\sigma_1); (*\sigma_2))$. However, $*((*\sigma_1); \sigma_2) = *(\sigma_1; \sigma_2)$, as argued above. Therefore, $*(\sigma_1; \sigma_2) = *((*\sigma_1); (*\sigma_2))$.

Theorem 1. *Let $\{\sigma_1, \dots, \sigma_k\}$, $k \geq 1$, be a finite set of reachability expressions. Then $*(\sigma_1 + \dots + \sigma_k) = *(\sigma_1; \dots; \sigma_k)$.*

Proof. We prove the theorem by induction on k .

Basis ($k = 1$): The result holds trivially.

Hypothesis: Assume the result holds for all k in 1 through m .

Induction step: Consider a set of $m + 1$ reachability expressions, and let $\sigma_Y = (\sigma_2 + \dots + \sigma_{m+1})$. We first show that $*(\sigma_1 + \sigma_Y) = *(\sigma_1; \sigma_Y)$.

Since $(\sigma_1 + \sigma_Y) \sqsubseteq (\sigma_1; \sigma_Y)$, by Lemma 5(c), $*(\sigma_1 + \sigma_Y) \sqsubseteq *(\sigma_1; \sigma_Y)$. However, $(\sigma_1; \sigma_Y) = \delta + \sigma_1 + \sigma_Y + (\sigma_1 \circ \sigma_Y) \sqsubseteq *(\sigma_1 + \sigma_Y)$. Therefore, by Lemmas 5(c) and (d), $*(\sigma_1; \sigma_Y) \sqsubseteq *(\sigma_1 + \sigma_Y)$. It follows that $*(\sigma_1 + \sigma_Y) = *(\sigma_1; \sigma_Y)$.

By Lemma 7, $*(\sigma_1; \sigma_Y) = *(\sigma_1; (*\sigma_Y))$. By the induction hypothesis, $(*\sigma_Y) = (*\sigma_Z)$, where $\sigma_Z = (\sigma_2; \dots;$

$\sigma_{m+1})$. Therefore, $*(\sigma_1 + \sigma_Y) = *(\sigma_1; (*\sigma_Z))$. Applying Lemma 7, we get $*(\sigma_1 + \sigma_Y) = *(\sigma_1; \sigma_Z)$. Since “+” and “;” are associative, the proof is complete.

Corollary 2. *The “;” operator is commutative within the scope of “*”.*

Proof. Follows from Theorem 1 and from the commutativity of “+”.

Corollary 3. *Let $B = (V, Q, Q_0, \mathcal{T})$ be a finite state transition system, and let $\{T_1, \dots, T_k\}$ be a set of extended clusters of B such that $\bigcup_{i=1}^k T_i = \mathcal{T}$. The set of reachable states of B is given by $\text{reach}(B) = \llbracket *(\mathbf{T}_1; \dots; \mathbf{T}_k) \rrbracket(Q_0)$.*

Proof. Follows from lemma 4 and Theorem 1.

Theorem 2. *Let $\{\sigma_1, \dots, \sigma_k\}$, $k \geq 2$, be a finite set of reachability expressions. Let $\sigma_Y = (\sigma_1 + \dots + \sigma_k)$ and $\sigma_Z = (\sigma_1; \dots; \sigma_k)$.*

- (a) *For all $n \geq 0$, $(\sum_{j=0}^n (\sigma_Y)^j) \sqsubseteq (\sigma_Z)^n$.*
- (b) *Furthermore, if $(\sigma_p)^2 \sqsubseteq \sigma_p$ for all p , then $(\sum_{j=0}^n (\sigma_Y)^j) \sqsubseteq (\sigma_Z)^{n - (\lfloor n/k \rfloor - 1)}$.*

Proof. (a) Since $\sigma_Y \sqsubseteq \sigma_Z$, by Lemma 5(c), $(\sigma_Y)^j \sqsubseteq (\sigma_Z)^j$ for all $j \geq 0$. Therefore, $(\sum_{j=0}^n (\sigma_Y)^j) \sqsubseteq (\sum_{j=0}^n (\sigma_Z)^j)$. Applying Lemma 5(b), it can now be shown that $(\sum_{j=0}^n (\sigma_Z)^j) = \sigma_Z^n$. Hence, $(\sum_{j=0}^n (\sigma_Y)^j) \sqsubseteq \sigma_Z^n$.

- (b) Let $n = k \cdot i + r$, where $i \geq 0$ and $0 \leq r < k$. We show by induction on k that $(\sum_{j=0}^{k \cdot i} (\sigma_Y)^j) \sqsubseteq (\sigma_Z)^{(k-1) \cdot i + 1}$ for $k \geq 2$.

Basis ($k = 2$): We first prove that $(\sum_{j=0}^{2i} (\sigma_1 + \sigma_2)^j) \sqsubseteq (\sigma_1; \sigma_2)^{i+1}$, for $i \geq 0$.

Consider a set S of states, and let $q \in \llbracket \sum_{j=0}^{2i} (\sigma_1 + \sigma_2)^j \rrbracket(S)$. Then, there exists a non-negative integer M , no greater than $2i$, such that $q \in \llbracket (\sigma_1 + \sigma_2)^M \rrbracket(S)$. If $M = 0$, $q \in \llbracket \delta \rrbracket(S)$. However, since $\delta \sqsubseteq (\sigma_1; \sigma_2)^{i+1}$, $q \in \llbracket (\sigma_1; \sigma_2)^{i+1} \rrbracket(S)$ as well. If $M > 0$, there exists a sequence of indices, (r_1, \dots, r_M) , such that each $r_j \in \{1, 2\}$ and $q \in \llbracket (\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) \rrbracket(S)$. We show below that $(\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) \sqsubseteq (\sigma_1 \circ \sigma_2)^{i+1}$.

Since $(\sigma_p)^2 \sqsubseteq \sigma_p$ for all $p \in \{1, 2\}$, it suffices to consider those sequences (r_1, \dots, r_M) in which all consecutive indices are distinct. However, since each $r_j \in \{1, 2\}$, indices 1 and 2 must alternate in (r_1, \dots, r_M) . We have the following cases:

- $M = 2q$, $q > 0$, $r_1 = 1$: $(\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) = (\sigma_1 \circ \sigma_2)^q \sqsubseteq (\sigma_1; \sigma_2)^q$.
- $M = 2q$, $q > 0$, $r_1 = 2$: $(\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) = (\sigma_2 \circ (\sigma_1 \circ \sigma_2)^{q-1} \circ \sigma_1) \sqsubseteq (\sigma_1; \sigma_2)^{q+1}$.

- $M = 2q + 1, q > 0, r_1 = 1: (\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) = ((\sigma_1 \circ \sigma_2)^q \circ \sigma_1) \subseteq (\sigma_1; \sigma_2)^{q+1}$.
- $M = 2q + 1, q > 0, r_1 = 2: (\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) = (\sigma_2 \circ (\sigma_1 \circ \sigma_2)^q) \subseteq (\sigma_1; \sigma_2)^{q+1}$.

Thus, in all cases, $(\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) \subseteq (\sigma_1; \sigma_2)^{\lfloor M/2 \rfloor + 1}$. Since M is bounded above by $2i$, $\lfloor M/2 \rfloor + 1 \leq i + 1$. Therefore, by Lemma 5(b) $(\sigma_{r_1} \circ \dots \circ \sigma_{r_M}) \subseteq (\sigma_1 \circ \sigma_2)^{i+1}$.

Hypothesis: Let $(\bigoplus_{j=0}^{k,i} (\sigma_Y)^j) \subseteq (\sigma_Z)^{(k-1),i+1}$ for $2 \leq k \leq m$ and $i \geq 0$.

Induction Step: Let $\sigma_W = \sigma_m + \sigma_{m+1}$. Then, $\sigma_Y = (\sigma_1 + \dots + \sigma_{m+1})$ can be rewritten as $(\sigma_1 + \dots + \sigma_{m-1} + \sigma_W)$. Note that the last expression is obtained by applying the “+” operator to m reachability expressions. Since “+” is idempotent and “o” distributes over “+”, $(\bigoplus_{j=0}^{(m+1),i} (\sigma_Y)^j) = (\bigoplus_{j=0}^{m,i} (\sigma_Y)^j) \circ (\bigoplus_{l=0}^i (\sigma_Y)^l)$. By the induction hypothesis, $(\bigoplus_{j=0}^{m,i} (\sigma_Y)^j) \subseteq (\sigma_1; \dots; \sigma_{m-1}; \sigma_W)^{(m-1),i+1}$. Since $\sigma_W = (\sigma_m + \sigma_{m+1}) \subseteq (\sigma_m; \sigma_{m+1})$, by Lemma 5(a) and 5(c), we get $(\sigma_1; \dots; \sigma_{m-1}; \sigma_W)^{(m-1),i+1} \subseteq (\sigma_1; \dots; \sigma_{m-1}; \sigma_m; \sigma_{m+1})^{(m-1),i+1} = \sigma_Z^{(m-1),i+1}$. Hence, by transitivity, $(\bigoplus_{j=0}^{m,i} (\sigma_Y)^j) \subseteq \sigma_Z^{(m-1),i+1}$.

We have also shown in Theorem 2(a) that $(\bigoplus_{l=0}^i (\sigma_Y)^l) \subseteq (\sigma_Z)^i$. By Lemma 5(a), we then get $(\bigoplus_{j=0}^{m,i} (\sigma_Y)^j) \circ (\bigoplus_{l=0}^i (\sigma_Y)^l) \subseteq (\sigma_Z^{(m-1),i+1} \circ \sigma_Z^i)$. In other words, $(\bigoplus_{j=0}^{(m+1),i} (\sigma_Y)^j) \subseteq \sigma_Z^{((m+1)-1),i+1}$. Hence, by induction, $(\bigoplus_{j=0}^{k,i} (\sigma_Y)^j) \subseteq \sigma_Z^{(k-1),i+1}$ for all $k \geq 2$.

Since $(\bigoplus_{j=0}^r (\sigma_Y)^j) \subseteq \sigma_Z^r$ (by Theorem 2(a)), applying Lemma 5(a), we get $(\bigoplus_{j=0}^{k,i+r} (\sigma_Y)^j) = (\bigoplus_{j=0}^{k,i} (\sigma_Y)^j) \circ (\bigoplus_{j=0}^r (\sigma_Y)^j) \subseteq \sigma_Z^{(k-1),i+1} \circ \sigma_Z^r = \sigma_Z^{(k,i+r-(i-1))}$.

The theorem is now proved by noting that $k \cdot i + r = n$ and $i = \lfloor n/k \rfloor$.

Let $B = (V, Q, Q_0, \Upsilon)$ be a finite state transition system. Let $\{\tau_1, \dots, \tau_k\}$ be a set of extended clusters, satisfying $\bigcup_{i=1}^k \tau_i = \Upsilon$. From Lemma 4 and Corollary 3, the set of reachable states of B , denoted $reach(B)$, is given by $\llbracket *(\mathbf{T}_1 + \dots + \mathbf{T}_k) \rrbracket(Q_0)$ and $\llbracket *(\mathbf{T}_1; \dots; \mathbf{T}_k) \rrbracket(Q_0)$. Furthermore, Theorem 2(a) guarantees that the number of image computation iterations using $*(\mathbf{T}_1; \dots; \mathbf{T}_k)$ never exceeds that required with $*(\mathbf{T}_1 + \dots + \mathbf{T}_k)$. Therefore, if computing $\llbracket (\sigma_Y)^{i+1} \rrbracket(S)$ and computing $\llbracket (\sigma_Z)^{i+1} \rrbracket(S)$ (using the terminology of Theorem 2) are of comparable complexity, it is advantageous to use $*(\mathbf{T}_1; \dots; \mathbf{T}_k)$. This advantage is also demonstrated by our experiments, as reported in Section 3. If $(\sigma_p)^2 \subseteq \sigma_p$ for all p , Theorem 2(b) improves the upper bound

of Theorem 2(a) even further. Note that we can have $(\sigma_p)^2 \subseteq \sigma_p$ under several circumstances, e.g. if σ_p is of the form $*\sigma_q$.

Theorem 3. Let $B = (V, Q, Q_0, \Upsilon)$ be a finite state transition system with extended clusters $\{\tau_1, \dots, \tau_k\}$, such that $\bigcup_{i=1}^k \tau_i = \Upsilon$ and $\mathbf{T}_k \not\subseteq *(\mathbf{T}_1 + \dots + \mathbf{T}_{k-1})$. Let σ_X denote $(\mathbf{T}_1 + \dots + \mathbf{T}_{k-1})$, and $\hat{\sigma}$ denote $(*\sigma_X) \circ *(\mathbf{T}_k ; (*\sigma_X))$.

- (a) $\llbracket \hat{\sigma} \rrbracket(Q_0) = reach(B)$.
- (b) Let σ be any reachability expression over $\{\tau_1, \dots, \tau_k\}$ such that $\llbracket \sigma \rrbracket(Q_0) = reach(B)$. Let $N_k(\sigma, Q_0)$ denote the number of times image under \mathbf{T}_k is computed until the complete set of states reachable from Q_0 is obtained during evaluation of $\llbracket \sigma \rrbracket(Q_0)$. Then $N_k(\hat{\sigma}, Q_0) \leq N_k(\sigma, Q_0) + 1$.

Proof. (a) From Lemma 7 and Theorem 1, $*(\mathbf{T}_k; (*\sigma_X)) = *(\mathbf{T}_k ; \sigma_X) = *(\mathbf{T}_k + \sigma_X)$. Since $\delta \subseteq *\sigma_X$, composing both sides with $*(\mathbf{T}_k + \sigma_X)$ (or, equivalently with $*(\mathbf{T}_k; (*\sigma_X))$) and by applying Lemma 5(a), we get $*(\mathbf{T}_k + \sigma_X) \subseteq (*\sigma_X) \circ *(\mathbf{T}_k ; (*\sigma_X))$. Therefore, $*(\mathbf{T}_k + \sigma_X) \subseteq \hat{\sigma}$. However, $\hat{\sigma} \subseteq *(\mathbf{T}_k + \sigma_X)$. Hence, $\hat{\sigma} = *(\mathbf{T}_k + \sigma_X)$. Since $\llbracket *(\mathbf{T}_k + \sigma_X) \rrbracket(Q_0) = reach(B)$, it follows that that $\llbracket \hat{\sigma} \rrbracket(Q_0) = reach(B)$.

- (b) Since $\hat{\sigma} = (*\sigma_X) ; *(\mathbf{T}_k ; (*\sigma_X))$, evaluation of $\llbracket \hat{\sigma} \rrbracket(Q_0)$ involves execution of the following steps:
 - (i) A variable W (of type “set of states”) is initialized to $\llbracket *\sigma_X \rrbracket(Q_0)$.
 - (ii) If no new states have been added to W in the previous step, we terminate and return W as the value of $\llbracket \hat{\sigma} \rrbracket(Q_0)$.
 - (iii) Else, W is updated to the union of W and $\llbracket \mathbf{T}_k ; (*\sigma_X) \rrbracket(W)$, and we go back to step(ii).

Since \mathbf{T}_k is evaluated once during each execution of step(iii) above, $N_k(\hat{\sigma}, Q_0)$ equals the number of executions of step(iii). Let σ be an arbitrary reachability expression over $\{T_1, \dots, T_k\}$ such that $\llbracket \sigma \rrbracket(Q_0) = reach(B)$. Lemma 6 implies all reachability expressions not involving \mathbf{T}_k , i.e. over $\{T_1, \dots, T_{k-1}\}$, are covered by $*\sigma_X$. Therefore, if M denotes the value of $N_k(\sigma, Q_0)$, we have $\llbracket \sigma \rrbracket(Q_0) \subseteq \llbracket *\sigma_X \rrbracket(Q_0) \cup \llbracket *\sigma_X \circ \mathbf{T}_k \circ *\sigma_X \rrbracket(Q_0) \cup \llbracket *\sigma_X \circ (\mathbf{T}_k \circ *\sigma_X)^2 \rrbracket(Q_0) \cup \dots \cup \llbracket *\sigma_X \circ (\mathbf{T}_k \circ *\sigma_X)^M \rrbracket(Q_0)$. Since $reach(B) = \llbracket \sigma \rrbracket(Q_0)$ (our premise), it follows that $reach(B) = \llbracket (*\sigma_X) \circ (\bigoplus_{j=0}^M (\mathbf{T}_k \circ (*\sigma_X))^j) \rrbracket(Q_0)$. Since $(\sigma_i \circ \sigma_j) \subseteq (\sigma_i; \sigma_j)$ for all σ_i and σ_j , we have $reach(B) \subseteq \llbracket (*\sigma_X) \circ (\bigoplus_{j=0}^M (\mathbf{T}_k; (*\sigma_X))^j) \rrbracket(Q_0)$. However, from

Theorem 3(a), we know that $\llbracket (*\sigma_X) \circ (\bigoplus_{j=0}^M (\mathbf{T}_k; (*\sigma_X))^j) \rrbracket(Q_0) \subseteq reach(B)$. It therefore follows that $reach(B) = \llbracket (*\sigma_X) \circ (\bigoplus_{j=0}^M (\mathbf{T}_k; (*\sigma_X))^j) \rrbracket(Q_0)$. This implies that when computing $\llbracket \hat{\sigma} \rrbracket(Q_0)$, the total number of iterations of step(iii) of the procedure outlined

above is no larger than $M + 1$. Hence, $N_k(\hat{\sigma}, Q_0) \leq N_k(\sigma, Q_0) + 1$.

Given a finite-state system, Theorem 3 gives us a reachability expression that guarantees that the number of image computations under τ_k is at worst 1 more than the minimum number needed to compute the reachable state space using any reachability expression. This is particularly useful when we have clusters with disparate image computation costs. For example, when performing reachability analysis of a network of timed automata, the discrete (or non-time-elapse) transitions of individual automata might be represented by τ_1 through τ_{k-1} , while a combined time-elapse transition for all automata might be represented by τ_k . Since clocks of all automata change synchronously, computing the image under τ_k requires synchronizing all the processes and updating the clocks of *all* automata, unlike computing the image under τ_1 through τ_{k-1} . Consequently, image computation under τ_k is expected to be more expensive (in terms of memory usage and CPU time) in general compared to image computation under τ_1 through τ_{k-1} . In such cases, it may be advantageous to minimize the number of expensive image computations by application of Theorem 3(b).

Often the set of extended clusters in a state transition system are related in such a way that starting from an initial set of states S_0 , if we compute the image under a cluster σ_j , no *new* states are reached unless the image under another cluster σ_i has already been computed. As an illustration, consider a combinational circuit in which the behaviour of each gate is modeled as a finite state transition system. Suppose the circuit contains a single-input gate g_1 that is fed by another gate g_2 . Suppose further that the circuit starts from a stable internal state (i.e., the output of no gate is scheduled to change). The inputs of the circuit are then changed after some delay, leading to a new state. If we compute the image of this new state under a cluster modeling the behaviour of g_1 , the set of reachable states cannot change unless the image under clusters corresponding to g_2 has already been computed. The following theorem shows that such dependencies can be exploited to simplify reachability expressions.

Theorem 4. *Let $\{\sigma_1, \dots, \sigma_k\}$ be a set of extended clusters and S be a set of states satisfying the following conditions:*

C1: $(\sigma_i \circ ((\ast\sigma_{i+1}) \circ \dots \circ (\ast\sigma_k))) \circ \sigma_i \sqsubseteq (\sigma_i \circ ((\ast\sigma_i) \circ \dots \circ (\ast\sigma_k)))$ for all $1 \leq i < k$.

C2: There exists m , $1 \leq m \leq k$ such that

C21: $(\sigma_i \circ \sigma_j) \sqsubseteq (\sigma_i + \sigma_j)$ for $1 \leq i, j \leq m$ and $i \neq j$.

C22: $\llbracket \ast\sigma_i \rrbracket(S) = S$, for all $i > m$.

Then $\llbracket \ast(\sigma_1; \dots; \sigma_k) \rrbracket(S) = \llbracket (\ast\sigma_1); \dots; (\ast\sigma_k) \rrbracket(S)$.

Proof. Let $\alpha_i = ((\ast\sigma_i) \circ \dots \circ (\ast\sigma_k))$ and $\beta_j = ((\ast\sigma_j) \circ \alpha_{j+1} \circ (\ast\sigma_j))$ for $1 \leq i \leq k$ and $1 \leq j < k$.

From condition C1 of Theorem 4, it is easy to show by induction on r and t that $((\sigma_i)^r \circ (\alpha_{i+1}) \circ (\sigma_i)^t) \sqsubseteq (\sigma_i \circ \alpha_i)$ for all $r \geq 1$ and $t \geq 0$. Therefore, $(\sigma_i \circ \beta_i) \sqsubseteq (\sigma_i \circ \alpha_i)$.

From the definition of the semantics, $(\ast\sigma_i); (\ast\sigma_j) = (\ast\sigma_i) \circ (\ast\sigma_j)$ for all $i, j \in \{1 \dots k\}$. Therefore, by Lemma 7, $\llbracket \ast(\sigma_1; \dots \sigma_k) \rrbracket(S) = \llbracket \ast((\ast\sigma_1); \dots (\ast\sigma_k)) \rrbracket(S) = \llbracket \ast\alpha_1 \rrbracket(S)$. In the following, we show by induction on r that $\llbracket (\alpha_1)^r \rrbracket(S) \subseteq \llbracket \alpha_1 \rrbracket(S)$ for all $r \geq 0$.

Basis ($r=0$): Follows from the definition of the semantics.

Hypothesis: Let $\llbracket (\alpha_1)^r \rrbracket(S) \subseteq \llbracket \alpha_1 \rrbracket(S)$ for $0 \leq r \leq m$.

Induction: By definition, $\llbracket (\alpha_1)^{m+1} \rrbracket(S) = \llbracket \alpha_1 \rrbracket(\llbracket (\alpha_1)^m \rrbracket(S))$. By the induction hypothesis and by Lemma 5(a), this is covered by $\llbracket (\alpha_1)^2 \rrbracket(S)$. To complete the induction, we show below that $\llbracket (\alpha_1)^2 \rrbracket(S) \subseteq \llbracket \alpha_1 \rrbracket(S)$.

By definition, $(\alpha_i)^2 = (\beta_i \circ \alpha_{i+1})$ for $1 \leq i < k$. We claim that $\llbracket \beta_i \rrbracket(S) \subseteq \llbracket \alpha_i \rrbracket(S)$. To see why this is so, we first note that $\llbracket \beta_i \rrbracket(S) = \llbracket \sigma_i \circ \beta_i \rrbracket(S) \cup \llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S)$. It has already been shown above that $(\sigma_i \circ \beta_i) \sqsubseteq (\sigma_i \circ \alpha_i) \sqsubseteq \alpha_i$. To see if $\llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S) \subseteq \llbracket \alpha_i \rrbracket(S)$, we consider the following cases:

1. $i > m$: Since $\llbracket \ast\sigma_i \rrbracket(S) = S$ for all $i > m$ (condition C22 in Theorem 4), $\llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S) = S$. Since $S \subseteq \llbracket \alpha_i \rrbracket(S)$, $\llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S) \subseteq \llbracket \alpha_i \rrbracket(S)$.
2. $i < m$: Since both i and $i + 1$ are no larger than m , by condition C21 in Theorem 4, $\llbracket \sigma_i \circ \sigma_{i+1} \rrbracket \sqsubseteq \llbracket \sigma_i + \sigma_{i+1} \rrbracket$. Therefore, $\llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S) \subseteq \llbracket \sigma_i \circ \beta_i \rrbracket(S)$. As has been shown above, this is contained in $\llbracket \alpha_i \rrbracket(S)$.
3. $i = m$: Since $\llbracket \ast\sigma_i \rrbracket(S) = S$ for all $i > m$ (condition C22 in Theorem 4), $\llbracket \alpha_{i+1} \rrbracket(S) = S$. Therefore, $\llbracket \alpha_{i+1} \circ (\ast\sigma_i) \rrbracket(S) = \llbracket \ast\sigma_i \rrbracket(S) \subseteq \llbracket \alpha_i \rrbracket(S)$.

Therefore, $\llbracket \beta_i \rrbracket(S) \subseteq \llbracket \alpha_i \rrbracket(S)$ and hence, $\llbracket (\alpha_i)^2 \rrbracket(S) = \llbracket (\beta_i \circ \alpha_{i+1}) \rrbracket(S) \subseteq \llbracket \alpha_i \circ \alpha_{i+1} \rrbracket(S)$. The last expression can be rewritten as $\llbracket (\ast\sigma_i) \circ (\alpha_{i+1})^2 \rrbracket(S)$. Therefore, $\llbracket (\alpha_i)^2 \rrbracket(S) \subseteq \llbracket (\ast\sigma_i) \circ (\alpha_{i+1})^2 \rrbracket(S)$ for all $1 \leq i < k$. By applying this result recursively, we obtain $\llbracket (\alpha_i)^2 \rrbracket(S) \subseteq \llbracket (\ast\sigma_i) \circ \dots \circ (\ast\sigma_{k-1}) \circ (\ast\sigma_k)^2 \rrbracket(S) = \llbracket \alpha_i \rrbracket(S)$.

The induction is completed by setting $i = 1$ in the above result. Since $\llbracket (\alpha_1)^r \rrbracket(S) \subseteq \llbracket \alpha_1 \rrbracket(S)$ for all $r \geq 0$, it follows that $\llbracket \ast(\sigma_1; \dots \sigma_k) \rrbracket(S) = \llbracket \ast\alpha_1 \rrbracket(S) \subseteq \llbracket \alpha_1 \rrbracket(S)$. The theorem is now proved by noting that by definition, $\llbracket \alpha_1 \rrbracket(S) \subseteq \llbracket \ast(\sigma_1; \dots \sigma_k) \rrbracket(S)$. Therefore, $\llbracket \alpha_1 \rrbracket(S) = \llbracket \ast(\sigma_1; \dots \sigma_k) \rrbracket(S)$.

Condition C1 in Theorem 4 formalizes an ordering of dependencies between the σ_i 's. Effectively, C1 states that the effect of computing the image under expressions $\{\sigma_{i+1}, \dots, \sigma_k\}$ does not affect the computation of image under σ_i for all $1 \leq i < k$. Condition C21 asserts that the first few expressions in the above ordering do not depend on any other expressions. Hence, computing the image under the composition of two such expressions gives the same result as computing the image under the expressions individually and then taking their union. Condition




Fig. 1. Fischer's mutual exclusion protocol

C22 states an additional ordering requirement: unless the image of S under one of $\{\sigma_1, \dots, \sigma_m\}$ is computed, the reachability expressions $\{\sigma_{m+1}, \dots, \sigma_k\}$ do not result in any new states being reached. If all three conditions are satisfied, Theorem 4 permits a simplification in the computation of reachable states. In particular, it allows us to obtain the entire set of reachable states by computing the reachable states under each σ_i only once.

3 Experimental Results and Their Analysis

In the previous section, we presented theorems on reachability expressions which embody heuristic strategies for improving the efficiency of symbolic search. These include strategies such as replacing symbolic breadth-first search by round-robin search, and minimizing the number of applications of costly transitions. In order to evaluate the effectiveness of these heuristics, we have implemented an interpreter for reachability expressions in a tool called NuSMV-DP. Our tool acts as a wrapper on top of the reachability analysis engine of NuSMV [9]. It takes as inputs: (a) a description of a finite state transition system as a collection of named clusters, (b) a reachability expression, and (c) an initial set of states. Our tool explores the reachable state space according to the reachability expression and reports performance statistics on termination of the search.

Brief overview of example suite: We have used two classes of examples for our experiments – Fischer protocol and gate-level circuits with delays. Our choice of examples is motivated by their popularity in the domain of timed system analysis, and also by the ease of scaling their sizes.

Fischer protocol: This is a distributed timed protocol used to ensure mutual exclusion when a number of processes access a shared resource. Each process P_i is modeled as a timed automaton, as shown in Figure 1, where x_i is the clock of P_i and k is a shared variable for communication between processes. In Figure 1, a and b are integer constants that bound the time spent by each process in the “Assign” and “Wait” states. For an n -process Fischer protocol, a network of timed automata is obtained by asynchronous parallel composition of n automata. De-

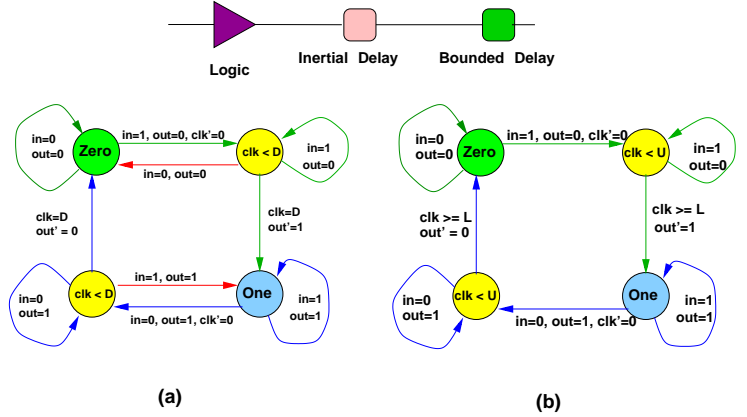


Fig. 2. (a) Inertial delay model (b) Bi-bounded pure delay model

tails of the model can be found in our technical report [14].

A natural clustering for an n -process Fischer protocol is to have one cluster per process, containing all *discrete* or non-time-elapse guarded actions of the process. Additionally, we must have one cluster containing the guarded action representing the synchronous advancement of time for all processes.

Circuits with inertial and bi-bounded delays: Our second set of examples consists of gate-level circuits. Each gate is modeled as consisting of three parts:

- A *boolean logic block* that gives the boolean value of the output as a function of the boolean values of the inputs.
- The output of the logic block is fed to an *inertial delay element* modeled as in Figure 2(a). If the inertial delay is D , the output of this element changes only if a change in its input persists for at least D units of time.
- The output of the inertial delay element is fed to a *bi-bounded pure delay element* which is modeled as shown in Figure 2(b). If the lower and upper bounds associated with this element are l and u respectively, this element delays each transition on its input by a non-deterministic delay between l and u units.

Given an interconnection of gates representing a circuit, we compose the state transition behaviour of the logic block, inertial delay element and bi-bounded delay element of each gate to form a network of timed automata. To simplify the model, we assume that D , l and u are identical for all gates. To ensure that every pure delay element causes its output to change once between two consecutive changes at its input, we also assume that $u < D$. When the output of a gate feeds the input of another gate, we ensure during composition that the corresponding output and input transitions occur simultaneously. Time is assumed to flow synchronously for all gates.

A natural clustering for an n -gate circuit modeled as above is to have a cluster for the discrete (non-time-elapse) transitions of each logic function, inertial and bi-bounded delay element, and an additional cluster for the synchronous advancement of time for all clocks. When the output of a gate feeds the input of another gate, we must combine the corresponding guarded actions at the output and input. For our experiments, the circuit inputs are modeled as signals that non-deterministically change their boolean values after a predefined delay Δ_{in} . A set of circuit layouts were chosen and the experiments were run with different delay parameters on this set. The details of circuits used in our experiments are given in [13]. Additional information regarding modeling of circuits as timed automata can be found in the technical report [14].

For both classes of examples, we assume that time is discrete, and model the timed behaviour using bounded-counter automata.

Performance comparisons: For the examples given above, let $\{\tau_1, \dots, \tau_k\}$ be the set of extended clusters representing non-time-elapse transitions, and let τ_t be the cluster representing synchronous advancement of time. Let Γ be the monolithic transition relation obtained by combining all transitions into a single cluster. Then, the reachability expression $S_0 = *\Gamma$ mimics symbolic breadth-first search using this monolithic transition relation, as in the original NuSMV tool. Γ can also be disjointively partitioned into its component clusters and the image computed using the reachability expression $S_1 = *(T_1 + \dots + T_k + T_t)$. While this reduces the effort for each image computation, the number of image computations increases significantly. To control this, we apply Theorem 2, and consider the reachability expression $S_2 = *(T_1; \dots; T_k; T_t)$ instead.

We have experimentally profiled the performance of reachability analysis using the expressions S_0 , S_1 and S_2 . All our experiments were run on a 3 GHz Intel Pentium 686 processor with 1 GB of main memory, and running Fedora Core Linux 3.4.3-6.fc3.

For the Fischer protocol examples, we computed the set of backward reachable states starting from a set of states in which mutual exclusion is violated. For simplicity, the parameters a and b were set to 1 and 2 respectively, for all processes. The results are shown as bar graphs in Figures 3 and 4. The total number of image computation iterations needed to compute the reachable states using S_0 , S_1 and S_2 respectively are shown as triples within parentheses along the abscissa in Figure 3. The missing data corresponds to experiments that did not terminate in 30 minutes. For the circuit examples, we computed the set of forward reachable states starting from a given set of initial states. The results are shown as bar graphs in Figures 5 and 6. In these figures, bar graphs corresponding to experiments on the same circuit but with different values of the parameters l, u, D




Fig. 3. Analysis of Fischer processes: Time and iteration counts for S_0 , S_1 , S_2 .

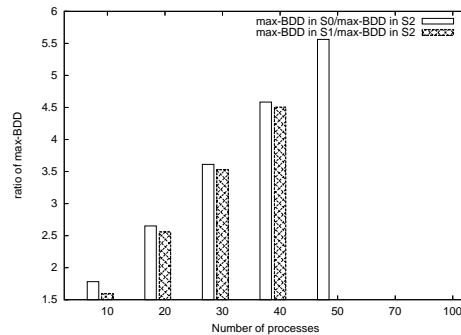


Fig. 4. Analysis of Fischer processes: Ratios of “max BDD”.

and Δ_{in} have been grouped together. The total number of image computation iterations needed to compute the reachable states using S_1 and S_2 respectively are shown as comma-separated pairs along the abscissa in Figure 5. For each circuit with r different combinations of l, u, D and Δ_{in} (r ranges from 2 to 4 in our experiments), there are r sets of bar graphs and r lines of comma-separated pairs above the circuit’s name. The i^{th} pair from the top and the i^{th} set of bar graphs from the left represent data obtained with the same set of parameters for a given circuit. The number of iterations using S_0 and S_1 were identical for all our circuit experiments. Details of the parameter values used for each circuit are available in the tables listed in the technical report [13].

In Figures 3 and 5, “Time (s)” denotes the time in seconds to compute the reachable state space. In Figures 4 and 6, “max BDD” denotes the maximum number of BDD nodes required to store the (partially computed) state space at any time during the state space search.

It can be seen that unguided disjunctive partitioning of the transition relation, as in S_1 , results in worse performance than reachability search using a monolithic transition relation. In the absence of guidance, disjunctive partitioning is, therefore, not an effective strategy. Theorem 2 guarantees that S_2 requires no more iterations of image computation than S_1 . This is clearly seen in the iteration counts in Figures 3 and 5. A reduction




Fig. 5. Analysis of circuits: Time and iteration counts for S_1, S_2 . Iteration counts for S_0 are identical to those for S_1 .




Fig. 6. Analysis of circuits: Ratios of “max BDD”.

in CPU time is expected from the combined effect of fewer iterations and operations on smaller BDDs. In the Fischer protocol examples, S_2 has the best performance. The only exception is the example with 10 processes. Here, the BDDs are small even when using S_0 , and no significant gains are obtained by decomposing the transition relation. Instead, iterating through the clusters incurs time overhead when using S_2 . For circuits, the ratios in Figure 6 are always greater than 1. Thus, S_2 results in the minimum “max BDD” value. However, as seen in Figure 5, circuits 6, 7 and 8 show better performance using S_0 with respect to time. These circuits were found to have very low “max BDD” values compared to the largest transition cluster size. This is in contrast to circuits 2, 3 and 4 where this ratio was much higher. Therefore, unlike in circuits 2, 3 and 4, BDD sizes of partially computed state sets do not significantly influence the performance of reachability analysis in circuits 6, 7 and 8. Since the largest transition cluster size is large compared to “max BDD”, reducing the total number of image computation iterations gives better performance. Thus S_0 performs better than S_2 for circuits 6, 7 and 8.

The BDD representation of the cluster “ τ_t ” is usually larger than that of other τ_i ’s since the transitions




Fig. 7. Analysis of circuits: Time for schedules S_3 and S_4 .

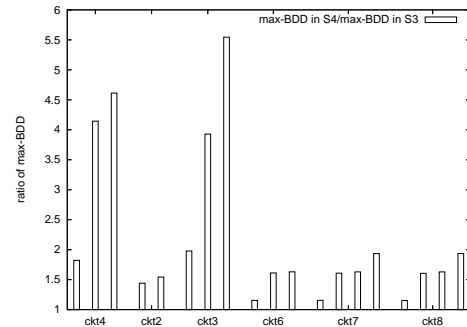


Fig. 8. Analysis of circuits: Ratios of “max BDD” using S_4 to that using S_3 .




Fig. 9. Analysis of circuits: Time for schedules S_3, S_5 and S_6 .

in τ_t involve clock variables of *all* processes. These large BDDs, in turn, lead to higher costs for computing the image under τ_t vis-a-vis the cost of computing the image under a τ_i . We have seen above that Theorem 3 gives us a way to reduce the number of costly image computations, potentially leading to performance improvements. To validate this experimentally, we measured and recorded the performance of computing the reachable state set using a reachability expression obtained from Theorem 3. In the Fischer protocol examples, the size of the BDD representation of τ_t is comparable to that of the

other τ_i clusters. Hence, the effect of minimizing applications of τ_t does not produce a significant performance difference for the Fischer examples, and we report results for only the circuits. Let $\sigma_x = (\mathbf{T}_1; \dots; \mathbf{T}_k)$ and let $S_3 = (*\sigma_x) \circ *(\mathbf{T}_t; (*\sigma_x))$. As seen in Theorem 3, S_3 minimizes the number of image computations under τ_t (up to 1 additional computation). We now compare its performance with an equivalent schedule S_4 , which is defined as: $S_4 = *(\mathbf{T}_1; \dots; \mathbf{T}_k; *\mathbf{T}_t)$.

Figure 8 shows the ratios of “max BDD” using S_4 to that using S_3 for the circuit examples. For each circuit, we used different sets of delay parameters, as in the earlier experiments. Within the set of experiments for each circuit, the number of image computations under τ_t until all reachable states are computed, increases from left to right for both S_4 and S_3 . As the number of computations increases, the effect of minimizing applications of τ_t becomes more pronounced, as can be seen from the rising ratio of “max BDD” using S_4 to that using S_3 .

Theorem 4 relates to the effect of applying a sequence of clusters consistently with the topological dependencies between them. To evaluate the effectiveness of this Theorem, we performed an additional set of experiments with the circuit examples. In these examples, one can obtain a topological ordering of the gates and circuit elements from the inputs to the outputs. When computing this order, a sub-circuit with a loop must be considered as a single circuit element without exposing the loop. Since the output of each gate/circuit element is fed to a gate or element with a higher topological index, there is an ordering of dependencies between the non-time-elapse clusters τ_i . By choosing the initial state S to be such that all gates are stable (i.e. no gate is scheduled to change its output), we ensure that condition C22 required in Theorem 4 is satisfied.

Let the input-to-output topological ordering of the non-time-elapse clusters be $\tau_1 < \dots < \tau_k$. Let $\sigma_x = (\mathbf{T}_1; \dots; \mathbf{T}_k)$ as before, and let $\sigma_y = (\mathbf{T}_k; \dots; \mathbf{T}_1)$ compute images of the clusters in reverse topological order. Moreover, let $\sigma_z = (*\mathbf{T}_1; \dots; *\mathbf{T}_k)$. Then, by Theorem 1, we have $*\sigma_x = *\sigma_y$ and by Theorem 4, we have $*\sigma_x = \sigma_z$. Earlier, we have considered the reachability expression $S_3 = (*\sigma_x) \circ *(\mathbf{T}_t; (*\sigma_x))$ that minimizes (up to 1 additional computation) the number of image computations under time transition cluster τ_t . We now consider the reachability expressions $S_5 = (*\sigma_y) \circ *(\mathbf{T}_t; (*\sigma_y))$ and $S_6 = \sigma_z \circ *(\mathbf{T}_t; \sigma_z)$. Using the above mentioned identities, it is easy to see that $S_3 = S_5 = S_6$, i.e. all three reachability expressions compute the same set of reachable states.

Using S_3 , S_5 , and S_6 in the circuit examples, “max BDD” was found to be nearly identical. However, the CPU times were different because of additional fix-point checks in S_5 and S_3 compared to those in S_6 . In Figure 9, the circuits are arranged from left to right in order of increasing topological depth. The increased number of fix-point checks due to an increase in the number of ordered

clusters results in the increased time difference between schedules S_5 and S_3 from circuit 4 to 8 as seen in Figure 9. For circuits with short topological depth, the performance of S_5 and S_3 are similar. However, as the topological depth increases, computing images in topologically sorted order leads to significant improvements compared to computing in reverse topological order. Furthermore, S_6 improves over S_3 , albeit marginally, in most cases. This is because all clusters τ_i other than the time transition cluster τ_t are self-disabling; hence computing the image under \mathbf{T}_i and $*\mathbf{T}_i$ require similar computational effort.

4 Discussion and conclusion

Reachability expressions give the user the ability to specify the heuristics of symbolic state space search. Semantically equivalent reachability expressions can have radically different costs of computation. In this paper, we presented a theory to reason about the equivalence and relative performance of alternative reachability expressions, and validated our predictions with experiments.

NuSMVDP is a novel tool that allows user specified scheduling of transition relations. SMART is another tool which has similar properties. SMART is the outcome of the saturation algorithm discussed in [8]. The saturation algorithm uses multi-way decision diagrams(MDD) and Kronecker operators to encode the next state function. The Kronecker encoding helps them exploit locality of effect that events have on state variables. This minimizes the peak memory consumption during reachability analysis. The algorithm is implemented in the tool named SMART. Both NuSMVDP and SMART have efficient implementation for identity transitions. In SMART, the identity transitions are ignored in the corresponding Kronecker matrix. In NuSMVDP all unspecified variables are assumed to undergo the identity transition. For example, consider a state variable s_1 which has the corresponding next state variable s'_1 . Let s_1 be a variable that is not affected by T_1 i.e. s'_1 is not in the support set of T_1 . In NuSMVDP, during the image computation using BDDs s_1 is not quantified and is renamed to s'_1 indicating that the variable underwent an identity transition. This significantly improves the computation time and memory usage because quantification of variables using BDDs is an expensive operation. In SMART, a transition relation is partitioned disjunctively along two dimensions: by events and by sub-models. This is called a Kronecker consistent decomposition, where it is possible to express the effect of firing an event α in a global state as the cross product of the local effects of α on each sub-model. The decomposition helps capture event locality and update sets of state variables in a manner that reduces peak memory usage. If there are M sub-models and E events then we need to consider a total of $E \times M$ transitions relations. A large number

of transition clusters implies increase in the number of iterations before fix-point is reached in reachability analysis. Hence a large number of transition clusters could increase the computation time because of scheduling of the transition relations in image computation. This problem can be overcome in NuSMVDP by selectively combining sub-models for certain events, which would reduce the number of transition relations to consider.

The motivation behind the saturation method in [8] is very similar to our motivation for taking fix-point with respect to local transition clusters in NuSMVDP. As we have shown in our comparison of schedules S_3 and S_4 , taking a local fix-point is not guaranteed to perform well. If the cluster size is large with a large support variable set, as the time transition cluster T_t in our examples, it may be better to avoid local fix-point computations with this transition cluster. This indicates that efficiency of the saturation method would also depend on the decomposition used for the transition relation.

Our experimental investigations indicate that when the absolute size of the BDD representation of a monolithic transition relation is small, it is advantageous to perform classical symbolic breadth-first search using the monolithic relation directly. Similarly, when the maximum BDD size encountered in representing (partially) computed state sets is small compared to the BDD size of the monolithic transition relation or the BDD size of the largest transition cluster, it helps to minimize the number of image computation iterations by using classical symbolic breadth-first search. However, if the maximum BDD size for representing (partially) computed state sets is large and there is a dominant cluster represented as a large BDD, it is advantageous to adopt a round-robin scheduling of clusters in a way that minimizes the image computations under the dominant cluster. Further, in round-robin scheduling, ordering the clusters in topological order is effective in examples such as deep circuits, which have considerable forward propagation of events.

The experiments and conclusions reported in this paper are promising, although preliminary. Much more data on a wider set of examples is needed before a comprehensive evaluation of the effectiveness of reachability expressions in improving the performance of state space search can be done. However, we believe that reachability expressions will serve as a useful addition to the existing repository of tools and techniques for making symbolic reachability analyzers more efficient.

Acknowledgments: We thank Varun Kanade for help with the experiments.

References

1. D. Beyer, C. Lewerentz, and A. Noack. Rabbit: A tool for BDD-based verification of real-time systems. In *Proceedings of CAV*, LNCS 2725, pages 122–125, 2003.
2. R. Bloem, K. Ravi, and F. Somenzi. Symbolic guided search for CTL model checking. In *Proceeding of ACM/IEEE DAC*, pages 29–34, 2000.
3. R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, 35(8):677–691, 1986.
4. J. R. Burch, E. M. Clarke, K. L. McMillan, and D. L. Dill. Sequential circuit verification using symbolic model checking. In *Proceeding of ACM/IEEE DAC*, pages 46–61, 1990.
5. G. Cabodi, P. Camurati, L. Lavagno, and S. Quer. Disjunctive partitioning and partial iterative squaring: An effective approach to symbolic traversal of large circuits. In *Proceeding of ACM/IEEE DAC*, pages 728–733, 1997.
6. P. Chauhan, E. M. Clarke, S. Jha, J. Kukula, T. Shiple, H. Veith, and D. Wang. Non-linear quantification scheduling in image computation. In *Proceedings of ACM/IEEE ICCAD*, pages 293–298, 2001.
7. P. Chauhan, E. M. Clarke, S. Jha, J. Kukula, H. Veith, and D. Wang. Using combinatorial optimization methods for quantification scheduling. In *Proceedings of CHARME*, LNCS 2144, pages 293–309, 2001.
8. Gianfranco Ciardo, Robert Marmorstein, and Radu Siminiceanu. The saturation algorithm for symbolic state-space exploration. *Int. J. Softw. Tools Technol. Transf.*, 8(1):4–25, 2006.
9. A. Cimatti, E. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. NuSMV version 2: An open source tool for symbolic model checkin. In *Proceedings of CAV*, LNCS 2404, pages 359–364, 2002.
10. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic Model Checking: 10^{20} States and Beyond. In *Proceedings of LICS*, pages 1–33, 1990.
11. A. Narayan, A. J. Isles, J. Jain, R. K. Brayton, and A. L. Sangiovanni-Vincentelli. Reachability analysis using partitioned-robdds. In *Proceedings of ACM/IEEE ICCAD*, pages 388–393, 1997.
12. P. K. Pandya and M. Raut. A kleene algebra of reachability expressions and its use in efficient symbolic guided search. Technical Report (in preparation), STCS, Tata Institute of Fundamental Research, 2005.
13. D. Thomas, S. Chakraborty, and P. K. Pandya. Efficient guided symbolic reachability using reachability expressions. Technical Report TR-06-1 (<http://www.cfdvs.iitb.ac.in/reports/techrep06.php3>), CFDVS, IIT Bombay, January 2006.
14. D. Thomas, P. K. Pandya, and S. Chakraborty. Scheduling clusters in model checking of real time systems. Technical Report TR-04-16 (<http://www.cfdvs.iitb.ac.in/reports/techrep04.php3>), CFDVS, IIT Bombay, September 2004.