Interpolant Strength

Vijay D'Silva[1,⋆], Daniel Kroening[1], Mitra Purandare[2,⋆⋆], and Georg Weissenbacher[1,2,⋆⋆⋆]

[1] Computing Laboratory, Oxford University
[2] Computer Systems Institute, ETH Zurich

Abstract. Interpolant-based model checking is an approximate method for computing invariants of transition systems. The performance of the model checker is contingent on the approximation computed, which in turn depends on the logical strength of the interpolants. A good approximation is coarse enough to enable rapid convergence but strong enough to be contained within the weakest inductive invariant. We present a system for constructing propositional interpolants of different strength from a resolution refutation. This system subsumes existing methods and allows interpolation systems to be ordered by the logical strength of the obtained interpolants. Interpolants of different strength can also be obtained by transforming a resolution proof. We analyse an existing proof transformation, generalise it, and characterise the interpolants obtained.

1 Introduction

Symbolic model checking techniques manipulate implicit representations of sets of states to verify correctness properties of transition systems. Image computation and fixed point detection, two essential steps in model checking, involve quantifier elimination, which is computationally expensive. Interpolant-based model checking of finite state systems uses approximate images to compute an inductive invariant that suffices to show correctness [11]. The approximate images are constructed from resolution refutations generated by a SAT solver, thereby avoiding quantifier elimination.

The performance of an interpolant-based model checker depends on the approximate images obtained. A coarse approximation typically contains spurious errors and causes the model checker to restart with a larger formula. Model checking with a larger formula is more resource intensive than with a smaller formula. On the other hand, a tight approximation delays convergence to a fixed point. If the property holds, the ideal approximate image is an inductive invariant that implies the property. If the property does not hold, the ideal approximation is one which enables the error to be detected efficiently. Thus, rather than strong

⋆ Supported by Microsoft Research's European PhD Scholarship Programme.
⋆⋆ Supported by the Semiconductor Research Corporation (SRC) under contract no. 2006-TJ-1539.
⋆⋆⋆ Supported by the EU FP7 STREP MOGENTES (project ID ICT-216679) and by Microsoft Research's European PhD Scholarship Programme.

G. Barthe and M. Hermenegildo (Eds.): VMCAI 2010, LNCS 5944, pp. 129–145, 2010.

or weak interpolants, procedures to compute interpolants of different strengths are required. A procedure for constructing interpolants from resolution refutations is called an *interpolation system* in this paper.

We study two orthogonal approaches to obtaining interpolants of different strengths. The first approach is to construct different interpolants from a refutation. This is a challenge because only two interpolation systems exist; a symmetric system, published independently by Huang [6], Krajíček [8] and Pudlák [13], and McMillan's system [11]. We are not aware of any results relating these two systems. The second approach, suggested by Jhala and McMillan [7], is to reorder the sequence of resolution steps in a proof to strengthen the interpolants obtained. Our implementation of their algorithm led us to find an error and was the motivation for much of this work. The effect of proof transformations has only been studied for McMillan's system [11]. It is not known if such transformations result in stronger interpolants in other systems.

Contributions. The contributions of this paper are as follows.

- An ordered family of linear-time interpolation systems. This family subsumes existing interpolation systems. An interpolation system Itp maps a resolution refutation R to an interpolant $\mathsf{Itp}(R)$. The order guarantees interpolant strength; if $\mathsf{Itp}_1 \preceq \mathsf{Itp}_2$ then $\mathsf{Itp}_1(R)$ implies $\mathsf{Itp}_2(R)$ for any refutation R.
- Operators for composing interpolation systems. The ordered family of interpolation systems with these operators forms a complete lattice with McMillan's systems being the strongest. Interpolation systems can be composed to obtain stronger and weaker interpolants as required.
- A study of the effect of pivot reordering on interpolant strength. A proof transformation due to Jhala and McMillan [7] is shown to produce invalid refutations and redundant interpolants in cases. These cases are analysed and characterised.

This paper is organised as follows. Background material on model checking and resolution proofs is covered in § 2. Existing interpolation systems are presented in § 3 and our parametrised interpolation system appears in § 4. Proof transformations that change interpolant strength are studied in § 5. We discuss related work in § 6 and conclude in § 7. The proofs of all statements in this paper are presented in the appendices of the supplemental technical report [5].

2 Preliminaries

2.1 Finite State Model Checking

A transition system $M = (S, T)$ is a finite set of states S and a transition relation $T \subseteq S \times S$. Fix the sets J and F, where $J \cap F = \emptyset$, as sets of initial and failure states respectively. A system is correct if no state in F is reachable from any state in J. The image operator $post : \wp(S) \to \wp(S)$ maps a set of states to its successors: $post(Q) = \{s' \in S | s \in Q \text{ and } (s, s') \in T\}$. Let $post^0(Q) = Q$ and