Counterexample-guided Abstraction Refinement

E. Clarke, O. Grumberg, S. Jha, Y. Lu, H. Veith

January 2000
CMU-CS-00-103

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

We present an automatic iterative abstraction-refinement methodology in which the initial abstract model is generated by an automatic analysis of the control structures in the program to be verified. Abstract models may admit erroneous (or "spurious") counterexamples. We devise new symbolic techniques which analyze such counterexamples and refine the abstract model correspondingly. The refinement algorithm keeps the size of the abstract state space small due to the use of abstraction functions which distinguish many degrees of abstraction for each program variable. We describe an implementation of our methodology in NuSMV. Practical experiments including a large Fujitsu IP core design with about 500 latches and 10000 lines of SMV code confirm the effectiveness of our approach.

This research is sponsored by the Semiconductor Research Corporation (SRC) under Contract No. 98-DJ-294, the National Science Foundation (NSF) under Grant No. CCR-9803774, and the Max Kade Foundation. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of SRC, NSF, or the United States Government.

Keywords: Abstraction, Model checking, and Counterexample.

1 Introduction

The state explosion problem remains a major hurdle in applying model checking to large industrial designs. Abstraction is certainly the most important technique for handling this problem. In fact, it is essential for verifying designs of industrial complexity. Currently, abstraction is typically a manual process, often requiring considerable creativity. In order for model checking to be used more widely in industry, automatic techniques are needed for generating abstractions. In this paper, we describe an automatic abstraction technique for ACTL* specifications which is based on an analysis of the structure of formulas appearing in the program. In general, our technique computes an upper approximation of the original program. Thus, when a specification is true in the abstract model, it will also be true in the concrete design. However, if the specification is false in the abstract model, the counterexample may be the result of some behavior in the approximation which is not present in the original model. When this happens, it is necessary to refine the abstraction so that the behavior which caused the erroneous counterexample is eliminated. The main contribution of this paper is an efficient automatic refinement technique which uses information obtained from erroneous counterexamples. The refinement algorithm keeps the size of the abstract state space small due to the use of abstraction functions which distinguish many degrees of abstraction for each program variable. Practical experiments including a large Fujitsu IP core design with about 500 latches and 10000 lines of SMV code confirm the competitiveness of our implementation. Although our current implementation is based on NuSMV, it is in principle not limited to the input language of SMV and can be applied to other languages.

Our paper follows the general framework established by Clarke, Grumberg, and Long [10]. We assume that the reader has some familiarity with that framework. In our methodology, *atomic formulas* are automatically extracted from the program that describes the model. The atomic formulas are similar to the *predicates* used for abstraction by Graf and Saidi [14] and later in [11, 20]. However, instead of using the atomic formulas to generate an abstract global transition system, we use them to construct an explicit *abstraction function*. The abstraction function preserves logical relationships among the atomic formulas instead of treating them as independent propositions. The initial abstract model is constructed by adapting the *existential abstraction* techniques proposed in [8, 10] to our framework. Then, a traditional model checker is used to determine whether ACTL* properties hold in the abstract model (ACTL* is a fragment of CTL* which only allows universal quantification over paths). If the answer is yes, then the concrete model also satisfies the property. If the answer is no, then the model checker generates a counterexample. Since the abstract model has more behaviors than the concrete one, the abstract counterexample might not be valid. We say that such a counterexample is *spurious*.

In our methodology, we provide a new symbolic algorithm to determine whether an abstract counterexample is *spurious*. If the counterexample is not spurious, we report it to the user and stop. If the counterexample is spurious, the abstraction function must be refined to eliminate it. In our methodology, we identify the shortest prefix of the abstract counterexample that does not correspond to an actual trace in the concrete model. The last abstract state in this prefix is split into less abstract states so that the spurious counterexample is eliminated. Thus, a more refined abstraction function is obtained. Note that there may be many ways of splitting the abstract state; each determines a different refinement of the abstraction function. It is desirable to obtain the coarsest refinement which eliminates the counterexample because this corresponds to the *smallest* abstract model that is suitable for verification. We prove, however, that finding the coarsest refinement is NP-hard. Because of this, we use a polynomial-time algorithm which gives a suboptimal but sufficiently good refinement of the abstraction function. The applicability of our heuristic algorithm is confirmed by our experiments. Using the refined abstraction function obtained in this manner, a new abstract model is built and the entire process is repeated. Our methodology is complete for the fragment of ACTL* which has counterexamples that are either paths or loops, i.e., we are guaranteed to either find a valid counterexample or prove that the system satisfies the desired property. In principle, our methodology can be extended to all of ACTL*.

Using counterexamples to refine abstract models has been investigated by a number of other researchers beginning with the *localization reduction* of Kurshan [15]. He models a concurrent system as a composition of L-processes $L_1, \ldots, L_n$ (L-processes are described in detail in [15]). The localization reduction is an iterative technique that starts with a small subset of relevant L-processes that are topologically close to the specification in the *variable dependency graph*. All other program variables are abstracted away with nondeterministic assignments. If the counterexample is found to be spurious, additional variables are added to eliminate the counterexample. The heuristic for selecting these variables also uses information from the variable dependency graph. Note that the localization reduction either leaves a variable unchanged or replaces it by a nondeterministic assignment. A similar approach has been described by Balarin in [2]. In our approach, the abstraction functions exploit logical relationships among variables appearing in atomic formulas that occur in the control structure of the program. Moreover, the way we use abstraction functions makes it possible to distinguish many degrees of abstraction for each variable. Therefore, in the refinement step only very small and local changes to the abstraction functions are necessary and the abstract model remains comparatively small.

Another refinement technique has recently been proposed by Lind-Nielson and Andersen [17]. Their model checker uses upper and lower approximations in order to handle all of CTL. Their approximation techniques enable them to avoid rechecking the entire model after each refinement step while guaranteeing completeness. As in [2, 15] the variable dependency graph is used both to obtain the initial abstraction and in the refinement process. Variable abstraction is also performed in a similar manner. Therefore, our abstraction-refinement methodology relates to their technique in essentially the same way as it relates to the classical localization reduction.

A number of other papers [16, 18, 19] have proposed abstraction-refinement techniques for CTL model checking. However, these papers do not use counterexamples to refine the abstraction. We believe that the methods described in these papers are orthogonal to our technique and may even be combined with ours in order to achieve better performance. A recent technique proposed by Govindaraju and Dill [13] may be a starting point in this direction, since it also tries to identify the first spurious state in an abstract counterexample. It randomly chooses a concrete state corresponding to the first spurious state and tries to construct a real counterexample starting with the image of this state under the transition relation. The paper only talks about safety properties and path counterexamples. It does not describe how to check liveness properties with cyclic counterexamples. Furthermore, our method does not use random choice to extend the counterexample; instead it analyzes the cause of the spurious counterexample and uses this information to guide the refinement process.

Summarizing, our technique has a number of advantages over previous work:

(i) The technique is complete for an important fragment of ACTL*.

(ii) The initial abstraction and the refinement steps are efficient and entirely automatic. All algorithms are symbolic.

(iii) In comparison to methods like the localization reduction, we distinguish more degrees of abstraction for each variable. Thus, the changes in the refinement are potentially finer in our approach.

(iv) The refinement procedure is guaranteed to eliminate spurious counterexamples while keeping the state space of the abstract model small.

We have implemented our new methodology in NuSMV [6] and applied it to a number of benchmark designs [6]. In addition we have used it to debug a large IP core being developed at Fujitsu [1]. The design has about 500 latches and 10000 lines of Verilog code. Before using our methodology, we implemented the *cone of influence* reduction [8] in NuSMV to enhance its ability to check large models. Neither our enhanced version of NuSMV nor the recent version of SMV developed by Yang [23] were able to verify the Fujitsu IP core design. However, by using our new technique, we were able to find a subtle error in the design. Our program automatically abstracted 144 symbolic variables and performed three refinement steps. Currently, we are evaluating the methodology on other complex industrial designs.

The paper is organized as follows: Section 2 gives the basic definitions and terminology used throughout the paper. A general overview of our methodology is given in Section 3. Detailed descriptions of our abstraction-refinement algorithms are provided in Section 4. Performance improvements for the implementation are described in Section 5. Experimental results are presented in Section 6. Future research is discussed in Section 7.

2 Preliminaries

A *program* P has a finite set of variables $V = \{v_1, \cdots, v_n\}$, where each variable v_i has an associated finite domain D_{v_i}. The set of all possible states for program P is $D_{v_1} \times \cdots D_{v_n}$ which we denote by D. *Expressions* are built from variables in V, constants in D_{v_i}, and function symbols in the usual way, e.g. $v_1 + 3$. *Atomic formulas* are constructed from expressions and relation symbols, e.g. $v_1 + 3 < 5$. Similarly, *predicates* are composed of atomic formulas using negation ($\neg$), conjunction ($\wedge$), and disjunction ($\vee$). Given a predicate p, Atoms(p) is the set of atomic formulas occurring in it. Let p be a predicate containing variables from V, and $d = (d_1, \ldots, d_n)$ be an element from D. Then we write $d \models p$ when the predicate obtained by replacing each occurrence of the variable v_i in p by the constant d_i evaluates to true.

Each variable v_i in the program has an associated *transition block*, which defines both the initial value and the transition relation for the variable v_i. An example of a transition block for the variable v_i is shown in Figure 1, where $I_i \subseteq D_{v_i}$ is the initial expression for

$\textbf{init}(v_i) := I_i;$	$\textbf{init}(x) := 0;$	$\textbf{init}(y) := 1;$
$\textbf{next}(v_i) := \textbf{case}$	$\textbf{next}(x) := \textbf{case}$	$\textbf{next}(y) := \textbf{case}$
$\quad C_i^1 : A_i^1;$	$\quad reset = \mathrm{TRUE} : 0;$	$\quad reset = \mathrm{TRUE} : 0;$
$\quad C_i^2 : A_i^2;$	$\quad x < y : x + 1;$	$\quad (x = y) \wedge \neg(y = 2) : y + 1;$
$\quad \cdots : \cdots;$	$\quad x = y : 0;$	$\quad (x = y) : 0;$
$\quad C_i^k : A_i^k;$	$\quad \textbf{else} : x;$	$\quad \textbf{else} : y;$
$\textbf{esac};$	$\textbf{esac};$	$\textbf{esac};$

Figure 1: A generic transition block and a typical example

the variable v_i, each condition C_i^j is a predicate, and A_i^j is an expression. The semantics of the transition block is similar to the semantics of the **case** statement in the modeling language of SMV, i.e., find the least j such that in the current state condition C_i^j is true and assign the value of the expression A_i^j to the variable v_i in the next state. Common hardware description languages like Verilog and VHDL can easily be compiled into this language.

We assume that the specifications are written in a fragment of CTL* called ACTL* (see [10]), where atomic formulas are used at the lowest level. ACTL* is the fragment of CTL*, where negation is restricted to the atomic level, and path quantification is restricted to universal path quantification. Assume that we are given an ACTL* specification φ, and a program P. For each transition block B_i let $\mathrm{Atoms}(B_i)$ be the set of atomic formulas that appear in the conditions. Let $\mathrm{Atoms}(\varphi)$ be the set of atomic formulas appearing in the specification φ. $\mathrm{Atoms}(P)$ is the set of atomic formulas that appear in the specification or in the conditions of the transition blocks.

Each program P naturally corresponds to a labeled *Kripke structure* $M = (S, I, R, L)$, where $S = D$ is the set of states, $I \subseteq S$ is a set of initial states, $R \subseteq S \times S$ is a transition relation, and $L : S \to 2^{\mathrm{Atoms}(\varphi)}$ is a labelling given by $L(d) = \{f \in \mathrm{Atoms}(\varphi) \mid d \models f\}$. Translating a program into a Kripke structure is straightforward and will not be described here.

An abstraction h for a program P is given by a surjection $h : D \to \widehat{D}$. Notice that the surjection h induces an equivalence relation $\equiv$ on the domain D in the following manner: let d, e be states in D, then

$$d \equiv e \text{ iff } h(d) = h(e).$$

Since an abstraction can be represented either by a surjection h or by an equivalence relation $\equiv$, we sometimes switch between these representations to avoid notational overhead.

Assume that we are given a program P and an abstraction function h for P. The *abstract Kripke structure* $\widehat{M} = (\widehat{S}, \widehat{I}, \widehat{R}, \widehat{L})$ corresponding to the abstraction function h is defined as follows:

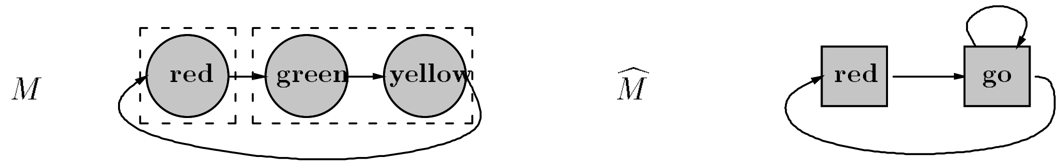

Figure 2: Abstraction of a Traffic Light.

1. $\widehat{S}$ is the abstract domain $\widehat{D}$.

2. $\widehat{I}(\widehat{d})$ iff $\exists d(h(d) = \widehat{d} \wedge I(d))$.

3. $\widehat{R}(\widehat{d_1}, \widehat{d_2})$ iff $\exists d_1 \exists d_2 (h(d_1) = \widehat{d_1} \wedge h(d_2) = \widehat{d_2} \wedge R(d_1, d_2))$.

4. $\widehat{L}(\widehat{d}) = \bigcup_{h(d)=\widehat{d}} L(d)$. (This definition will be justified in Theorem 2.1.)

This abstraction technique is called *existential abstraction* [8]. An atomic formula f *respects* an abstraction function h if for all d and d' in the domain D, $(d \equiv d') \Rightarrow (d \models f \Leftrightarrow d' \models f)$. Let $\widehat{d}$ be an abstract state. $\widehat{L}(\widehat{d})$ is *consistent*, if all concrete states corresponding to $\widehat{d}$ satisfy all labels in $\widehat{L}(\widehat{d})$, i.e., for all $d \in h^{-1}(\widehat{d})$ it holds that $d \models \bigwedge_{f \in \widehat{L}(\widehat{d})} f$.

Theorem 2.1 *Let h be an abstraction and φ be an* ACTL* *specification where the atomic subformulas respect h. Then the following holds: (i) $\widehat{L}(\widehat{d})$ is consistent for all abstract states $\widehat{d}$ in $\widehat{M}$; (ii) $\widehat{M} \models \varphi \;\Rightarrow\; M \models \varphi$.*

In other words, correctness of the abstract model implies correctness of the concrete model. On the other hand, if the abstract model invalidates an ACTL* specification, i.e., $\widehat{M} \not\models \varphi$, *the actual model may still satisfy the specification.*

Example 2.1 *Assume that for a traffic light controller (see Figure 2), we want to prove $\psi = \mathbf{AG}\,\mathbf{AF}(state = red)$ using the abstraction function $h(red) = red$ and $h(green) = h(yellow) = go$. It is easy to see that $M \models \psi$ while $\widehat{M} \not\models \psi$. There exists an infinite trace $\langle red, go, go, \dots \rangle$ that invalidates the specification.*

If an abstract counterexample does not correspond to some concrete counterexample, we call it *spurious*. For example, $\langle red, go, go, \dots \rangle$ in the above example is a spurious counterexample.

When the set of possible states is given as the product $D_1 \times \cdots D_n$ of smaller domains, an abstraction h can be described by surjections $h_i : D_i \to \widehat{D_i}$, such that $h(d_1, \dots, d_n)$ is equal to $(h_1(d_1), \dots, h_n(d_n))$, and $\widehat{D}$ is equal to $\widehat{D_1} \times \cdots \widehat{D_n}$. In this case, we write $h = (h_1, \dots, h_n)$. The equivalence relations $\equiv_i$ corresponding to the individual surjections h_i induce an equivalence relation $\equiv$ over the entire domain $D = D_1 \times \cdots \times D_n$ in the obvious manner:

$$(d_1, \cdots, d_n) \equiv (e_1, \cdots, e_n) \text{ iff } d_1 \equiv_1 e_1 \wedge \cdots \wedge d_n \equiv_n e_n$$

In previous work on existential abstraction [10], abstractions were defined for each variable domain, i.e., D_i in the above paragraph was chosen to be D_{v_i}, where D_{v_i} is the set of possible values for variable v_i. Unfortunately, many abstraction functions h can not be described in this simple manner. For example, let $D = \{0,1,2\} \times \{0,1,2\}$, and $\widehat{D} = \{0,1\} \times \{0,1\}$. Then there are $4^9 = 262144$ functions h from D to $\widehat{D}$. Next, consider $h = (h_1, h_2)$. Since there are $2^3 = 8$ functions from $\{0,1,2\}$ to $\{0,1\}$, there are only 64 functions of this form from D to $\widehat{D}$.

In this paper, we define abstraction functions in a different way. We partition the set V of variables into sets of related variables called *variable clusters* $VC_1, \ldots, VC_m$, where each variable cluster VC_i has an associated domain $D_{VC_i} := \prod_{v \in VC_i} D_v$. Consequently, $D = D_{VC_1} \times \cdots D_{VC_m}$. We define abstraction functions as surjections on the domains D_{VC_i}, i.e., D_i in the above paragraph is equal to D_{VC_i}. Thus, the notion of abstraction used in this paper is more general than the one used in [10].

3 Overview

For a program P and an ACTL* formula φ, our goal is to check whether the Kripke structure M corresponding to P satisfies φ. Our methodology consists of the following steps.

1. *Generate the initial abstraction:* We generate an initial abstraction h by examining the transition blocks corresponding to the variables of the program. We consider the conditions used in the **case** statements and construct variable clusters for variables which interfere with each other via these conditions. Details can be found in Section 4.1.

2. *Model-check the abstract structure:* Let $\widehat{M}$ be the abstract Kripke structure corresponding to the abstraction h. We check whether $\widehat{M} \models \varphi$. If the check is affirmative, then we can conclude that $M \models \varphi$ (see Theorem 2.1). Suppose the check reveals that there is a counterexample $\widehat{T}$. We ascertain whether $\widehat{T}$ is an actual counterexample, i.e., a counterexample in the unabstracted structure M. If $\widehat{T}$ turns out to be an actual counterexample, we report it to the user, otherwise $\widehat{T}$ is a spurious counterexample, and we proceed to step 3.

3. *Refine the abstraction:* We refine the abstraction function h by partitioning a *single equivalence class* of $\equiv$ so that after the refinement the abstract structure $\widehat{M}$ corresponding to the refined abstraction function does not admit the spurious counterexample $\widehat{T}$. We will discuss partitioning algorithms for this purpose in Section 4.3. After refining the abstraction function, we return to step 2.

4 The Abstraction-Refinement Framework

4.1 Generating The Initial Abstraction

Assume that we are given a program P with n variables $\{v_1, \cdots, v_n\}$. Given an atomic formula f, let $var(f)$ be the set of variables appearing in f, e.g., $var(x = y)$ is $\{x, y\}$. Given a set of atomic formulas U, $var(U)$ equals $\bigcup_{f \in U} var(f)$. In general, for any syntactic entity X, $var(X)$ will be the set of variables appearing in X. We say that two atomic formulas f_1 and f_2 *interfere* iff $var(f_1) \cap var(f_2) \neq \emptyset$. Let $\equiv_I$ be the equivalence relation on Atoms(P) that is the reflexive, transitive closure of the interference relation. The equivalence class of an atomic formula $f \in$ Atoms(P) is called the *formula cluster* of f and is denoted by $[f]$. Let f_1 and f_2 be two atomic formulas. Then $var(f_1) \cap var(f_2) \neq \emptyset$ implies that $[f_1] = [f_2]$. In other words, a variable v_i cannot appear in formulas that belong to two different formula clusters. Moreover, the formula clusters induce an equivalence relation $\equiv_V$ on the set of variables V in the following way:

> $v_i \equiv_V v_j$ if and only if v_i and v_j appear in atomic formulas that belong to the same formula cluster.

The equivalence classes of $\equiv_V$ are called *variable clusters*. For instance, consider a formula cluster $FC_i = \{v_1 > 3, v_1 = v_2\}$. The corresponding variable cluster is $VC_i = \{v_1, v_2\}$. Let $\{FC_1, \ldots, FC_m\}$ be the set of formula clusters and $\{VC_1, \ldots, VC_m\}$ the set of corresponding variable clusters. We construct the initial abstraction $h = (h_1, \ldots, h_m)$ as follows. For each h_i, we set $D_{VC_i} = \prod_{v \in VC_i} D_v$, i.e., D_{VC_i} is the domain corresponding to the variable cluster VC_i. Since the variable clusters form a partition of the set of variables V, it follows that $D = D_{VC_1} \times \cdots D_{VC_m}$. For each variable cluster $VC_i = \{v_{i_1}, \ldots, v_{i_k}\}$, the corresponding abstraction h_i is defined on D_{VC_i} as follows. $h_i(d_1, \cdots, d_k) = h_i(e_1, \cdots, e_k)$ iff for all atomic formulas $f \in FC_i$,

$$(d_1, \cdots, d_k) \models f \Leftrightarrow (e_1, \cdots, e_k) \models f.$$

In other words two values are in the same equivalence class if they cannot be "distinguished" by atomic formulas appearing in the formula cluster FC_i. The following example illustrates how we construct the initial abstraction h.

Example 4.1 *Consider the program P with three variables $x, y \in \{0, 1, 2\}$, and reset $\in$ $\{$TRUE, FALSE$\}$ shown in Figure 1. The set of atomic formulas is* Atoms(P) $= \{(reset =$ TRUE$), (x = y), (x < y), (y = 2)\}$. *There are two formula clusters, $FC_1 = \{(x = y), (x < y), (y = 2)\}$ and $FC_2 = \{(reset =$ TRUE$)\}$. The corresponding variable clusters are $\{x, y\}$ and $\{reset\}$, respectively. Consider the formula cluster FC_1. Values $(0, 0)$ and $(1, 1)$ are in the same equivalence class because for all the atomic formulas f in the formula cluster FC_1 it holds that $(0, 0) \models f$ iff $(1, 1) \models f$. It can be shown that the domain $\{0, 1, 2\} \times \{0, 1, 2\}$ is partitioned into a total of five equivalence classes by this criterion. We denote these classes by the natural numbers $0, 1, 2, 3, 4$, and list them below:*

$$1 = \{(0,0), (1,1)\},\ 2 = \{(0,1)\},\ 3 = \{(0,2), (1,2)\},\ 4 = \{(1,0), (2,0), (2,1)\},\ 5 = \{(2,2)\}$$

The domain $\{$TRUE, FALSE$\}$ has two equivalence classes – one containing FALSE *and the other* TRUE. *Therefore, we define two abstraction functions $h_1 : \{0, 1, 2\}^2 \to \{0, 1, 2, 3, 4\}$ and $h_2 : \{$TRUE, FALSE$\} \to \{$TRUE, FALSE$\}$. The first function h_1 is given by $h_1(0, 0) = h_1(1, 1) = 0$, $h_1(0, 1) = 1$, $h_1(0, 2) = h_1(1, 2) = 2$, $h_1(1, 0) = h_1(2, 0) = h_1(2, 1) = 3$, $h_1(2, 2) = 4$. The second function h_2 is just the identity function, i.e., $h_2(reset) = reset$.*

4.2 Model Checking The Abstract Model

Given an ACTL* specification φ, an abstraction function h (assume that φ respects h), and a program P with a finite set of variables $V = \{v_1, \cdots, v_n\}$, let $\widehat{M}$ be the abstract Kripke structure corresponding to the abstraction function h. We use standard symbolic model checking procedures to determine whether $\widehat{M}$ satisfies the specification φ. If it does, then by Theorem 2.1 we can conclude that the original Kripke structure also satisfies φ. Otherwise, assume that the model checker produces a counterexample $\widehat{T}$ corresponding to the abstract model $\widehat{M}$. In the rest of this section, we will focus on counterexamples which are either *(finite) paths* or *loops*.

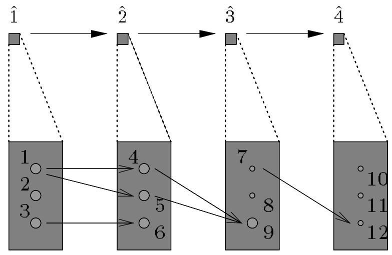

Figure 3: An abstract counterexample

4.2.1 Identification Of Spurious Path Counterexamples

First, we will tackle the case when the counterexample $\widehat{T}$ is a path $\langle \widehat{s_1}, \cdots, \widehat{s_n} \rangle$. Given an abstract state $\widehat{s}$, the set of concrete states s such that $h(s) = \widehat{s}$ is denoted by $h^{-1}(\widehat{s})$, i.e., $h^{-1}(\widehat{s}) = \{s | h(s) = \widehat{s}\}$. We extend h^{-1} to sequences in the following way: $h^{-1}(\widehat{T})$ is the set of concrete paths given by the following expression

$$\{\langle s_1, \cdots, s_n \rangle | \bigwedge_{i=1}^{n} h(s_i) = \widehat{s_i} \wedge I(s_1) \wedge \bigwedge_{i=1}^{n-1} R(s_i, s_{i+1})\}.$$

We will occasionally write h^{-1}_{path} to emphasize the fact that h^{-1} is applied to a sequence. Next, we give a *symbolic* algorithm to compute $h^{-1}(\widehat{T})$. Let $S_1 = h^{-1}(\widehat{s_1}) \cap I$ and R be the transition relation corresponding to the unabstracted Kripke structure M. For $1 < i \leq n$, we define S_i in the following manner: $S_i := Img(S_{i-1}, R) \cap h^{-1}(\widehat{s_i})$. In the definition of S_i, $Img(S_{i-1}, R)$ is the forward image of S_{i-1} with respect to the transition relation R. The sequence of sets S_i is computed symbolically using OBDDs and the standard image computation algorithm. The following lemma establishes the correctness of this procedure.

Lemma 4.1 *The following are equivalent:*

(i) The path $\widehat{T}$ corresponds to a concrete counterexample.

(ii) The set of concrete paths $h^{-1}(\widehat{T})$ is non-empty.

(iii) For all $1 \le i \le n$, $S_i \neq \emptyset$.

Suppose that condition (iii) of Lemma 4.1 is violated, and let i be the largest index such that $S_i \neq \emptyset$. Then $\widehat{s_i}$ is called the *failure state* of the spurious counterexample $\widehat{T}$.

Example 4.2 *Consider a program with only one variable with domain* $D = \{1, \cdots, 12\}$. *Assume that the abstraction function* h *maps* $x \in D$ *to* $\lfloor (x-1)/3 \rfloor + 1$. *There are four abstract states corresponding to the equivalence classes* $\{1,2,3\}$, $\{4,5,6\}$, $\{7,8,9\}$, *and* $\{10,11,12\}$. *We call these abstract states* $\widehat{1}$, $\widehat{2}$, $\widehat{3}$, *and* $\widehat{4}$. *The transitions between states in the concrete model are indicated by the arrows in Figure 3; small dots denote non-reachable states. Suppose that we obtain an abstract counterexample* $\widehat{T} = \langle \widehat{1}, \widehat{2}, \widehat{3}, \widehat{4} \rangle$. *It is easy to see that* $\widehat{T}$ *is spurious. Using the terminology of Lemma 4.1, we have* $S_1 = \{1,2,3\}$, $S_2 = \{4,5,6\}$, $S_3 = \{9\}$, *and* $S_4 = \emptyset$. *Notice that* S_4 *and therefore* $Img(S_3, R)$ *are both empty. Thus,* $\widehat{s_3}$ *is the failure state.*

Algorithm SplitPATH($\widehat{T}$)

$S := h^{-1}(\widehat{s_1}) \cap I$
$j := 1$
while ($S \neq \emptyset$ and $j < n$) {
  $j := j + 1$
  $S_{\text{prev}} := S$
  $S := Img(S, R) \cap h^{-1}(\widehat{s_j})$ }
if $S \neq \emptyset$ **then** output "counterexample exists"
else output j, S_{prev}

Figure 4: SplitPATH checks if an abstract path is spurious.

It follows from Lemma 4.1 that if $h^{-1}(\widehat{T})$ is empty (i.e., if the counterexample $\widehat{T}$ is spurious), then there exists a minimal i $(2 \le i \le n)$ such that $S_i = \emptyset$. The symbolic Algorithm **SplitPATH** in Figure 4 computes this number and the set of states S_{i-1}; the states in S_{i-1} are called *dead-end* states. After the detection of the dead-end states, we proceed to the refinement step (see Section 4.3). On the other hand, if the conditions stated in Lemma 4.1 are true, then **SplitPATH** will report a "real" counterexample and we can stop.

4.2.2 Identification of Spurious Loop Counterexamples

Now we consider the case when the counterexample $\widehat{T}$ includes a loop, which we write as $\langle \widehat{s_1}, \cdots, \widehat{s_i} \rangle \langle \widehat{s_{i+1}}, \cdots, \widehat{s_n} \rangle^{\omega}$. The loop starts at the abstract state $\widehat{s_{i+1}}$ and ends at $\widehat{s_n}$. Since this case is more complicated than the path counterexamples, we first present an example in which some of the typical situations occur.

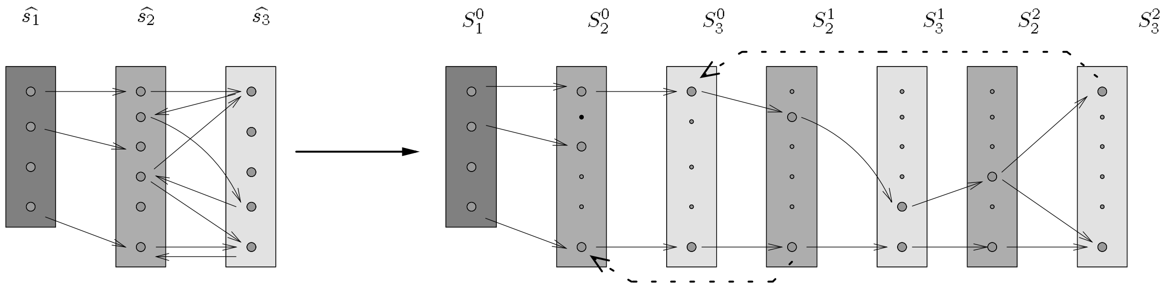

Figure 5: A loop counterexample, and its unwinding.

Example 4.3 *We consider a loop $\langle \widehat{s}_1 \rangle \langle \widehat{s}_2, \widehat{s}_3 \rangle^\omega$ as shown in Figure 5. In order to find out if the abstract loop corresponds to concrete loops, we unwind the counterexample as demonstrated in the figure. There are two situations where cycles occur. In the figure, for each of these situations, an example cycle (the first one occurring) is indicated by a fat dashed arrow. We make the following important observations: (i) A given abstract loop may correspond to several concrete loops of* different size. *(ii) Each of these loops may start at different stages of the unwinding. (iii) The unwinding eventually becomes periodic (in our case $S_3^0 = S_3^2$), but only after several stages of the unwinding. The size of the period is the least common multiple of the size of the individual loops, and thus, in general* exponential.

We conclude from the example that a naive algorithm may have exponential time complexity due to an exponential number of loop unwindings. The following surprising theorem however shows that a polynomial number of unwindings is sufficient. Let min be the minimum size of all abstract states in the loop, i.e., $min = \min_{i+1 \le j \le n} |h^{-1}(\widehat{s}_j)|$. $\widehat{T}_{\text{unwind}}$ denotes the the finite abstract path $\langle \widehat{s}_1, \ldots, \widehat{s}_i \rangle \langle \widehat{s_{i+1}}, \ldots, \widehat{s}_n \rangle^{min+1}$, i.e., the path obtained by unwinding the loop part of $\widehat{T}$ $min + 1$ times.

Theorem 4.1 *The following are equivalent:*

(i) $\widehat{T}$ corresponds to a concrete counterexample.

(ii) $h^{-1}_{\text{path}}(\widehat{T}_{\text{unwind}})$ is not empty.

We conclude that loop counterexamples can be reduced to path counterexamples. In Figure 6, we describe the algorithm **SplitLOOP** which is an extension of **SplitPATH**. In the algorithm, $\widehat{T}_{\text{unwind}}$ is computed by the subprogram **unwind**. The subprogram **LoopIndex**(j) computes the index of the abstract state at position j in the unwound counterexample $\widehat{T}_{\text{unwind}}$, i.e.,

$$\textbf{LoopIndex}(j) = \begin{cases} j & \text{if } j \le n \\ ((j - i - 1) \bmod (n - i)) + (i + 1) & \text{otherwise} \end{cases}$$

If the abstract counterexample is spurious, then the algorithm **SplitLOOP** outputs a set S_{prev} and indices k, p, such that the following conditions hold:

Algorithm SplitLOOP$(\widehat{T})$

$min = \min\{|h^{-1}(\widehat{s_{i+1}})|, \ldots, |h^{-1}(\widehat{s_n})|\}$
$\widehat{T}_{\text{unwind}} = \mathbf{unwind}(\widehat{T}, min + 1)$
Compute j and S_{prev} as in **SplitPATH**$(\widehat{T}_{\text{unwind}})$
$k :=$ **LoopIndex**(j)
$p :=$ **LoopIndex**$(j + 1)$
output S_{prev}, k, p

Figure 6: SplitLOOP checks if an abstract loop is spurious

1. The states in S_{prev} correspond to the abstract state $\widehat{s_p}$, i.e., $S_{\text{prev}} \subseteq h^{-1}(\widehat{s_p})$
2. All states in S_{prev} are reachable from $h^{-1}(\widehat{s_1}) \cap I$.
3. k is the successor index of p within the loop, i.e., if $p = n$ then $k = i+1$, and otherwise $k = p + 1$.
4. There is no transition from a state in S_{prev} to $h^{-1}(\widehat{s_k})$, i.e., $Img(S_{\text{prev}}, R) \cap h^{-1}(\widehat{s_k})$ is empty.
5. Therefore, $\widehat{s_p}$ is the failure state of the loop counterexample.

Thus, the final situation encountered is indeed very similar as in the case of path counterexamples. Note that the nontrivial feature of the algorithm **SplitLOOP** is the fact that only *min* unwindings of the loop are necessary. The correctness of this approach is not trivial, and details are deferred to the appendix.

4.3 Refining The Abstraction

First, we will consider the case when the counterexample $\widehat{T} = \langle \widehat{s_1}, \cdots, \widehat{s_n} \rangle$ is a path. Let us return to a previous example for a closer investigation of failure states.

Example 4.4 *Recall that in the spurious counterexample of Figure 3, the abstract state* $\widehat{3}$ *was the* **failure state**. *There are three types of concrete states in the failure state* $\widehat{3}$*:*

(i) The **dead-end state** 9 *is reachable, but there are no outgoing transitions to the next state in the counterexample.*

(ii) The **bad state** 7 *is not reachable but outgoing transitions cause the spurious counterexample. The spurious counterexamples is caused by the bad state.*

(iii) The **irrelevant state** 8 *is neither reachable nor bad.*

The goal of the refinement methodology described in this section is to refine h so that the dead-end states and bad states do not belong to *the same* abstract state. Then the spurious counterexample will be eliminated.

	3	*4*	*5*
7	*1*	*x*	*x*
8	*0*	*x*	*1*
9	*x*	*0*	*0*

Equivalence Class

	3/4	*5*
7	*1*	*x*
8	*0*	*1*
9	*0*	*0*

Refinement (a)

	3	**4/5**
7/9	*1*	*0*
8	*0*	*1*

Refinement (b)

Figure 7: Two possible refinements of an Equivalence Class.

If $\widehat{T}$ does not correspond to a real counterexample, by Lemma 4.1 (iii) there always exists a set S_i of dead-end states, i.e., $S_i \subseteq h^{-1}(\widehat{s_i})$ with $1 \leq i < n$ such that $Img(S_i, R) \cap h^{-1}(\widehat{s_{i+1}}) = \emptyset$ and S_i is reachable from initial state set $h^{-1}(\widehat{s_1}) \cap I$. Moreover, the set S_i of dead-end states can be obtained as the output S_{prev} of **SplitPATH** or **SplitLOOP**. Since there is a transition from $\widehat{s_i}$ to $\widehat{s_{i+1}}$ in the abstract model, there is at least one transition from a *bad* state in $h^{-1}(\widehat{s_i})$ to a state in $h^{-1}(\widehat{s_{i+1}})$ even though there is no transition from S_i to $h^{-1}(\widehat{s_{i+1}})$, and thus the set of bad states is not empty. We partition $h^{-1}(\widehat{s_i})$ into three subsets $S_{i,0}$, $S_{i,1}$, and $S_{i,x}$ as follows:

Name	*Partition*	*Definition*
dead-end states	$S_{i,0}$	S_i
bad states	$S_{i,1}$	$\{s \in h^{-1}(\widehat{s_i}) \mid \exists s' \in h^{-1}(\widehat{s_{i+1}}).R(s, s')\}$
irrelevant states	$S_{i,x}$	$h^{-1}(\widehat{s_i}) \setminus (S_{i,0} \cup S_{i,1})$

Intuitively, $S_{i,0}$ denotes the set of dead-end states, i.e., states in $h^{-1}(\widehat{s_i})$ that are reachable from initial states. $S_{i,1}$ denotes the set of bad states,i.e., those states in $h^{-1}(\widehat{s_i})$ that are not reachable from initial states, but have at least one transition to some state in $h^{-1}(\widehat{s_{i+1}})$. The set $S_{i,1}$ cannot be empty since we know that there is a transition from $h^{-1}(\widehat{s_i})$ to $h^{-1}(\widehat{s_{i+1}})$. $S_{i,x}$ denotes the set of irrelevant states, i.e., states that are not reachable from initial states, and do not have a transition to a state in $h^{-1}(\widehat{s_{i+1}})$. Since $S_{i,1}$ is not empty, there is a spurious transition $\widehat{s_i} \rightarrow \widehat{s_{i+1}}$. This causes the spurious counterexample $\widehat{T}$. Hence in order to refine the abstraction h so that the new model does not allow $\widehat{T}$, we need a refined abstraction function which separates the two sets $S_{i,0}$ and $S_{i,1}$, i.e., we need an abstraction function, in which no abstract state simultaneously contains states from $S_{i,0}$ and from $S_{i,1}$.

It is natural to describe the needed refinement in terms of equivalence relations: Recall that $h^{-1}(\widehat{s})$ is an equivalence class of $\equiv$ which has the form $E_1 \times \cdots \times E_m$, where each E_i is an equivalence class of $\equiv_i$. Thus, the refinement $\equiv'$ of $\equiv$ is obtained by partitioning the equivalence classes E_j into subclasses, which amounts to refining the equivalence relations $\equiv_j$. The *size of the refinement* is the number of new equivalence classes. Ideally, we would like to find the coarsest refinement that separates the two sets, i.e., the separating refinement with the smallest size.

Example 4.5 *Assume that we have two variables v_1, v_2. The failure state corresponds to* one equivalence class $E_1 \times E_2$, *where* $E_1 = \{3, 4, 5\}$ *and* $E_2 = \{7, 8, 9\}$. *In Figure 7, dead-end states $S_{i,0}$ are denoted by 0, bad states $S_{i,1}$ by 1, and irrelevant states by x.*

Let us consider two possible partitions of $E_1 \times E_2$:

- *Case (a) :* $\{(3,4),(5)\} \times \{(7),(8),(9)\}$ *(6 classes)*
- *Case (b) :* $\{(3),(4,5)\} \times \{(7,9),(8)\}$ *(4 classes)*

Clearly, case (b) generates a coarser refinement than case (a). It can be easily checked that no other refinement is coarser than (b).

In general, the problem of finding the coarsest refinement problem is computationally intractable.

Theorem 4.2 *The problem of finding the coarsest refinement is NP-hard.*

The proof is provided in Appendix B.

We therefore need to obtain a good heuristics for abstraction refinement. When $S_{i,x}$ is empty, there is a polynomial algorithm which can find the coarsest refinement. The algorithm **PolyRefine** (see Figure 8) corresponds to this case. Let P_j^+, P_j^- be two projection functions, such that for $s = (d_1, \ldots, d_m)$, $P_j^+(s) = d_j$ and $P_j^-(s) = (d_1, \ldots, d_{j-1}, d_{j+1}, \ldots, d_m)$. Then $proj(S_{i,0}, j, a)$ denotes the *projection* set $\{P_j^-(s) | P_j^+(s) = a, s \in S_{i,0}\}$. Intuitively, the condition $proj(S_{i,0}, j, a) \neq proj(S_{i,0}, j, b)$ in the algorithm means that there exists $(d_1, \ldots, d_{j-1}, d_{j+1}, \ldots, d_m) \in proj(S_{i,0}, j, a)$ and $(d_1, \ldots, d_{j-1}, d_{j+1}, \ldots, d_m) \notin proj(S_{i,0}, j, b)$. According to the definition of $proj(S_{i,0}, j, a)$, $s_1 = (d_1, \ldots, d_{j-1}, a, d_{j+1}, \ldots, d_m) \in S_{i,0}$ and $s_2 = (d_1, \ldots, d_{j-1}, b, d_{j+1}, \ldots, d_m) \notin S_{i,0}$, i.e., $s_2 \in S_{i,1}$. The only way to separate s_1 and s_2 into different equivalence classes is that a and b have to be in different equivalence classes of $\equiv'_j$, i.e., $a \not\equiv'_j b$.

Algorithm PolyRefine

> **for** j := 1 **to** m {
> $\equiv'_j := \equiv_j$
>   **for** every $a, b \in E_j$ {
>   **if** $proj(S_{i,0}, j, a) \neq proj(S_{i,0}, j, b)$
>   **then** $\equiv'_j := \equiv'_j \setminus \{(a,b)\}$ }}

Figure 8: The algorithm **PolyRefine**

Lemma 4.2 *When $S_{i,x} = \emptyset$, the relation $\equiv'_j$ computed by* **PolyRefine** *is an equivalence relation which refines $\equiv_j$ and separates $S_{i,0}$ and $S_{i,1}$. Furthermore, the equivalence relation $\equiv'_j$ is the coarsest refinement of $\equiv_j$.*

The proof of this lemma is provided in Appendix B.

Note that in symbolic presentation, the projection operation $proj(S_{i,0}, j, a)$ amounts to computing a generalized cofactor, which can be easily done by standard BDD methods.

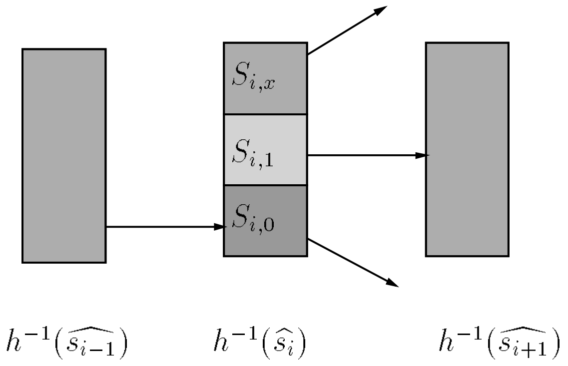

Figure 9: Three sets $S_{i,0}, S_{i,1}$, and $S_{i,x}$

Given a function $f : D \rightarrow \{0,1\}$, a generalized cofactor of f with respect to $g = (\bigwedge_{k=p}^{q} x_k = d_k)$ is the function $f_g = f(x_1, \ldots, x_{p-1}, d_p, \ldots, d_q, x_{q+1}, \ldots, x_n)$. In other words, f_g is the projection of f with respect to g. Symbolically, the set $S_{i,0}$ is represented by a function $f_{S_{i,0}} : D \rightarrow \{0,1\}$, and therefore, the projection $proj(S_{i,0}, j, a)$ of $S_{i,0}$ to value a of the jth component corresponds to a cofactor of $f_{S_{i,0}}$.

In our implementation, we use an heuristics which is based on the following corollary to the proof of Lemma 4.2.

Corollary 4.1 *Even if $S_{i,x}$ is not empty, the relation $\equiv_j'$ computed by* **PolyRefine** *is an equivalence relation which refines $\equiv_j$ and separates $S_{i,0}$ and $S_{i,1}$.*

Refinement Heuristics *We merge the states in $S_{i,x}$ into $S_{i,1}$, and use the algorithm* **PolyRefine** *to find the coarsest refinement that separates the sets $S_{i,0}$ and $S_{i,1} \cup S_{i,x}$. The equivalence relation computed by* **PolyRefine** *in this manner is in general not optimal, but it is a correct refinement which separates $S_{i,0}$ and $S_{i,1}$, and eliminates the spurious counterexample. This heuristic has given good results in our practical experiments.*

Since according to Theorem 4.1, the algorithm **SplitLOOP** for loop counterexamples works analogously as **SplitPATH**, the refinement procedure for spurious loop counterexamples works analogously, i.e., it uses **SplitLOOP** to identify the failure state, and **PolyRefine** to obtain a heurisitc refinement.

Our refinement procedure continues to refine the abstraction function by partitioning equivalence classes until a real counterexample is found, or the ACTL* property is verified. The partitioning procedure is guaranteed to terminate since each equivalence class must contain at least one element. Thus, our method is complete.

Theorem 4.3 *Given a model M and an* ACTL* *specification φ whose counterexample is either path or loop, our algorithm will find a model $\widehat{M}$ such that $\widehat{M} \models \varphi \Leftrightarrow M \models \varphi$.*

5 Performance Improvements

The symbolic methods described in Section 4 can be directly implemented using BDDs. Our implementation uses additional heuristics which are outlined in this section. For details, we refer to our technical report [7].

Two-phase Refinement Algorithms. Consider the spurious loop counterexample $\widehat{T} = \langle \widehat{1}, \widehat{2} \rangle^\omega$ of Figure 10. Although $\widehat{T}$ is spurious, the concrete states involved in the example contain an infinite path $\langle 1, 1, \dots \rangle$ which is a potential counterexample. Since we know that our method is complete, such cases could be ignored. Due to practical performance considerations, however, we came to the conclusion that the relatively small effort to detect additional counterexamples is justified as a valuable heuristic. For a general loop counterexample $\widehat{T} = \langle \widehat{s}_1, \dots, \widehat{s}_i \rangle \langle \widehat{s}_{i+1}, \dots, \widehat{s}_n \rangle^\omega$, we therefore proceed in two phases:
(i) We restrict the model to the state space $S_{\text{local}} := (\bigcup_{1 \leq i \leq n} h^{-1}(\widehat{s}_i))$ of the counterexample and use the standard fixpoint computation for temporal formulas (see e.g. [8]) to check the property on the Kripke structure restricted to S_{local}. If a concrete counterexample is found, then the algorithm terminates.
(ii) If no counterexample is found, we use **SplitLOOP** and **PolyRefine** to compute a refinement as described above.
This two-phase algorithm is slightly slower than the original one if we do not find a concrete counterexample; in many cases however, it can speed up the search for a concrete counterexample. An analogous two phase approach is used for finite path counterexamples.

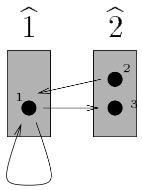

Figure 10: A spurious loop counterexample $\langle \widehat{1}, \widehat{2} \rangle^\omega$

Approximation. Despite the use of partitioned transition relations it is often infeasible to compute the total transition relation of the model M [8]. Therefore, the abstract model $\widehat{M}$ cannot be computed from M directly. In previous work [2, 10], a method which we call *early approximation* has been introduced: first, abstraction is applied to the BDD representation of each transition block and then the BDDs for the partitioned transition relation are built from the already abstracted BDDs for the transition blocks. The disadvantage of early approximation is that it *over-approximates* the abstract model $\widehat{M}$ [9]. In our approach, a heuristic individually determines for each variable cluster VC_i, if early approximation should be applied or if the abstraction function should be applied in an exact manner. Our method has the advantage that it balances overapproximation and memory usage. Moreover, the overall method presented in our paper remains complete with this approximation.

Lemma 5.1 *Let $\widehat{R}$ be the abstract transition relation obtained from existential abstraction. Let $\{R_i^{\text{early}}\}$ be a partitioned transition relation obtained from early approximation. Let*

$\{R_i^{\text{combined}}\}$ *be the final partitioned transition relation which we obtain in our approach. Then* $\widehat{R} \rightarrow \bigwedge_i R_i^{\text{combined}}$ *and* $\bigwedge_i R_i^{\text{combined}} \rightarrow \bigwedge_i R_i^{\text{early}}$.

Thus, the approximation in our approach indeed is intermediate between early approximation and exact existential abstraction. Our method remains complete, because during the symbolic simulation of the counterexample the algorithms **SplitPATH** and **SplitLOOP** treat both forms of overapproximations, i.e., virtual transitions and spurious transitions, in the same way.

Abstractions For Distant Variables. In addition to the methods of Section 4.1, we completely abstract variables whose distance from the specification in the *variable dependency graph* is greater than a user-defined constant. Note that the variable dependency graph is also used for this purpose in the localization reduction [2, 15, 17] in a similar way. However, the refinement process of the localization reduction [15] can only turn a completely abstracted variable into a completely unabstracted variable, while our method uses intermediate abstraction functions.

A user-defined integer constant `far` determines which variables are close to the specification φ. The set `NEAR` of near variables contains those variables whose distance from the specification in the dependency graph is at most `far`, and `FAR` $= var(P) -$ `NEAR` is the set of far variables. For variable clusters without far variables, the abstraction function remains unchanged. For variable clusters with far variables their far variables are completely abstracted away, and their near variables remain unabstracted. Note that the initial abstraction for variable clusters with far variables looks similar as in the localization reduction. However, the refinement process of the localization reduction [15] can only turn a completely abstracted variable into a completely unabstracted variable, while our method uses intermediate abstraction functions.

6 Experimental Results

We have implemented our methodology in NuSMV [6] which uses the CUDD package [21] for symbolic representation. We performed two sets of experiments. One set is on five benchmark designs. The other was performed on an industrial design of a multimedia processor from Fujitsu [1]. All the experiments were carried out on a 200MHz PentiumPro PC with 1GB RAM memory using Linux.

The first benchmark designs are publicly available. The PCI example is extracted from [5]. The results for these designs are listed in the table.

Design	#Var	#Prop	**NuSMV+COI**				**NuSMV+ABS**			
			#COI	Time	$\|TR\|$	$\|MC\|$	#ABS	Time	$\|TR\|$	$\|MC\|$
gigamax	10(16)	1	0	0.3	8346	1822	9	0.2	13151	816
guidance	40(55)	8	30	35	140409	30467	34-39	30	147823	10670
p-queue	12(37)	1	4	0.5	51651	1155	5	0.4	52472	1114
waterpress	6(21)	4	0-1	273	34838	129595	4	170	38715	3335
PCI bus	50(89)	10	4	2343	121803	926443	12-13	546	160129	350226

In the table, the performance for an enhanced version of NuSMV with cone of influence reduction (**NuSMV + COI**) and our implementation (**NuSMV + ABS**) are compared.

#Var and #Prop are properties of the designs: #Var = $x(y)$ means that x is the number of symbolic variables, and y the number of Boolean variables in the design. #Prop is the number of verified properties. The columns #COI and #ABS contain the number of symbolic variables which have been abstracted using the cone of influence reduction (#COI), and our initial abstraction (#ABS). The column "Time" denotes the accumulated running time to verify all #Prop properties of the design. $|TR|$ denotes the maximum number of BDD nodes used for building the transition relation. $|MC|$ denotes the maximum number of *additional* BDD nodes used during the verification of the properties. Thus, $|TR| + |MC|$ is the maximum BDD size during the total model checking process. For the larger examples, we use partitioned transition relations by setting the BDD size limit to 10000.

Although our approach in one case uses 50% more memory than the traditional cone of influence reduction to *build* the abstract transition relation, it requires one magnitude less memory during *model checking*. This is an important achievement since the model checking process is the most difficult task in verifying large designs. More significant improvement is further demonstrated by the Fujitsu IP core design.

The Fujitsu IP core design is a multimedia assist (MMA-ASIC) processor [1]. The design is a system-on-a-chip that consists of a co-processor for multimedia instructions, a graphic display controller, peripheral I/O units, and five bus bridges. The RTL implementation of MM-ASIC is described in about 61,500 lines of Verilog-HDL code. After manual abstraction by engineers from Fujitsu in [22], there still remain about 10,600 lines of code with roughly 500 registers. We translated this abstracted Verilog code into 9,500 lines of SMV code. In [22], the authors verified this design using a "navigated" model checking algorithm in which state traversal is restricted by navigation conditions provided by the user. Therefore, their methodology is not complete, i.e., it may fail to prove the correctness even if the property is true. Moreover, the navigation conditions are usually not automatically generated.

In order to compare our model checker to others, we tried to verify this design using two state-of-the-art model checkers - Yang's SMV [23] and NuSMV [6]. We implemented the cone of influence reduction for NuSMV, but not for Yang's SMV. Both NuSMV+COI and Yang's SMV failed to verify the design. On the other hand, our system abstracted 144 symbolic variables and with three refinement steps, successfully verified the design, and found a bug which has not been discovered before.

7 Conclusion and Future Work

We have presented a novel abstraction refinement methodology for symbolic model checking. The advantages of our methodology have been demonstrated by experimental results. We believe that our technique is general enough to be adapted for other forms of abstraction. There are many interesting avenues for future research. First, we want to find efficient approximation algorithms for the NP-complete separation problem encountered during the refinement step. Moreover, in a recent paper [4], the fragment of ACTL* that admits "trace"-like counterexamples (of a potentially more complicated structure than paths and loops) has been characterized; we plan to extend our refinement algorithm to this language. Since the symbolic methods described in this paper are not tied to representation by BDDs, we will also investigate how they can be applied to recent work on symbolic model checking without BDDs [3]. We are currently applying our technique to verify other large examples.

References

[1] Fujitsu aims media processor at DVD. *MicroProcessor Report*, pages 11–13, 1996.

[2] F. Balarin and A. L. Sangiovanni-Vincentelli. An iterative approach to language containment. In *Computer-Aided Verification*, volume 697 of *LNCS*, pages 29–40, 1993.

[3] A. Biere, A. Cimatti, E. Clarke, M. Fujita, and Y. Zhu. Symbolic model checking using SAT procedures instead of BDDs. In *Design Automation Conference*, pages 317–320, 1999.

[4] F. Buccafurri, T. Eiter, G. Gottlob, and N. Leone. On ACTL formulas having deterministic counterexamples. Technical report, Vienna University of Technology, 1999. available at http://www.kr.tuwien.ac.at/research/reports/index.html.

[5] P. Chauhan, E. Clarke, Y. Lu, and D. Wang. Verifying IP-core based System-On-Chip design. In *IEEE ASIC*, September 1999.

[6] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri. NuSMV: a new symbolic model checker. *Software Tools for Technology Transfer*, 1998.

[7] E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. Technical Report CMU-CS-00-103, Computer Science, Carnegie Mellon University, 2000.

[8] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Publishers, 1999.

[9] E. Clarke, S. Jha, Y. Lu, and D. Wang. Abstract BDDs : a technique for using abstraction in model checking. In *Correct Hardware Design and Verification Methods*, volume 1703 of *LNCS*, pages 172–186, 1999.

[10] E. M. Clarke, O. Grumberg, and D. E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and System (TOPLAS)*, 16(5):1512–1542, September 1994.

[11] S. Das, D. L. Dill, and S. Park. Experience with predicate abstraction. In *Computer-Aided Verification*, volume 1633 of *LNCS*, pages 160–171. Springer Verlag, July 1999.

[12] M. R. Garey and D. S. Johnson. *Computers and interactability: a guide to the theory of NP-Completeness*. W. H. Freeman And Company, 1979.

[13] Shankar G. Govindaraju and David L. Dill. Verification by approximate forward and backward reachability. In *Proceedings of International Conference on Computer-Aided Design*, November 1998.

[14] S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *Computer-Aided Verification*, volume 1254 of *LNCS*, pages 72–83, June 1997.

[15] R. P. Kurshan. *Computer-Aided Verification of Coordinating Processes.* Princeton University Press, 1994.

[16] W. Lee, A. Pardo, J. Jang, G. Hachtel, and F. Somenzi. Tearing based abstraction for CTL model checking. In *Proceedings of the International Conference on Computer-Aided Design*, pages 76–81, November 1996.

[17] J. Lind-Nielsen and H. R. Andersen. Stepwise CTL model checking of state/event systems. In *Computer-Aided Verification*, volume 1633 of *LNCS*, pages 316–327. Springer Verlag, 1999.

[18] A. Pardo. *Automatic Abstraction Techniques for Formal Verification of Digital Systems.* PhD thesis, University of Colorado at Boulder, Dept. of Computer Science, August 1997.

[19] A. Pardo and G.D. Hachtel. Incremental CTL model checking using BDD subsetting. In *Design Automation Conference*, pages 457–462, 1998.

[20] H. Saidi and N. Shankar. Abstract and model checking while you prove. In *Computer-Aided Verification*, number 1633 in LNCS, pages 443–454, July 1999.

[21] F. Somenzi. CUDD: CU decision diagram package. Technical report, University of Colorado at Boulder, 1997.

[22] K. Takayama, T. Satoh, T. Nakata, and F. Hirose. An approach to verify a large scale system-on-chip using symbolic model checking. In *International Conference of Computer Design*, pages 308–313, 1998.

[23] B. Yang et al. A performance study of BDD-based model checking. In *Formal Methods in Computer-Aided Design*, volume 1522 of *LNCS*. Springer Verlag, 1998.

APPENDIX

A Identification of Spurious Loop Counterexamples

Let $\widehat{T} = \langle \widehat{s_1}, \cdots, \widehat{s_i} \rangle \langle \widehat{s_{i+1}}, \cdots, \widehat{s_n} \rangle^\omega$ be an abstract loop counterexample. For an index j, let j^+ denote its successor index in the counterexample, i.e., $n^+ = i + 1$, and for $j < n$, $j^+ = j + 1$.

Theorem 4.1 *The following are equivalent:*

(i) $\widehat{T}$ *corresponds to a concrete counterexample.*

(ii) $h_{\text{path}}^{-1}(\widehat{T}_{\text{unwind}})$ *is not empty.*

Proof. Let us first start with some easy observations. Recall that R is the transition relation of the Kripke structure. By definition, the elements of $h_{\text{path}}^{-1}(\widehat{T}_{\text{unwind}})$ are all the finite R-paths P of the form

$$\langle a_1, \ldots, a_i, \quad b_{i+1}^1, \ldots, b_n^1, \quad \ldots, \quad b_{i+1}^{min+1}, \ldots, b_n^{min+1} \rangle \tag{$*$}$$

for which the following two properties hold:

1. $a_j \in h^{-1}(\widehat{s_j})$ for all a_j in P, and
2. $b_j^k \in h^{-1}(\widehat{s_j})$ for all b_j^k in P.

Each such path P has length $L := i + (min + 1) \times (n - i)$, and we can equivalently write P in the form

$$\langle d_1, \ldots, d_L \rangle \tag{$**$}$$

with the properties

1. $d_1 \in h^{-1}(\widehat{s_1})$, and
2. for all $j < n$, if $d_j \in h^{-1}(\widehat{s_k})$ then $d_{j+1} \in h^{-1}(\widehat{s_{k+}})$.

Recall that min was defined to be the size of the smallest abstract state in the loop, i.e., $\min\{|h^{-1}(\widehat{s_{i+1}})|, \ldots, |h^{-1}(\widehat{s_n})|\}$, and let M be the index of an abstract state $\widehat{s_M}$ s.t. $|h^{-1}(\widehat{s_M})| = min$. (Such a state must exist, because the minimum must be obtained somewhere.)

(i) $\to$ (ii) Suppose there exists a concrete counterexample. Since the counterexample contains a loop, there exists an *infinite R-path* $I = \langle c_1, \ldots \rangle$ such that $c_1 \in h^{-1}(\widehat{s_1})$, and for all j, if $c_j \in h^{-1}(\widehat{s_k})$, then $c_{j+1} \in h^{-1}(\widehat{s_{k+}})$. According to $(**)$, the finite prefix $\langle c_1, \ldots, c_L \rangle$ of I is contained in $h_{\text{path}}^{-1}(\widehat{T}_{\text{unwind}})$, and thus $h_{\text{path}}^{-1}(\widehat{T}_{\text{unwind}})$ is not empty.
(ii) $\to$ (i) Suppose that $h_{\text{path}}^{-1}(\widehat{T}_{\text{unwind}})$ contains a finite R-path P.
Claim: *There exists a state which appears at least twice in P.*
Proof of Claim: Suppose P is in form $(*)$. Consider the states $b_M^1, b_M^2, \ldots, b_M^{min+1}$. By $(*)$, all b_M^k are contained in $h^{-1}(\widehat{s_M})$. By definition of M, however, $h^{-1}(\widehat{s_M})$ contains only min elements, and thus there must be at least one repetition in the sequence $b_M^1, b_M^2, \ldots, b_M^{min+1}$. Therefore, there exists a repetition in the finite R-path P, and the claim is proved. $\Box$ (Claim)

Let us now write P in form $(**)$, i.e., $P = \langle d_1, \ldots, d_L \rangle$, and let a repetition be given by two indices $\alpha < \beta$, s.t. $d_\alpha = d_\beta$. Because of the repetition, there must be a transition from $d_{\beta-1}$ to d_α, and therefore, d_α is the successor state of $d_{\beta-1}$ in a cycle. We conclude that

$$\langle d_1, \ldots, d_{\alpha-1} \rangle \langle d_\alpha, \ldots, d_{\beta-1} \rangle^\omega$$

is a concrete counterexample. $\Box$

A.1 Optimal Abstraction Refinement is NP-complete

Recall that in figure 7, we have visualized the special case of two variables and two equivalence relations in terms of matrices:

	3	4	5
7	1	x	x
8	0	x	1
9	x	0	0

Equivalence Class

	3/4	5
7	1	x
8	0	1
9	0	0

Refinement (a)

	3	**4/5**
7/9	1	0
8	0	1

Refinement (b)

In order to formally capture this visualization, let us define the **Matrix Squeezing** problem.

Definition A.1 Matrix Squeezing
Given an integer constant Γ and a finite (n, m) matrix with entries $0, 1, x$, is it possible to obtain a matrix with $\leq \Gamma$ entries by iterating the following operations:

1. *Merging two compatible rows.*

2. *Merging two compatible columns.*

Two rows are compatible*, if there is no position, where one row contains* 1 *and the other row contains* 0*. All other combinations are allowed, i.e., x does not affect compatibility.* Merging *two rows means replacing the rows by a new one which contains* 1 *at those positions where at least one of the two columns contained* 1*, and* 0 *at those positions, where at least one of the two columns contained* 0*.*

For columns, the definitions are analogous.

Since **Matrix Squeezing** is a special case of the refinement problem, it is sufficient to show NP-hardness for **Matrix Squeezing**. Then it follows that the refinement problem is NP-hard, too, and thus Theorem 4.2 is proved.

As mentioned **Matrix Squeezing** is easy to visualize: If we imagine the symbol x to be transparent, then merging two columns can be thought of as putting the two (transparent) columns on top of each other. **Column Squeezing** is a variant of **Matrix Squeezing**, where only columns can be merged, and the number of rows is left unchanged. We will first show NP-completeness of **Column Squeezing**, and then show NP-completeness of **Matrix Squeezing** by a reduction from **Column Squeezing**.

Definition A.2 Column Squeezing
Given an integer constant Δ and a finite (n, m) matrix with entries $0, 1, x$, is it possible to obtain a matrix with $\leq \Delta$ columns *by iterated merging of* columns *?*

The proof will be by reduction from problem GT15 in [12]:

Definition A.3 Partition Into Cliques
Given an undirected graph (V, E) and and a number $K \geq 3$, is there a partition of V into $k \leq K$ classes, such that each class induces a clique on (V, E) ?

Theorem A.1 (Karp 72) *Partition Into Cliques is NP-complete.*

Theorem A.2 Column Squeezing *is NP-complete.*

Proof: Membership is trivial. Let us consider hardness. We reduce **Partition Into Cliques** to **Column Squeezing**. Given a graph (V, E) and a number K, we have to construct a matrix M and a number Δ such that M can be squeezed to size $\leq \Delta$ iff (V, E) can be partitioned in $\leq K$ cliques.

We construct a $(|V|, |V|)$ matrix $(a_{i,j})$ which is very similar to the adjacency matrix of (V, E):

$$a_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } (i,j) \notin E, i \neq j \\ x & \text{if } (i,j) \in E, i \neq j \end{cases}$$

Assume w.l.o.g. that $V = \{1, \ldots, n\}$. Then it is not hard to see that for all $i, j \in V$, columns i and j are compatible iff $(i,j) \in E$, since the 0 entries in the matrix were chosen in such a way that the columns corresponding to two non-adjacent edges cannot be merged.

By construction, (V, E) contains a clique C with vertices $c_1, \ldots, c_l$ iff the columns $c_1, \ldots, c_l$ can all be merged into one. (Note however that compatibility is not a transitive relation.)

Thus, (V, E) can be partitioned into $\leq K$ cliques, iff the columns of $(a_{i,j})$ can be merged into $\leq K$ columns. Setting $\Delta = K$ concludes the proof. □

Theorem A.3 Matrix Squeezing *is NP-complete.*

Proof: Membership is trivial. We show hardness by reducing **Column Squeezing** to **Matrix Squeezing**. For an integer n, let $|bin(n)|$ denote the size of the binary representation of n. Given an (n, m) matrix M and a number Δ, it is easy to construct an $(n + 1, m + |bin(m - 1)|)$ matrix $B(M)$ by adding additional columns to A in such a way that

(i) all rows of $B(M)$ become incompatible, and

(ii) no new column is compatible with any other (new or old) column.

An easy construction to obtain this is to concatenate the rows of M with the binary encodings of the numbers $0, \ldots, m-1$ over alphabet $\{0, 1\}$, such that the ith row is concatenated with the binary encoding of the number $i-1$. Since any two different binary encodings are distinguished by at least one position, no two rows are compatible. In addition, we add an $n+1$st row which contains 1 on positions in the original columns, and 0 on positions in the new columns. Thus, in matrices of the form $B(M)$, only columns which already appeared in M (with an additional 0 symbol below) can be compatible.

It remains to determine Γ. We set $\Gamma := (\Delta + |bin(m - 1)|) \times (n + 1)$.□ The summand $|bin(m - 1)|$ takes into account that we have added $|bin(m - 1)|$ columns, and the factor $(n + 1)$ takes into account that Δ is counting columns, while Γ is counting matrix entries.□

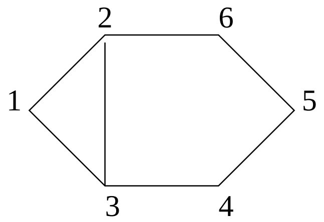

	1	2	3	4	5	6
1	1	x	x	0	0	0
2	x	1	x	0	0	x
3	x	x	1	x	0	0
4	0	0	x	1	x	0
5	0	0	0	x	1	x
6	0	x	0	0	x	1

Column Squeezing

	1	2	3	4	5	6	7	8	9
1	1	x	x	0	0	0	0	0	0
2	x	1	x	0	0	x	0	0	1
3	x	x	1	x	0	0	0	1	0
4	0	0	x	1	x	0	0	1	1
5	0	0	0	x	1	x	1	0	0
6	0	x	0	0	x	1	1	0	1
7	0	0	0	0	0	0	1	1	1

Matrix Squeezing

Figure 11: An instance of **Partition into Cliques**, and its reduction images.

Example A.1 *Figure 11 demonstrates how a graph instance is reduced to a matrix instance. Note for example that* $\{1,2,3\}$ *is a clique in the graph, and therefore, the columns* $1,2,3$ *of the* **Column Squeezing** *problem are compatible. In the* **Matrix Squeezing Instance**, *Columns* $7,8,9$ *enforce that no rows can be merged. Row* 7 *guarantees that columns* $7,8,9$ *can not be merged with columns* $1,\dots,6$.

B Proofs about Path Counterexamples

Lemma 4.1. *The following are equivalent:*

(i) The path $\widehat{T}$ *corresponds to a concrete counterexample.*

(ii) The set of concrete paths $h^{-1}(\widehat{T})$ *is non-empty.*

(iii) For all $1 \le i \le n$, $S_i \neq \emptyset$.

Proof: (i) $\rightarrow$ **(ii)** Assume that $\widehat{T}$ corresponds to a concrete counterexample $T = \langle s_1, \dots, s_n \rangle$. From the definition of $\widehat{T}$, $h(s_i) = \widehat{s_i}$ and $s_i \in h^{-1}(\widehat{s_i})$. Since T is a trace in the concrete model, it has to satisfy the transition relation and start from initial state, i.e. $R(s_i, s_{i+1})$ and $s_1 \in I$. From the definition of $h^{-1}(\widehat{T})$, it follows that $T \in h^{-1}(\widehat{T})$.
(ii) $\rightarrow$ **(i)** Assume that $h^{-1}(\widehat{T})$ is non-empty. We pick a trace $\langle s_1, \dots, s_n \rangle$ from $h^{-1}(\widehat{T})$. Then $\langle h(s_1), \dots, h_(s_n) \rangle = \widehat{T}$, and therefore $\widehat{T}$ corresponds to a concrete counterexample.
(ii) $\rightarrow$ **(iii)** Assume that $h^{-1}(\widehat{T})$ is not empty. Then there exists a path $\langle s_1, \dots, s_n \rangle$ where $h(s_i) = \widehat{s_i}$ and $s_1 \in I$. Therefore, we have $s_1 \in S_1$. Let us assume that $s_i \in S_i$. By the definition of $h^{-1}(\widehat{T})$, $s_{i+1} \in Img(s_i, R)$ and $s_{i+1} \in h^{-1}(\widehat{s_{i+1}})$. Therefore, $s_{i+1} \in S_{i+1}$, since $S_{i+1} = Img(S_i, R) \cap h^{-1}(\widehat{s_{i+1}})$. By induction, $S_i \neq \emptyset$, for $i \le n$.
(iii) $\rightarrow$ **(ii)** Assume that $S_i \neq \emptyset$ for $1 \le i \le n$. We choose a state $s_n \in S_n$ and inductively construct a trace backward. Assume that $s_i \in S_i$. From the definition of S_i, if follows that $s_i \in Img(S_{i-1}, R) \cap h^{-1}(\widehat{s_i})$ and S_{i-1} is not empty. Select s_{i-1} from S_{i-1}. From the definition of S_{i-1}, $S_{i-1} \subseteq h^{-1}(\widehat{s_{i-1}})$. Hence, $s_{i-1} \in h^{-1}(\widehat{s_{i-1}})$. By induction, $s_1 \in S_1 = h^{-1}(\widehat{s_1}) \cap I$. Therefore, the trace $\langle s_1, \dots, s_n \rangle$ that we have constructed satisfies the definition of $h^{-1}(\widehat{T})$. Thus, $h^{-1}(\widehat{T})$ is not empty. $\square$

In the following, we prove that when $S_{i,x}$ is empty, there exists a polynomial algorithm to find the coarsest refinement. Let $s \in h^{-1}(\widehat{s_i})$ be a state and P_j^+, P_j^- be two projection functions, such that for $s = (d_1, \ldots, d_m)$, $P_j^+(s) = d_j$ and $P_j^-(s) = (d_1, \ldots, d_{j-1}, d_{j+1}, \ldots, d_m)$. Note that this definition is consistent to the definition in Section 4.3. Since $S_{i,x}$ is empty, $S_{i,0}$ and $S_{i,1}$ form a partition of $h^{-1}(\widehat{s_i})$. A refinement of $h^{-1}(\widehat{s_i})$ can be achieved by refining each equivalence relations $\equiv_j$ (and thus, simultaneously, the abstraction functions h_j).

We will replace each equivalence relation $\equiv_j$ by the equivalence relation $\equiv'_j$ in the following way: We put two elements a, b of D_{VC_j} in the same equivalence class (symbolically, $a \equiv'_j b$) if and only if the *projection* sets $P_{j,a} = \{P_j^-(s) | P_j^+(s) = a, s \in S_{i,1}\}$ and $P_{j,b} = \{P_j^-(s) | P_j^+(s) = b, s \in S_{i,1}\}$ are equal. Intuitively, this means that any two states which only differ in the jth component are either both in $S_{i,1}$ or both not in $S_{i,1}$. As shown in Section 2, the equivalence relations $\equiv'_j$ $(1 \leq j \leq m)$ define an equivalence relation $\equiv'$ on D.

Lemma 4.2 *When $S_{i,x} = \emptyset$, the relation $\equiv'_j$ computed by* **PolyRefine** *is an equivalence relation which refines $\equiv_j$ and separates $S_{i,0}$ and $S_{i,1}$. Furthermore, the equivalence relation $\equiv'_j$ is the coarsest refinement of $\equiv_j$.*

Proof: First, we argue that $\equiv'_j$ is an equivalence relation:

- Reflexivity: for any $a \in E_j$, (a, a) is not removed from $\equiv_j$, therefore, $a \equiv'_j a$;
- Symmetry: $a \equiv'_j b$ implies that $proj(S_{i,0}, j, a) = proj(S_{i,0}, j, b)$. According to **PolyRefine**, (b, a) is not removed from $\equiv_j$. Therefore, $b \equiv'_j a$;
- Transitivity: assume that $a \equiv'_j b$ and $b \equiv'_j c$, Then $proj(S_{i,0}, j, a) = proj(S_{i,0}, j, b)$ and $proj(S_{i,0}, j, b) = proj(S_{i,0}, j, c)$. Hence, $proj(S_{i,0}, j, a) = proj(S_{i,0}, j, c)$. This implies that $a \equiv'_j c$.

Secondly, we show that $\equiv'$ is a correct refinement, i.e., for any two states $s_1 \in S_{i,1}$ and $s_2 \in S_{i,0}$, $s_1 \not\equiv' s_2$. Assume that there are two states $s_1 \in S_{i,1}$ and $s_2 \in S_{i,0}$ where $s_1 \equiv' s_2$. Also assume that $s_1 = (d_1, \ldots, d_m)$ and $s_2 = (e_1, \ldots, e_m)$ where $d_j \equiv'_j e_j$. Without loss of generality, we assume that $d_j \neq e_j$ for $1 \leq j \leq k$ and $d_j = e_j$ for $k < j \leq m$ where $1 < k \leq m$. Consider another state $s_3 = (e_1, d_2, \ldots, d_m)$. Since $e_1 \in E_1$, $d_j \in E_j$ for $1 < j \leq m$, $s_3 \in h^{-1}(\widehat{s_i})$. On the other hand, $s_1 \equiv' s_3$ because $d_1 \equiv'_1 e_1$ and $d_j \equiv'_j d_j$ for all j. According to our definition of $\equiv'_1$, any two states which only differ in the jth component are either both in $S_{i,1}$ or both not in $S_{i,1}$. Since $s_1 \in S_{i,1}$, it follows that $s_3 \in S_{i,1}$. Furthermore, we consider $s_4 = (e_1, e_2, d_3, \ldots, d_m)$. Following the same argument, $s_3 \equiv' s_4$ and $s_4 \in S_{i,1}$. Therefore, $s_1 \equiv' s_4$. By repeating this step k times, we will obtain that $s_1 \equiv' s_2$ and $s_2 \in S_{i,1}$. Hence, $S_{i,1} \cap S_{i,0} \neq \emptyset$. This contradicts our definition of $S_{i,1}$ and $S_{i,0}$. Therefore, the equivalence relation $\equiv'$ partitions $S_{i,1}$ and $S_{i,0}$ into different equivalence classes.

Finally, we prove that the equivalence relation $\equiv'$ defines the coarsest refinement. Towards contradiction, we assume that there is another equivalence relation $\equiv''$ which defines a coarser refinement than $\equiv'$ and it eliminates the counterexample. Note that a coarser refinement implies that there are a fewer number of equivalence classes generated by $\equiv''$ than $\equiv'$. This implies that there exists a j such that $\equiv''_j$ generates fewer equivalence classes than $\equiv'_j$. Therefore, there must exist two elements $a, b \in D_{VC_j}$ where $a \not\equiv'_j b$ but $a \equiv''_j b$. According to the definition of $\equiv'_j$, $a \not\equiv'_j b$ if and only if there exist two states s_1 and s_2, s.t. $P_j^+(s_1) = a$, $P_j^+(s_2) = b$ and $P_j^-(s_1) = P_j^-(s_2)$, however, either $s_1 \in S_{i,1} \wedge s_2 \notin S_{i,1}$ or $s_1 \notin S_{i,1} \wedge s_2 \in S_{i,1}$. We will first consider the case of $s_1 \in S_{i,1} \wedge s_2 \notin S_{i,1}$. The second case will follow the same argument. Because $S_{i,x}$ is empty, $s_2 \notin S_{i,1}$ implies that $s_2 \in S_{i,0}$. On the other hand, $a \equiv''_j b$ implies that $s_1 \equiv'' s_2$ according to the definition of $\equiv''$. Therefore, $\equiv''$ cannot partition $S_{i,1}$ and $S_{i,0}$ into different equivalence classes, i.e., it cannot eliminate the counterexample. Hence, $\equiv'$ defines the coarsest refinement. $\square$

Theorem 4.3 *Given a model M and an* ACTL* *specification φ whose counterexample is either path or loop, our algorithm will find a model $\widehat{M}$ such that $\widehat{M} \models \varphi \Leftrightarrow M \models \varphi$.*

Proof: There are three cases to consider.

(i) If $\widehat{M} \models \varphi$, then $M \models \varphi$ according to Theorem 2.1

(ii) If $\widehat{M} \not\models \varphi$, and the generated abstract counterexample is not spurious, then there exists a concrete counterexample, and hence, $M \not\models \varphi$.

(iii) If $\widehat{M} \not\models \varphi$, and the generated abstract counterexample is spurious, then **PolyRefine** will refine the abstraction. Since each refinement step partitions an existing equivalence classe into *strictly* smaller equivalence classes, after a finite number of steps the equivalence relation will become the *equality* relation, and therefore $\widehat{M} = M$. Hence $M \not\models \varphi$.

$\square$