Proof-Like Counter-Examples

Arie Gurfinkel and Marsha Chechik

Department of Computer Science, University of Toronto,
Toronto, ON M5S 3G4, Canada.

{arie, chechik}@cs.toronto.edu

Abstract. Counter-examples explain why a desired temporal logic property fails to hold, and as such considered to be the most useful form of output from model-checkers. Reported explanations are typically short and described in terms of states and transitions of the model; as a result, they can be effectively used for debugging. However, counter-examples are not available for every CTL property and are often inadequate for explaining exactly what the answer means [CLJV02].

In this paper, we present the approach of annotating counter-examples with additional proof steps. This approach does not sacrifice any of the advantages of traditional counter-examples, yet allows the user to understand and navigate through the counter-example better. We describe our proof system, discuss how to connect it with counter-example generators, and present KEGVis – a tool for visualizing and browsing the annotated counter-examples.

1 Introduction

A model-checker can tell the user not only whether a desired temporal property holds, but also generate a counter-example, explaining the reasons why this property failed. Typically, counter-examples are fairly small and are given in terms of states and transitions of the model; thus, they are readily understood by engineers and can be effectively used for debugging the model. The counter-example generation ability has been one of the major advantages of model-checking in comparison with other verification methods.

Counter-examples are a form of mathematical proof: to disprove that some property φ holds on all elements of some set S , it is sufficient to produce a single element $s \in S$ such that $\neg\varphi$ holds on s . For model-checking, this has two ramifications. First, counter-examples are restricted to universally-quantified formulas, and second, counter-examples have been viewed as infinite or finite *paths*, starting from the initial state, that illustrate failure of a given property. This notion of path-like counter-examples has been implemented in SMV [McM93,CGMZ95]. Yet only a subset of universally-quantified CTL (ACTL), namely $\text{ACTL} \cap \text{LTL}$, has linear counter-examples [CLJV02]. Recent work by Clarke et. al. [CLJV02] has extended this notion to *tree-like* counter-examples. This method generates trees instead of paths as counter-examples and is complete for ACTL. For example, a counter-example for $AF(\neg y \wedge AX\neg x)$, where shaded areas indicate which subformula is being disproved, is shown in Figure 1. This example is taken from [CLJV02]. Note that tree-like counter-examples are fairly hard to understand: different parts of the counter-example correspond to different parts of the property; thus, local information is insufficient to understand what each branch is attempting to disprove.

Fig. 1. Counter-example for $AF(\neg y \wedge AX\neg x)$ (from [CLJV02]).

This makes it difficult to navigate to “interesting” parts of the counter-example, so the user has to understand the counter-example in its entirety.

The typical approach of existing counter-example generators is to give a complete explanation, if it is available. Yet the model-checker cannot explain why an existential property is false or why a universal one is true, and gives no feedback to the user in these cases. Such an explanation could include the entire model! Thus, in general, we say that a counter-example is not available if it is too large to be practical. Certainly, the problem gets even more difficult when temporal quantifiers are nested, e.g., $EFAX\varphi$ (is there a reachable state from which φ holds in all successors?). Counter-example generators, conventional or tree-like, do not give the user any feedback for such properties, even though counter-examples are available for parts of the formula.

The goals of the work reported in this paper are as follows: (1) to preserve the desired usability aspects of counter-examples, i.e., their short length and the close correspondence to the model; (2) to provide some feedback even if in general the counter-example is not available; (3) to help the user in *understanding* complex counter-examples.

Our results follow from the primary observation that counter-examples are simply proofs by example; yet they are the *coarsest* type of proofs which skips all steps except those that result in the transition to the next state. Additional proof steps that explain *why* the model-checker chose this sequence of states for presentation can be added as means of annotating the counter-examples. With this approach, short linear counter-examples remain intuitive, but long and bushy counter-examples become significantly easier to understand because the user can refer to the proof for the explanation. Further, when counter-examples are not available, they can be replaced by proof obligations which are either discharged by a theorem-prover or taken on faith. For example, if the model-checker determines that a property $AX\varphi$ holds, we (1) show which states are successors of the current state; (2) indicate that φ holds in these states; and (3) generate a proof obligation that the current state has no additional successors. The user might check the last claim by using a theorem-prover or simply by eye-balling the model. We call the resulting construction *proof-like counter-examples*.

In this paper, we present and formalize the concept of proof-like counter-examples and show how to generate them from a run of a model-checker. We also describe a system for visualizing and browsing such counter-examples. Here, we concentrate on the problem of computing witnesses to existential properties. This problem is dual to the one of computing counter-examples. We choose witnesses for our presentation because

we find it more natural to talk about why a property is true as opposed to why a negation of a property is false [CGMZ95].

The rest of this paper is organized as follows: Section 2 introduces our notation and gives some background on CTL model-checking. Before we can describe the generation and the presentation of proof-like witnesses, we have to develop a proof system that allows us to generate proofs for fair CTL, which we do in Section 3. We show how to extract proof-like witnesses from these proofs in Section 4. Section 4 also discusses tool support for browsing these witnesses. Our primary goal is not to produce proofs for all possible temporal properties. In that, we differ from the work that uses model-checking for proof generation [PZ01, PPZ01, Nam01, TC02]. Instead, we concentrate on annotating counter-examples with proof steps. We compare our results with related work in Section 5. Section 6 concludes the paper with the summary of our approach and venues for future work.

2 Background

We assume that the reader is familiar with the basics of model checking and CTL; this information is available in [CGP99]. Below, we recall some specific concepts and fix the notation.

We use $\top$ and $\perp$ to represent *True* and *False*, respectively. $\mathbb{B}$ and $\mathbb{N}$ are the set of boolean values $\{\top, \perp\}$ and all natural numbers, respectively. An unnamed function over the domain D is denoted by $\lambda x \in D \cdot F\text{-}n \text{ Body}$. $\mu Z \cdot f(Z)$ and $\nu Z \cdot f(Z)$ are the least and the greatest fixpoints of a function f , respectively.

A Kripke structure is a tuple $M = (S, R, s_0, A, I)$ where S is a set of states; $R : S \times S \rightarrow \mathbb{B}$ is a (total) transition function; $s_0 \in S$ is an initial state, A is a set of atomic propositions; and $I : S \rightarrow 2^A$ is a (total) labeling function. CTL formulas are evaluated over infinite trees of computations produced by M . We write $\llbracket \varphi \rrbracket^M(s)$ to indicate the value of φ in state s of M . If M is clear from the context, we omit it from our notation. If a formula φ holds in the initial state, i.e. $\llbracket \varphi \rrbracket(s_0) = \top$, it is considered to hold in the model.

We use EG , EX and EU as our adequate set for CTL [CGP99]. These operators are defined as follows:

$$\begin{aligned} EX\varphi &\triangleq \lambda s \cdot \exists t \in S \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t) && \text{def. of } EX \\ EG\varphi &\triangleq \nu Z \cdot \varphi \wedge EXZ && \text{def. of } EG \\ E[\varphi U \psi] &\triangleq \mu Z \cdot \psi \vee (\varphi \wedge EXZ) && \text{def. of } EU \end{aligned}$$

We also explicitly define the bounded versions of EU :

$$\begin{aligned} E[\varphi U_0 \psi] &\triangleq \psi && \text{def. of } EU_0 \\ E[\varphi U_i \psi] &\triangleq \psi \vee (\varphi \wedge EXE[\varphi U_{i-1} \psi]) && \text{def. of } EU_i \end{aligned}$$

Note that

$$E[\varphi U \psi] = E[\varphi U_\infty \psi]$$

The remaining operators can be computed from these, as shown in [CGP99].

Fig. 2. A simple Kripke structure.

A computation of M on which all of the given fairness conditions $C = \{c_1, \dots, c_k\}$ occur infinitely often is called a *fair computation*. Model-checking restricted to fair computations of M is called *fair*, and the resulting language is called FCTL. This language can be obtained by providing a fair version of EG ($E_C G$) as follows:

$$E_C G \varphi \triangleq \nu Z \cdot \varphi \wedge \bigwedge_{i=1}^k EXE[\varphi U \varphi \wedge c_i \wedge Z] \quad \text{def. of } E_C G$$

Thus, $\llbracket E_C G \top \rrbracket(s) = \top$ if and only if s is a starting state of some fair computation of M . Computing $E_C X \varphi$ in s amounts to restricting successors of s to the ones at the start of some fair computation, and evaluating $EX \varphi$ using only these successors. Computing $E_C U$ is similar:

$$\begin{aligned} E_C X \varphi &\triangleq EX(\varphi \wedge E_C G \top) && \text{def. of } E_C X \\ E_C[\varphi U \psi] &\triangleq E[\varphi U (\psi \wedge E_C G \top)] && \text{def. of } E_C U \end{aligned}$$

ECTL is the universal fragment of CTL where every path is existentially quantified and negation is restricted to atomic propositions. ACTL is the dual universal fragment of CTL, i.e., only universal quantification is allowed. Their fair counterparts are referred to as FECTL and FACTL, respectively.

In this paper, we use p and q to stand for arbitrary atomic propositions, s and t to represent states, and φ and ψ to represent CTL formulas. We also use the notation $\{s\}$ to express a formula that evaluates to $\top$ at state s and $\perp$ otherwise, i.e. $\llbracket \{s\} \rrbracket(t) \triangleq (s = t)$. We use $\{s\}$ for the negation of $\{s\}$.

3 Generating Proofs

In this section, we develop a proof system that allows us to generate proofs for fair CTL. We start with ECTL and then extend our framework to deal with fairness, universal quantification, and negation, resulting in fair CTL. We show how to extract proof-like witnesses from these proofs in Section 4.

3.1 Proof Rules for ECTL

Our initial goal is to develop a sound and complete proof system that allows us to prove validity of sentences of the form $\llbracket \varphi \rrbracket(s)$, where φ is an ECTL formula, and s is a state of a given Kripke structure M .

We assume that our proof system includes all axioms of boolean and propositional logic. Several of such axioms are shown in Figure 3. For example, in the one-point rule, also known as the $\exists$ introduction rule, f is a predicate and d is some element of D . Intuitively, the one-point rule states that to justify an existential statement $\exists x \in D \cdot f(x)$,

$$\begin{array}{c}
\frac{a}{a \vee b} \quad \vee\text{-intro} \qquad \frac{a \quad b}{a \wedge b} \quad \wedge\text{-intro} \\
\\
\frac{b}{a \vee b} \quad \vee\text{-intro} \qquad \frac{f(d)}{\exists x \in D \cdot f(x)} \quad \text{one-point rule} \\
\\
\frac{f(d_1) \quad f(d_2) \quad \dots \quad f(d_n)}{\forall d \in D \cdot f(d)} \quad \text{finite quantification, with } \bigcup_{i=1}^n \{d_i\} = D \\
\\
\frac{\forall x \in D_1 \cdot f(x) \wedge g(x) \quad \forall x \in D \setminus D_1 \cdot \neg f(x)}{\forall x \in D \cdot f(x) \rightarrow g(x)} \quad \text{universal case splitting}
\end{array}$$

Fig. 3. Some axioms of propositional and boolean logic.

$$\begin{array}{c}
\frac{\top}{\llbracket \top \rrbracket(s)} \quad \text{value-rule} \qquad \frac{\llbracket \varphi \rrbracket(s) \quad \llbracket \psi \rrbracket(s)}{\llbracket \varphi \wedge \psi \rrbracket(s)} \quad \wedge\text{-rule} \\
\\
\frac{I(s, p)}{\llbracket p \rrbracket(s)} \quad \text{atomic-rule} \qquad \frac{\llbracket \varphi \rrbracket(s) \vee \llbracket \psi \rrbracket(s)}{\llbracket \varphi \vee \psi \rrbracket(s)} \quad \vee\text{-rule} \\
\\
\frac{\neg I(s, p)}{\llbracket \neg p \rrbracket(s)} \quad \text{neg-atomic-rule} \qquad \frac{\exists t \in S \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t)}{\llbracket EX \varphi \rrbracket(s)} \quad EX
\end{array}$$

Fig. 4. Proof rules for non-temporal operators and EX .

one simply needs to exhibit an element $d \in D$ for which $f(d)$ holds. Note also we only consider quantification over finite domains.

In addition, we assume that all axioms of the theory of Kripke structures and the axiomatization of a particular Kripke structure M are available. The latter includes statements about M 's transition relation R and its labeling function I . For example, some of the axioms describing the Kripke structure in Figure 2 are:

$$\begin{array}{ll}
R(s_0, s_1) = \top & I(s_0, p) = \top \\
R(s_0, s_0) = \perp & I(s_0, q) = \perp
\end{array}$$

The proof rules for non-temporal operators and EX are shown in Figure 4. They follow directly from the definition of the corresponding operators. The EX -rule introduces an existential quantifier, which is typically eliminated by the one-point rule shown in Figure 3.

The proof rules for the bounded EU are given in Figure 5 and follow directly from the definition of this operator. To derive the rule for the unbounded EU , we start by noting the monotonicity of EU_i :

$$\begin{array}{c}
\frac{\llbracket \psi \rrbracket(s)}{\llbracket E[\varphi \ U_0 \ \psi] \rrbracket(s)} \ EU_0 \\
\frac{\llbracket \psi \vee (\varphi \wedge EXE[\varphi \ U_{i-1} \ \psi]) \rrbracket(s)}{\llbracket E[\varphi \ U_i \ \psi] \rrbracket(s)} \ EU_i \\
\frac{\llbracket (\varphi \wedge EXE[\varphi \ U \ \varphi \wedge \{s\}]) \vee (\varphi \wedge EXEG(\varphi \wedge \overline{\{s\}})) \rrbracket(s)}{\llbracket EG\varphi \rrbracket(s)} \ EG
\end{array}
\quad
\frac{\exists n \in \mathbb{N} \cdot \llbracket E[\varphi \ U_n \ \psi] \rrbracket(s)}{\llbracket E[\varphi \ U \ \psi] \rrbracket(s)} \ EU$$

Fig. 5. Proof rules for EU and EG .

Proposition 1 *Let φ, ψ be ECTL formulas and $i, j \in \mathbb{N}$. Then,*

$$i \geq j \Rightarrow \forall s \in S \cdot (\llbracket E[\varphi \ U_i \ \psi] \rrbracket(s) \Leftarrow \llbracket E[\varphi \ U_j \ \psi] \rrbracket(s))$$

Since we assume that the state space is finite, the rule is actually bi-directional. That is, for a given Kripke structure M , there always exists a natural number n , which depends on the diameter of the directed graph induced by M , such that $E[\varphi \ U \ \psi] = E[\varphi \ U_n \ \psi]$. The proof rule for the unbounded EU is given in Figure 5.

To complete our proof system, we need to find a proof rule for EG . Unfortunately, we cannot proceed in the same manner as before and use the ECTL equivalence $EG\varphi = \varphi \wedge EXEG\varphi$. Doing so would result in a proof system which is not complete, since this proof rule can potentially be applied an infinite number of times.

Instead, note that $\llbracket EG\varphi \rrbracket(s)$ is the result of evaluating $G\varphi$ on all infinite paths emanating from the state s . Moreover, since we are dealing with finite state systems, every infinite path can be decomposed into a finite (possibly empty) prefix and a finite repeating suffix. Thus, we can decompose $\llbracket EG\varphi \rrbracket(s)$ into EG restricted to all non-trivial cycles around s , and EG restricted to all infinite paths that do not contain s in the future.

First, we consider the restriction of $\llbracket EG\varphi \rrbracket(s)$ to all non-trivial cycles around s . Essentially, this is simply a fair- EG , where the fairness condition is given by a single formula $\{s\}$. That is, the set of non-trivial cycles around s is exactly the set of paths along which s occurs infinitely often. Furthermore, since our starting state is s , any infinite path along which s occurs infinitely often is equivalent to a finite path from s to itself. Thus, to evaluate $\llbracket EG\varphi \rrbracket(s)$ restricted to cycles around s , it is sufficient to consider only finite paths from s to s . This intuition is formalized in the following theorem, the proof of which is available in [Gur02]:

Theorem 1 *Let φ be an ECTL formula and s be a state of a Kripke structure. Then,*

$$\llbracket EG\varphi \rrbracket(s) = \llbracket (\varphi \wedge EXE[\varphi \ U \ \varphi \wedge \{s\}]) \vee (\varphi \wedge EXEG(\varphi \wedge \overline{\{s\}})) \rrbracket(s)$$

A proof rule for EG is given in Figure 5.

Theorem 2 *The proof system for ECTL is sound and complete.*

Proof:

The proof of soundness comes from the fact that our proof rules have been derived using definitions and equivalences between ECTL operators.

To prove completeness, we show that any valid statement of the form $\llbracket \varphi \rrbracket(s)$ can be proven by a finite number of applications of our proof rules. The proof proceeds on the structure of the formula φ . A proof sketch follows:

(1) Let φ be a propositional temporal formula, that is, φ does not contain EU and EG . Each rule given in Figure 4 reduces φ to its subformulas. Thus, the rest of the proof for this case proceeds by induction on the number of subformulas of φ .

(2) If in addition to (1), φ can also contain EU_i , EU_i can be removed by expanding it using rules for EU_0 and EU_i given in Figure 5.

(3) If φ can contain bounded or unbounded EU , we note that for a given Kripke structure, there exists an equivalent formula in which all unbounded EU operators are replaced by their bounded versions, reducing the resulting formula to case (2) considered above.

(4) If φ can also contain EG , i.e., $\varphi \in \text{ECTL}$, the EG -rule reduces a formula with EG to two formulas: (a) the one with EU , handled by the above case, and (b) a new formula containing EG , where the EG operator is restricted to a subset of the state space which does not contain the current state s . Therefore, this rule can only be applied up to $|S|$ times, ensuring that a valid statement $\llbracket \varphi \rrbracket(s)$ can be proven by a finite number of applications of our proof rules. $\square$

3.2 Automatic Proof Generation

Given a statement $\llbracket \varphi \rrbracket(s)$, we are interested in an automated proof of its validity. We can achieve this by embedding the proof system of Section 3.1 into an automated theorem prover, such as PVS [OSR93], and use its facilities for generating the proof. Yet we can do so more efficiently if we use the model-checker as a decision procedure for (a) deciding the validity of a given subformula (so that our proof generator avoids exploring irrelevant proof branches) and for (b) applying the one-point rule. We call this decision procedure `modelCheck` and assume that `modelCheck`(φ, s) computes $\llbracket \varphi \rrbracket(s)$.

We start with the boolean connective $\vee$. Given a statement of the form $\llbracket \varphi \vee \psi \rrbracket(s)$, we apply the $\vee$ -rule:

$$\frac{\llbracket \varphi \rrbracket(s) \vee \llbracket \psi \rrbracket(s)}{\llbracket \varphi \vee \psi \rrbracket(s)} \quad \vee\text{-rule}$$

Next, we must apply the $\vee$ -introduction rule, using either φ or ψ . Since the choice depends on the validity of $\llbracket \varphi \rrbracket(s)$ and $\llbracket \psi \rrbracket(s)$, this suggests a simple proof strategy: if $\llbracket \varphi \rrbracket(s)$ is valid, apply the $\vee$ -introduction rule with φ ; otherwise, if $\llbracket \psi \rrbracket(s)$ is valid, apply the $\vee$ -introduction rule with ψ ; otherwise, terminate declaring that the statement is invalid. This proof strategy is implemented by an algorithm shown in Figure 6(a).

We now examine the case of the unbounded until (EU) operator. Given the statement $\llbracket E[\varphi U \psi] \rrbracket(s)$, we first apply the EU -rule:

$$\frac{\exists n \in \mathbb{N} \cdot \llbracket E[\varphi U_n \psi] \rrbracket(s)}{\llbracket E[\varphi U \psi] \rrbracket(s)} \quad EU$$

```

1: proc orRule( $\varphi, \psi, s$ )
2:    $k_\varphi := \text{modelCheck}(\varphi, s)$ 
3:    $k_\psi := \text{modelCheck}(\psi, s)$ 
4:   if  $k_\varphi$  then
5:     apply  $\vee$ -introduction with  $\varphi$ 
(a) 6:   else if  $k_\psi$  then
7:     apply  $\vee$ -introduction with  $\psi$ 
8:   else
9:     terminate with invalid
10:   end if
11: end proc

1: proc exOnePoint( $\varphi, s, \ell$ )
2:    $k := \text{modelCheck}(EX\varphi, s)$ 
3:   if not  $k$  then
4:     terminate with invalid
(c) 5:   end if
6:    $\hat{t} := \text{exWitness}(\varphi, s)$ 
7:   apply the one-point rule substituting  $\hat{t}$  for  $t$ 
8: end proc

1: proc euOnePoint( $\varphi, \psi, s, \ell$ )
2:    $i := 0$ 
3:    $eu := \text{modelCheck}(E[\varphi \ U \ \psi], s)$ 
4:    $ewi := \perp$ 
5:   while  $ewi \neq eu$  do
6:      $ewi := \text{modelCheck}(E[\varphi \ U_i \ \psi], s)$ 
(b) 7:    $i := i + 1$ 
8:   end while
9:   if  $ewi$  then
10:    apply the one-point rule substituting  $i$  for  $n$ 
11:   else
12:    terminate with invalid
13:   end if
14: end proc

```

Fig. 6. Algorithms for generating proof-like witnesses.

The next step is to find an instantiation of n for the one-point rule. Recall that the bounded EU_i is monotone when viewed as a function of i (by Proposition 1). Moreover, it is bounded above by the unbounded EU . Therefore, we can find the instantiation of n by a linear search, starting from $n = 0$. The algorithm for the application of the one-point rule is given in Figure 6(b). The intermediate computations performed by this algorithm are exactly the same as the ones done by a symbolic model-checking algorithm. Thus, if the results of the intermediate computations performed by $\text{modelCheck}(E[\varphi \ U \ \psi], s)$ are available, a more efficient binary search can be used instead of the linear one.

For example, consider the Kripke structure in the Figure 2 and assume that we want to prove that $\llbracket E[p \ U \ q] \rrbracket(s_0)$. After the application of the EU -rule, we get

$$\exists n \in \mathbb{N} \cdot \llbracket E[p \ U_n \ q] \rrbracket(s_0)$$

To apply the one-point rule, we first try $\llbracket E[p \ U_0 \ q] \rrbracket(s_0) = \llbracket q \rrbracket(s_0)$. This yields $\perp$. Increasing the bound, we get $\llbracket E[p \ U_1 \ q] \rrbracket(s_0) = \top$ and therefore can apply the one-point rule by instantiating n to 1.

Finally, given the statement $\llbracket EX\varphi \rrbracket(s)$, we start by applying the EX rule:

$$\frac{\exists t \in S \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t)}{\llbracket EX\varphi \rrbracket(s)} \quad EX$$

Then, we need to eliminate the existential quantifier by applying the one-point rule. First, let us define a function $\text{img} : S \rightarrow \mathbb{B}$ as

$$\text{img}(x) \triangleq R(s, x) \wedge \llbracket \varphi \rrbracket(x)$$

Clearly,

$$\exists t \in S \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t) \Leftrightarrow \exists t \in S \cdot \text{img}(t) = \top$$

Thus, to apply the one-point rule, we need to find an element $\hat{t} \in \text{img}^{-1}(\top)$. Unfortunately, unlike the previous cases, we cannot obtain this element by only using

the model-checker as a black box. Note, however, that if a path $s, \hat{t}$ is a witness to $\llbracket EX\varphi \rrbracket(s)$, then the element $\hat{t} \in \text{img}^{-1}(\top)$. Thus, we assume the availability of a function $\text{exWitness}(\varphi, s)$ which computes a witness to $\llbracket EX\varphi \rrbracket(s)$. A version of this function can be easily obtained from any symbolic model-checker. The algorithm for the application of the one-point rule in the case of EX is given in Figure 6(c).

3.3 Extention to FCTL

Here we extend the results presented earlier in this section: to fair ECTL (FECTL), to ACTL, and finally to full FCTL.

Extention to ECTL with fairness Let a set of fairness conditions $C = \{c_1, \dots, c_k\}$ be given. $E_C G$, a fair version of EG , is defined as a restriction of EG to paths where each fairness condition occurs infinitely often (see Section 2). This intuition is formalized below:

Theorem 3 *Let $C = \{c_1, \dots, c_k\}$ be a set of fairness conditions, and $F_i(Z) = EX E[\varphi U \varphi \wedge c_i \wedge Z]$. Then,*

$$E_C G\varphi = \nu Z \cdot \varphi \wedge (F_1 \circ \dots \circ F_k)(Z)$$

For example, if $C = \{c_1, c_2\}$, then

$$E_C G\varphi = \nu Z \cdot \varphi \wedge EX E[\varphi U \varphi \wedge c_1 \wedge EX E[\varphi U \varphi \wedge c_2 \wedge Z]]$$

The proofs of this and the remaining theorems in this paper are available in [Gur02].

The proof rule for $E_C G$ is obtained similarly to the EG -rule. We show how $\llbracket E_C G\varphi \rrbracket(s)$ can be decomposed into (1) $E_C G$ restricted to fair cycles around s and (2) fair paths that do not contain s , and then use this result to define the proof rule.

Theorem 4 *Let φ be a FECTL formula, s be a state of a Kripke structure, $C = \{c_1, \dots, c_k\}$ be a set of fairness conditions, and F_i be defined as in Theorem 3. Then*

$$\llbracket E_C G\varphi \rrbracket(s) = \llbracket (\varphi \wedge (F_1 \circ \dots \circ F_k)(EX E[\varphi U \varphi \wedge \{s\}])) \vee (\varphi \wedge EX E_C G(\varphi \wedge \overline{\{s\}})) \rrbracket(s)$$

This theorem gives rise to the proof rule for $E_C G$, shown in Figure 7.

Recall from Section 2 that other FECTL operators are defined through $E_C G$, so no additional proof rules are required.

Theorem 5 *The proof system for FECTL is sound and complete.*

The proof is similar to the one for ECTL.

$$\frac{\llbracket (\varphi \wedge (F_1 \circ \dots \circ F_k)(EXE[\varphi U \varphi \wedge \{s\}])) \vee (\varphi \wedge EXE_C G(\varphi \wedge \overline{\{s\}})) \rrbracket(s)}{\llbracket E_C G \varphi \rrbracket(s)} E_C G$$

Fig. 7. A proof rule for $E_C G$.

$$\begin{array}{c} \frac{\forall t \in S \cdot R(s, t) \rightarrow \llbracket \varphi \rrbracket(t)}{\llbracket AX \varphi \rrbracket(s)} AX \\ \frac{\llbracket \varphi \wedge AX AG(\{s\} \vee \varphi) \rrbracket(s)}{\llbracket AG \varphi \rrbracket(s)} AG \end{array} \quad \begin{array}{c} \frac{\llbracket \psi \rrbracket(s)}{\llbracket A[\varphi U_0 \psi] \rrbracket(s)} AU_0 \\ \frac{\llbracket \psi \vee (\varphi \wedge AX A[\varphi U_{i-1} \psi]) \rrbracket(s)}{\llbracket A[\varphi U_i \psi] \rrbracket(s)} AU_i \\ \frac{\exists n \in \mathbb{N} \cdot \llbracket A[\varphi U_n \psi] \rrbracket(s)}{\llbracket A[\varphi U \psi] \rrbracket(s)} AU \end{array}$$

Fig. 8. Proof rules for ACTL.

Extension to ACTL The ACTL subset of CTL is very similar to ECTL. The essential difference is that the EX operator in the fixpoint definition of the ECTL temporal operators is replaced by its dual AX . The proof rules are derived similarly to the ones for ECTL, and are summarized in Figure 8.

The AX -rule introduces universal quantification over the state space. Although the state-space is finite, the result is too large to present to the user explicitly and thus forms the basis of a proof obligation. Yet sometimes it is possible to reduce the complexity of such proof obligations. The application of the universal case splitting rule

$$\frac{\forall t \in S_1 \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t) \quad \forall t \in S \setminus S_1 \cdot \neg R(s, t)}{\forall t \in S \cdot R(s, t) \rightarrow \llbracket \varphi \rrbracket(t)} \text{universal case splitting}$$

tells us that we can show that φ holds for all successors of s , and no other states are successors of s . If the set S_1 is relatively small, it is demonstrated to the user via the finite quantification rule:

$$\frac{R(s, t_1) \wedge \llbracket \varphi \rrbracket(t_1) \quad \dots \quad R(s, t_n) \wedge \llbracket \varphi \rrbracket(t_n)}{\forall t \in S_1 \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t)} \text{finite quantification, with } \bigcup_{i=1}^n \{t_i\} = S_1$$

and the fact that the elements of $S \setminus S_1$ are not successors of s is retained as a proof obligation.

ACTL proofs that do not involve the use of AX , such as proving $\llbracket A[\varphi U \psi] \rrbracket(s)$ when $\llbracket \psi \rrbracket(s)$ holds, are generated automatically.

Extension to FCTL To extend the proof rules to full CTL, we need to extend the negation rule to temporal operators. Using the well-known dualities between ECTL and ACTL operators [CGP99], we obtain the negation rules summarized in Figure 9. Note that the negation applied to the EU operator (neg- EU rule) is handled differently. Since

$$\begin{array}{c}
\frac{\llbracket AX \neg \varphi \rrbracket(s)}{\llbracket \neg EX \varphi \rrbracket(s)} \text{ neg-}EX \qquad \frac{\llbracket A[\top U \neg \varphi] \rrbracket(s)}{\llbracket \neg EG \varphi \rrbracket(s)} \text{ neg-}EG \\
\\
\frac{\llbracket EX \neg \varphi \rrbracket(s)}{\llbracket \neg AX \varphi \rrbracket(s)} \text{ neg-}AX \qquad \frac{\llbracket E[\top U \neg \varphi] \rrbracket(s)}{\llbracket \neg AG \varphi \rrbracket(s)} \text{ neg-}AG \\
\\
\frac{\llbracket (\neg \psi \wedge \neg \varphi) \vee (\neg \psi \wedge \neg EX E[\varphi \wedge \overline{\{s\}} U \psi \wedge \neg \overline{\{s\}}]) \rrbracket(s)}{\llbracket \neg E[\varphi U \psi] \rrbracket(s)} \text{ neg-}EU \\
\\
\frac{\llbracket EG \neg \psi \vee E[\neg \varphi U \neg \varphi \wedge \neg \psi] \rrbracket(s)}{\llbracket \neg A[\varphi U \psi] \rrbracket(s)} \text{ neg-}AU
\end{array}$$

Fig. 9. Proof rules for negation.

$$\begin{array}{ll}
\text{(a)} \quad \frac{\frac{R(s, \hat{t}) \quad \llbracket \varphi \rrbracket(\hat{t})}{\exists t \in S \cdot R(s, t) \wedge \llbracket \varphi \rrbracket(t)} \text{ one-point rule}}{\llbracket EX \varphi \rrbracket(s)} EX & \text{(c)} \quad \begin{array}{c} \llbracket \varphi \rrbracket(\hat{t}) \\ \uparrow \\ \llbracket EX \varphi \rrbracket(s) \end{array} \\
\\
\text{(b)} \quad \frac{\llbracket \varphi \rrbracket(\hat{t})}{\llbracket EX \varphi \rrbracket(s)} & \text{(d)} \quad \begin{array}{c} \hat{t} \\ \uparrow \\ s \end{array}
\end{array}$$

Fig. 10. From proofs to witnesses.

any adequate set for CTL must contain the EU operator [Lar95], it is not possible to express its negation as a combination of only ACTL operators.

Finally, to yield the proof system for FCTL, we combine the above proof rules with the one for $E_C G$.

4 Generating Proof-Like Counter-Examples

In this section we describe how to use the proof system introduced in Section 3 to generate annotated witnesses and counter-examples. We also discuss and illustrate the tool support for this approach.

4.1 From Proofs to Counter-Examples

There is a clear one-to-one correspondence between witnesses and proofs for CTL. Note, however, that proofs are finite, whereas witnesses can be infinite and are typically

Fig. 11. Proof-like witness for $\llbracket E[\top U q] \rrbracket(s_0)$.

represented using a finite prefix and a finite repeating suffix. We return to this subject below.

Consider the proof for $\llbracket EX\varphi \rrbracket(s)$ shown in Figure 10(a). Here, $\hat{t}$ is some state in S obtained by the model-checking procedure $\text{exWitness}(\varphi, s)$ as part of the one-point rule for EX (see line 6 of the algorithm in Figure 6(c)). The witness can be visually obtained from the proof of EX via the following steps:

1. remove all nodes from the proof tree except for: (a) the root node, and (b) nodes that are the result of the application of the one-point rule to EX and that do not correspond to the axioms of the transition relation (see Figure 10(b));
2. replace horizontal bars by directed edges (see Figure 10(c));
3. relabel the root node by s and each node of the form $\llbracket \varphi \rrbracket(\hat{t})$ by $\hat{t}$ (see Figure 10(d)).

Clearly, the result of the application of the above procedure to the proof of $\llbracket EX\varphi \rrbracket(s)$ is a path through the Kripke structure that is a witness to $\llbracket EX\varphi \rrbracket(s)$. A similar procedure that uses finite quantification instead of the one-point rule also exists for AX .

Note that a sequence of steps in a Kripke structure is the only information given by the witness generators. We referred to it as the *coarsest* type of proof in Section 1. All of such steps are results of the application of the one-point rule to EX or the finite quantification rule to AX . Other proof steps do not result in state transitions and are not shown by conventional witness generators. When extracting the witness from our proofs, we label each state in a witness with these missing proof steps. Branches in the witness result from applications of $\wedge$ -intro and finite quantification rules. Loops can be identified by merging states with the same label, allowing the representation of nested loops. The resulting structures are called *proof-like witnesses* or *counter-examples*.

For example, a proof-like witness for $\llbracket E[\top U q] \rrbracket(s_0)$ on the Kripke structure in Figure 2 is shown in Figure 11. In this figure, the circled nodes represent states and solid arrows between them represent state transitions. Thus, the witness for $\llbracket E[\top U q] \rrbracket(s_0)$ consists of two states: s_0 followed by s_1 . Each state is labeled with the part of the proof that directly depends on it. This is indicated by the dashed *proof* arrows. Finally, each step in the proof that adds new states to the witness is labeled with the states it introduces, indicated by the dotted *one-point rule* (finite quantification rule) arrows.

As can be seen from this example, a proof-like witness is essentially a composition of two trees: the witness state tree, and the proof tree. This allows the user to either ignore the proof part and only explore the witness, or ignore the witness part and only explore the proof, or switch between the two representations at will.

In Section 1, we set out three goals for our work. The first goal, preserving usability aspects of witnesses, is achieved by our ability to generate conventional witnesses. The second goal, providing some feedback even if the complete counter-example is not available, is satisfied by generating proof obligations for universal properties and combining them with other proof steps so that users can understand the reasons behind these. The last goal, helping users in understanding complex counter-examples, is achieved by providing *complete* information necessary to justify the result of the model-checker. Yet proof-like counter-examples often end up being quite large. Automated support for *browsing* these counter-examples is a key for the effective use of this information for understanding and debugging the model. We address the tool support and browsing strategies below.

4.2 Tool Support for Browsing Counter-Examples

We have developed a prototype proof/witness browser tool called KEGVis (Kounter-example generator and visualizer). The tool uses the symbolic model-checker NuSMV [CCGR99] to generate proof-like witnesses for CTL formulas, and the daVinci [FW94] graph visualization package for their visual presentation. KEGVis can give the user a high-level overview of proof-like witnesses and allow him/her to skip certain steps, fast forward to “interesting” parts of the witness, determine whether the current state is part of a loop, etc.

Figure 12 lists several examples of strategies for efficient navigation through proof-like counter-examples. We distinguish between static and dynamic modes of navigation. In the former, the witness is generated prior to browsing, while in the latter, it is generated in response to user actions.

In the static mode, the exploration can proceed either in the forward direction, convenient for discovering why a particular path in a witness justifies a given subformula, or in the backward direction, used for identifying why a particular “suspicious” state is part of the witness. In both cases, the user can restrict the part of the witness of property φ to be displayed by selecting subformulas or operators of φ . This restriction is implemented in KEGVis by letting the user specify the initial condition, i.e. the one that should hold in the beginning of a path to explore, and the final condition, i.e. the one that tells KEGVis when to stop the exploration. For example, to restrict the witness to the EF operator of property $P = EGEF(x \wedge EXx)$, we set the initial and the final conditions to $EF(x \wedge EXx)$ and $x \wedge EXx$, respectively. Note that there are several paths satisfying the initial condition; all of them will be displayed. If a chosen subformula or operator φ occurs in multiple places in the temporal formula (e.g., x in property P above), then paths satisfying φ can be found by the following algorithm: (1) find the smallest unique subformula ψ that contains φ ; (2) find all states that justify ψ ; (3) for each state t found in step (2), follow the proof to a state labeled with φ . For example, the smallest unique subformula that contains the second occurrence of x in P is EXx . Thus, this x can be justified by successors of states labelled with EXx .

Fig. 12. Browsing strategies for proof-like witnesses/counter-examples.

Witness generation is usually quite expensive. The dynamic mode of witness exploration allows the user to control which path of the witness is to be generated, or what should be the granularity of the logical step before further user input is solicited. These choices are based on the information available from the proof part of the proof-like witness, and can be taken manually or by a user-supplied heuristic function. Consider exploring the witness for the formula $\llbracket EFx \wedge EFy \rrbracket(s)$. After a few steps of the exploration, the proof labeling state s may indicate that fixpoints for EFx and EFy have been achieved in 3 and 5 iterations, respectively:

$$\frac{\frac{\llbracket E[\top U_3 x] \rrbracket(s)}{\llbracket EFx \rrbracket(s)} \quad \frac{\llbracket E[\top U_5 y] \rrbracket(s)}{\llbracket EFy \rrbracket(s)}}{\llbracket EFx \wedge EFy \rrbracket(s)} \wedge\text{-intro}$$

The user can then use this information to guide the witness generation in the direction of the shorter path.

Current implementation of KEGVis supports static exploration only; the dynamic exploration mode is still under development.

5 Related Work

The proof system developed in this paper is similar to the tableaux used for local model-checking [SW91]. In fact, the automated proof-generation technique can be seen as a simulation of a run of a local model-checker, where the information collected from the run of a global model-checker is used to guide the construction of the proof.

Several other researches have explored the idea of generating proofs from the model-checking runs [PZ01, PPZ01, Nam01, TC02]. Their common motivation is to use proofs to provide *complete* information justifying the result of the model-checker, in particular, in cases where a witness, or a counter-example, is not feasible.

The work of Namjoshi [Nam01] concentrates on generating a proof of validity of a run of a global μ -calculus model-checker. This is accomplished by augmenting the

model-checker to record the set of states satisfying each subformula and the convergence bound for each fixpoint operator. In cases where the formula has a finite witness, this information can also be used to produce proof-like witnesses, using techniques similar to the one described in Section 3.2. Tan and Cleaveland [TC02] extend Namjoshi's work to local model-checking.

An approach of Peled et. al. [PZ01,PPZ01] is similar to ours in spirit. The main goal is to generate proofs that are as easy to understand as possible, and use them to communicate the reasons behind the model-checking result to the user. The technique developed in [PZ01,PPZ01] generates proofs of satisfaction for LTL properties. Note that this is similar to providing witnesses for ACTL; thus, the approach is complimentary to ours, as it addresses the problem of discharging proof obligations that we generate in the case of AX .

In this paper, we have only explored how proofs can be used to aid in understanding witnesses, yet other applications are also possible. For example, Namjoshi [Nam01] suggests that one can use proofs extracted from the model-checker to debug the model-checker itself, or use them as a basis for integration between model-checking and theorem proving. Since proof-like witnesses are effectively proofs, they can be used in any application that calls for a proof of satisfaction of a temporal property.

6 Conclusion

The main contribution of this paper is the concept of *proof-like witnesses* (and counterexamples) that bridges the gap between proofs and witnesses. Our proof-like witnesses are expressed in terms of computations of the model (thus making them easy to understand), yet allow us to provide some feedback to the user even if in general the witness is not available. Overall, we have shown that expressing witnesses as proofs, in combination with an intelligent browser tool, provides a significant support for understanding complex witnesses.

In order to create proof-like witnesses, we have developed a sound and complete proof system for CTL. This makes our approach a logical, rather than an algorithmic, one. The main advantage here is that a proof of correctness of our witness generator is significantly simplified; furthermore, the technique can be easily extended to other, non-traditional, applications such as witness generation for temporal-logic queries [GDC02].

The major limitation of our approach is that it provides only limited information in the case of witnesses for ACTL and properties that use both universal and existential quantifiers. In the future, we plan to explore possible solutions to this problem along the lines of [PPZ01]. We also plan to enhance our tool, KEGVis, to support dynamic witness exploration.

References

- [CCGR99] A. Cimatti, E.M. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: a new Symbolic Model Verifier. In N. Halbwachs and D. Peled, editors, *Proceedings of 11th Conference on Computer-Aided Verification (CAV'99)*, number 1633 in Lecture Notes in Computer Science, pages 495–499, Trento, Italy, July 1999. Springer.

- [CGMZ95] E.M. Clarke, O. Grumberg, K.L. McMillan, and X. Zhao. Efficient Generation of Counterexamples and Witnesses in Symbolic Model Checking. In *Proceedings of 32nd Design Automation Conference (DAC 95)*, pages 427–432, San Francisco, CA, USA, 1995.
- [CGP99] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [CLJV02] E.M. Clarke, Y. Lu, S. Jha, and H. Veith. Tree-Like Counterexamples in Model Checking. In *Proceedings of the Seventeenth Annual IEEE Symposium on Logic in Computer Science (LICS'02)*, pages 19–29, Copenhagen, Denmark, July 2002. IEEE Computer Society.
- [FW94] M. Fröhlich and M. Werner. The Graph Visualization System daVinci – A user interface for applications. Technical Report 5/94, Department of Computer Science, Bremen University, 1994.
- [GDC02] A. Gurfinkel, B. Devereux, and M. Chechik. “Model Exploration with Temporal Logic Query Checking”. In *Proceedings of SIGSOFT Conference on Foundations of Software Engineering (FSE'02)*, Charleston, South Carolina, November 2002. ACM Press.
- [Gur02] A. Gurfinkel. Multi-valued symbolic model-checking: Fairness, counter-examples, running time. Master’s thesis, University of Toronto, Department of Computer Science, October 2002.
- [Lar95] F. Laroussinie. “About the Expressive Power of CTL Combinators”. *Information Processing Letters*, 54:343–345, 1995.
- [McM93] K.L. McMillan. *Symbolic Model Checking*. Kluwer Academic, 1993.
- [Nam01] K. Namjoshi. Certifying Model Checkers. In *Proceedings of 13th International Conference on Computer-Aided Verification (CAV'01)*, volume 2102 of LNCS. Springer-Verlag, 2001.
- [OSR93] S. Owre, N. Shankar, and J. Rushby. “User Guide for the PVS Specification and Verification System (Draft)”. Technical report, Computer Science Lab, SRI International, Menlo Park, CA, 1993.
- [PPZ01] D. Peled, A. Pnueli, and L. Zuck. From falsification to verification. In *FST&TCS*, volume 2245 of LNCS. Springer-Verlag, 2001.
- [PZ01] D. Peled and L. Zuck. From model checking to a temporal proof. In *Proceedings of the 8th International SPIN Workshop (SPIN'2001)*, volume 2057 of LNCS, pages 1–14, Toronto, Canada, May 2001. Springer.
- [SW91] C. Stirling and D. Walker. Local model-checking in the modal μ -calculus. *Theoretical Computer Science*, 89, 1991.
- [TC02] L. Tan and R. Cleaveland. Evidence-Based Model Checking. In *Proceedings of 14th Conference on Computer-Aided Verification (CAV'02)*, volume 2404 of LNCS, pages 455–470, Copenhagen, Denmark, July 2002. Springer-Verlag.