Efficient Generation of Counterexamples and Witnesses in Symbolic Model Checking

E. Clarke O. Grumberg* K. McMillan† X. Zhao

October 1994
CMU-CS-94-204

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

* Computer Science Department
The Technion
Haifa, 32000 Israel

† Cadence Berkeley Laboratories
1919 Addison Street, Ste. 303
Berkeley, CA 94704-1144

This research was sponsored in part by the National Science Foundation under grant no. CCR-8722633, by the Semiconductor Research Corporation under contract 92-DJ-294, and by The Defense Advanced Research Projects Agency, Information Science and Technology Office, under the title "Research on Parallel Computing", ARPA Order No. 7330, issued by DARPA/CMO under Contract MDA972-90-C-0035.

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 NSF, SRC, ARPA, or the U.S. government.

Keywords: automatic verification, temporal logic, model checking, binary decision diagrams, counterexamples

Abstract

Model checking is an automatic technique for verifying sequential circuit designs and protocols. An efficient search procedure is used to determine whether or not the specification is satisfied. If it is not satisfied, our technique will produce a counterexample execution trace that shows the cause of the problem. Although finding counterexamples is extremely important, there is no description of how to do this in the literature on model checking. We describe an efficient algorithm to produce counterexamples and witnesses for symbolic model checking algorithms. This algorithm is used in the SMV model checker and works quite well in practice. We also discuss how to extend our technique to more complicated specifications. This extension makes it possible to find counterexamples for verification procedures based on showing language containment between various types of ω-automata.

1. Introduction

Complex state-transition systems occur frequently in the design of sequential circuits and protocols. During the past ten years, researchers at Carnegie Mellon University have developed an alternative approach to verification called *temporal logic model checking* [5, 6]. In this approach specifications are expressed in a propositional temporal logic, and circuit designs and protocols are modeled as state-transition systems. An efficient search procedure is used to determine automatically if the specifications are satisfied by the transition systems.

One of the most important advantages of model checking over mechanical theorem provers or proof checkers for verification of circuits and protocols is its *counterexample facility*. Typically, the user provides a high level representation of the model and the specification to be checked. The model checking algorithm either terminates with the answer *true*, indicating that the model satisfies the specification, or gives a counterexample execution that shows why the formula is not satisfied. The counterexamples can be essential in finding subtle errors in complex designs.

The main disadvantage of model checking is the state explosion which can occur if the system being verified has many components that can make transitions in parallel. Recently, the size of the transition systems that can be verified by model checking techniques has increased dramatically after the introduction of *ordered binary decision diagrams* (OBDDs) [2]. By applying this technique, verification of systems that have more than 10^{100} states has become possible [3, 11]. However, finding counterexamples is significantly more difficult when OBDDs are used in model checking instead of explicit state enumeration techniques, especially when fairness constraints are involved.

Although finding counterexamples is extremely important, as far as we know, there is no description of how to do this in the literature on model checking. In this paper, we describe an efficient algorithm to produce counterexamples and witnesses for model checking algorithms. The algorithm is, in fact, the one that is used in the SMV model checker developed at Carnegie Mellon [11] and works quite well in practice. We show how the counterexample facility can be used to debug a subtle asynchronous circuit design. We also discuss how to extend our technique to more complicated temporal formulas. This extension makes it possible to find counterexamples for verification procedures based on showing language containment between various types of ω-automata.

This paper is organized as follows: The properties of OBDDs that we need are briefly discussed in Section 2. The next section describes the temporal logic CTL that we use for specifying properties of sequential circuits and protocols. Section 4 explains the symbolic model checking algorithm for CTL, and Section 5 shows how fairness constraints can be handled. Section 6 is the main section of the paper. We describe how counterexamples and witnesses are generated. We also give an example that shows how this facility can be used in sequential circuit verification. In the next section we extend the counterexample facility to a wider class of temporal properties. Section 8 describes how our techniques can be used to generate counterexamples for verification procedures that are based on showing inclusion between ω-automata. The paper concludes in Section 9 with a discussion of possible directions for future research.

2. Binary Decision Diagrams

Ordered binary decision diagrams (OBDDs) are a canonical form representation for boolean formulas [2]. They are often substantially more compact than traditional normal forms such as conjunctive normal form and disjunctive normal form, and they can be manipulated very efficiently. An OBDD is similar to a binary decision tree, but has the following properties:

- Its structure is a directed acyclic graph rather than a tree.
- Variables occur in the same order on every path from root to leaf.
- No two subgraphs in the graph represent the same function.

For a given variable ordering, the OBDD representation of a boolean formula is unique [2].

We can perform most logical operations efficiently using OBDDs. The function that restricts some argument x_i of a boolean function f to a constant value b, denoted by $f\mid_{x_i \leftarrow b}$, can be performed in time linear in the size of the original OBDD [2]. The restriction algorithm allows us to compute the OBDD for the formula $\exists x f$ as $f\mid_{x\leftarrow 0} \vee f\mid_{x\leftarrow 1}$. All 16 two-argument logical operations can be implemented efficiently on boolean functions that are represented as OBDDs. The complexity of these operations is linear in the size of the argument OBDDs [2]. Moreover, equivalence of two boolean functions can be decided in constant time [1].

OBDDs are used in this paper for obtaining concise representations of relations over finite domains [3]. If R is n-ary relation over $\{0,1\}$ then R can be represented by the OBDD of its *characteristic function*

$$f_R(x_1,\ldots,x_n) = 1 \text{ iff } R(x_1,\ldots,x_n).$$

If R is an n-ary relation over the finite domain D with $|D| > 2$, R can still be represented as an OBDD if an appropriate binary encoding is used for D.

3. The temporal logic CTL

The logic that we use to specify circuits is a propositional temporal logic of branching time, called CTL or Computation Tree Logic [6]. In this logic each of the usual forward-time operators of linear temporal logic (**G** *globally* or *invariantly*, **F** *sometime in the future*, **X** *nexttime* and **U** *until*) must be directly preceded by a *path quantifier*. The path quantifier can either be an **A** (for all computation paths) or an **E** (for some computation path). Thus, some typical CTL formulas are $\mathbf{AG}\, f$, which holds in a state provided that f holds globally along all possible computation paths starting from that state, and $\mathbf{EF}\, f$, which holds in a state provided that there is a computation path such that f holds in the future on the path.

In order to explain our verification procedure, it is convenient to express the CTL operators with universal path quantifiers in terms of the operators with existential path quantifiers, taking advantage of the duality between universal and existential quantification. Consequently, in our description of the syntax and semantics of CTL, we specify the existential

path quantifiers directly and treat the universal path quantifiers as syntactic abbreviations. Let P be the set of *atomic propositions,* then:

1. Every atomic proposition p in P is a formula in CTL.
2. If f and g are CTL formulas, then so are $\neg f$, $f \vee g$, $\mathbf{EX}\, f$, $\mathbf{E}[f\, \mathbf{U}\, g]$ and $\mathbf{EG}\, f$.

The semantics of a CTL formula is defined with respect to a *labeled state-transition graph.* A labeled state-transition graph is a 5-tuple $\mathbf{M} = (AP, S, L, N, S_0)$ where AP is a set of atomic propositions, S is a finite set of states, L is a function labeling each state with a set of atomic propositions, $N \subseteq S \times S$ is a transition relation, and S_0 is a set of initial states. A *computation path* is an infinite sequence of states $s_0, s_1, s_2, \ldots$ such that $N(s_i, s_{i+1})$ is true for every i.

The propositional connectives $\neg$ and $\vee$ have their usual meanings of negation and disjunction. The other propositional operators can be defined in terms of these. $\mathbf{X}$ is the *nexttime* operator: $\mathbf{EX}\, f$ will be true in a state s of $\mathbf{M}$ if and only if s has a successor s' such that f is true at s'. $\mathbf{U}$ is the *until* operator: $\mathbf{E}[f\, \mathbf{U}\, g]$ will be true in a state s of $\mathbf{M}$ if and only if there exists a computation path starting in s and an initial prefix of the path such that g holds at the last state of the prefix and f holds at all other states along the prefix. The operator $\mathbf{G}$ is used to express the *invariance* of some property over time: $\mathbf{EG}\, f$ will be true at a state s if there is a path starting at s such that f holds at each state on the path. If f is true in state s of structure $\mathbf{M}$, we write $\mathbf{M}, s \models f$. A CTL formula f is identified with the set $\{s | \mathbf{M}, s \models f\}$ of states that make f true. We use the following syntactic abbreviations for CTL formulas:

- $\mathbf{AX}\, f \equiv \neg\, \mathbf{EX}\, \neg f$ which means that f holds at all successor states of the current state (f must hold at the *next* state).
- $\mathbf{EF}\, f \equiv \mathbf{E}[\text{true}\, \mathbf{U}\, \text{f}]$ which means that for some path, there exists a state on the path at which f holds (f is *possible* in the future).
- $\mathbf{AF}\, f \equiv \neg\, \mathbf{EG}\, \neg f$ which means that for every path, there exists a state on the path at which f holds (f is *inevitable* in the future).
- $\mathbf{AG}\, f \equiv \neg\, \mathbf{EF}\, \neg f$ which means that for every path, f holds in each state on the path (f holds *globally* along all paths).
- $\mathbf{A}[f\, \mathbf{U}\, g] \equiv \neg\, \mathbf{E}[\neg g\, \mathbf{U}\, \neg f \wedge \neg g] \wedge \neg\, \mathbf{EG}\, \neg g$ which means that for every path, there exists an initial prefix of the path such that g holds at the last state of the prefix and f holds at all other states along the prefix (f holds *until* g holds, along all paths).

4. Symbolic Model Checking

Model checking is the problem of finding the set of states in a state-transition graph where a given CTL formula is true. There is a program called EMC (Extended Model Checker) that solves this problem using efficient graph-traversal techniques. If the model is represented as

a state-transition graph, the complexity of the algorithm is linear in the size of the graph and in the length of the formula. The algorithm is quite fast in practice [5, 6]. However, an explosion in the size of the model may occur when the state-transition graph is extracted from a finite state concurrent system that has many processes or components.

In this section, we describe a *symbolic model checking* algorithm for CTL which uses OBDDs to represent the state-transition graph. Assume that the behavior of the concurrent system is determined by n boolean state variables $v_1, v_2, \ldots, v_n$. The transition relation $R(\bar{v}, \bar{v}')$ for the concurrent system is given as a boolean formula in terms of two copies of the state variables: $\bar{v} = (v_1, \ldots, v_n)$ which represents the current state and $\bar{v}' = (v_1', \ldots, v_n')$ which represents the next state. The formula $R(\bar{v}, \bar{v}')$ is now converted to an OBDD. This usually results in a very concise representation of the transition relation.

Our model checking algorithm is based on the standard fixpoint characterizations of the basic CTL operators. A *fixpoint* of $\tau : 2^S \rightarrow 2^S$ is a set $S' \subseteq S$ such that $\tau(S') = S'$. If τ is monotonic, it has a fixpoint S_0 that is a subset of every other fixpoint of τ. S_0 is called the *least fixpoint* of τ and is denoted by $\mathbf{lfp}\, f\, [\tau(f)]$. The *greatest fixpoint* of τ, $\mathbf{gfp}\, f\, [\tau(f)]$, can be defined similarly as the fixpoint of τ that is a superset of all other fixpoints. It can be shown that the least fixpoint $\mathbf{lfp}\, f\, [\tau(f)]$ is the limit of the sequence of approximations

$$\text{False}, \tau(\text{False}), \tau^2(\text{False}), \ldots, \tau^i(\text{False}), \ldots$$

and the greatest fixpoint $\mathbf{gfp}\, f\, [\tau(f)]$ is the limit of the sequence of approximations

$$\text{True}, \tau(\text{True}), \tau^2(\text{True}), \ldots, \tau^i(\text{True}), \ldots$$

When the state-transition graph is finite, both of these sequences are guaranteed to converge in a finite number of steps.

Each of the basic CTL operators can be characterized as a least or greatest fixpoint of some functional $\tau : 2^S \rightarrow 2^S$. In particular, it is shown in [5] that

- $\mathbf{E}[f\ \mathbf{U}\ g] = \mathbf{lfp}\, Z\, [g \vee (f \wedge \mathbf{EX}\, Z)]$, and
- $\mathbf{EG}\, f = \mathbf{gfp}\, Z\, [f \wedge \mathbf{EX}\, Z]$.

The symbolic model checking algorithm is implemented by a procedure *Check* that takes the CTL formula to be checked as its argument and returns an OBDD that represents exactly those states of the system that satisfy the formula. Of course, the output of *Check* depends on the system being checked; this parameter is implicit in the discussion below. We define *Check* inductively over the structure of CTL formulas. If f is an atomic proposition v_i, then $Check(f)$ is simply the OBDD for v_i. Formulas of the form $\mathbf{EX}\, f$, $\mathbf{E}[f\ \mathbf{U}\ g]$, and $\mathbf{EG}\, f$ are handled by the procedures:

$$\begin{aligned} &Check(\mathbf{EX}\, f) = CheckEX\,(Check(f)), \\ &Check(\mathbf{E}[f\ \mathbf{U}\ g]) = CheckEU\,(Check(f), Check(g)), \\ &Check(\mathbf{EG}\, f) = CheckEG(Check(f)). \end{aligned}$$

Notice that these intermediate procedures take boolean formulas as their arguments, while *Check* takes a CTL formula as its argument. CTL formulas of the form $f \vee g$ or $\neg f$ are

handled using the standard algorithms for computing boolean connectives with OBDDs. Since $\mathbf{AX}\, f$, $\mathbf{A}[f\ \mathbf{U}\ g]$ and $\mathbf{AG}\, f$ can all be rewritten using just the above operators, this definition of *Check* covers all CTL formulas.

The procedure for *CheckEX* is straightforward since the formula $\mathbf{EX}\, f$ is true in a state if the state has a successor in which f is true.

$$\mathit{CheckEX}\,(f(\bar{v})) = \exists \bar{v}'\,[f(\bar{v}') \wedge R(\bar{v}, \bar{v}')].$$

If we have OBDDs for f and R, then we can compute an OBDD for

$$\exists \bar{v}'\,[f(\bar{v}') \wedge R(\bar{v}, \bar{v}')].$$

using the BDD operations given in Section 2.

The procedure for *CheckEU* is based on the least fixpoint characterization for the CTL operator $\mathbf{EU}$.

$$\mathit{CheckEU}\,(f(\bar{v}), g(\bar{v})) = \mathbf{lfp}\, Z(\bar{v})\,[g(\bar{v}) \vee (f(\bar{v}) \wedge \mathit{CheckEX}\,(Z(\bar{v})))].$$

In this case we can compute the sequence of approximations

$$Q_0, Q_1, \ldots, Q_i, \ldots$$

for the least fixpoint as described above. If we have OBDDs for f, g, and the current approximation Q_i, then we can compute an OBDD for the next approximation Q_{i+1}. Since OBDDs provide a canonical form of boolean functions, it is easy to test for convergence by comparing consecutive approximations. When $Q_i = Q_{i+1}$, this process terminates. The set of states corresponding to $\mathbf{E}[f\ \mathbf{U}\ g]$ will be represented by the OBDD for Q_i.

CheckEG is similar. In this case the procedure is based on the greatest fixpont characterization for the CTL operator $\mathbf{EG}$

$$\mathit{CheckEG}(f(\bar{v})) = \mathbf{gfp}\, Z(\bar{v})\,[f(\bar{v}) \wedge \mathit{CheckEX}\,(Z(\bar{v}))].$$

If the OBDD for f is given, then the sequence of approximations for the greatest fixpoint can be used to compute the OBDD representation for the set of states that satisfy $\mathbf{EG}\, f$.

5. Fairness Constraints

Next, we consider the issue of *fairness*. In many cases, we are only interested in the correctness along fair computation paths. For example, if we are verifying an asynchronous circuit with an arbiter, we may wish to consider only those executions in which the arbiter does not ignore one of its request inputs forever. This type of property cannot be expressed directly in CTL. In order to handle such properties we must modify the semantics of CTL slightly. A *fairness constraint* can be an arbitrary set of states, usually described by a formula of the logic. A path is said to be *fair* with respect to a set of fairness constraints if each constraint holds *infinitely often* along the path. The path quantifiers in CTL formulas are then restricted to fair paths. In the remainder of this section we describe how to modify

the algorithm above to handle fairness constraints. We assume the fairness constraints are given by a set of CTL formulas $H = \{h_1, \ldots, h_n\}$. We define a new procedure *CheckFair* for checking CTL formulas relative to the fairness constraints in H. We do this by giving definitions for new intermediate procedures *CheckFairEX*, *CheckFairEU*, and *CheckFairEG* which correspond to the intermediate procedures used to define *Check*.

Consider the formula $\mathbf{EG}\, f$ given fairness constraints H. The formula means that there exists a path beginning with the current state on which f holds globally (invariantly) and each formula in H holds infinitely often on the path. The set of such states S is the largest set with the following two properties:

1. all of the states in S satisfy f, and
2. for all fairness constraints $h_k \in H$ and all states $s \in S$, there is a sequence of states of length one or greater from s to a state in S satisfying h_k such that all states on the path satisfy f.

It is easy to show that if these conditions hold, each state in the set is the beginning of an infinite computation path on which f is always true, and for which every formula in H holds infinitely often. Thus, the procedure $\mathit{CheckFairEG}(f(\bar{v}))$ will compute the greatest fixpoint

$$\mathbf{gfp}\, Z(\bar{v})\,[f(\bar{v}) \wedge \bigwedge_{k=1}^{n} \mathit{CheckEX}\,(\mathit{CheckEU}(f(\bar{v}), Z(\bar{v}) \wedge \mathit{Check}(h_k)))].$$

The fixed point can be evaluated in the same manner as before. The main difference is that each time the above expression is evaluated, several nested fixed point computations are done (inside *CheckEU*).

Checking $\mathbf{EX}\, f$ and $\mathbf{E}[f\ \mathbf{U}\ g]$ under fairness constraints is simpler. The set of all states which are the start of some fair computation is

$$\mathit{fair}(\bar{v}) = \mathit{CheckFair}(\mathbf{EG}\, \mathit{True}).$$

The formula $\mathbf{EX}\, f$ is true under fairness constraints in a state s if and only if there is a successor state s' such that s' satisfies f and s' is at the beginning of some fair computation path. It follows that the formula $\mathbf{EX}\, f$ (under fairness constraints) is equivalent to the formula $\mathbf{EX}(f \wedge \mathit{fair})$ (without fairness constraints). Therefore, we define

$$\mathit{CheckFairEX}\,(f(\bar{v})) = \mathit{CheckEX}\,(f(\bar{v}) \wedge \mathit{fair}(\bar{v})).$$

Similarly, the formula $\mathbf{E}[f\ \mathbf{U}\ g]$ (under fairness constraints) is equivalent to the formula $\mathbf{E}[f\ \mathbf{U}\ (g \wedge \mathit{fair})]$ (without fairness constraints). Hence, we define

$$\mathit{CheckFairEU}(f(\bar{v}), g(\bar{v})) = \mathit{CheckEU}(f(\bar{v}), g(\bar{v}) \wedge \mathit{fair}(\bar{v})).$$

6. Counterexamples and Witnesses

One of the most important features of CTL model checking algorithms is the ability to find *counterexamples* and *witnesses*. When this feature is enabled and the model checker

determines that a formula with a universal path quantifier is false, it will find a computation path which demonstrates that the negation of the formula is true. Likewise, when the model checker determines that a formula with an existential path quantifier is true, it will find a computation path that demonstrates why the formula is true. For example, if the model checker discovers that the formula $\mathbf{AG}\, f$ is false, it will produce a path to a state in which $\neg f$ holds. Similarly, if it discovers that the formula $\mathbf{EF}\, f$ is true, it will produce a path to a state in which f holds. Note that the counterexample for a universally quantified formula is the witness for the dual existentially quantified formula. By exploiting this observation we can restrict our discussion of this feature to finding witnesses for the three basic CTL operators **EX**, **EG**, and **EU**.

We start by considering the complexity of finding a good witness for the formula $\mathbf{EG}\, f$ under the set of fairness constraints $H = \{h_1, \ldots, h_n\}$. We will identify each h_i with the set of states that make it true. Given a state s satisfying $\mathbf{EG}\, f$, we must exhibit a path π starting with s, such that f holds at each state, and every fairness constraint $h \in H$ is satisfied infinitely often along the path π. Since the witness is an infinite path, we must find a finite representation for it. It is easy to see that a witness can always be found that consists of a finite prefix followed by a repeating cycle. Each fairness constraint h_i is satisfied at least once on the cycle. Such a path is called a *finite witness*. The length of a finite witness is defined as the total length of the prefix and the cycle. It is desirable to find a finite witness with minimal length; however, this problem is NP-complete.

Theorem 1 *If fairness constraints are permitted, finding a finite witness with minimal length for the formula* $\mathbf{EG}\,\mathrm{True}$ *is NP-complete.*

Proof: It is relatively easy to see that this problem in NP. The prefix of a minimal finite witness cannot contain a cycle, so its length is bounded by the number of states. The cycle of a minimal finite witness can be decomposed into several simple cycles. Each simple cycle must contain a state that satisfies a fairness constraint that does not hold in any other simple cycle. Otherwise, we can eliminate this simple cycle from the witness. The length of the complete cycle is therefore bounded by the product of the number of fairness constraints and the number of states. Consequently, it is possible to guess a prefix and cycle and check to see whether they constitute a minimal finite witness in polynomial time in the size of the graph.

Finding a Hamiltonian cycle for a directed graph is known to be an NP-complete problem. Thus, it is sufficient to prove that the Hamiltonian cycle problem can be reduced to the minimal finite witness problem. Consider an instance of the Hamiltonian cycle problem for a directed graph with n nodes. This graph is treated as a state-transition graph and the set of fairness constraints $H = \{h_1, \ldots, h_n\}$ is selected so that each state satisfies a distinct fairness constraint. On any finite witness, each state must appear at least once on the cycle; hence, the length of the finite witness must be at least n. The length of the minimal finite witness is n if and only if the n states on the path form a Hamiltonian cycle. Thus, the Hamiltonian cycle problem reduces to finding a minimal finite witness and checking if this path has length n. This reduction can be performed in polynomial time. Consequently, the minimal finite witness problem is also NP-complete. □

Although we are unable to find the minimal finite witness easily, we still want to obtain a finite witness that is as short as possible. In order to accomplish this task, we will need to examine the strongly connected components of the transition graph determined by the Kripke structure. We will say that two states s_1 and s_2 are *equivalent* if there is a path from s_1 to s_2 and also from s_2 to s_1. We will call the equivalence classes of this relation *strongly connected components*. We can form a new graph in which the nodes are the strongly connected components and there is an edge from one strongly connected component to another if and only if there is an edge from a state in one to a state in the other. It is easy to see that the new graph does not contain any proper cycles, i.e., each cycle in the graph is contained in one of the strongly connected components. Moreover, since we only consider finite Kripke structures, each infinite path must have a suffix that is entirely contained within a strongly connected component of the transition graph.

Recall that the set of states that satisfy the formula $\mathbf{EG}\, f$ with the fairness constraints H is given by the formula

$$\mathbf{gfp}\, Z\, [f \wedge \bigwedge_{k=1}^{n} \mathbf{EX}(\mathbf{E}[f\ \mathbf{U}\ Z \wedge h_k])] \tag{1}$$

For brevity, we will use $\mathbf{EG}\, f$ to denote the set of states that satisfy $\mathbf{EG}\, f$ under the fairness constraints H. We construct the witness path incrementally by giving a sequence of prefixes of the path of increasing length until a cycle is found. At each step in the construction we must ensure that the current prefix can be extended to a fair path along which each state satisfies f. This invariant is guaranteed by making sure that each time we add a state to the current prefix, the state satisfies $\mathbf{EG}\, f$.

First, we evaluate the above fixpoint formula. In every iteration of the outer fixpoint computation, we compute a collection of least fixpoints associated with the formulas $\mathbf{E}[f\ \mathbf{U}\ Z \wedge h]$, for each fairness constraint $h \in H$. For every constraint h, we obtain an increasing sequence of approximations $Q_0^h, Q_1^h, Q_2^h, \ldots$, where Q_i^h is the set of states from which a state in $Z \wedge h$ can be reached in i or fewer steps, while satisfying f. In the last iteration of the outer fixpoint when $Z = \mathbf{EG}\, f$, we save the sequence of approximations Q^h for each h in H.

Now, suppose we are given an initial state s satisfying $\mathbf{EG}\, f$. Then s belongs to the set of states computed in equation (1), so it must have a successor in $\mathbf{E}[f\ \mathbf{U}\ (\mathbf{EG}\, f) \wedge h]$ for each $h \in H$. In order to minimize the length of the witness path, we choose the first fairness constraint that can be reached from s. This is accomplished by testing the saved sets Q_i^h for increasing values of i until one is found that contains some successor t of s. Note that since $t \in Q_i^h$, it has a path to a state in $(\mathbf{EG}\, f) \wedge h$ and therefore t is in $\mathbf{EG}\, f$. If $i > 0$, we find a successor of t in Q_{i-1}^h. This is done by finding the set of successors of t, intersecting it with Q_{i-1}^h, and then choosing an arbitrary element of the resulting set. Continuing until $i = 0$, we obtain a path from the initial state s to some state in $(\mathbf{EG}\, f) \wedge h$. We then eliminate h from further consideration, and repeat the above procedure until all of the fairness constraints have been visited. Let s' be the final state of the path obtained thus far.

To complete a cycle, we need a non-trivial path from s' to the state t along which each state satisfies f. In other words, we need a witness for the formula $\{s'\} \wedge \mathbf{EX}\, \mathbf{E}[f\ \mathbf{U}\ \{t\}]$. If this formula is true, we have found the witness path for s. This case is illustrated in Figure 1. If the formula is false, there are several possible strategies. The simplest is to restart the

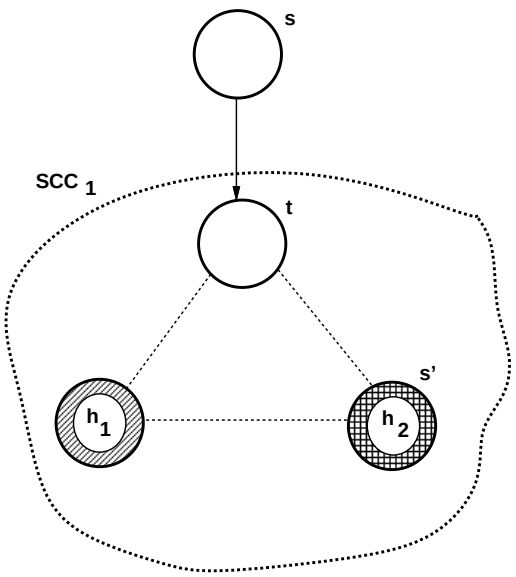

Figure 1: Witness is in first strongly connected component

procedure from the final state s'. Since $\{s'\} \wedge \mathbf{EX}\,\mathbf{E}[f\,\mathbf{U}\,\{t\}]$ is false, we know that s' is not in the strongly connected component of f containing t, however s' is in $\mathbf{EG}\,f$. Thus, if we continue this strategy, we must descend in the directed acyclic graph of strongly connected components, eventually either finding a cycle π, or reaching a terminal strongly connected component of f. In the latter case, we are guaranteed to find a cycle, since we cannot exit a terminal strongly connected component. This case is illustrated in Figure 2.

A slightly more sophisticated approach would be to precompute $\mathbf{E}[(\mathbf{EG}\,f)\,\mathbf{U}\,\{t\}]$. The first time we exit this set, we know the cycle cannot be completed, so we restart from that state. Heuristically, these approaches tend to find short counterexamples (probably because the number of strongly connected components tends to be small), so no attempt is made to find the shortest cycle.

The witness procedure for $\mathbf{EG}\,f$ under fairness constraints H can be used to extend witnesses for $\mathbf{E}[f\,\mathbf{U}\,g]$ and $\mathbf{EX}\,f$ to infinite fair paths. Let $fair$ be the set of states that satisfy $\mathbf{EG}\,True$ under the fairness constraints H. We can compute $\mathbf{E}[f\,\mathbf{U}\,g]$ under H by using the standard CTL model checking algorithm (without fairness constraints) to compute $\mathbf{E}[f\,\mathbf{U}\,(g \wedge fair)]$. Similarly, We can compute $\mathbf{EX}\,f$ by using the standard CTL model checking algorithm to compute $\mathbf{EX}(f \wedge fair)$.

In order to test the procedure for finding counterexamples when fairness constraints are used, we have examined an error in an arbiter design originally developed by Seitz [12]. The circuit is shown in Figure 3; it is designed to be *speed independent*, which means that each gate can take an arbitrarily long time to respond to its inputs. Fairness constraints are used to ensure that every gate eventually responds.

An attempt was made to verify the circuit using an explicit state model checker [7]. However, the attempt failed because the number of states was too large. In order to complete the verification, one of the input devices had to be disabled. By using symbolic model checking techniques, we are able to verify the original circuit without using any simplifying

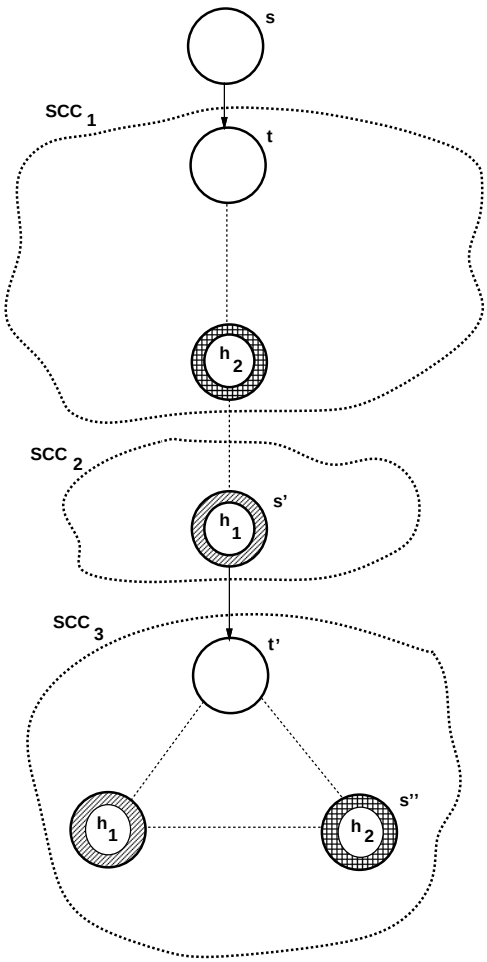

Figure 2: Witness spans three strongly connected components

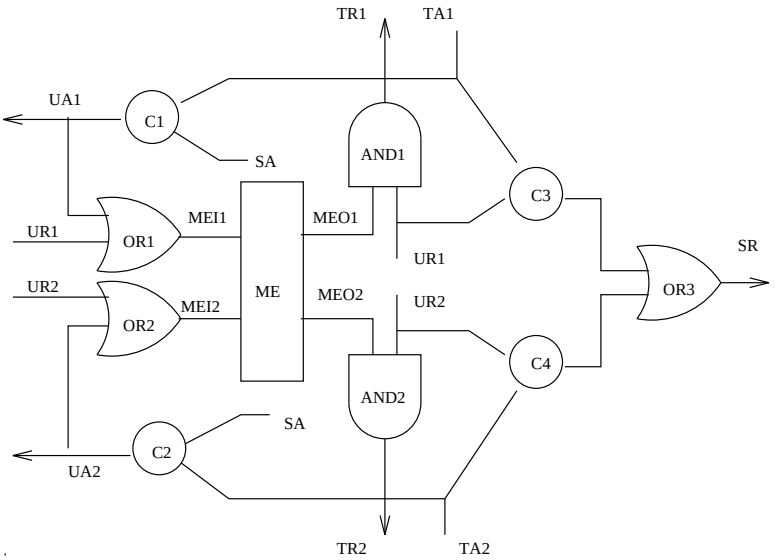

Figure 3: An asynchronous arbiter

assumptions. The model contains 33,633 reachable states, and the entire verification takes only a few minutes.

We have verified several liveness properties which require that each request signal inevitably leads to an acknowledgement signal. Such properties can be easily represented by CTL formulas with the form $\mathbf{AG}(r \rightarrow \mathbf{AF}\, a)$, where r represents a *request* and a represents an *acknowledgment*. An error was discovered when the specification $\mathbf{AG}(tr1 \rightarrow \mathbf{AF}\, ta1)$ was checked. The algorithm given earlier in this section found a counterexample that was seventy eight states long and had a cycle with length thirty. The counterexample showed that the following execution sequence was possible. The circuit could reach a state where every node was low except *meo1* if the *ME* element took a long time to respond. When *ur1* was issued, *tr1, ta1, sr, sa* and *ua1* became true consecutively. Because of the long delay of the *OR1* gate, *mei1* remained low. Eventually, the *ME* element responded to its inputs and set *meo1* low. This caused *tr1* and *ta1* to become low. Next, *OR1* responded and *mei1* became high. Then, the *ME* element and the *AND1* gate caused *tr1* to become high again while *ta1* continued to be low. In this state, the formula $tr1 \rightarrow \mathbf{AF}\, ta1$ was false. Since *ua1* was already high, *ur1* could become low. This caused *tr1* to become low. The counterexample showed that *ur1* was always low. Therefore, *ta1* remained low as well. A correction for the error was proposed in [7], but will not be discussed here.

7. Counterexamples and Witnesses for CTL* Formulas

In the previous sections, we described how to perform model checking and find counterexamples or witnesses for CTL formulas. However, some temporal properties that are important for reasoning about sequential circuit designs and protocols cannot be expressed by CTL formulas. In these cases, an extension of CTL, called CTL*, is often used. There are two types of formulas in CTL*: *state formulas* (which are true in a specific state) and *path formulas* (which are true along a specific path). As before, let AP be the set of atomic propositions. The syntax of state formulas is given by the following rules:

- If $p \in AP$, then p is a state formula.
- If f and g are state formulas, then $\neg f$ and $f \vee g$ are state formulas.
- If f is a path formula, then $\mathbf{E}(f)$ is a state formula.

Two additional rules are needed to specify the syntax of path formulas:

- If f is a state formula, then f is also a path formula.
- If f and g are path formulas, then $\neg f$, $f \vee g$, $\mathbf{X}\, f$, and $f\, \mathbf{U}\, g$ are path formulas.

CTL* is the set of state formulas generated by the above rules. The logical connectives $\neg$ and $\vee$ have their usual meaning. The formula $\mathbf{E}(f)$ is true in a state when there exists a path from the state such that f holds along the path. Let $\pi = s_0, s_1, \ldots$ be a path. We use π^i to denote the *suffix* of π starting at s_i. A state formula holds along π when it is true in

the first state s_0. $\mathbf{X}\, f$ holds along π when f holds along π^1. Finally, the formula $f\,\mathbf{U}\,g$ holds along π when there exists a $k \geq 0$ such that g holds on π^k and f holds along every π^j where $0 \leq j < k$. The following abbreviations are used in writing CTL* formulas:

- $f \wedge g \equiv \neg(\neg f \vee \neg g)$
- $\mathbf{A}(f) \equiv \neg\, \mathbf{E}(\neg f)$
- $\mathbf{F}\, f \equiv true\, \mathbf{U}\, f$
- $\mathbf{G}\, f \equiv \neg\, \mathbf{F}\, \neg f$

In general, model checking is very expensive for CTL* formulas. However, for a large class of formulas which have the form $\mathbf{E} \bigvee_{i=1}^{n} \bigwedge_{j=1}^{n_i} (\mathbf{GF}\, p_{ij} \vee \mathbf{FG}\, q_{ij})$, efficient model checking algorithms exist [8]. Because

$$\mathbf{E} \bigvee_{i=1}^{n} \bigwedge_{j=1}^{n_i} (\mathbf{GF}\, p_{ij} \vee \mathbf{FG}\, q_{ij}) = \bigvee_{i=1}^{n} \mathbf{E} \bigwedge_{j=1}^{n_i} (\mathbf{GF}\, p_{ij} \vee \mathbf{FG}\, q_{ij}),$$

it is sufficient to check formulas having the form $\mathbf{E} \bigwedge_{j=1}^{n} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j)$. A fixed point characterization for these formulas is given in [8]

$$\mathbf{E} \bigwedge_{j=1}^{n} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) = \mathbf{EF}\, \mathbf{gfp}\, Y\, [\bigwedge_{j=0}^{n} ((q_j \wedge \mathbf{EX}\, Y) \vee \mathbf{EX}\, \mathbf{E}[Y\, \mathbf{U}\, (p_j \wedge Y)])].$$

By performing a computation that is similar to the one described in Section 5, we are able to check the restricted class of CTL* formulas mentioned above. The problem of finding witnesses for these formulas is more complicated. Suppose that we want find a witness for $s_0 \models \mathbf{E} \bigwedge_{j=1}^{n} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j)$. It is easy to see that

$$\begin{aligned}
& \mathbf{E} \bigwedge_{j=1}^{n} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \\
= \;& \mathbf{E} \bigwedge_{j=1}^{n-1} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \wedge (\mathbf{GF}\, p_n \vee \mathbf{FG}\, q_n) \\
= \;& \left(\mathbf{E} \bigwedge_{j=1}^{n-1} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \wedge \mathbf{GF}\, p_n \right) \vee \left(\mathbf{E} \bigwedge_{j=1}^{n-1} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \wedge \mathbf{FG}\, q_n \right).
\end{aligned}$$

Consequently, if $s_0 \models \mathbf{E} \bigwedge_{j=1}^{n-1} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \wedge \mathbf{FG}\, q_n$, it is sufficient to find a witness for this formula; otherwise, a witness must exist for $\mathbf{E} \bigwedge_{j=1}^{n-1} (\mathbf{GF}\, p_j \vee \mathbf{FG}\, q_j) \wedge \mathbf{GF}\, p_n$. If we continue this process for the remainder of the formula, we will eventually obtain a formula which has the form $\mathbf{E}\, \mathbf{FG}\, q_{i_1} \wedge \ldots \wedge \mathbf{FG}\, q_{i_k} \wedge \mathbf{GF}\, p_{j_1} \wedge \ldots \wedge \mathbf{GF}\, p_{j_{n-k}}$. Because

$$\mathbf{E}(\mathbf{FG}\, q_{i_1} \wedge \ldots \wedge \mathbf{FG}\, q_{i_k} \wedge \mathbf{GF}\, p_{j_1} \wedge \ldots \wedge \mathbf{GF}\, p_{j_{n-k}}) = \mathbf{EF}\, \mathbf{EG}(q_{i_1} \wedge \ldots \wedge q_{i_k} \wedge \mathbf{F}\, p_{j_1} \wedge \ldots \wedge \mathbf{F}\, p_{j_{n-k}}),$$

this formula is true if and only if the CTL formula $\mathbf{EG}(q_{i_1} \wedge \ldots \wedge q_{i_k})$ is true under the fairness constraints $p_{j_1}, \ldots, p_{j_{n-k}}$. A witness can be computed in exactly the same manner as in the last section.

8. Counterexamples for Language Containment Problems

An alternative technique for verifying finite-state systems is based on showing language inclusion between finite ω-automata [9, 10, 13]. We model the system to be verified by an ω-automaton K_{sys}. The specification to be checked is given by a second ω-automaton K_{spec}. The system will satisfy its specification if the language accepted by K_{sys} is contained in the language accepted by K_{spec}, i.e. $\mathcal{L}(K_{sys}) \subseteq \mathcal{L}(K_{spec})$. In this section we show how the techniques described in last section can be used to find counterexamples for language containment problems. Although there are many types of ω-automata, in this paper we only consider *Streett automata*. These automata are particularly useful for modeling systems with complicated fairness constraints that cannot be handled using the technique described in Section 5. Counterexamples for other types of ω-automata can be determined in a similar manner by using results from [4]. In general, checking language inclusion between two nondeterministic ω-automata is PSPACE-hard. For this reason we require that the specification automaton be deterministic. We require that both automata be complete.

A (*nondeterministic*) ω-automaton is a 5-tuple $K = \langle S, s_0, \Sigma, \Delta, F \rangle$, where

- S is a finite set of *states*
- $s_0 \in S$ is the *initial state*
- Σ is a finite *alphabet*
- $\Delta \subseteq S \times \Sigma \times S$ is the *transition relation*
- F is the *acceptance condition.*

The automaton is *deterministic* if for all states $s, t_1, t_2 \in S$ and input symbols $\sigma \in \Sigma$, if $\langle s, \sigma, t_1 \rangle$ and $\langle s, \sigma, t_2 \rangle$ are two transitions in Δ, then $t_1 = t_2$. The automaton is *complete* if for every state $s \in S$ and for every symbol $\sigma \in \Sigma$, there is a state $s' \in S$ such that $(s, \sigma, s') \in \Delta$. An infinite sequence of states $s_0 s_1 s_2 \ldots \in S^\omega$ is a *run* of an ω-automaton if there exists an infinite sequence $\sigma_0 \sigma_1 \sigma_2 \ldots \in \Sigma^\omega$ such that $\forall i \geq 0, (s_i, \sigma_i, s_{i+1}) \in \Delta$. The *infinitary set* of a sequence $s_0 s_1 s_2 \ldots \in S^\omega$, denoted by $\inf(s_0 s_1 \ldots)$, is the set of all the states that appear infinitely many times in the sequence. The *Streett acceptance condition* has the form $F = \{(U_1, V_1), \ldots, (U_n, V_n)\}$, where $U_i, V_i \subseteq S$. A sequence $\sigma_0 \sigma_1 \sigma_2 \ldots \in \Sigma^\omega$ is *accepted* by a Streett automaton if there is a corresponding run $s_0 s_1 s_2 \ldots \in S^\omega$ with the property that for every $i \in \{1, \ldots, n\}$, $\inf(r) \subseteq U_i$ or $\inf(r) \cap V_i \neq \emptyset$. The set of sequences accepted by an automaton M is called the *language* of M and is denoted by $\mathcal{L}(M)$.

Let $K = \langle S, s_0, \Sigma, \Delta, F \rangle$, $K' = \langle S', s'_0, \Sigma, \Delta', F' \rangle$ be a pair of Streett automata over the same alphabet. It is shown in [4] that the path formula ϕ_F expresses the acceptance condition of K:

$$\phi_F = \bigwedge_{(U,V) \in F} (\mathbf{FG}(\bigvee_{s \in U} s) \vee \mathbf{GF}(\bigvee_{s \in V} s)),$$

and that $\neg\phi_{F'}$ expresses the negation of the acceptance condition for K':

$$\neg\phi_{F'} = \bigvee_{(U',V') \in F'} (\mathbf{GF}(\bigvee_{s \in \overline{U'}} s) \wedge \mathbf{FG}(\bigvee_{s \in \overline{V'}} s))$$

Let $M(K,K') = (S \times S', (s_0, s_0'), \mathcal{L}, \mathcal{R})$ be a state-transition system such that $\mathcal{L}(s,s') = \{s,s'\}$ and $(s,s')\mathcal{R}(t,t') \Leftrightarrow (\exists a \in \Sigma : (s,a,t) \in \Delta$ and $(s',a,t') \in \Delta)$. If K is a nondeterministic Streett automaton and K' is a deterministic Streett automaton, then

$$\mathcal{L}(K) \subseteq \mathcal{L}(K') \Leftrightarrow M(K,K') \models \neg\mathbf{E}(\phi_F \wedge \neg\phi_{F'})$$

where ϕ_F and $\neg\phi_{F'}$ are the formulas given above. Note that the above equivalence does not hold if K' is a nondeterministic automaton. The formula $\mathbf{E}(\phi_F \wedge \neg\phi_{F'})$ is equivalent to

$$\left(\bigwedge_{(U,V)\in F} (\mathbf{FG}(\bigvee_{s\in U} s) \vee \mathbf{GF}(\bigvee_{s\in V} s))\right) \wedge \left(\bigvee_{(U',V')\in F'} (\mathbf{GF}(\bigvee_{s\in \overline{U'}} s) \wedge \mathbf{FG}(\bigvee_{s\in \overline{V'}} s))\right)$$

$$\equiv \bigvee_{(U',V')\in F'} \left(\left(\bigwedge_{(U,V)\in F} (\mathbf{FG}(\bigvee_{s\in U} s) \vee \mathbf{GF}(\bigvee_{s\in V} s))\right) \wedge \mathbf{GF}(\bigvee_{s\in \overline{U'}} s) \wedge \mathbf{FG}(\bigvee_{s\in \overline{V'}} s)\right).$$

This formula is an instance of the CTL* formulas discussed in last section. Thus, the technique given in last section for finding witnesses can be used to find a counterexample when $\mathcal{L}(K)$ is not contained in $\mathcal{L}(K')$. Counterexamples for the language inclusion problems of Büchi, Muller, Rabin, and L automata can be found in essentially the same way.

9. Directions for Future Research

In this paper, we have described an efficient technique for generating counterexamples and witnesses for symbolic model checking algorithms. However, when the number of reachable states is very large, the counterexample can still be very long. Techniques for generating even shorter counterexamples will make symbolic model checking more useful in practice.

Finding a counterexample can sometimes take most of the execution time required for model checking. Additional research is needed to develop more efficient algorithms. This is particularly important because the model checking algorithm may need to be invoked several times in order to find the witness for a CTL* formula.

Another problem with the counterexample generated by the model checker is that it is sometimes hard to read. A more readable form will be helpful to engineers who are not familiar with model checking.

References

[1] K. S. Brace, R. L. Rudell, and R. E. Bryant. Efficient implementation of a BDD package. In *Proceedings of the 27th ACM/IEEE Design Automation Conference*. IEEE Computer Society Press, June 1990.

[2] R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, C-35(8), 1986.

[3] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98(2):142–170, June 1992.

[4] 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.

[5] E. M. Clarke and E. A. Emerson. Synthesis of synchronization skeletons for branching time temporal logic. In *Logic of Programs: Workshop, Yorktown Heights, NY, May 1981*, volume 131 of *Lecture Notes in Computer Science*. Springer-Verlag, 1981.

[6] 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.

[7] D. L. Dill and E. M. Clarke. Automatic verification of asynchronous circuits using temporal logic. *IEE Proceedings*, Part E 133(5), 1986.

[8] E. A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional mu-calculus. In *Proceedings of the First Annual Symposium on Logic in Computer Science*. IEEE Computer Society Press, June 1986.

[9] Z. Har'El and R. P. Kurshan. Software for analytical development of communications protocols. *AT&T Technical Journal*, 69(1):45–59, Jan.–Feb. 1990.

[10] R. P. Kurshan. Analysis of discrete event coordination. In 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.

[11] K. L. McMillan. *Symbolic Model Checking: An Approach to the State Explosion Problem*. PhD thesis, Carnegie Mellon University, 1992.

[12] C. L. Seitz. Ideas about arbiters. *Lambda*, 10(4), 1980.

[13] H. J. Touati, R. K. Brayton, and R. P. Kurshan. Testing language containment for ω-automata using BDD's. In *Proceedings of the 1991 International Workshop on Formal Methods in VLSI Design*, January 1991.