

Model Checking and Modular Verification*

Orna Grumberg
Computer Science Department
The Technion
Haifa 32000, Israel
orna@techsel (BITNET)

David E. Long
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
long+@cs.cmu.edu (Internet)

August 15, 1991

Abstract

We describe a framework for compositional verification of finite state processes. The framework is based on two ideas: a subset of the logic CTL for which satisfaction is preserved under composition; and a preorder on structures which captures the relation between a component and a system containing the component. Satisfaction of a formula in the logic corresponds to being below a particular structure (a tableau for the formula) in the preorder. We show how to do assume-guarantee style reasoning within this framework. In addition, we demonstrate efficient methods for model checking in the logic and for checking the preorder in several special cases. We have implemented a system based on these methods, and we use it to give a compositional verification of a CPU controller.

1 Introduction

Temporal logic model checking procedures are useful tools for the verification of finite state systems [3, 12, 20]. However, these procedures have traditionally suffered from the state explosion problem. This problem arises in systems which are composed of many parallel processes; in general, the size of the state space grows exponentially with the number of processes. By introducing symbolic representations for sets of states and transition relations and using a symbolic model checking procedure, systems with very large state spaces (10^{100} or more states) can be verified [1, 8]. Further, the time and space requirements with these techniques may in practice be polynomial in the number of *components* of the system.

*This research was sponsored in part by the Avionics Laboratory, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U.S. Air Force, Wright-Patterson AFB, Ohio 45433-6543 under Contract F33615-90-C-1465, ARPA Order No. 7597 and in part by the National Science Foundation under Contract No. CCR-9005992 and in part by the U.S.-Israeli Binational Science Foundation.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

Unfortunately, the symbolic procedures still have limits, and many realistic problems are not tractable due to their size. Thus, we are motivated to search for additional methods of handling the state explosion problem, methods which work well in conjunction with the symbolic techniques.

An obvious method for trying to avoid the state explosion problem is to use the natural decomposition of the system. The goal is to verify properties of individual components, infer that these hold in the complete system, and use them to deduce additional properties of the system. When verifying properties of the components, it may also be necessary to make assumptions about the environment. This approach is exemplified by Pnueli’s assume-guarantee paradigm [23]. A formula in his logic is a triple $\langle\varphi\rangle M \langle\psi\rangle$ where φ and ψ are temporal formulas and M is a program. The formula is true if whenever M is part of a system satisfying φ , the system must also satisfy ψ . A typical proof shows that $\langle\varphi\rangle M \langle\psi\rangle$ and $\langle true\rangle M' \langle\varphi\rangle$ hold and concludes that $\langle true\rangle M \parallel M' \langle\psi\rangle$ is true.

In order to automate this approach, a model checker must have several properties. It must be able to check that a property is true of *all* systems which can be built using a given component. More generally, it must be able to restrict to a given class of environments when doing this check. It must also provide facilities for performing temporal reasoning. Most existing model checkers were not designed to provide these facilities. Instead, they typically assume that they are given complete systems.

An elegant way to obtain a system with the above properties is to provide a preorder on the finite state models that captures the notion of “more behaviors” and to use a logic whose semantics relate to the preorder. The preorder should preserve satisfaction of formulas of the logic, i.e., if a formula is true for a model, it should also be true for any model which is smaller in the preorder. In addition, composition should preserve the preorder, and a system should be smaller in the preorder than its individual components. Finally, satisfaction of a formula should correspond to being smaller than a particular model (a tableau for the formula) in the preorder. In such a framework, the above reasoning sequence might be expressed as: T is the tableau of φ , $M \parallel T \models \psi$, $M' \preceq T$, and hence $M \parallel M' \models \psi$. Note that assumptions may be given either as formulas or directly as finite state models, whichever is more concise or convenient. More complex forms of reasoning such as induction [18] are also possible within this framework.

In choosing a computational model, a logic and a preorder to obtain a system such as this, we are guided by the following considerations. First, we must be able to realistically model physical systems such as circuits. Second, there should be efficient procedures for model checking and for checking the preorder. Finally, it should be possible to implement these procedures effectively using symbolic techniques.

In this paper, we propose a preorder for use with a subset of the logic CTL* [11]. This subset is strictly more expressive than LTL. Further, the induced subset of CTL is expressive enough for most verification tasks and has an efficient model checking algorithm. We also give a tableau construction for this CTL subset. The construction provides a means of temporal reasoning and makes it possible to use formulas as assumptions. Our preorder and the semantics of our logics both include a notion of fairness. This is essential for modeling systems such as communication protocols. We show how to use our results to verify systems composed of Moore machines. Moore machines have an explicit notion of input and output

and are particularly suitable for modeling synchronous circuits. Finally, we suggest efficient methods for checking the preorder in several interesting cases. We have implemented a system based on these results; the system supports efficient compositional verification and temporal reasoning.

Our paper is organized as follows. Section 2 surveys some related work. In section 3, we present the logic and its semantics (for Kripke structures). The preorder and some of its properties are given in section 4. The next section defines the semantics of the logic for Moore machines. Given a Moore machine and a formula, we show how to efficiently check whether for all environments, the Moore machine in the environment satisfies the formula. Section 6 presents the tableau construction and demonstrates how to use it for temporal reasoning. Methods for checking the preorder are discussed in section 7. Section 8 gives a compositional verification of a simple CPU controller. We conclude with a summary and some directions for future work.

2 Related work

Much of the work on reducing the complexity of automatic verification can be grouped into two classes. The first class includes methods to build a reduced global state graph or to expand only the needed portion of the global state graph.

Local model checking algorithms [6, 26, 29] based on logics like the μ -calculus use a tableau-based procedure to deduce that a specific state (the initial state of the system) satisfies a given logical formula. The state space can be generated as needed in such an algorithm, and for some formulas, only a small portion of the space may have to be examined. The main drawback of these algorithms is that often the entire space is generated (for example, when checking that a property holds globally). It is also not clear whether the algorithms can take good advantage of symbolic representations.

Graf and Steffen [13] describe a method for generating a reduced version of the global state space given a description of how the system is structured and specifications of how the components interact. Clarke, Long and McMillan [4] describe a similar attempt. Both methods will still produce large state graphs if most of the states in the system are not equivalent, and much of the verification must be redone if part of the system changes. Shtadler and Grumberg [24] show how to verify networks of processes whose structure is described by grammars. In this approach, which involves finding the global behavior of each component, networks of arbitrary complexity can be verified by checking one representative system. For many systems, however, the number of states may still be prohibitive, and it is not clear whether the method can use symbolic representations.

The second class of methods are compositional; properties of the individual components are verified, and properties of the global system are deduced from these. A representation of the global state space is not built.

Josko [15] gives an algorithm for checking whether a system satisfies a CTL specification in all environments. His algorithm also allows assumptions about the environment to be specified in a restricted linear-time logic. The system is able to handle assume-guarantee reasoning. The method is fairly *ad hoc* however, and more complex forms of reasoning such as induction cannot be easily incorporated into the system.

Within the framework of CCS [22], there have been a number of suggestions for compositional reasoning. Larsen [19] investigates the expressive power of formalisms for specifying the behavior of a process in a system. He suggests equivalence, refinement and satisfaction (of a formula) as three interesting relations between an implementation and its specification. However, he does not discuss the applicability of these ideas to verification, nor does he suggest how they can be implemented. Walker [27] demonstrates how to use a preorder plus knowledge of how a system should operate to simplify the verification of bisimulation equivalence. Cleaveland and Steffen [7] use a similar idea. Winskel [28] proposes a method for decomposing specifications into properties which the components of a system must satisfy for the specification to hold. The approach is very appealing, but unfortunately, dealing with parallel composition is difficult. It is not apparent whether any of these methods will work well with symbolic representations.

Kurshan [16] describes a verification methodology based on testing containment of ω -regular languages. Homomorphic reductions are used to map implementations to specifications, and the specifications may be used as implementations at the next level of abstraction. Dill [10] proposes an elegant form of trace theory which can be used in a similar manner, but the framework does not handle liveness properties well. Both approaches depend on specifications being deterministic for efficiency, and neither approach makes provisions for using logical formulas as specifications or assumptions.

Shurek and Grumberg [25] describe criteria for obtaining a modular framework, and illustrate the idea using CTL* with only universal path quantifiers. This system is closest to the work presented here, but they give no provisions for handling fairness efficiently, using formulas as assumptions, or supporting temporal reasoning. Models in their system are also associated with a fixed decomposition into components; hence it is unclear how to perform inductive reasoning in the framework.

3 Temporal logic

The logics presented in this section are branching-time temporal logics. In order to be able to efficiently decide whether a formula is true in all systems containing a given component, we eliminate the existential path quantifier from the logics. Thus, a formula may include only the universal quantifier over paths, but unlike in linear-time temporal logic, nesting of path quantifiers is allowed. To ensure that existential path quantifiers do not arise via negation, we will assume that formulas are expressed in negation normal form. In other words, negations are applied only to atomic propositions. The logics are interpreted over a form of Kripke structure with fairness constraints. Path quantifiers range over the fair paths in the structures.

Definition 1 (\forall CTL*) *The logic \forall CTL* is the set of state formulas given by the following inductive definition.*

1. *The constants true and false are state formulas. For every atomic proposition p , p and $\neg p$ are state formulas.*
2. *If φ and ψ are state formulas, then $\varphi \wedge \psi$ and $\varphi \vee \psi$ are state formulas.*

3. If φ is a path formula, then $\forall(\varphi)$ is a state formula.
4. If φ is a state formula, then φ is a path formula.
5. If φ and ψ are path formulas, then so are $\varphi \wedge \psi$, $\varphi \vee \psi$.
6. If φ and ψ are path formulas, then so are
 - (a) $\mathbf{X}\varphi$,
 - (b) $\varphi \mathbf{U}\psi$, and
 - (c) $\varphi \mathbf{V}\psi$.

We also use the following abbreviations: $\mathbf{F}\varphi$ and $\mathbf{G}\varphi$, where φ is a path formula, denote (true $\mathbf{U}\varphi$) and (false $\mathbf{V}\varphi$) respectively.

$\forall\text{CTL}$ is a restricted subset of $\forall\text{CTL}^*$ in which the \forall path quantifier may only precede a restricted set of path formulas. More precisely, $\forall\text{CTL}$ is the logic obtained by eliminating rules 3 through 6 above and adding the following rule.

- 3'. If φ and ψ are state formulas, then $\forall\mathbf{X}\varphi$, $\forall(\varphi \mathbf{U}\psi)$, and $\forall(\varphi \mathbf{V}\psi)$ are state formulas.

In practice, we have found that many of the formulas which are used in specifying and verifying systems are expressible in $\forall\text{CTL}$, and almost all are expressible in $\forall\text{CTL}^*$. An example formula which is not expressible in $\forall\text{CTL}^*$ is a weak form of absence of deadlock: $\forall\mathbf{G}\exists\mathbf{F}p$ states that it should always be possible to reach a state where p holds.

We will give the semantics of the logic using a form of Kripke structure with fairness constraints.

Definition 2 (structure) A structure $M = \langle S, S_0, \mathcal{A}, \mathcal{L}, R, \mathcal{F} \rangle$ is a tuple of the following form.

1. S is a finite set of states.
2. $S_0 \subseteq S$ is a set of initial states.
3. \mathcal{A} is a finite set of atomic propositions.
4. \mathcal{L} is a function that maps each state to the set of atomic propositions true in that state.
5. $R \subseteq S \times S$ is a transition relation.
6. \mathcal{F} is a Streett acceptance condition, represented by pairs of sets of states.

Definition 3 A path in M is an infinite sequence of states $\pi = s_0s_1s_2\dots$ such that for all $i \in \mathbb{N}$, $R(s_i, s_{i+1})$.

Definition 4 Define $\text{inf}(\pi) = \{s \mid s = s_i \text{ for infinitely many } i\}$. π is a fair path in M iff for every $(P, Q) \in \mathcal{F}$, if $\text{inf}(\pi) \cap P \neq \emptyset$, then $\text{inf}(\pi) \cap Q \neq \emptyset$.

The notation π^n will denote the suffix of π which begins at s_n . We now consider the semantics of the logic $\forall\text{CTL}^*$ with atomic propositions drawn from the set \mathcal{A} .

Definition 5 (satisfaction of a formula) *Satisfaction of a state formula φ by a state s ($s \models \varphi$) and of a path formula ψ by a fair path π ($\pi \models \psi$) is defined inductively as follows.*

1. $s \models \text{true}$, and $s \not\models \text{false}$. $s \models p$ iff $p \in \mathcal{L}(s)$. $s \models \neg p$ iff $p \notin \mathcal{L}(s)$.
2. $s \models \varphi \wedge \psi$ iff $s \models \varphi$ and $s \models \psi$. $s \models \varphi \vee \psi$ iff $s \models \varphi$ or $s \models \psi$.
3. $s \models \forall(\varphi)$ iff for every fair path π starting at s , $\pi \models \varphi$.
4. $\pi \models \varphi$, where φ is a state formula, iff the first state of π satisfies the state formula.
5. $\pi \models \varphi \wedge \psi$ iff $\pi \models \varphi$ and $\pi \models \psi$. $\pi \models \varphi \vee \psi$ iff $\pi \models \varphi$ or $\pi \models \psi$.
6. (a) $\pi \models \mathbf{X}\varphi$ iff $\pi^1 \models \varphi$.
 (b) $\pi \models \varphi \mathbf{U} \psi$ iff there exists $n \in \mathbb{N}$ such that $\pi^n \models \psi$ and for all $i < n$, $\pi^i \models \varphi$.
 (c) $\pi \models \varphi \mathbf{V} \psi$ iff for all $n \in \mathbb{N}$, if $\pi^i \not\models \varphi$ for all $i < n$, then $\pi^n \models \psi$.

$M \models \varphi$ indicates that for every $s_0 \in S_0$, $s_0 \models \varphi$.

Emerson and Halpern [11] compared the expressive power of the three logics LTL, CTL and CTL^* . They showed that LTL and CTL have incomparable expressive power, while CTL^* is strictly more expressive than either of the others. Eliminating the existential path quantifier from CTL and CTL^* does not affect the relative expressive power of the logics. $\forall\text{CTL}^*$ trivially encompasses LTL and $\forall\text{CTL}$. The formula $\forall\mathbf{F}\forall\mathbf{G}p$ is a formula of $\forall\text{CTL}$ that does not have an equivalent LTL formula. On the other hand, there is no $\forall\text{CTL}$ formula that is equivalent to the LTL formula $\forall\mathbf{F}\mathbf{G}p$. Thus, LTL and $\forall\text{CTL}$ are incomparable, and both are strictly less expressive than $\forall\text{CTL}^*$.

4 Homomorphisms and composition of structures

In this section, we define the preorder which we use and examine some of its properties. We also show how these properties make assume-guarantee style reasoning possible.

Definition 6 (structure homomorphism) *Let M and M' be two structures with $\mathcal{A} \supseteq \mathcal{A}'$, and let t and t' be states in S and S' , respectively. A relation $H \subseteq S \times S'$ is a homomorphism from (M, t) to (M', t') iff the following conditions hold.*

1. $H(t, t')$.
2. For all s and s' , $H(s, s')$ implies
 - (a) $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$; and
 - (b) for every fair path $\pi = s_0s_1 \dots$ from $s = s_0$ in M there exists a fair path $\pi' = s'_0s'_1 \dots$ from $s' = s'_0$ in M' such that for every $i \in \mathbb{N}$, $H(s_i, s'_i)$.

When H satisfies property 2, we say H is a homomorphism. H is a homomorphism from M to M' iff for every $s_0 \in S_0$ there is $s'_0 \in S'_0$ such that $H(s_0, s'_0)$. To indicate that two paths correspond as in item 2b above, we write $H(\pi, \pi')$.

Definition 7 For $s \in S$ and $s' \in S'$, $(M, s) \preceq (M', s')$ iff there is a homomorphism from (M, s) to (M', s') . $M \preceq M'$ iff there exists a homomorphism from M to M' .

When M and M' are understood, we sometimes write $s \preceq s'$. Intuitively, two states are homomorphic if their labels agree on the atomic propositions of the second structure and if for every fair path from the first state there is a corresponding fair path from the second state. Two structures are homomorphic if for every initial state of the first, there is a corresponding initial state of the second. One may view the second structure as a specification and the first as its implementation. Since a specification may hide some of the implementation details, it may have a smaller set of atomic propositions.

Definition 8 (composition of structures) Let M and M' be two structures. The composition of M and M' , denoted $M \parallel M'$, is the structure M'' defined as follows.

1. $S'' = \{ (s, s') \mid \mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s') \cap \mathcal{A} \}$.
2. $S''_0 = (S_0 \times S'_0) \cap S''$.
3. $\mathcal{A}'' = \mathcal{A} \cup \mathcal{A}'$.
4. $\mathcal{L}''((s, s')) = \mathcal{L}(s) \cup \mathcal{L}'(s')$.
5. $R''((s, s'), (t, t'))$ iff $R(s, t)$ and $R'(s', t')$.
6. $\mathcal{F}'' = \{ ((P \times S') \cap S'', (Q \times S') \cap S'') \mid (P, Q) \in \mathcal{F} \}$
 $\cup \{ ((S \times P') \cap S'', (S \times Q') \cap S'') \mid (P', Q') \in \mathcal{F}' \}$.

The choice of this definition of composition is motivated by its correspondence with composition of Moore machines. Each transition of the composition is a joint transition of the components, and states of the composition are pairs of component states that agree on their common atomic propositions. We first note that this composition operator has the usual properties.

Theorem 1 *Composition of structures is commutative and associative (up to isomorphism).*

Proof Straightforward but tedious. □

We now turn to the connections between the relation \preceq and composition. To begin, we note that a path in $M \parallel M'$ is fair iff its restriction to each component results in a fair path.

Lemma 1 *Let $M'' = M \parallel M'$. The following conditions are equivalent.*

1. $\pi'' = (s_0, s'_0)(s_1, s'_1) \dots$ is a fair path in M'' .

2. $\pi = s_0 s_1 \dots$ and $\pi' = s'_0 s'_1 \dots$ are fair paths in M and M' respectively, and (s_i, s'_i) is a state of M'' for all $i \in \mathbb{N}$.

Proof Assume condition 1 above. By the definition of composition, $\pi = s_0 s_1 \dots$ is a path in M . Let $(P, Q) \in \mathcal{F}$, and suppose $\inf(\pi) \cap P \neq \emptyset$. Now $(P'', Q'') = ((P \times S') \cap S'', (Q \times S') \cap S'') \in \mathcal{F}''$, and $\inf(\pi'') \cap P'' \neq \emptyset$. By the definition of a fair path, $\inf(\pi'') \cap Q'' \neq \emptyset$. Hence $\inf(\pi) \cap Q \neq \emptyset$, and so π is a fair path in M . Similarly, $\pi' = s'_0 s'_1 \dots$ is a fair path in M' .

Assume condition 2 above. From the definition of composition, $\pi'' = (s_0, s'_0)(s_1, s'_1) \dots$ is a path in M'' . Suppose $(P'', Q'') \in \mathcal{F}''$ and $\inf(\pi'') \cap P'' \neq \emptyset$. Either $(P'', Q'') = ((P \times S') \cap S'', (Q \times S') \cap S'')$ for some $(P, Q) \in \mathcal{F}$, or $(P'', Q'') = ((S \times P') \cap S'', (S \times Q') \cap S'')$ for some $(P', Q') \in \mathcal{F}'$. In the first case, we have $\inf(\pi) \cap P \neq \emptyset$, and so $\inf(\pi) \cap Q \neq \emptyset$. This implies $\inf(\pi'') \cap Q'' \neq \emptyset$. The second case is similar. Hence π'' is a fair path in M'' . \square

Theorem 2

1. \preceq is a preorder.
2. For all M and M' , $M \parallel M' \preceq M$.
3. For all M, M' and M'' , if $M \preceq M'$ then $M \parallel M'' \preceq M' \parallel M''$.
4. For all M , $M \preceq M \parallel M$.

Proof

1. The relation $H = \{ (s, s) \mid s \in S \}$ is a homomorphism from M to M , so \preceq is reflexive. Thus it only remains to show that \preceq is transitive. Assume $M \preceq M'$ and $M' \preceq M''$. Let H_0 be a homomorphism from M to M' , and let H_1 be a homomorphism from M' to M'' . Define H_2 as the relational product of H_0 and H_1 , i.e.,

$$H_2 = \{ (s, s'') \mid \exists s' [H_0(s, s') \wedge H_1(s', s'')] \}.$$

If $s_0 \in S_0$, then by the definition of homomorphism, there exists $s'_0 \in S'_0$ such that $H_0(s_0, s'_0)$. Similarly, there exists $s''_0 \in S''_0$ such that $H_1(s'_0, s''_0)$, and hence $H_2(s_0, s''_0)$.

Suppose $H_2(s, s'')$, and let s' be such that $H_0(s, s')$ and $H_1(s', s'')$. By the definition of homomorphism, $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$ and $\mathcal{L}'(s') \cap \mathcal{A}'' = \mathcal{L}''(s'')$. Then since $\mathcal{A}' \supseteq \mathcal{A}''$, we have $\mathcal{L}(s) \cap \mathcal{A}'' = \mathcal{L}''(s'')$. If π is a fair path in M from s , then there exists a fair path π' from s' in M' such that $H_0(\pi, \pi')$. Since H_1 is a homomorphism, there exists a fair path π'' from s'' in M'' such that $H_1(\pi', \pi'')$. But then $H_2(\pi, \pi'')$, and hence H_2 is a homomorphism from M to M'' . Thus $M \preceq M''$.

2. Define H by

$$H = \left\{ \left((s, s'), s \right) \mid (s, s') \in S^{M \parallel M'} \right\}.$$

If (s_0, s'_0) is an initial state of $M \parallel M'$, then $s_0 \in S_0$. The label of (s, s') is $\mathcal{L}(s) \cup \mathcal{L}'(s')$, and $(\mathcal{L}(s) \cup \mathcal{L}'(s')) \cap \mathcal{A} = \mathcal{L}(s)$. If $(s_0, s'_0)(s_1, s'_1) \dots$ is a fair path in $M \parallel M'$, then by the previous lemma, $s_0 s_1 \dots$ is a fair path in M . By the definition of H , $H((s_i, s'_i), s_i)$ for every i . Hence H is a homomorphism and $M \parallel M' \preceq M$.

3. Let H_0 be a homomorphism from M to M' . Define H_1 by

$$H_1 = \left\{ \left((s, s''), (s', s'') \right) \mid H_0(s, s') \right\}.$$

We show that H_1 is a homomorphism. Let (s_0, s''_0) be an initial state of $M \parallel M''$. By the definition of composition, $s_0 \in S_0$ and $s''_0 \in S''_0$. Since $M \preceq M'$, there exists $s'_0 \in S'_0$ such that $H_0(s_0, s'_0)$. Now (s'_0, s''_0) is a state of $M' \parallel M''$ since

$$\begin{aligned} \mathcal{L}'(s'_0) \cap \mathcal{A}'' &= (\mathcal{L}(s_0) \cap \mathcal{A}') \cap \mathcal{A}'' \\ &= (\mathcal{L}(s_0) \cap \mathcal{A}'') \cap \mathcal{A}' \\ &= (\mathcal{L}''(s''_0) \cap \mathcal{A}) \cap \mathcal{A}' \\ &= \mathcal{L}''(s''_0) \cap \mathcal{A}'. \end{aligned}$$

Further, (s'_0, s''_0) is an initial state of $M' \parallel M''$ by the definition of composition. By definition of H_1 , we have $H_1((s_0, s''_0), (s'_0, s''_0))$.

Suppose $H_1((s, s''), (s', s''))$. First note that

$$\begin{aligned} (\mathcal{L}(s) \cup \mathcal{L}''(s'')) \cap (\mathcal{A}' \cup \mathcal{A}'') &= (\mathcal{L}(s) \cap \mathcal{A}') \cup (\mathcal{L}(s) \cap \mathcal{A}'') \\ &\quad \cup (\mathcal{L}''(s'') \cap (\mathcal{A}' \cup \mathcal{A}'')) \\ &= \mathcal{L}'(s') \cup (\mathcal{L}''(s'') \cap \mathcal{A}) \cup \mathcal{L}''(s'') \\ &= \mathcal{L}'(s') \cup \mathcal{L}''(s''). \end{aligned}$$

Let $(s_0, s''_0)(s_1, s''_1) \dots$ be a fair path in $M \parallel M''$ from $(s, s'') = (s_0, s''_0)$. Then for every $i \in \mathbb{N}$, we have $\mathcal{L}(s_i) \cap \mathcal{A}'' = \mathcal{L}''(s''_i) \cap \mathcal{A}$. By the previous lemma, $\pi = s_0 s_1 \dots$ is a fair path in M starting at s , and $\pi'' = s''_0 s''_1 \dots$ is a fair path in M'' from s'' . Since $H_0(s, s')$, there is a path $\pi' = s'_0 s'_1 \dots$ from $s' = s'_0$ in M' such that for every $i \in \mathbb{N}$, $H_0(s_i, s'_i)$. By the definition of homomorphism, $\mathcal{L}(s_i) \cap \mathcal{A}' = \mathcal{L}'(s'_i)$ for all i . Arguing as above, we then have $\mathcal{L}'(s'_i) \cap \mathcal{A}'' = \mathcal{L}''(s''_i) \cap \mathcal{A}'$ for each i , and so each (s'_i, s''_i) is a state in $M' \parallel M''$. Now $H_1((s_i, s''_i), (s'_i, s''_i))$ by the definition of H_1 . Applying the previous lemma, we find that $(s'_0, s''_0)(s'_1, s''_1) \dots$ is a fair path starting in (s', s'') and corresponding to the path $(s_0, s''_0)(s_1, s''_1) \dots$.

4. First note that for every state s of M , (s, s) is a state of $M \parallel M$. Define $H = \left\{ (s, (s, s)) \mid s \in S \right\}$. If $s_0 \in S_0$, then by the definition of composition, (s_0, s_0) is an

initial state of $M \parallel M$. (s, s) trivially has the same label as s . Using the previous lemma and the definition of composition, we find that if $s_0 s_1 \dots$ is a fair path in M , then $(s_0, s_0)(s_1, s_1) \dots$ is a fair path in $M \parallel M$. By the definition of H , we have $H(s_i, (s_i, s_i))$ for all i . Hence H is a homomorphism and $M \preceq M \parallel M$. \square

Theorem 3 *Let s and s' be states of M and M' , and let H be a homomorphism such that $H(s, s')$. Let π and π' be fair paths such that $H(\pi, \pi')$. Then*

1. *for every $\forall CTL^*$ (state) formula φ (with all atomic propositions in \mathcal{A}'), if $s' \models \varphi$ then $s \models \varphi$; and*
2. *for every $\forall CTL^*$ path formula φ (with all atomic propositions in \mathcal{A}'), if $\pi' \models \varphi$, then $\pi \models \varphi$.*

Proof The proof proceeds by induction on the structure of the formula.

1. If $\varphi = \text{true}$ or $\varphi = \text{false}$, the result is trivial. If $\varphi = p$, an atomic proposition, then $s' \models \varphi$ if and only if $p \in \mathcal{L}'(s')$. By the definition of homomorphism, $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$, and so $p \in \mathcal{L}(s)$ iff $p \in \mathcal{L}'(s')$. Thus $s \models \varphi$. The case where $\varphi = \neg p$ is similar.
2. If $\varphi = \varphi_1 \wedge \varphi_2$, then $s' \models \varphi$ iff $s' \models \varphi_1$ and $s' \models \varphi_2$. The induction hypothesis implies $s \models \varphi_1$ and $s \models \varphi_2$; hence $s \models \varphi$. The case where $\varphi = \varphi_1 \vee \varphi_2$ is similar.
3. If $\varphi = \forall(\varphi_1)$, then $s \models \varphi$ iff for every fair path π from s , $\pi \models \varphi_1$. Let π be any fair path from s . By the definition of homomorphism, there exists a fair path π' from s' such that $H(\pi, \pi')$. If $s' \models \varphi$, then $\pi' \models \varphi_1$ for any π' from s' . The induction hypothesis then implies $\pi \models \varphi_1$, and hence $s \models \varphi$.
4. If φ is a path formula consisting of only a state formula and $\pi' \models \varphi$, then the initial state s' of π' satisfies φ . By the induction hypothesis, $s \models \varphi$, and since s is the initial state of π , $\pi \models \varphi$.
5. The cases for the conjunction and disjunction of path formulas are similar to case 2.
6. (a) If $\varphi = \mathbf{X} \varphi_1$, then $\pi' \models \varphi$ implies $\pi'^1 \models \varphi$. Now since $H(\pi, \pi')$, we also have $H(\pi^1, \pi^1)$. Then the induction hypothesis implies $\pi^1 \models \varphi_1$. Thus $\pi \models \varphi$.
(b) If $\varphi = \varphi_1 \mathbf{U} \varphi_2$, then $\pi' \models \varphi$ implies there exists n such that $\pi'^n \models \varphi_2$ and for all $i < n$, $\pi'^i \models \varphi_1$. Since $H(\pi, \pi')$, we have $H(\pi^j, \pi'^j)$ for any j . Applying the induction hypothesis, $\pi^n \models \varphi_2$ and $\pi^i \models \varphi_1$ for all $i < n$. Hence $\pi \models \varphi$.
(c) The case where $\varphi = \varphi_1 \mathbf{V} \varphi_2$ is similar to the previous two cases. \square

Corollary 1 *Suppose $M \preceq M'$. Then for every $\forall CTL^*$ formula φ (with atomic propositions in \mathcal{A}'), $M' \models \varphi$ implies $M \models \varphi$.*

Proof Immediate. \square

Using theorem 2 and this corollary, we see that a standard CTL (CTL*) model checking algorithm [3], when restricted to \forall CTL (\forall CTL*), can be viewed as determining whether a formula is true of all systems containing a given component. This is the key to compositional verification. With the theorem and corollary, it is also straightforward to justify the soundness of the assume-guarantee paradigm when assumptions are given as structures. (The connection between structures and formulas will be examined in section 6.) Discharging an assumption involves checking for the relation \preceq . Suppose that we wish to check that $M \parallel M' \models \varphi$ and that we have verified the following relationships:

$$\begin{aligned} M &\preceq A \\ A \parallel M' &\preceq A' \\ M \parallel A' &\models \varphi. \end{aligned}$$

In other words, M discharges assumption A , M' under assumption A discharges assumption A' , and M under assumption A' satisfies the desired formula. From theorem 2, we have

$$\begin{aligned} M \parallel M' &\preceq M \parallel M \parallel M' \\ &\preceq M \parallel A \parallel M' \\ &\preceq M \parallel A'. \end{aligned}$$

Then corollary 1 implies that $M \parallel M' \models \varphi$. The theorem and corollary also show that any system containing $M \parallel M'$ will satisfy φ . Note that φ is not necessarily true in either M or M' and may involve atomic propositions from both M and M' .

5 Moore machines

We have seen that the structures defined earlier (definition 2) can be used for compositional reasoning about synchronous systems. However, such systems are typically given using a more common finite state model such as Moore machines [14]. Moore machines are models of computation with an explicit notion of inputs and outputs. Since the inputs originate from an external, uncontrolled environment, the machine can always receive any combination of input values. Moore machines are synchronous; in a composition of Moore machines, each machine makes a single step at every point. Thus, they are most suitable for modeling synchronous circuits. In this section, we show a natural correspondence between Moore machines with an empty set of inputs and the structures defined earlier. We use this correspondence to define the semantics of \forall CTL* with respect to Moore machines, and we show how to use compositional reasoning to verify a system composed of Moore machines.

Definition 9 (Moore machine) *A Moore machine $M = \langle S, S_0, I, O, \mathcal{L}, R \rangle$ is a tuple of the following form*

1. S is a finite set of states.
2. $S_0 \subseteq S$ is a set of initial states.

3. I is a finite set of input propositions.
4. O is a finite set of output propositions.
5. \mathcal{L} is a function that maps each state to the set of output propositions true in that state.
6. $R \subseteq S \times 2^I \times S$ is the transition relation.

We require that $I \cap O = \emptyset$ and that for every $s \in S$ and $v \subseteq I$, there exists some $t \in S$ such that $R(s, v, t)$. We also let \mathcal{A} denote $I \cup O$.

Definition 10 (composition of Moore machines) Let M and M' be Moore machines with $O \cap O' = \emptyset$. The composition of M and M' , denoted $M \parallel M'$, is the Moore machine M'' defined as follows.

1. $S'' = S \times S'$.
2. $S''_0 = S_0 \times S'_0$.
3. $I'' = (I \cup I') \setminus (O \cup O')$.
4. $O'' = O \cup O'$.
5. $\mathcal{L}''((s, s')) = \mathcal{L}(s) \cup \mathcal{L}'(s')$.
6. $R''((s, s'), v, (t, t'))$ iff $R(s, (v \cup \mathcal{L}'(s')) \cap I, t)$ and $R'(s', (v \cup \mathcal{L}(s)) \cap I', t')$.

We now turn to the question of how to define satisfaction of a specification by a Moore machine M . The key consideration is that we wish to have a compositional method of reasoning. Thus, *M satisfying a specification should mean that M plus any environment satisfies that specification.* We will achieve this by considering the behavior of complete systems involving M .

Definition 11 A Moore machine M is called closed if $I = \emptyset$.

Intuitively, the behavior of a closed machine cannot be altered. For such a machine, there is a structure which naturally corresponds to it. We define this structure precisely now. The definition here is actually slightly more general in that it assigns a structure to non-closed machines as well.

Definition 12 (structure for a Moore machine) The structure M' corresponding to a Moore machine M , denoted by $K(M)$, is defined as follows.

1. $S' = S \times 2^I$.
2. $S'_0 = S_0 \times 2^I$.
3. $\mathcal{A}' = \mathcal{A} = I \cup O$.
4. $\mathcal{L}'((s, v)) = \mathcal{L}(s) \cup v$.

5. $R'((s, v_1), (t, v_2))$ iff $R(s, v_1, t)$.
6. $\mathcal{F}' = \emptyset$.

Definition 13 A Moore machine M' is called a closing environment for M if $O \cap O' = \emptyset$, $I \subseteq O'$ and $I' \subseteq O$.

If M' is a closing environment for M , then M and M' can be composed, and the resulting Moore machine will be closed. We now define satisfaction of a formula by a Moore machine.

Definition 14 (satisfaction in a Moore machine) If M is a Moore machine and φ is a \forall CTL* formula with atomic propositions over \mathcal{A} , then $M \models \varphi$ iff for every closing environment M' for M , $K(M \parallel M') \models \varphi$.

We must now demonstrate how to efficiently check whether $M \models \varphi$.

Lemma 2 If M and M' are Moore machines with $O \cap O' = \emptyset$, then $K(M \parallel M')$ is isomorphic to $K(M) \parallel K(M')$.

Proof Define ϕ mapping the states of $K(M \parallel M')$ to the states of $K(M) \parallel K(M')$ as follows.

$$\phi(((s, s'), v)) = ((s, (v \cup \mathcal{L}'(s')) \cap I), (s', (v \cup \mathcal{L}(s)) \cap I'))$$

Suppose $((s, s'), v)$ and $((t, t'), u)$ both map to the same state of $K(M) \parallel K(M')$. Then from the definition of ϕ , we immediately have $s = t$ and $s' = t'$. Also, $(v \cup \mathcal{L}'(s')) \cap I = (u \cup \mathcal{L}'(t')) \cap I$ and $(v \cup \mathcal{L}(s)) \cap I' = (u \cup \mathcal{L}(t)) \cap I'$. By the definition of Moore machine composition, v and u are disjoint from $O \cup O'$. Hence $v \cap I = u \cap I$ and $v \cap I' = u \cap I'$. This implies $v \cap (I \cup I') = u \cap (I \cup I')$, i.e., $v = u$. Hence ϕ is an injection.

To argue that ϕ is surjective, we consider the cardinalities of the two sets of states. First, we have

$$|S^{K(M \parallel M')}| = |S| \cdot |S'| \cdot 2^{|(I \cup I') \setminus (O \cup O')|}.$$

Now consider $|S^{K(M) \parallel K(M')}|$. This is the number of states in the cross product $S^{K(M)} \times S^{K(M')}$ which have compatible labelings. Fix a pair of states s and s' . There are $2^{|I|}$ states in $K(M)$ with s as their first component and $2^{|I'|}$ in $K(M')$ with s' as the first component. Thus there are potentially $2^{|I|} \cdot 2^{|I'|}$ states in $K(M) \parallel K(M')$ corresponding to s and s' . However, each pair must correspond on the atomic propositions in $I \cap O'$, $I' \cap O$, and $I \cap I'$. Thus there are exactly

$$\frac{2^{|I|} \cdot 2^{|I'|}}{2^{|I \cap O'|} \cdot 2^{|I' \cap O|} \cdot 2^{|I \cap I'|}}$$

states in $K(M) \parallel K(M')$ corresponding to s and s' . Thus we have

$$\begin{aligned}
|S^{K(M) \parallel K(M')}| &= \frac{|S| \cdot |S'| \cdot 2^{|I|} \cdot 2^{|I'|}}{2^{|I \cap O'|} \cdot 2^{|I' \cap O|} \cdot 2^{|I \cap I'|}} \\
&= \frac{|S| \cdot |S'| \cdot 2^{|I \cup I'|}}{2^{|I \cap O'|} \cdot 2^{|I' \cap O|}} \\
&= \frac{|S| \cdot |S'| \cdot 2^{|I \cup I'|}}{2^{|I \cap O'|} \cdot 2^{|I \cap O|} \cdot 2^{|I' \cap O|} \cdot 2^{|I' \cap O'|}} \\
&= \frac{|S| \cdot |S'| \cdot 2^{|I \cup I'|}}{2^{|(I \cup I') \cap (O \cup O')|}} \\
&= |S| \cdot |S'| \cdot 2^{|(I \cup I') \setminus (O \cup O')|} \\
&= |S^{K(M) \parallel M'}|.
\end{aligned}$$

Hence ϕ is a bijection.

If $((s_0, s'_0), v)$ is an initial state of $K(M \parallel M')$, then $s_0 \in S_0$ and $s'_0 \in S'_0$. Then $\phi((s_0, s'_0), v)$ is an initial state of $K(M) \parallel K(M')$ since $s_0 \in S_0$ implies $(s_0, (v \cup \mathcal{L}'(s'_0)) \cap I)$ is an initial state of $K(M)$ and $s'_0 \in S'_0$ implies $(s'_0, (v \cup \mathcal{L}(s_0)) \cap I')$ is an initial state of $K(M')$. Similarly, if $\phi((s, s'), v)$ is an initial state of $K(M) \parallel K(M')$, then $((s, s'), v)$ is an initial state of $K(M \parallel M')$.

The sets of atomic propositions of the two structures are clearly identical. The labeling of $((s, s'), v)$ is $\mathcal{L}(s) \cup \mathcal{L}'(s') \cup v$. The labeling of $\phi((s, s'), v)$ is

$$\begin{aligned}
&\mathcal{L}^{K(M)}((s, (v \cup \mathcal{L}'(s')) \cap I)) \cup \mathcal{L}^{K(M')}((s', (v \cup \mathcal{L}(s)) \cap I')) \\
&= \mathcal{L}(s) \cup ((v \cup \mathcal{L}'(s')) \cap I) \cup \mathcal{L}'(s') \cup ((v \cup \mathcal{L}(s)) \cap I') \\
&= \mathcal{L}(s) \cup \mathcal{L}'(s') \cup (v \cap (I \cup I')) \\
&= \mathcal{L}(s) \cup \mathcal{L}'(s') \cup v \\
&= \mathcal{L}^{K(M) \parallel M'}(((s, s'), v)).
\end{aligned}$$

$R^{K(M) \parallel M'}(((s, s'), v), ((t, t'), u))$ iff $R^{M \parallel M'}((s, s'), v, (t, t'))$ iff $R(s, (v \cup \mathcal{L}'(s')) \cap I, t)$ and $R'(s', (v \cup \mathcal{L}(s)) \cap I', t')$ iff $R^{K(M)}((s, (v \cup \mathcal{L}'(s')) \cap I), (t, (u \cup \mathcal{L}'(t')) \cap I))$ and $R^{K(M')}((s', (v \cup \mathcal{L}(s)) \cap I'), (t', (u \cup \mathcal{L}(t)) \cap I'))$ iff $R^{K(M) \parallel K(M')}(\phi(((s, s'), v)), \phi((t, t'), u))$. The fairness sets of both structures are empty. \square

Definition 15 *If M is a Moore machine, the maximal closing environment for M , denoted $E(M)$, is the Moore machine M' defined as follows.*

1. $S' = 2^I$.
2. $S'_0 = S'$.

3. $I' = \emptyset$.
4. $O' = I$.
5. $\mathcal{L}'(s') = s'$.
6. $R'(s', \emptyset, t')$ is identically true.

The maximal environment (for M) represents an environment which can do anything at each step. Intuitively, a possible behavior of M in an arbitrary environment must also be a possible behavior of M in the maximal environment. The logics we use specify properties that should hold for every possible behavior of a system. Hence, if M plus its maximal environment satisfies a formula, then M in any environment should satisfy that formula.

Lemma 3 *Suppose M' is a closing environment for M , and suppose $M'' = E(M)$. Then $K(M') \preceq K(M'')$.*

Proof Define

$$H = \{ (s', s'') \mid \mathcal{L}'(s') \cap \mathcal{A}'' = \mathcal{L}''(s'') \}.$$

Note that for every $s' \in S'$, there is some $s'' \in S''$ such that $H(s', s'')$ (in particular, the state $\mathcal{L}'(s') \cap \mathcal{A}''$ in M''). Thus, if $s'_0 \in S'_0$, there is s''_0 which is related to it by H , and every state in M'' is an initial state.

If $H(s', s'')$, then by the definition of H , we have $\mathcal{L}'(s') \cap \mathcal{A}'' = \mathcal{L}''(s'')$. If π' is a fair path in M' , then the fact that every state in M' is related to some state in M'' plus the fact that R'' is identically true implies that there is a path π'' in M'' such that $H(\pi', \pi'')$. Further, every path in M'' is fair. Thus H is a homomorphism. \square

Lemma 4 *Let M be a Moore machine. Then $K(M)$ is isomorphic to $K(M \parallel E(M))$.*

Proof Let $M' = K(M)$ and $M'' = K(M \parallel E(M))$. Define ϕ mapping the states of M'' to the states of M' by $\phi(((s, v), \emptyset)) = (s, v)$. ϕ is obviously an injection, and ϕ is a surjection since each subset of 2^I is a state of $E(M)$.

If $((s_0, v), \emptyset) \in S''_0$, then s_0 must be in S_0 . Hence $(s_0, v) \in S'_0$. Similarly, if $(s_0, v) \in S'_0$, then $s_0 \in S_0$ and so $((s_0, v), \emptyset) \in S''_0$. \mathcal{A}'' and \mathcal{A}' are trivially equal. We also have

$$\begin{aligned} \mathcal{L}''(((s, v), \emptyset)) &= \mathcal{L}^{M \parallel E(M)}((s, v)) \cup \emptyset \\ &= \mathcal{L}(s) \cup \mathcal{L}^{E(M)}(v) \\ &= \mathcal{L}(s) \cup v \\ &= \mathcal{L}'((s, v)). \end{aligned}$$

Finally, we have $R''(((s, v_1), \emptyset), ((t, v_2), \emptyset))$ iff $R^{M \parallel E(M)}(((s, v_1), \emptyset), (t, v_2))$ iff $R(s, v_1, t)$ iff $R'((s, v_1), (t, v_2))$. \mathcal{F}'' and \mathcal{F}' are both empty. \square

Theorem 4 *If M is a Moore machine, then $M \models \varphi$ iff $K(M) \models \varphi$.*

Proof Suppose $K(M) \models \varphi$. By lemma 4, we find $K(M \parallel E(M)) \models \varphi$, and then by lemma 2, $K(M) \parallel K(E(M)) \models \varphi$. Let M' be any closing environment for M . By lemma 3, $K(M') \preceq K(E(M))$. Hence by theorem 2, $K(M) \parallel K(M') \preceq K(M) \parallel K(E(M))$. Applying corollary 1, we have $K(M) \parallel K(M') \models \varphi$. By lemma 2, $K(M) \parallel K(M')$ is isomorphic to $K(M \parallel M')$, and thus $K(M \parallel M') \models \varphi$. Since M' was arbitrary, $M \models \varphi$.

If $M \models \varphi$, then $K(M \parallel E(M)) \models \varphi$, and hence by lemma 4, $K(M) \models \varphi$. \square

Thus, to determine if a system $M_1 \parallel M_2 \parallel \dots \parallel M_n$ satisfies a formula φ , we instead check that $K(M_1 \parallel M_2 \parallel \dots \parallel M_n)$ satisfies φ . By lemma 2, this is equivalent to checking that $K(M_1) \parallel K(M_2) \parallel \dots \parallel K(M_n)$ satisfies the formula. As illustrated in the previous section, we can use the assume-guarantee paradigm to try to verify this latter relation. Thus, during an actual verification we will be working with structures even though the thing we want to verify is a property of a composition of Moore machines.

6 The tableau construction

In this section, we give a tableau construction for \forall CTL formulas (for a similar construction for LTL, see Burch *et al.* [1]). We show that the tableau of a formula is a maximal model for the formula under the relation \preceq . Thus, the structure generated in the construction can be used as an assumption by composing the structure with the desired system before applying the model checking algorithm. Discharging the assumption is simply a matter of checking that the environment satisfies the formula. We also indicate how the tableau can be used to do temporal reasoning. For the remainder of this section, fix a \forall CTL formula ψ .

Definition 16 *The set $\text{sub}(\varphi)$ of subformulas of the formula φ is defined by the following equations.*

1. If $\varphi = \text{true}$ or $\varphi = \text{false}$ or $\varphi = p$, an atomic proposition, then $\text{sub}(\varphi) = \{\varphi\}$. If $\varphi = \neg p$, a negated atomic proposition, then $\text{sub}(\varphi) = \{\varphi, p\}$.
2. If $\varphi = \varphi_1 \wedge \varphi_2$ or $\varphi = \varphi_1 \vee \varphi_2$, then $\text{sub}(\varphi) = \{\varphi\} \cup \text{sub}(\varphi_1) \cup \text{sub}(\varphi_2)$.
3. (a) If $\varphi = \forall \mathbf{X} \varphi_1$, then $\text{sub}(\varphi) = \{\varphi\} \cup \text{sub}(\varphi_1)$.
 (b) If $\varphi = \forall(\varphi_1 \mathbf{U} \varphi_2)$, then $\text{sub}(\varphi) = \{\varphi\} \cup \text{sub}(\varphi_1) \cup \text{sub}(\varphi_2)$.
 (c) If $\varphi = \forall(\varphi_1 \mathbf{V} \varphi_2)$, then $\text{sub}(\varphi) = \{\varphi\} \cup \text{sub}(\varphi_1) \cup \text{sub}(\varphi_2)$.

Definition 17 *The set $\text{el}(\varphi)$ of elementary formulas of the formula φ is defined by the following equations.*

1. If $\varphi = \text{true}$ or $\varphi = \text{false}$, then $\text{el}(\varphi) = \emptyset$. If $\varphi = p$, an atomic proposition, or $\varphi = \neg p$, then $\text{el}(\varphi) = \{p\}$.
2. If $\varphi = \varphi_1 \wedge \varphi_2$ or $\varphi = \varphi_1 \vee \varphi_2$, then $\text{el}(\varphi) = \text{el}(\varphi_1) \cup \text{el}(\varphi_2)$.

3. (a) If $\varphi = \forall \mathbf{X} \varphi_1$, then $\text{el}(\varphi) = \{\forall \mathbf{X} \varphi_1\} \cup \text{el}(\varphi_1)$.
- (b) If $\varphi = \forall(\varphi_1 \mathbf{U} \varphi_2)$, then $\text{el}(\varphi) = \{\forall \mathbf{X} \text{false}, \forall \mathbf{X} \forall(\varphi_1 \mathbf{U} \varphi_2)\} \cup \text{el}(\varphi_1) \cup \text{el}(\varphi_2)$.
- (c) If $\varphi = \forall(\varphi_1 \mathbf{V} \varphi_2)$, then $\text{el}(\varphi) = \{\forall \mathbf{X} \text{false}, \forall \mathbf{X} \forall(\varphi_1 \mathbf{V} \varphi_2)\} \cup \text{el}(\varphi_1) \cup \text{el}(\varphi_2)$.

The special elementary subformula $\forall \mathbf{X} \text{false}$ denotes the nonexistence of a fair path; $s \models \forall \mathbf{X} \text{false}$ indicates that no fair path begins at s .

Definition 18 (tableau of a formula) *The tableau of ψ , denoted $\mathcal{T}(\psi)$, is the structure $\langle S, S_0, \mathcal{A}, \mathcal{L}, R, \mathcal{F} \rangle$ defined as follows.*

1. $S = 2^{\text{el}(\psi)}$.
2. $S_0 = \Phi(\psi)$, where Φ is the map from $\text{el}(\psi) \cup \text{sub}(\psi) \cup \{\text{true}, \text{false}\}$ to S defined by the following equations.
 - (a) $\Phi(\text{true}) = S$, $\Phi(\text{false}) = \emptyset$. If $\varphi \in \text{el}(\psi)$, then $\Phi(\varphi) = \{s \mid \varphi \in s\}$. If $\varphi = \neg\varphi_1$, then $\Phi(\varphi) = S \setminus \Phi(\varphi_1)$.
 - (b) If $\varphi = \varphi_1 \wedge \varphi_2$, then $\Phi(\varphi) = \Phi(\varphi_1) \cap \Phi(\varphi_2)$. If $\varphi = \varphi_1 \vee \varphi_2$, then $\Phi(\varphi) = \Phi(\varphi_1) \cup \Phi(\varphi_2)$.
 - (c) i. If $\varphi = \forall(\varphi_1 \mathbf{U} \varphi_2)$, then $\Phi(\varphi) = \left(\Phi(\varphi_2) \cup \left(\Phi(\varphi_1) \cap \Phi(\forall \mathbf{X} \varphi) \right) \right) \cup \Phi(\forall \mathbf{X} \text{false})$.
 - ii. If $\varphi = \forall(\varphi_1 \mathbf{V} \varphi_2)$, then $\Phi(\varphi) = \left(\Phi(\varphi_2) \cap \left(\Phi(\varphi_1) \cup \Phi(\forall \mathbf{X} \varphi) \right) \right) \cup \Phi(\forall \mathbf{X} \text{false})$.
3. $\mathcal{A} = \{p \mid p \in \text{el}(\psi)\}$.
4. $\mathcal{L}(s) = \{p \mid p \in s\}$.
5. $R(s, t)$ iff for each formula $\forall \mathbf{X} \varphi$ in $\text{el}(\psi)$, $\forall \mathbf{X} \varphi \in s$ implies $t \in \Phi(\varphi)$.
6. $\mathcal{F} = \left\{ \left(\Phi(\forall \mathbf{X} \forall(\varphi_1 \mathbf{U} \varphi_2)), \Phi(\varphi_2) \right) \mid \forall \mathbf{X} \forall(\varphi_1 \mathbf{U} \varphi_2) \in \text{el}(\psi) \right\}$.

Lemma 5 *For all subformulas φ of ψ , if $s \in \Phi(\varphi)$, then $s \models \varphi$.*

Proof The proof proceeds by induction on the structure of φ .

1. If $\varphi = \text{true}$, then $\Phi(\varphi) = S$, and every state satisfies *true*. If $\varphi = \text{false}$, then $\Phi(\varphi) = \emptyset$, so the result is trivial. If $\varphi = p$, an atomic proposition, then $\Phi(\varphi) = \{s \mid p \in s\}$. But $\mathcal{L}(s) = \{q \mid q \in s\}$, and so $p \in \mathcal{L}(s)$ and $s \models p$. If $\varphi = \neg p$, a negated atomic proposition, then $\Phi(\varphi) = S \setminus \{s \mid p \in s\}$. Since $\mathcal{L}(s) = \{q \mid q \in s\}$, we have that $p \notin \mathcal{L}(s)$ and so $s \models \neg p$.
2. If $\varphi = \varphi_1 \wedge \varphi_2$, then $\Phi(\varphi) = \Phi(\varphi_1) \cap \Phi(\varphi_2)$, and hence $s \in \Phi(\varphi_1)$ and $s \in \Phi(\varphi_2)$. By the induction hypothesis, $s \models \varphi_1$ and $s \models \varphi_2$, which implies $s \models \varphi_1 \wedge \varphi_2$. The case where $\varphi = \varphi_1 \vee \varphi_2$ is similar.
3. (a) If $\varphi = \forall \mathbf{X} \varphi_1$, then $\Phi(\varphi) = \{s \mid \forall \mathbf{X}(\varphi_1) \in s\}$, and so $\forall \mathbf{X}(\varphi_1) \in s$. Suppose $R(s, t)$. By the definition of R , we have $t \in \Phi(\varphi_1)$, and then the induction hypothesis implies $t \models \varphi_1$. Since t was chosen arbitrarily, any fair path from s satisfies φ_1 at its second state, and hence $s \models \forall \mathbf{X}(\varphi_1)$.

- (b) If $\varphi = \forall(\varphi_1 \mathbf{U} \varphi_2)$, then $\Phi(\varphi) = \left(\Phi(\varphi_2) \cup \left(\Phi(\varphi_1) \cap \Phi(\forall \mathbf{X} \varphi) \right) \right) \cup \Phi(\forall \mathbf{X} \text{false})$. Let t be any state in $\Phi(\varphi)$. Then either
- i. $t \in \Phi(\forall \mathbf{X} \text{false})$, in which case t has no successors and $t \models \varphi$ trivially, or
 - ii. $t \in \Phi(\varphi_2)$, in which case the induction hypothesis implies $t \models \varphi_2$, or
 - iii. $t \in \Phi(\varphi_1) \cap \Phi(\forall \mathbf{X} \varphi)$. In this case, the induction hypothesis implies $t \models \varphi_1$.
By the definition of R , we also know that if $R(t, u)$, then $u \in \Phi(\varphi)$.

Let $s = s_0$, and consider a fair path $\pi = s_0 s_1 s_2 \dots$ from s . Note that no state on this path can satisfy the first condition above. There are two cases to consider.

- i. There is some j such that $s_j \models \varphi_2$. Let s_i be the first such state on the path.
By the above, for every $j < i$, $s_j \models \varphi_1$. Hence the path satisfies $\varphi_1 \mathbf{U} \varphi_2$.
- ii. For every j , $s_j \not\models \varphi_2$. Then the above implies that for every j , $s_j \in \Phi(\forall \mathbf{X} \varphi)$.
By the induction hypothesis, we know that each s_j is not in $\Phi(\varphi_2)$. But then $\text{inf}(\pi) \cap \Phi(\forall \mathbf{X} \varphi) \neq \emptyset$ and $\text{inf}(\pi) \cap \Phi(\varphi_2) = \emptyset$. By the definition of \mathcal{F} , this contradicts the fact that π is fair, and so this case is impossible.

Thus $s \models \forall(\varphi_1 \mathbf{U} \varphi_2)$.

- (c) If $\varphi = \forall(\varphi_1 \mathbf{V} \varphi_2)$, then $\Phi(\varphi) = \left(\Phi(\varphi_2) \cap \left(\Phi(\varphi_1) \cup \Phi(\forall \mathbf{X} \varphi) \right) \right) \cup \Phi(\forall \mathbf{X} \text{false})$. If t is any state in $\Phi(\varphi)$, then either
- i. $t \in \Phi(\forall \mathbf{X} \text{false})$, in which case t has no successors and $t \models \varphi$ trivially, or
 - ii. $t \in \Phi(\varphi_2)$. In this case, we also have either $t \in \Phi(\varphi_1)$ or for every u such that $R(t, u)$, $u \in \Phi(\varphi)$.

Let $s = s_0$, and let $\pi = s_0 s_1 s_2 \dots$ be a fair path from s . Note that no s_i can satisfy the first condition above. If s_i is such that for all $j < i$, $s_j \not\models \varphi_1$, then the induction hypothesis implies that $s_j \notin \Phi(\varphi_1)$. Hence $s_j \in \Phi(\varphi_2)$, and also $s_i \in \Phi(\varphi_2)$; then the induction hypothesis implies for all $j \leq i$, $s_j \models \varphi_2$. Thus the path satisfies $\varphi_1 \mathbf{V} \varphi_2$, and hence we have $s \models \forall(\varphi_1 \mathbf{V} \varphi_2)$. \square

Now let $M = \mathcal{T}(\psi)$, and fix a structure M' .

Lemma 6 Define a relation $H \subseteq S' \times S$ by

$$H = \left\{ (s', s) \mid s = \{ \varphi \mid \varphi \in \text{el}(\psi), s' \models \varphi \} \right\}.$$

If $H(s', s)$, then for every subformula or elementary formula φ of ψ , $s' \models \varphi$ implies $s \in \Phi(\varphi)$.

Proof The proof proceeds by induction on the structure of φ , where the base cases for the induction are the elementary subformulas of ψ , plus *true* and *false*.

1. If $\varphi = \text{true}$, then $\Phi(\varphi) = S$, so the result is trivial. If $\varphi = \text{false}$, then s' cannot satisfy φ . If $\varphi \in \text{el}(\psi)$, then by the definition of H , $s' \models \varphi$ implies $\varphi \in s$. Now $\Phi(\varphi) = \{ s \mid \varphi \in s \}$, so $s \in \Phi(\varphi)$.
2. If $\varphi = \neg p$, a negated atomic proposition, then $s' \models \varphi$ implies $p \notin s$. Since $\Phi(\varphi) = S \setminus \{ s \mid p \in s \}$, $s \in \Phi(\varphi)$.

3. If $\varphi = \varphi_1 \wedge \varphi_2$, then $\Phi(\varphi) = \Phi(\varphi_1) \cap \Phi(\varphi_2)$. We have $s' \models \varphi$ implies $s' \models \varphi_1$ and $s' \models \varphi_2$. By the induction hypothesis, $s \in \Phi(\varphi_1)$ and $s \in \Phi(\varphi_2)$, and so $s \in \Phi(\varphi_1) \cap \Phi(\varphi_2)$. The case when $\varphi = \varphi_1 \vee \varphi_2$ is similar.
4. If $\varphi = \forall(\varphi_1 \mathbf{U} \varphi_2)$, then $\Phi(\varphi) = (\Phi(\varphi_2) \cup (\Phi(\varphi_1) \cap \Phi(\forall \mathbf{X} \varphi))) \cup \Phi(\forall \mathbf{X} \text{false})$. Given $s' \models \varphi$, there are three cases.
 - (a) If no fair paths start at s' , then $s' \models \forall \mathbf{X} \text{false}$. The induction hypothesis implies $s \in \Phi(\forall \mathbf{X} \text{false})$, and so $s \in \Phi(\varphi)$.
 - (b) If $s' \models \varphi_2$, then by the induction hypothesis, $s \in \Phi(\varphi_2)$, and so $s \in \Phi(\varphi)$.
 - (c) Otherwise, $s' \models \varphi_1$ and $s' \models \forall \mathbf{X} \varphi$. By the induction hypothesis, $s \in \Phi(\varphi_1)$ and $s \in \Phi(\forall \mathbf{X} \varphi)$ (since $\forall \mathbf{X} \varphi \in \text{el}(\psi)$). Hence $s \in \Phi(\varphi)$.

In all cases, $s \in \Phi(\forall(\varphi_1 \mathbf{U} \varphi_2))$.

5. If $\varphi = \forall(\varphi_1 \mathbf{V} \varphi_2)$, then $\Phi(\varphi) = (\Phi(\varphi_2) \cap (\Phi(\varphi_1) \cup \Phi(\forall \mathbf{X} \varphi))) \cup \Phi(\forall \mathbf{X} \text{false})$. Since $s' \models \varphi$, either
 - (a) no fair paths start at s' , in which case $s' \models \forall \mathbf{X} \text{false}$ and the induction hypothesis implies $s \in \Phi(\varphi)$, or
 - (b) $s' \models \varphi_2$, and so by the induction hypothesis, $s \in \Phi(\varphi_2)$. Also, either $s' \models \varphi_1$ or $s' \models \forall \mathbf{X} \varphi$. Applying the induction hypothesis again, either $s \in \Phi(\varphi_1)$ or $s \in \Phi(\forall \mathbf{X} \varphi)$. In both cases, $s \in \Phi(\varphi)$.

Thus in all cases, $s \in \Phi(\forall(\varphi_1 \mathbf{V} \varphi_2))$. □

Lemma 7 *The relation H given above is a homomorphism.*

Proof Note that for every state s' of M' , there is a (single) state s of M such that $H(s', s)$. Let A be the set of atomic propositions for M , and assume $H(s', s)$. We have $\mathcal{L}(s) = \{p \mid p \in s\}$. From the definition of H , $p \in s$ implies $s' \models p$. Further, if $s' \models p$ and $p \in A$, then $p \in \text{el}(\psi)$, and hence $p \in s$, $p \in \mathcal{L}(s)$. Thus we find $\mathcal{L}'(s') \cap A = \mathcal{L}(s)$.

Let $s'_0 = s'$, and suppose $\pi' = s'_0 s'_1 s'_2 \dots$ is a fair path from s' . Let $\forall \mathbf{X} \varphi_1, \forall \mathbf{X} \varphi_2, \dots, \forall \mathbf{X} \varphi_n$ be all the formulas of the form $\forall \mathbf{X} \varphi$ in $\text{el}(\psi)$ which s' satisfies. Then we have $s'_1 \models \varphi_1, s'_1 \models \varphi_2, \dots, s'_1 \models \varphi_n$. Let s_1 be the state of M related to s'_1 by H . By the previous lemma, $s_1 \in \Phi(\varphi_1), s_1 \in \Phi(\varphi_2), \dots, s_1 \in \Phi(\varphi_n)$. Now by the definition of H , the formulas of the form $\forall \mathbf{X} \varphi$ in s must be exactly $\forall \mathbf{X} \varphi_1, \forall \mathbf{X} \varphi_2, \dots, \forall \mathbf{X} \varphi_n$. Then from the definition of R , we see that $R(s, s_1)$. Since $H(s'_1, s_1)$, we can continue the process. Defining $s_0 = s$, we get a sequence of states $\pi = s_0 s_1 s_2 \dots$ starting at s such that $H(s'_i, s_i)$ for all i . To complete the proof, we must show that this sequence is fair.

Assume that π is not fair. Looking at \mathcal{F} , we see that there must be some elementary subformula $\forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b)$ such that $\text{inf}(\pi) \cap \Phi(\forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b)) \neq \emptyset$ and $\text{inf}(\pi) \cap \Phi(\varphi_b) = \emptyset$. Consider one of the states s_i . $s_i \in \Phi(\forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b))$ iff $\forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b) \in s_i$, and then the definition of H implies $s'_i \models \forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b)$. In addition, the previous lemma implies that if $s_i \notin \Phi(\varphi_b)$, then $s'_i \not\models \varphi_b$. Choose i so that $s_i \in \Phi(\forall \mathbf{X} \forall(\varphi_a \mathbf{U} \varphi_b))$ and so that for all $j \geq i$,

$s_j \notin \Phi(\varphi_b)$. Then $s'_i s'_{i+1} \dots$ is a fair path in M' starting at s'_i , and every state on this path satisfies $\neg\varphi_b$. But $s'_i \models \forall \mathbf{X} \forall (\varphi_a \mathbf{U} \varphi_b)$, a contradiction. Hence π is in fact a fair path in M . \square

Theorem 5 $M' \models \psi$ iff $M' \preceq \mathcal{T}(\psi)$.

Proof Suppose $M' \preceq \mathcal{T}(\psi)$. By lemma 5 and the definition of the tableau, every initial state of $\mathcal{T}(\psi)$ satisfies ψ , i.e., $\mathcal{T}(\psi) \models \psi$. Then since $M' \preceq \mathcal{T}(\psi)$, $M' \models \psi$.

If $M' \models \psi$, then by definition, every $s'_0 \in S'_0$ satisfies ψ . By the definition of H , every such s'_0 is paired with a (unique) s_0 . Lemma 6 implies that $s_0 \in \Phi(\psi)$, and by the definition of the tableau, $s_0 \in S_0$. By lemma 7, H is a homomorphism, so $M' \preceq \mathcal{T}(\psi)$. \square

The tableau construction can also be used to reason about formulas. We are typically interested in whether every model of a formula φ is also a model of some other formula ψ . Let $\varphi \models \psi$ denote this semantic relation.

Proposition 1 $\varphi \models \psi$ iff $\mathcal{T}(\varphi) \models \psi$.

Proof If $\varphi \models \psi$, then every model of φ , in particular $\mathcal{T}(\varphi)$, is also a model of ψ . Assume $\mathcal{T}(\varphi) \models \psi$, and let $M \models \varphi$. By the previous corollary, $M \preceq \mathcal{T}(\varphi)$. Since $\mathcal{T}(\varphi) \models \psi$, $\mathcal{T}(\varphi) \preceq \mathcal{T}(\psi)$. Hence $M \preceq \mathcal{T}(\psi)$, i.e., $M \models \psi$. \square

We will sometimes extend the set of elementary formulas of a formula by adding additional atomic propositions. For example, if we wished to check whether *true* implied p , we would extend the set of atomic propositions for *true* to include p (another way to view this is to imagine rewriting *true* as $true \wedge (p \vee \neg p)$). The formula ψ has a nontrivial model iff it is not the case that $\psi \models \forall \mathbf{X} \text{false}$. ψ is true in every model iff $true \models \psi$.

7 Checking for homomorphism

In this section, we discuss the problem of determining whether there exists a homomorphism between two structures M and M' . Our goal is to efficiently determine if $M \preceq M'$. First note that if H_1 and H_2 are homomorphisms, then $H_1 \cup H_2$ is a homomorphism. Also, \emptyset is trivially a homomorphism. These facts imply that there is a maximal homomorphism under set inclusion. We will actually give an algorithm for computing this maximal homomorphism.

We also note the following facts.

1. If s is a state of M and no fair paths start at s , then s is homomorphic to exactly those states s' in M' for which $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$.
2. If s' is a state of M' and no fair paths start at s' , then s' is homomorphic exactly to those states s in M which are the start of no fair path and for which $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$.

States which are the start of no fair path can be detected in polynomial time [12] and eliminated in a preprocessing step. Hence, without loss of generality, we can assume that every state in M and M' is the start of some fair path. We now describe polynomial time algorithms for checking the preorder in several important special cases.

Suppose that M' has a trivial acceptance condition, i.e., $\mathcal{F}' = \emptyset$.

Definition 19 Define a sequence of relations H_i as follows.

1. $H_0 = \{ (s, s') \mid \mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s') \}$
2. $H_{i+1} = H_i \cap \{ (s, s') \mid \forall t [R(s, t) \rightarrow \exists t' (R'(s', t') \wedge H_i(t, t'))] \}$

Define H_ω to be the first H_i such that $H_i = H_{i+1}$ (such an i exists since $H_{j+1} \subseteq H_j$ for all j and each H_j is finite).

Theorem 6 For every $s \in S$ and $s' \in S'$, $s \preceq s'$ iff $H_\omega(s, s')$.

Proof We first note that H_ω is the greatest fixed point of the equation

$$H = H \cap \{ (s, s') \mid \mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s') \wedge \forall t [R(s, t) \rightarrow \exists t' (R'(s', t') \wedge H(t, t'))] \}.$$

Suppose s and s' are states such that $H_\omega(s, s')$. We have $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$ immediately. Let $s_0 = s$ and $s'_0 = s'$, and assume $\pi = s_0 s_1 \dots$ is a fair path starting from s . From the above equation, there exists a state s'_1 such that $R'(s'_0, s'_1)$ and $H_\omega(s_1, s'_1)$. Continuing in this fashion, we find a path $s'_0 s'_1 \dots$ starting from s' such that $H_\omega(s_i, s'_i)$ for all i . Since $\mathcal{F}' = \emptyset$, this path is fair. Hence H_ω is a homomorphism from s to s' , i.e., $s \preceq s'$.

To show that $s \preceq s'$ implies $H_\omega(s, s')$, we show that any homomorphism H is a fixed point of the above equation. Since H_ω is the greatest fixed point, we will have $H \subseteq H_\omega$. Hence if there is some homomorphism H such that $H(s, s')$, then $H_\omega(s, s')$. It is enough to show that H is a subset of the set

$$\{ (s, s') \mid \mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s') \wedge \forall t [R(s, t) \rightarrow \exists t' (R'(s', t') \wedge H(t, t'))] \}.$$

If $H(s, s')$, then we have $\mathcal{L}(s) \cap \mathcal{A}' = \mathcal{L}'(s')$. If $R(s, t)$, then by our earlier assumption, there exists a fair path from t . Hence, letting $s_0 = s$ and $s_1 = t$, there is some fair path $s_0 s_1 \dots$ from s through t . Since $H(s, s')$, there exists a fair path $s'_0 s'_1 \dots$ from $s' = s'_0$ such that $H(s_i, s'_i)$ for all i . Now if we take $t' = s'_1$, we see that (s, s') is in the above set. \square

We note that $H_\omega = H_i$ for some i which is at most $|S| \cdot |S'|$. Each H_{j+1} can also be computed in polynomial time from H_j ; hence H_ω can be computed in polynomial time.

Another important case is when M' is deterministic, i.e., if $R'(s', t')$ and $R'(s', u')$, then $\mathcal{L}'(t') \neq \mathcal{L}'(u')$. For this case, $s \preceq s'$ iff the language of s is contained in the language of s' (the language for a state s is the set of sequences of labelings which occur along the fair paths starting at s). This relation can be checked in polynomial time using the techniques of Clarke, Draghicescu and Kurshan [2].

Finally, if M' is the result of a tableau construction, say $M' = \mathcal{T}(\psi)$, then as shown in the previous section, checking whether $M \preceq M'$ reduces to the problem of checking whether $M \models \psi$.

8 An example

We have implemented a BDD-based model checker based on the theory developed in the previous sections. The model checker is written in a combination of T (Yale’s dialect of Scheme) and C. It includes facilities for model checking, temporal reasoning (via the tableau construction), and checking for homomorphism. To illustrate the system, we use the controller of a simple CPU as an example. The controller is written in a state machine description language called CSML [5] which is compiled into Moore machines. We give only a brief description of the CPU here; Clarke, Long and McMillan [5] give details. The CPU is a simple stack-based machine, i.e., part of the CPU’s memory contains a stack from which instruction operands are popped and onto which results are pushed. There are two parts to the CPU controller. The first part is called the access unit and is responsible for all the CPU’s memory references. The second part, called the execution unit, interprets the instructions and controls the arithmetic unit, shifter, etc. These two parts operate in parallel. The access unit and execution unit communicate via a small number of signals. Three of the signals, *push*, *pop* and *fetch*, are inputs of the access unit and indicate that the execution unit wants to push or pop something from the stack or to get the next instruction. For each of these signals there is a corresponding ready output from the access unit. The execution unit must wait for the appropriate ready signal before proceeding. One additional signal, *branch*, is asserted by the execution unit when it wants to jump to a new program location.

In order to increase performance, the access unit attempts to keep the value on the top of the stack in a special register called the TS register. The goal is to keep the execution unit from having to wait for the memory. For example, when the TS register contains valid data, a pop operation can proceed immediately. In addition, when a value is pushed on the stack, it is moved into this register and copied to memory at some later point. The access unit also loads instructions into a queue when possible so that fetches do not require waiting for the memory. This queue is flushed whenever the CPU branches.

Clarke, Long and McMillan gave a number of correctness conditions for the controller. We demonstrate here how these formulas can be verified in a compositional fashion. From the form of the conditions, we divide them into three classes. The first class consists of simple safety properties of the access unit. For example, one of these formulas is

$$\forall \mathbf{G}(sptomema \rightarrow tsload \vee tsstore),$$


which states that if the access unit outputs the top-of-stack pointer as a memory address, then it is either reading or writing the TS register. The model checker verified that each of these properties held for the access unit alone. Hence, they hold in any system containing the access unit.

The conditions in the second class are slightly more complex. These properties are safety properties which specify what sequences of operations are allowed. For example, one condition is

$$\forall \mathbf{G}\left(\textit{pushed} \rightarrow \forall \mathbf{X}\forall (tsstored \vee popped \mathbf{V} \neg(\textit{pushed} \vee tsload))\right).$$

Here *pushed* is an abbreviation for *push* \wedge *pushrdy*, and *popped* abbreviates *pop* \wedge *poprdy*. The formula asserts that if a push operation is completed, then another push cannot be

completed and the access unit cannot attempt to load the TS register from memory until either a pop occurs or the TS register is stored on the stack. In other words, once the TS register contains a value which needs to be pushed on the stack, the CPU cannot do anything that would destroy this value until the value is either used or successfully stored in memory. Since all of the properties in this class essentially specify when the access unit may assert its ready signals, it is tempting to check whether they hold for the access unit alone as well. This is not possible, however, because the properties also depend on how the memory acknowledgment signal behaves. To verify these properties, we made a simple model of the memory (see figure 1). For conciseness, the figure shows a Moore machine; the actual model used is obtained by adding the fairness constraint shown in the figure to the structure corresponding to this Moore machine. All of the properties in this class except for one turn out to be true in the system composed of the access unit and this model of the memory. The exception is an analog to the previous formula that deals with what occurs after a pop. The counterexample produced by the model checker for this formula showed that the formula was false because a push and a pop could occur simultaneously. When we examined the access unit, we saw that it had been designed assuming that these operations would be mutually exclusive. The formula turns out to be true with the additional assumption $\forall \mathbf{G}(\neg push \vee \neg pop)$. The model checker verified this by building the tableau for this assumption, composing it with the access unit and memory model, and checking the formula.



\mathcal{F} is defined by
 $\mathbf{GF}(memrd \vee memwr \rightarrow memack)$

Figure 1: Memory abstraction

The final class of criteria consists of a single liveness property: $\forall \mathbf{G} \forall \mathbf{F}(fetch \wedge fetchrdy)$. This formula states that the CPU always fetches another instruction. We demonstrate two different ways of verifying this property.

One way is to observe that for this formula to be true, it must obviously be the case that the memory responds to requests eventually and that the execution unit does not execute infinite sequences of pushes, pops and branches. The memory model already has a fairness constraint ensuring the first of these, but there is nothing to guarantee the second. We can take care of this by using a simple model of the execution unit (see figure 2). Again, the actual model is the structure derived from the Moore machine, plus the indicated fairness constraint. The output *idle* in this figure is an abbreviation for $\neg(push \vee pop \vee fetch \vee branch)$.

The model checker verified that the access unit plus the models of the execution unit and the memory satisfied the above formula. It also verified that there was a homomorphism between the (structure for the) actual execution unit and the model. Thus, we can conclude that this formula holds in the final system provided there is a homomorphism from the actual memory to our model. We also checked that the execution unit model satisfied the assumption $\forall \mathbf{G}(\neg \text{push} \vee \neg \text{pop})$ used above. Since there is a homomorphism from the execution unit to the model, we know that the execution unit must satisfy this assumption as well. This final step allows us to conclude that the composition of the access and execution units satisfies the entire specification provided the memory is homomorphic to the model we used.

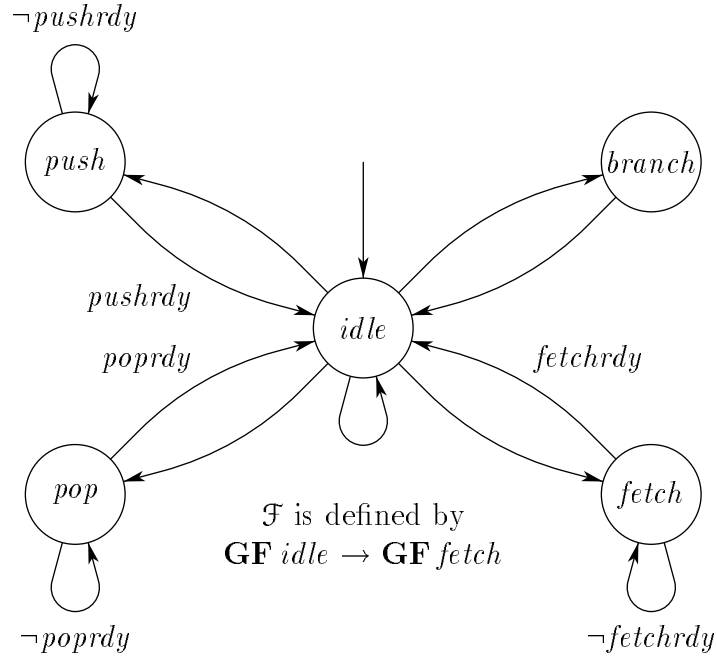


Figure 2: Execution unit abstraction

We can also verify the final property using a series of \forall CTL assumptions. The idea will be to check the property for the execution unit. In order for the formula to be true, the access unit must eventually respond to push and pop requests and must fill the instruction queue when appropriate. We can only guarantee that the access unit meets these conditions if we know that the execution unit does not try to do two operations at once and that it will not remove a request before the corresponding operation can complete. We begin with these properties.

$$\forall \mathbf{G}(\neg(\text{fetch} \wedge \text{push}) \wedge \neg(\text{fetch} \wedge \text{pop}) \wedge \dots \wedge \neg(\text{pop} \wedge \text{branch})) \quad (1)$$

$$\forall \mathbf{G}(\text{push} \rightarrow \forall(\text{pushed} \mathbf{V} \text{push})) \quad (2)$$

$$\forall \mathbf{G}(\text{pop} \rightarrow \forall(\text{popped} \mathbf{V} \text{pop})) \quad (3)$$

The first of these specifies that every pair of operations the execution unit can perform are mutually exclusive. The other two formulas state that if the execution unit makes a push or pop request, then it does not deassert the request until the operation completes. The model checker verified that these properties hold in the execution unit alone, and (using the tableau construction) that the first property implies the assumption $\forall \mathbf{G}(\neg \textit{push} \vee \neg \textit{pop})$ used above. Now using formulas 1 and 2 as assumptions, we checked that the system composed of the access unit and the memory model satisfied the formula

$$\forall \mathbf{G}(\textit{push} \rightarrow \forall(\textit{push} \mathbf{U} \textit{pushed})). \quad (4)$$

This specification states that every push operation will be completed. Similarly, using formulas 1 and 3 as assumptions, we verified

$$\forall \mathbf{G}(\textit{pop} \rightarrow \forall(\textit{pop} \mathbf{U} \textit{popped})). \quad (5)$$

The system composed of the access unit and the memory model also satisfies the formula $\forall \mathbf{G} \forall \mathbf{F}(\textit{fetchrdy} \vee \textit{branch})$ (at any point, either the access unit will eventually fill the instruction queue or a branch will occur). Finally, using this formula and formulas 4 and 5 as assumptions, the model checker verified that the execution unit satisfies $\forall \mathbf{G} \forall \mathbf{F}(\textit{fetch} \wedge \textit{fetchrdy})$. (Again, to complete the verification we would have to demonstrate a homomorphism between the actual memory and our model of it.)

9 Conclusion

We have identified a subset, $\forall \text{CTL}^*$, of CTL^* which is appropriate for compositional reasoning. For this subset, satisfaction is preserved under composition; hence a standard model checking algorithm can be used to answer the question: Is a formula true for all systems containing a specified component? We have also proposed a preorder \preceq which is appropriate for $\forall \text{CTL}^*$. The preorder captures the relation between a component and a system containing that component. It provides the basis for using an assume-guarantee style of reasoning with the logic. Assumptions which are given as structures are discharged by checking the preorder. We have given a tableau construction for the $\forall \text{CTL}$ subset of $\forall \text{CTL}^*$. Satisfaction of a $\forall \text{CTL}$ formula corresponds to being below the tableau of the formula in the preorder. The construction makes it possible to use $\forall \text{CTL}$ formulas as assumptions and to do temporal reasoning. $\forall \text{CTL}$ also has an efficient model checking algorithm. We have implemented a symbolic verification system based on these results and have used it to verify some nontrivial systems in a compositional fashion.

There are several directions for future work. Intuitively, the $\forall \text{CTL}^*$ subset of CTL^* should be maximal in the sense that any formula for which satisfaction is preserved under composition should be equivalent to a formula of $\forall \text{CTL}^*$, but we have not proved this. Another idea is to look at different logics with the same flavor, such as $\forall \text{CTL}^*$ extended with automata operators or the μ -calculus with only $[\cdot]$ modalities. It would also be interesting to try to extend the tableau construction of section 6 to all of $\forall \text{CTL}^*$. In order to accomplish this however, it will almost certainly be necessary to use a more complex type of structure than that given in definition 2. Another question is whether it is possible to apply our ideas to branching-time logics with existential path quantifiers. For example, is there a reasonable

algorithm which will determine whether a CTL formula is true in all systems containing a given component? It is fairly easy to come up with algorithms which are sound, but completeness seems more difficult to achieve. We also wish to examine the problem of efficiently checking the preorder for arbitrary structures. Finally, it is essential to try to apply the compositional reasoning methods we have considered to more complex systems in order to evaluate the techniques.

References

- [1] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and J. Hwang. Symbolic model checking: 10^{20} states and beyond. In LICS90 [21].
- [2] E. M. Clarke, I. A. Draghicescu, and R. P. Kurshan. A unified approach for showing language containment and equivalence between various types of ω -automata. In A. Arnold and N. D. Jones, editors, *Proceedings of the 15th Colloquium on Trees in Algebra and Programming*, volume 407 of *Lecture Notes in Computer Science*. Springer-Verlag, May 1990.
- [3] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
- [4] E. M. Clarke, D. E. Long, and K. L. McMillan. Compositional model checking. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1989.
- [5] E. M. Clarke, D. E. Long, and K. L. McMillan. A language for compositional specification and verification of finite state hardware controllers. In J. A. Darringer and F. J. Rammig, editors, *Proceedings of the Ninth International Symposium on Computer Hardware Description Languages and their Applications*. North-Holland, June 1989.
- [6] R. Cleaveland. Tableau-based model checking in the propositional mu-calculus. *Acta Informatica*, 27:725–747, 1990.
- [7] R. Cleaveland and B. Steffen. When is “partial” adequate? a logic-based proof technique using partial specifications. In LICS90 [21].
- [8] O. Coudert, C. Berthet, and J. C. Madre. Verifying temporal properties of sequential machines without building their state diagrams. In Kurshan and Clarke [17].
- [9] J. W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors. *Proceedings of the REX Workshop on Stepwise Refinement of Distributed Systems, Models, Formalisms, Correctness*, volume 430 of *Lecture Notes in Computer Science*. Springer-Verlag, May 1989.
- [10] D. L. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-Independent Circuits*. ACM Distinguished Dissertations. MIT Press, 1989.
- [11] E. A. Emerson and J. Y. Halpern. “Sometimes” and “Not Never” revisited: On branching time versus linear time. *Journal of the ACM*, 33:151–178, 1986.
- [12] E. A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional mu-calculus. In *Proceedings of the Second Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1986.

- [13] S. Graf and B. Steffen. Compositional minimization of finite state processes. In Kurshan and Clarke [17].
- [14] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [15] B. Josko. Verifying the correctness of AADL-modules using model checking. In de Bakker et al. [9].
- [16] R. P. Kurshan. Analysis of discrete event coordination. In de Bakker et al. [9].
- [17] R. P. Kurshan and E. M. Clarke, editors. *Proceedings of the 1990 Workshop on Computer-Aided Verification*, June 1990.
- [18] R. P. Kurshan and K. L. McMillan. A structural induction theorem for processes. In *Proceedings of the Eighth Annual ACM Symposium on Principles of Distributed Computing*. ACM Press, August 1989.
- [19] K. G. Larsen. The expressive power of implicit specifications. To appear in Proceedings of the Eighteenth International Colloquium on Automata, Languages, and Programming.
- [20] O. Lichtenstein and A. Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Proceedings of the Twelfth Annual ACM Symposium on Principles of Programming Languages*, January 1985.
- [21] *Proceedings of the Fifth Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1990.
- [22] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, 1980.
- [23] A. Pnueli. In transition for global to modular temporal reasoning about programs. In K. R. Apt, editor, *Logics and Models of Concurrent Systems*, volume 13 of *NATO ASI series. Series F, Computer and system sciences*. Springer-Verlag, 1984.
- [24] Z. Shtadler and O. Grumberg. Network grammars, communication behaviors and automatic verification. In J. Sifakis, editor, *Proceedings of the 1989 International Workshop on Automatic Verification Methods for Finite State Systems, Grenoble, France*, volume 407 of *Lecture Notes in Computer Science*. Springer-Verlag, June 1989.
- [25] G. Shurek and O. Grumberg. The modular framework of computer-aided verification: Motivation, solutions and evaluation criteria. In Kurshan and Clarke [17].
- [26] C. Stirling and D. J. Walker. Local model checking in the modal mu-calculus. In J. Diaz and F. Orejas, editors, *Proceedings of the 1989 International Joint Conference on Theory and Practice of Software Development*, volume 351–352 of *Lecture Notes in Computer Science*. Springer-Verlag, March 1989.
- [27] D. Walker. Bisimulations and divergence. In *Proceedings of the Third Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1988.
- [28] G. Winskel. Compositional checking of validity on finite state processes. Draft copy.

- [29] G. Winskel. Model checking in the modal ν -calculus. In *Proceedings of the Sixteenth International Colloquium on Automata, Languages, and Programming*, 1989.