

Efficient Generation of Craig Interpolants in Satisfiability Modulo Theories

ALESSANDRO CIMATTI

FBK-IRST

and

ALBERTO GRIGGIO and ROBERTO SEBASTIANI

Università di Trento

The problem of computing Craig interpolants has recently received a lot of interest. In this article, we address the problem of efficient generation of interpolants for some important fragments of first-order logic, which are amenable for effective decision procedures, called *satisfiability modulo theory* (SMT) solvers.

We make the following contributions. First, we provide interpolation procedures for several basic theories of interest: the theories of linear arithmetic over the rationals, difference logic over rationals and integers, and UTVPI over rationals and integers. Second, we define a novel approach to interpolate combinations of theories that applies to the delayed theory combination approach.

Efficiency is ensured by the fact that the proposed interpolation algorithms extend state-of-the-art algorithms for satisfiability modulo theories. Our experimental evaluation shows that the MathSAT SMT solver can produce interpolants with minor overhead in search, and much more efficiently than other competitor solvers.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—*Mechanical theorem proving*

General Terms: Theory, Algorithms

Additional Key Words and Phrases: Craig interpolation, decision procedures, SMT

ACM Reference Format:

Cimatti, A., Griggio, A., and Sebastiani, R. 2010. Efficient generation of craig interpolants in satisfiability modulo theories. *ACM Trans. Comput. Logic*, 12, 1, Article 7 (October 2010), 54 pages. DOI = 10.1145/1838552.1838559 <http://doi.acm.org/10.1145/1838552.1838559>

A. Cimatti is supported in part by the European Commission FP7-2007-IST-1-217069 COCONUT. A. Griggio and R. Sebastiani are supported in part by SRC under GRC Custom Research Project 2009-TJ-1880 WOLFLING, and by MIUR under PRIN project 20079E5KM8.002.

Authors' addresses: A. Cimatti, FBK-IRST, Via Sommarive 18, 38050 Povo, Trento, Italy; email: cimatti@fbk.eu; A. Griggio and Roberto Sebastiani, DISI, Università di Trento, Via Sommarive 14, 38050 Povo, Trento, Italy; email: {griggio, rseba}@disi.unitn.it.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2010 ACM 1529-3785/2010/10-ART7 \$10.00 DOI 10.1145/1838552.1838559 <http://doi.acm.org/10.1145/1838552.1838559>

1. INTRODUCTION

One of the most successful applications of computational logic is formal verification, which aims at proving (or disproving) certain properties of the behaviors of a reactive system. In recent years, also thanks to the impressive improvements of SAT solvers, a wide variety of verification methods based on SAT solving have been proposed. These methods proved effective for discrete state systems, most notably hardware components. The approach is made practical by the fact that SAT solvers, in addition to proving efficiently the satisfiability of huge propositional formulas, provide several functionalities, such as model generation, proof production, extraction of unsatisfiable cores, and generation of Craig interpolants (interpolation). In particular, since the seminal paper of McMillan [2003], interpolation has been recognized to be a substantial tool for verification in the case of Boolean systems (e.g., Cabodi et al. [2006]; Li and Somenzi [2006]; Marques-Silva [2007]).

One of the main limitations of SAT-based approaches is in their expressive power. Many systems of practical interest, containing integer or real valued variables, such as software and timed and hybrid systems, cannot be represented directly within propositional logic. This has prompted research in the analysis of fragments of first-order logic: given a formula referring to variables, the problem is to find a satisfying assignment in a theory of interest (e.g., linear arithmetic). This field, referred to as *Satisfiability Modulo Theory* (SMT), has resulted in substantial theoretical results, and in very effective decision procedures, known as SMT solvers. State-of-the-art SMT solvers complement the Boolean SAT algorithms with specialized decision procedures for conjunctions of literals in some given theory (*theory solvers*). In addition to checking satisfiability, SMT solvers are able to generate models, produce proofs, and extract unsatisfiable cores. This has allowed the lifting of many SAT-based verification algorithms to SMT-based verification, as well as opening up the way to abstraction refinement with SMT.

Quite surprisingly, however, the research on interpolation for SMT has not kept the pace of SMT solving. In fact, most of the approaches to producing interpolants for fragments of first-order theories proposed in the last few years [McMillan 2005; Yorsh and Musuvathi 2005; Rybalchenko and Sofronie-Stokkermans 2007; Kroening and Weissenbacher 2007; Kapur et al. 2006; Jain et al. 2009] suffer from a number of problems. Some of the approaches are severely limited in terms of their expressiveness. For instance, the tool described in Rybalchenko and Sofronie-Stokkermans [2007] can only deal with conjunctions of literals, while the work described in Kroening and Weissenbacher [2007] cannot deal with many useful theories. Furthermore, very few tools are available [Rybalchenko and Sofronie-Stokkermans 2007; McMillan 2005], and these tools do not seem to scale particularly well. More than to naïve implementation, this appears to be due to the underlying algorithms, which substantially deviate from or ignore choices common in state-of-the-art SMT. For instance, in the domain of linear arithmetic over the rationals ($\mathcal{LA}(\mathbb{Q})$), strict inequalities were encoded in McMillan [2005] as the conjunction of a weak inequality and a disequality; although sound, this choice

destroys the structure of the constraints, forces reasoning in the combination of theories $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{EUF}$, requires additional splitting, and ultimately results in a larger search space. Similarly, the fragment of difference logic ($\mathcal{DL}(\mathbb{Q})$) is dealt with by means of a general-purpose algorithm for full $\mathcal{LA}(\mathbb{Q})$, rather than one of the well-known and much faster specialized algorithms. An even more fundamental example is the fact that state-of-the-art SMT reasoners use dedicated algorithms for linear arithmetic [Dutertre and de Moura 2006].

In this article, we tackle the problem of generating interpolants for SMT problems, fully leveraging the algorithms used in a state-of-the-art SMT solver. In particular, our main contributions are the following.

- (1) An interpolation algorithm for $\mathcal{LA}(\mathbb{Q})$ that exploits a variant of the algorithm presented in Dutertre and de Moura [2006], and that is capable of handling the full $\mathcal{LA}(\mathbb{Q})$ —including strict inequalities and disequalities—without the need of theory combination.
- (2) An algorithm for computing interpolants in \mathcal{DL} —both over the rationals and over the integers—that builds on top of the efficient graph-based decision algorithms given in Cotton and Maler [2006] and Nieuwenhuis and Oliveras [2005], which ensures that the generated interpolants are still in the \mathcal{DL} fragment of linear arithmetic, and that allows for computing stronger interpolants than the existing algorithms for the full linear arithmetic.
- (3) An algorithm for computing interpolants in $UTVPI$ —both over the rationals and over the integers—that builds on an encoding of \mathcal{DL} . The algorithm ensures that the generated interpolants are still in the $UTVPI$ fragment of linear arithmetic, and that allows for computing stronger interpolants than the existing algorithms for the full linear arithmetic.
- (4) An algorithm for computing interpolants in a combination $\mathcal{T}_1 \cup \mathcal{T}_2$ of theories based on the Delayed Theory Combination (DTC) method [Bozzano et al. 2006; Bruttomesso et al. 2009] (as an alternative to the traditional Nelson-Oppen method), which does not require ad hoc interpolant combination methods, but exploits the propositional interpolation algorithm for performing the combination of theories.
- (5) An efficient implementation of all the proposed techniques within the MATHSAT 4 SMT solver [Bruttomesso et al. 2008], and an extensive experimental evaluation on a wide range of benchmarks.

This comprehensive approach advances the state of the art in two main directions: on the one side, we show how to extend efficient SMT solving techniques to SMT interpolation, for a wide class of important theories, without paying a substantial price in performance; on the other, we present an interpolating SMT solver that is able to produce interpolants for a much wider class of problems than its competitors, and, on problems that can be dealt with by other tools, shows dramatic improvements in performance, often by orders of magnitude.

This article is structured as follows. In Section 2 we present some background on interpolation in SMT. In Sections 3, 4, and 5 we show how to efficiently

interpolate $\mathcal{LA}(\mathbb{Q})$, \mathcal{DL} , and $UTVPI$, respectively. In Section 6 we discuss interpolation for combined theories. The proposed techniques are experimentally evaluated in Section 7. In Section 8 we draw some conclusions. The discussion of related work is distributed in the technical sections (Sections 3–6).

2. BACKGROUND AND STATE-OF-THE-ART

2.1 Satisfiability Modulo Theory

Our setting is standard first-order logic. A 0-ary function symbol is called a *constant*. A *term* is a first-order term built out of function symbols and variables. We write $t_1 \equiv t_2$ when the two terms t_1 and t_2 are syntactically identical. If t_1, \dots, t_n are terms and p is a predicate symbol, then $p(t_1, \dots, t_n)$ is an *atom*. A *literal* is either an atom or its negation. A *formula* ϕ is built in the usual way out of the universal and existential quantifiers, Boolean connectives, and atoms. We call a formula *quantifier-free* if it does not contain quantifiers, and *ground* if it does not contain free variables. A *clause* is a disjunction of literals. A formula is said to be in *conjunctive normal form* (CNF) if it is a conjunction of clauses. For every non-CNF T -formula ϕ , an equisatisfiable CNF formula ψ can be generated in polynomial time [Tseitin 1968].

We also assume the usual first-order notions of interpretation, satisfiability, validity, logical consequence, and theory, as given, for example, in Enderton [2001]. A *first-order theory*, T , is a set of first-order sentences. In this article, we consider only theories with equality. A structure \mathcal{A} is a model of a theory T if \mathcal{A} satisfies every sentence in T . A formula is *satisfiable in T* (or *T -satisfiable*) if it is satisfiable in a model of T .

We call *Satisfiability Modulo (the) Theory T* , $\text{SMT}(T)$, the problem of deciding the satisfiability of quantifier-free formulas¹ with respect to a background theory T . We denote formulas with ϕ, ψ, A, B, C, I , T -variables with x, y, z , Boolean variables with p, q , and numeric constants with a, b, c, l, u . Given a theory T , we write $\phi \models_T \psi$ (or simply $\phi \models \psi$) to denote that the formula ψ is a logical consequence of ϕ in the theory T . With $\phi \leq \psi$ we denote that all uninterpreted (in T) symbols of ϕ appear in ψ . If C is a clause, $C \downarrow B$ is the clause obtained by removing all the literals whose atoms do not occur in B , and $C \setminus B$ that obtained by removing all the literals whose atoms do occur in B . With a little abuse of notation, we might sometimes denote conjunctions of literals $l_1 \wedge \dots \wedge l_n$ as sets $\{l_1, \dots, l_n\}$ and vice versa. If $\eta \stackrel{\text{def}}{=} \{l_1, \dots, l_n\}$, we might write $\neg\eta$ to mean $\neg l_1 \vee \dots \vee \neg l_n$. A theory T is *stably-infinite* iff every quantifier-free T -satisfiable formula is satisfiable in an infinite model of T . A theory T is *convex* iff, for every collection $l_1, \dots, l_k, e_1, \dots, e_n$ of literals in T such that e_1, \dots, e_n are of the form $(x = y)$, x, y being variables, we have that $\{l_1, \dots, l_k\} \models_T \bigvee_{i=1}^n e_i$ if and only if $\{l_1, \dots, l_k\} \models_T e_i$ for some $1 \leq i \leq n$.

Given a decidable first-order theory T , we call a *theory solver for T* , T -*solver*, any tool able to decide the satisfiability in T of sets/conjunctions of ground atomic formulas and their negations—*theory literals* or *T -literals*—in

¹The general definition of SMT deals also with quantified formulas. Nevertheless, in this article we restrict our interest to quantifier-free formulas.

the language of \mathcal{T} . If $S \stackrel{\text{def}}{=} \{l_1, \dots, l_n\}$ is a set of literals in \mathcal{T} , we call (\mathcal{T})-*conflict set* any subset η of S which is inconsistent in \mathcal{T} .² We call $\neg\eta$ a \mathcal{T} -*lemma*. (Notice that $\neg\eta$ is a \mathcal{T} -valid clause.)

Definition 2.1 (Resolution Proof). Given a set of clauses $S \stackrel{\text{def}}{=} \{C_1, \dots, C_n\}$ and a clause C , we call a *resolution proof* of the deduction $\bigwedge_i C_i \models_{\mathcal{T}} C$ a DAG \mathcal{P} such that

- (1) C is the root of \mathcal{P} ;
- (2) the leaves of \mathcal{P} are either elements of S or \mathcal{T} -lemmas;
- (3) each nonleaf node C' has two premises C_{p_1} and C_{p_2} such that $C_{p_1} \stackrel{\text{def}}{=} p \vee \phi_1$, $C_{p_2} \stackrel{\text{def}}{=} \neg p \vee \phi_2$, and $C' \stackrel{\text{def}}{=} \phi_1 \vee \phi_2$. The atom p is called the *pivot* of C_{p_1} and C_{p_2} .

If C is the empty clause (denoted with \perp), then \mathcal{P} is a *resolution proof of (\mathcal{T})-unsatisfiability* for $\bigwedge_i C_i$.

We consider the $\text{SMT}(\mathcal{T})$ problem for some background theory \mathcal{T} .

Definition 2.2 (Craig Interpolant). Given an ordered pair (A, B) of formulas such that $A \wedge B \models_{\mathcal{T}} \perp$, a *Craig interpolant* (simply “interpolant” hereafter) is a formula I such that

- (i) $A \models_{\mathcal{T}} I$,
- (ii) $I \wedge B \models_{\mathcal{T}} \perp$,
- (iii) $I \leq A$ and $I \leq B$.

2.2 Algorithms for SMT

A standard technique for solving the $\text{SMT}(\mathcal{T})$ problem is to integrate a DPLL-based SAT solver and a \mathcal{T} -solver in a “lazy” manner. The idea underlying every lazy $\text{SMT}(\mathcal{T})$ procedure is that (a complete set of) the truth assignments for the propositional abstraction of ϕ are enumerated and checked for satisfiability in \mathcal{T} ; the procedure either returns sat if one \mathcal{T} -satisfiable truth assignment is found, or it returns unsat otherwise.

Figure 1 presents a simplified schema of a lazy $\text{SMT}(\mathcal{T})$ procedure, called the *offline schema*. The bijective function $\mathcal{T}2\mathcal{P}$ (“Theory-to-Boolean”), called *Boolean abstraction*, maps Boolean atoms into themselves and non-Boolean \mathcal{T} -atoms into fresh Boolean atoms—so that two atom instances in ϕ are mapped into the same Boolean atom iff they are syntactically identical—and extends to \mathcal{T} -formulas and sets of \mathcal{T} -formulas in the obvious way—that is, $\mathcal{T}2\mathcal{P}(\neg\phi_1) \stackrel{\text{def}}{=} \neg\mathcal{T}2\mathcal{P}(\phi_1)$, $\mathcal{T}2\mathcal{P}(\phi_1 \bowtie \phi_2) \stackrel{\text{def}}{=} \mathcal{T}2\mathcal{P}(\phi_1) \bowtie \mathcal{T}2\mathcal{P}(\phi_2)$ for each Boolean connective \bowtie , $\mathcal{T}2\mathcal{P}(\{\phi_i\}_i) \stackrel{\text{def}}{=} \{\mathcal{T}2\mathcal{P}(\phi_i)\}_i$. The function $\mathcal{P}2\mathcal{T}$ (“propositional-to-theory”), called *refinement*, is the inverse of $\mathcal{T}2\mathcal{P}$. The propositional abstraction ϕ^p of the input formula ϕ is given as input to an SAT solver based on the DPLL algorithm [Davis et al. 1962; Zhang and Malik 2002], which either decides that ϕ^p is unsatisfiable, and hence ϕ is \mathcal{T} -unsatisfiable, or returns a satisfying assignment μ^p ; in the latter case, $\mathcal{P}2\mathcal{T}(\mu^p)$ is given as input to \mathcal{T} -solver. If

²In the next sections, as we are in an $\text{SMT}(\mathcal{T})$ context, we often omit specifying “in the theory \mathcal{T} ” when speaking of consistency, validity, etc.

```

1. SatValue Lazy-SMT-Solver ( $\mathcal{T}$ -formula  $\phi$ ) {
2.    $\phi' = \text{convert\_to\_cnf}(\phi)$ 
3.    $\phi^p = T2P(\phi')$ 
4.   while (DPLL( $\phi^p, \mu^p$ ) == sat) {
5.      $(\rho, \eta) = T\text{-solver}(\mathcal{P}2T(\mu^p))$ 
6.     if ( $\rho == \text{sat}$ ) then return sat
7.      $\phi^p = \phi^p \wedge T2P(\neg\eta)$ 
8.   }
9.   return unsat
10. }
```

Fig. 1. A simplified schema for lazy SMT(\mathcal{T}) procedures.

$\mathcal{P}2T(\mu^p)$ is found \mathcal{T} -consistent, then ϕ is \mathcal{T} -consistent. If not, \mathcal{T} -solver returns the conflict set η which caused the \mathcal{T} -inconsistency of $\mathcal{P}2T(\mu^p)$; the abstraction of the \mathcal{T} -lemma $\neg\eta$, $T2P(\neg\eta)$, is then added as a clause to ϕ^p . Then the DPLL solver is restarted from scratch on the resulting formula.

Practical implementations follow a more elaborated schema, called the *on-line schema* (see Sebastiani [2007]; Barrett et al. [2009]). As before, ϕ^p is given as input to a modified version of DPLL, and when a satisfying assignment μ^p is found, the refinement $\mathcal{P}2T(\mu^p)$ of μ^p is fed to the \mathcal{T} -solver; if $\mathcal{P}2T(\mu^p)$ is found \mathcal{T} -consistent, then ϕ is \mathcal{T} -consistent; otherwise, \mathcal{T} -solver returns the conflict set η which caused the \mathcal{T} -inconsistency of $\mathcal{P}2T(\mu^p)$. Then the clause $\neg\eta^p$ is added in conjunction to ϕ^p , either temporarily or permanently (*\mathcal{T} -learning*), and the algorithm backtracks up to the highest point in the search where one of the literals in $\neg\eta^p$ is unassigned (*\mathcal{T} -backjumping*), and therefore its value is (propositionally) implied by the others in $\neg\eta^p$. Another important improvement is *early pruning (EP)*: before every literal selection, intermediate assignments are checked for \mathcal{T} -satisfiability and, if not \mathcal{T} -satisfiable, they are pruned (since no refinement can be \mathcal{T} -satisfiable). Finally, *theory propagation* can be used to reduce the search space by allowing the \mathcal{T} -solvers to explicitly return truth values for unassigned literals, which can be unit-propagated by the SAT solver. The interested reader is pointed to, for example, Sebastiani [2007] and Barrett et al. [2009] for details and further references.

With a small modification of the embedded DPLL engine, a lazy SMT solver can also be used to generate a resolution proof of unsatisfiability (see, e.g., van Gelder [2007]).

2.3 Interpolation in SMT

The use of interpolation in formal verification was introduced by McMillan [2003] for purely propositional formulas, and it was subsequently extended to handle SMT($\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$) formulas in McMillan [2005], \mathcal{EUF} being the theory of equality and uninterpreted functions. The technique is based on earlier work by Pudlák [1997], where two interpolant-generation algorithms were described: one for computing interpolants for propositional formulas from resolution proofs of unsatisfiability, and one for generating interpolants for conjunctions of (weak) linear inequalities in $\mathcal{LA}(\mathbb{Q})$. An interpolant for a pair (A, B)

Algorithm 1. Interpolant generation for $\text{SMT}(\mathcal{T})$

- (1) Generate a resolution proof of unsatisfiability \mathcal{P} for $A \wedge B$.
- (2) For every \mathcal{T} -lemma $\neg\eta$ occurring in \mathcal{P} , generate an interpolant $I_{\neg\eta}$ for $(\eta \setminus B, \eta \downarrow B)$.
- (3) For every input clause C in \mathcal{P} , set $I_C \stackrel{\text{def}}{=} C \downarrow B$ if $C \in A$, and $I_C \stackrel{\text{def}}{=} \top$ if $C \in B$.
- (4) For every inner node C of \mathcal{P} obtained by resolution from $C_1 \stackrel{\text{def}}{=} p \vee \phi_1$ and $C_2 \stackrel{\text{def}}{=} \neg p \vee \phi_2$, set $I_C \stackrel{\text{def}}{=} I_{C_1} \vee I_{C_2}$ if p does not occur in B , and $I_C \stackrel{\text{def}}{=} I_{C_1} \wedge I_{C_2}$ otherwise.
- (5) Output I_{\perp} as an interpolant for (A, B) .

of CNF formulas is constructed from a resolution proof of unsatisfiability of $A \wedge B$, generated as outlined in Section 2.1. The algorithm works by computing a formula I_C for each clause in the resolution refutation, such that the formula I_{\perp} associated with the empty root clause is the computed interpolant. In particular, the following invariant holds for every clause C and corresponding formula I_C in the resolution refutation [McMillan 2005]:

- (a) $A \models_{\mathcal{T}} I_C \vee (C \setminus B)$;
- (b) $B \wedge I_C \models (C \downarrow B)$; and
- (c) $I_C \preceq A$ and $I_C \preceq B$.

The algorithm can be described as follows.

Example 2.1. Consider the following two formulas in $\mathcal{LA}(\mathbb{Q})$:

$$A \stackrel{\text{def}}{=} (p \vee (0 \leq x_1 - 3x_2 + 1)) \wedge (0 \leq x_1 + x_2) \wedge (\neg q \vee \neg(0 \leq x_1 + x_2)),$$

$$B \stackrel{\text{def}}{=} (\neg(0 \leq x_3 - 2x_1 - 3) \vee (0 \leq 1 - 2x_3)) \wedge (\neg p \vee q) \wedge (p \vee (0 \leq x_3 - 2x_1 - 3)).$$

Figure 2(a) shows a resolution proof of unsatisfiability for $A \wedge B$, in which the clauses from A have been underlined. The proof contains the following $\mathcal{LA}(\mathbb{Q})$ lemma (displayed in boldface):

$$\neg(0 \leq x_1 - 3x_2 + 1) \vee \neg(0 \leq x_1 + x_2) \vee \neg(0 \leq x_3 - 2x_1 - 3) \vee \neg(0 \leq 1 - 2x_3).$$

Figure 2(b) shows, for each clause Θ_i in the proof, the formula I_{Θ_i} generated by Algorithm 1. For the $\mathcal{LA}(\mathbb{Q})$ lemma, it is easy to see that $(0 \leq 4x_1 + 1)$ is an interpolant for $((0 \leq x_1 - 3x_2 + 1) \wedge (0 \leq x_1 + x_2), (0 \leq x_3 - 2x_1 - 3) \wedge (0 \leq 1 - 2x_3))$ as required by Step 2 of the algorithm. (We will show how to obtain this interpolant in Example 2.2.) Therefore, $I_{\perp} \stackrel{\text{def}}{=} (p \vee (0 \leq 4x_1 + 1)) \wedge \neg q$ is an interpolant for (A, B) .

Algorithm 1 can be applied also when A and B are not in CNF. In this case, it suffices to preconvert them into CNF by using disjoint sets of auxiliary Boolean atoms in the usual way [McMillan 2005].

Notice that Step 2 of the algorithm is the only part which depends on the theory \mathcal{T} , so that the problem of interpolant generation in $\text{SMT}(\mathcal{T})$ reduces to that of finding interpolants for \mathcal{T} -lemmas. To this extent, McMillan [2005] gave a set of rules for constructing interpolants for \mathcal{T} lemmas in the theory of \mathcal{EUF} , that of weak linear inequalities $(0 \leq t)$ in $\mathcal{LA}(\mathbb{Q})$, and their combination. Linear equalities $(0 = t)$ can be reduced to conjunctions $(0 \leq t) \wedge (0 \leq -t)$ of inequalities. Thanks to the combination of theories, also strict linear inequalities $(0 < t)$ can

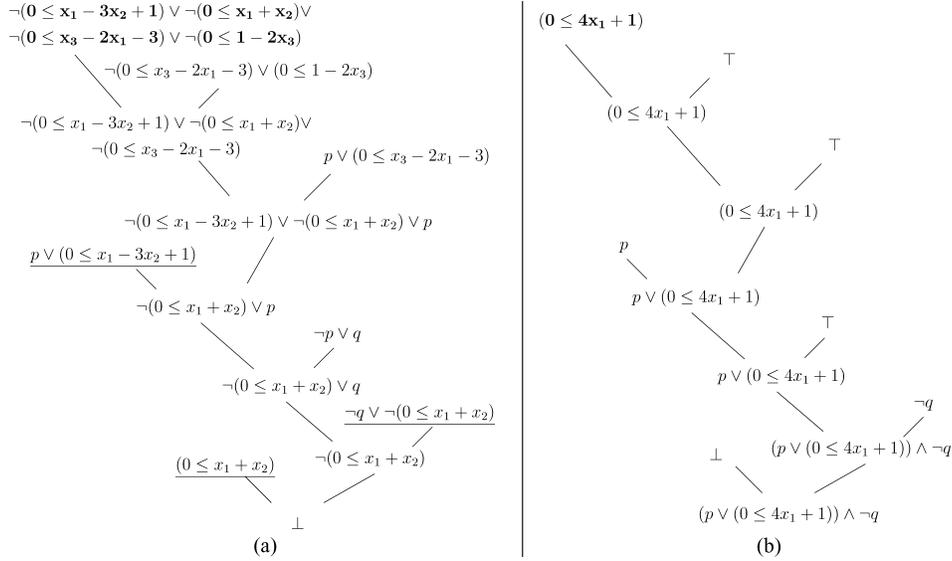


Fig. 2. Resolution proof of unsatisfiability (a) and interpolant (b) for the pair (A, B) of formulas of Example 2.1. In the tree on the left, \mathcal{T} -lemmas are displayed in boldface, and clauses from A are underlined.

be handled in $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ by replacing them with the conjunction $(0 \leq t) \wedge (0 \neq t)$,³ but this solution can be very inefficient.

The combination $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ can also be used to compute interpolants for other theories, such as those of lists, arrays, sets and multisets [Kapur et al. 2006].

In McMillan [2005], interpolants in the combined theory $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ were obtained by means of ad hoc combination rules. Yorsh and Musuvathi [2005], instead, presented a method for generating interpolants for $\mathcal{T}_1 \cup \mathcal{T}_2$ using the interpolant-generation procedures of \mathcal{T}_1 and \mathcal{T}_2 as black boxes, using the Nelson-Oppen approach [Nelson and Oppen 1979].

Also the method of Rybalchenko and Sofronie-Stokkermans [2007] allows the computation interpolants in $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$. Its peculiarity is that it is not based on unsatisfiability proofs. Instead, it generates interpolants in $\mathcal{LA}(\mathbb{Q})$ by solving a system of constraints using an off-the-shelf linear programming (LP) solver. The method allows both weak and strict inequalities. Extension to uninterpreted functions is achieved by means of reduction to $\mathcal{LA}(\mathbb{Q})$ using a hierarchical calculus [Sofronie-Stokkermans 2006]. The algorithm works only with conjunctions of atoms, although in principle it could be integrated in Algorithm 1 to generate interpolants for \mathcal{T} lemmas in $\mathcal{LA}(\mathbb{Q})$. As an alternative, Rybalchenko and Sofronie-Stokkermans [2007] showed how to generate interpolants for formulas that are in disjunctive normal form (DNF).

Another different approach was explored in Kroening and Weissenbacher [2007]. There the authors used the *eager* SMT approach to encode the

³The details were not given in McMillan [2005]. One possible way of doing this is to rewrite $(0 \neq t)$ as $(y = t) \wedge (z = 0) \wedge (z \neq y)$, z and y being fresh variables.

original SMT problem into an equisatisfiable propositional problem, for which a propositional proof of unsatisfiability is generated. This proof is later “lifted” to the original theory, and used to generate an interpolant in a way similar to Algorithm 1. At the moment, the approach is, however, limited to the theory of equality only (without uninterpreted functions).

The problem of *efficient* generation of interpolants in SMT has been addressed only recently. To the best of our knowledge, our work in Cimatti et al. [2008] was the first to go in this direction, with the introduction of efficient SMT-based interpolation algorithms for $\mathcal{LA}(\mathbb{Q})$, \mathcal{DL} , and combinations of convex theories. In Cimatti et al. [2009] covered interpolation for $UTVPI(\mathbb{Q})$ and $UTVPI(\mathbb{Z})$. Lynch and Tang [2008] presented interpolation algorithms for $\mathcal{LA}(\mathbb{Q})$ and some fragments of $\mathcal{LA}(\mathbb{Z})$. An efficient interpolation algorithm for \mathcal{EUF} was given in Fuchs et al. [2009]. Finally, another SMT-based method for interpolation in combined theories was recently proposed in Goel et al. [2009].

All the above techniques construct *one* interpolant for (A, B) . In general, however, interpolants are not unique. In particular, some of them can be better than others, depending on the particular application domain. Jhala and McMillan [2005] showed how to manipulate proofs in order to obtain stronger⁴ interpolants. Jhala and McMillan [2006, 2007], instead, presented a technique to restrict the language used in interpolants and showed it to be useful in preventing divergence of techniques based on predicate abstraction.

One of the most important applications of interpolation in formal verification is abstraction refinement [Henzinger et al. 2004; McMillan 2006]. In such a setting, every input problem ϕ has the form $\phi \stackrel{\text{def}}{=} \phi_1 \wedge \dots \wedge \phi_n$, and the interpolating solver is asked to compute several interpolants I_1, \dots, I_{n-1} corresponding to different partitions of ϕ into A_i and B_i , such that

$$\forall i. A_i \stackrel{\text{def}}{=} \phi_1 \wedge \dots \wedge \phi_i, \text{ and } B_i \stackrel{\text{def}}{=} \phi_{i+1} \wedge \dots \wedge \phi_n. \quad (1)$$

Moreover, I_1, \dots, I_{n-1} should be related by the following:

$$I_i \wedge \phi_{i+1} \models I_{i+1}. \quad (2)$$

A sufficient condition for (2) to hold is that all the I_i 's are computed from the same proof of unsatisfiability Π for ϕ [Henzinger et al. 2004].

2.3.1 Interpolants for Conjunctions of $\mathcal{LA}(\mathbb{Q})$ Literals. We recall the algorithm of McMillan [2005] for computing interpolants from $\mathcal{LA}(\mathbb{Q})$ proofs of unsatisfiability, for conjunctions of equalities and weak inequalities in $\mathcal{LA}(\mathbb{Q})$.

An $\mathcal{LA}(\mathbb{Q})$ *proof rule* R for a conjunction Γ of equalities and weak inequalities is either an element of Γ , or it has the form $\frac{P}{\phi}$, where ϕ is an equality or a weak inequality and P is a sequence of proof rules, called the *premises* of R . An $\mathcal{LA}(\mathbb{Q})$ *proof of unsatisfiability* for a conjunction of equalities and weak inequalities Γ is simply a rule in which $\phi \equiv 0 \leq c$ and where c is a negative numerical constant.⁵

⁴In the implication ordering (see Section 3.3).

⁵In the following, we might sometimes write \perp as a synonym of an atom “ $0 \leq c$ ” when c is a negative numerical constant.

$$\text{LEQEQ} \frac{0 = t}{0 \leq t} \quad \text{COMB} \frac{0 \leq t_1 \quad 0 \leq t_2}{0 \leq c_1 t_1 + c_2 t_2} \quad c_1, c_2 > 0$$

Fig. 3. $\mathcal{LA}(\mathbb{Q})$ -proof rules for a conjunction Γ of equalities and weak inequalities.

Similarly to McMillan [2005], we use the proof rules of Figure 3: LEQEQ for deriving inequalities from equalities, and COMB for performing linear combinations.⁶

Given an $\mathcal{LA}(\mathbb{Q})$ proof of unsatisfiability P for a conjunction Γ of equalities and weak inequalities partitioned into (A, B) , an interpolant I can be computed simply by replacing every atom $0 \leq t$ occurring in B (respectively, $0 = t$) with $0 \leq 0$ (respectively, $0 = 0$) in each leaf subrule of P , and propagating the results: the interpolant is then the single weak inequality $0 \leq t$ at the root of P [McMillan 2005].

Example 2.2. Consider the following sets of $\mathcal{LA}(\mathbb{Q})$ atoms:

$$A \stackrel{\text{def}}{=} \{(0 \leq x_1 - 3x_2 + 1), (0 \leq x_1 + x_2)\},$$

$$B \stackrel{\text{def}}{=} \{(0 \leq x_3 - 2x_1 - 3), (0 \leq 1 - 2x_3)\}.$$

An $\mathcal{LA}(\mathbb{Q})$ proof of unsatisfiability P for $A \wedge B$ is the following:

$$\frac{\frac{1 * (0 \leq x_1 - 3x_2 + 1) \quad 4 * (0 \leq x_1 + x_2)}{1 * (0 \leq 4x_1 + 1)} \quad \frac{2 * (0 \leq x_3 - 2x_1 - 3) \quad 1 * (0 \leq 1 - 2x_3)}{1 * (0 \leq -4x_1 - 5)}}{(0 \leq -4)}$$

By replacing inequalities in B with $(0 \leq 0)$, we obtain the proof P' :

$$\frac{\frac{1 * (0 \leq x_1 - 3x_2 + 1) \quad 4 * (0 \leq x_1 + x_2)}{1 * (0 \leq 4x_1 + 1)} \quad \frac{2 * (0 \leq 0) \quad 1 * (0 \leq 0)}{1 * (0 \leq 0)}}{(0 \leq 4x_1 + 1)}$$

Thus, the interpolant obtained is $(0 \leq 4x_1 + 1)$.

3. FROM SMT($\mathcal{LA}(\mathbb{Q})$) SOLVING TO SMT($\mathcal{LA}(\mathbb{Q})$) INTERPOLATION

Traditionally, SMT solvers used some kind of incremental simplex algorithm [Vanderbei 2001] as a T -solver for the $\mathcal{LA}(\mathbb{Q})$ theory. Recently, Dutertre and de Moura [2006] have proposed a new simplex-based algorithm, specifically designed for integration in a lazy SMT solver. The algorithm is extremely suitable for SMT, and SMT solvers embedding it were shown to significantly outperform (often by orders of magnitude) the ones based on other simplex variants. It has now been integrated in several SMT solvers, including ARGOLIB, CVC3, MATHSAT, YICES, and Z3. Remarkably, this algorithm also allows the handling of strict inequalities.

In this section, we shall deal only with conjunctions of constraints in $\mathcal{LA}(\mathbb{Q})$. Interpolation for the SMT ($\mathcal{LA}(\mathbb{Q})$) formula, is then obtained by combining this

⁶In McMillan [2005] the LEQEQ rule is not used in $\mathcal{LA}(\mathbb{Q})$, because the input is assumed to consist only of inequalities.

algorithm with the general procedure for SM (τ) described in Section 2.3. In Section 3.1 we begin by considering the case in which the input atoms are only equalities and nonstrict inequalities. In this case, we only need to show how to generate a proof of unsatisfiability, since then we can use the interpolation rules defined in McMillan [2005]. Then, in Section 3.2 we show how to generate interpolants for problems containing also strict inequalities and disequalities.

3.1 Interpolation with Nonstrict Inequalities

3.1.1 The Original Dutertre-de Moura Algorithm. In its original formulation, the Dutertre-de Moura algorithm assumes that the variables x_i are partitioned a priori into two sets, hereafter denoted as $\hat{\mathcal{B}}$ (“initially basic” or “dependent”) and $\hat{\mathcal{N}}$ (“initially nonbasic” or “independent”), and that the algorithm receives as inputs two kinds of atomic formulas⁷:

- a set of *equations* eq_i , one for each $x_i \in \hat{\mathcal{B}}$, of the form $\sum_{x_j \in \hat{\mathcal{N}}} \hat{a}_{ij}x_j + \hat{a}_{ii}x_i = 0$ such that all \hat{a}_{ij} ’s are numerical constants;
- elementary atoms* of the form $x_j \geq l_j$ or $x_j \leq u_j$ such that $x_j \in \hat{\mathcal{B}} \cup \hat{\mathcal{N}}$ and l_j, u_j are numerical constants.

In order to handle problems that are not in the above form, a satisfiability-preserving preprocessing step is applied upfront, before invoking the algorithm.

The initial equations eq_i are then used to build a tableau T :

$$\{x_i = \sum_{x_j \in \hat{\mathcal{N}}} a_{ij}x_j \mid x_i \in \mathcal{B}\}, \quad (3)$$

where \mathcal{B} (“basic” or “dependent”), \mathcal{N} (“nonbasic” or “independent”) and a_{ij} are such that initially $\mathcal{B} \equiv \hat{\mathcal{B}}$, $\mathcal{N} \equiv \hat{\mathcal{N}}$ and $a_{ij} \equiv -\hat{a}_{ij}/\hat{a}_{ii}$.

In order to decide the satisfiability of the input problem, the algorithm performs manipulations of the tableau that change the sets \mathcal{B} and \mathcal{N} and the values of the coefficients a_{ij} , always keeping the tableau T in (3) equivalent to its initial version. In particular, the algorithm maintains a mapping $\beta : \mathcal{B} \cup \mathcal{N} \mapsto \mathbb{Q}$ representing a candidate model which, at every step, satisfies the following invariants:

$$\forall x_j \in \mathcal{N}, \quad l_j \leq \beta(x_j) \leq u_j, \quad \forall x_i \in \mathcal{B}, \quad \beta(x_i) = \sum_{j \in \mathcal{N}} a_{ij}\beta(x_j). \quad (4)$$

The algorithm tries to adjust the values of β and the sets \mathcal{B} and \mathcal{N} , and hence the coefficients a_{ij} of the tableau, such that $l_i \leq \beta(x_i) \leq u_i$ holds also for all the x_i ’s in \mathcal{B} . Inconsistency is detected when this is not possible without violating any constraint in (4): as the bounds on the variables in \mathcal{N} are always satisfied by β , then there is a variable $x_i \in \mathcal{B}$ such that the inconsistency is caused either by the elementary atom $x_i \geq l_i$ or by the atom $x_i \leq u_i$ [Dutertre and de Moura 2006]; in the first case,⁸ a conflict set η is generated as follows:

$$\eta = \{x_j \leq u_j \mid x_j \in \mathcal{N}^+\} \cup \{x_j \geq l_j \mid x_j \in \mathcal{N}^-\} \cup \{x_i \geq l_i\}, \quad (5)$$

where $(x_i = \sum_{x_j \in \mathcal{N}} a_{ij}x_j)$ is the row of the current version of the tableau T (3) corresponding to x_i , \mathcal{N}^+ is $\{x_j \in \mathcal{N} \mid a_{ij} > 0\}$ and \mathcal{N}^- is $\{x_j \in \mathcal{N} \mid a_{ij} < 0\}$.

⁷Notationally, we use the hat symbol $\hat{\cdot}$ to denote the initial value of the generic symbol.

⁸Here we do not consider the second case $x_i \leq u_i$ as it is analogous to the first one.

Notice that η is a conflict set in the sense that it is made inconsistent by (some of) the equations in the tableau T (3), that is, $T \cup \eta \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$. In general, however, $\eta \not\models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$.

3.1.2 Our Proof-Producing Variant. In order to make it suitable for interpolant generation, we have conceived the following variant of the Dutertre-de Moura algorithm.

We take as input an arbitrary set of inequalities $l_k \leq \sum_h \hat{a}_{kh} y_h$ or $u_k \geq \sum_h \hat{a}_{kh} y_h$, and apply an internal preprocessing step to obtain a set of equations and a set of elementary bounds. In particular, we introduce a “slack” variable s_k for each distinct term $\sum_h \hat{a}_{kh} y_h$ occurring in the input inequalities. Then we replace such term with s_k (thus obtaining $l_k \leq s_k$ or $u_k \geq s_k$) and add an equation $s_k = \sum_h \hat{a}_{kh} y_h$. Notice that we introduce a slack variable even for “elementary” inequalities ($l_k \leq y_k$). With this transformation, the initial tableau T (3) is

$$\{s_k = \sum_h \hat{a}_{kh} y_h\}_k, \quad (6)$$

such that $\hat{\mathcal{B}}$ is made of all the slack variables s_k 's, $\hat{\mathcal{N}}$ is made of all the original variables y_h 's, and the elementary atoms contain only slack variables s_k 's.

Then the algorithm proceeds as described above, producing a set η (5) in case of inconsistency. In our variant of the algorithm, we can use η to generate a conflict set η' , thanks to the following theorem.

THEOREM 3.1. *In the set η of (5), x_i and all the x_j 's are slack variables introduced by our preprocessing step. Moreover, the set $\eta' \stackrel{\text{def}}{=} \eta_{\mathcal{N}^+} \cup \eta_{\mathcal{N}^-} \cup \eta_i$ is a conflict set, where*

$$\begin{aligned} \eta_{\mathcal{N}^+} &\stackrel{\text{def}}{=} \{u_k \geq \sum_h \hat{a}_{kh} y_h \mid s_k \equiv x_j \text{ and } x_j \in \mathcal{N}^+\}, \\ \eta_{\mathcal{N}^-} &\stackrel{\text{def}}{=} \{l_k \leq \sum_h \hat{a}_{kh} y_h \mid s_k \equiv x_j \text{ and } x_j \in \mathcal{N}^-\}, \\ \eta_i &\stackrel{\text{def}}{=} \{l_k \leq \sum_h \hat{a}_{kh} y_h \mid s_k \equiv x_i\}. \end{aligned}$$

PROOF. We consider the case in which η (5) is generated from a row $x_i = \sum_{x_j \in \mathcal{N}} a_{ij} x_j$ in the tableau T (3) such that $\beta(x_i) < l_i$. Dutertre and de Moura [2006] showed that in this case the following facts hold:

$$\forall x_j \in \mathcal{N}^+, \beta(x_j) = u_j, \quad \text{and} \quad \forall x_j \in \mathcal{N}^-, \beta(x_j) = l_j. \quad (7)$$

(We recall that $\mathcal{N}^+ = \{x_j \in \mathcal{N} \mid a_{ij} > 0\}$ and $\mathcal{N}^- = \{x_j \in \mathcal{N} \mid a_{ij} < 0\}$.) The bounds u_j and l_j can be introduced only by elementary atoms. Since in our variant the elementary atoms contain only slack variables, each x_j must be a slack variable (namely, s_k). The same holds for x_i (since its value is bounded by l_i).

Now consider η again. Dutertre and de Moura [2006] showed that when a conflict is detected because $\beta(x_i) < l_i$, then the following fact holds:

$$\beta(x_i) = \sum_{x_j \in \mathcal{N}^+} a_{ij} u_j + \sum_{x_j \in \mathcal{N}^-} a_{ij} l_j. \quad (8)$$

From the i th row of tableau T (3) we can derive

$$0 \leq \sum_{x_j \in \mathcal{N}} a_{ij} x_j - x_i. \quad (9)$$

If we take each inequality $0 \leq u_j - x_j$ multiplied by the coefficient a_{ij} for all $x_j \in \mathcal{N}^+$, each inequality $0 \leq x_j - l_j$ multiplied by coefficient $-a_{ij}$ for all

$x_j \in \mathcal{N}^-$, and the inequality $(0 \leq x_i - l_i)$ multiplied by 1, and we add them to (9), we obtain

$$0 \leq \sum_{\mathcal{N}^+} a_{ij} u_j + \sum_{\mathcal{N}^-} a_{ij} l_j - l_i, \quad (10)$$

which by (8) is equivalent to $0 \leq \beta(x_i) - l_i$. Thus we have obtained $0 \leq c$ with $c \equiv \beta(x_i) - l_i$, which is strictly lower than zero. Therefore, η is inconsistent under the definitions in T . Since we know that x_i and all the x_j 's in η are slack variables, we can replace every x_j (i.e., every s_k) with its corresponding term $\sum_h \hat{a}_{kh} y_h$, thus obtaining η' , which is thus inconsistent. \square

When our variant of the algorithm detects an inconsistency, we construct a proof of unsatisfiability as follows. From the set η of (5) we build a conflict set η' by replacing each elementary atom in it with the corresponding original atom, as shown in Theorem 3.1. Using the HYP rule, we introduce all the atoms in $\eta_{\mathcal{N}^+}$, and combine them with repeated applications of the COMB rule: if $u_k \geq \sum_h \hat{a}_{kh} y_h$ is the atom corresponding to s_k , we use as coefficient for the COMB the a_{ij} (in the i th row of the current tableau) such that $s_k \equiv x_j$. Then we introduce each of the atoms in $\eta_{\mathcal{N}^-}$ with HYP, and add them to the previous combination, again using COMB. In this case, the coefficient to use is $-a_{ij}$. Finally, we introduce the atom in η_i and add it to the combination with coefficient 1.

COROLLARY 3.2. *The result of the linear combination described above is the atom $0 \leq c$, such that c is a numerical constant strictly lower than zero.*

PROOF. Follows immediately by the proof of Theorem 3.1. \square

Besides the case just described (and its dual when the inconsistency is due to an elementary atom $x_i \leq u_i$), another case in which an inconsistency can be detected is when two contradictory atoms are asserted: $l_k \leq \sum_h \hat{a}_{kh} y_h$ and $u_k \geq \sum_h \hat{a}_{kh} y_h$, with $l_k > u_k$. In this case, the proof is simply the combination of the two atoms with coefficient 1.

The extension for handling also equalities like $b_k = \sum_h \hat{a}_{kh} y_h$ is straightforward: we simply introduce two elementary atoms $b_k \leq s_k$ and $b_k \geq s_k$ and, in the construction of the proof, we use the LEQEQ rule to introduce the proper inequality.

Finally, notice that the current implementation in MATHSAT (see Section 7) is slightly different from what presented here, and significantly more efficient. In practice, η , η' are not constructed in sequence; rather, they are built simultaneously. Moreover, some optimizations are applied to eliminate some slack variables when they are not needed. (In particular, in an actual implementation there is no need of introducing slack variables for elementary inequalities of the form $l_k \leq x_k$. We refer the reader to de Moura and Bjørner [2008] for more details about this fact.)

Example 3.1. Consider again the two sets of $\mathcal{L}\mathcal{A}(\mathbb{Q})$ atoms of Example 2.2:

$$\begin{aligned} A &\stackrel{\text{def}}{=} \{(0 \leq x_1 - 3x_2 + 1), (0 \leq x_1 + x_2)\}, \\ B &\stackrel{\text{def}}{=} \{(0 \leq x_3 - 2x_1 - 3), (0 \leq 1 - 2x_3)\}. \end{aligned}$$

With our variant of the Dutertre-de Moura algorithm, four “slack” variables are introduced, resulting in the following tableau and elementary constraints:

$$T \stackrel{\text{def}}{=} \begin{cases} s_1 = x_1 - 3x_2, & -1 \leq s_1, \\ s_2 = x_1 + x_2, & 0 \leq s_2, \\ s_3 = x_3 - 2x_1, & 3 \leq s_3, \\ s_4 = -2x_3, & -1 \leq s_4. \end{cases}$$

To detect the inconsistency, the algorithm performs some pivoting steps, resulting in the final tableau T' :

$$T' \stackrel{\text{def}}{=} \begin{cases} x_2 = -\frac{1}{12}s_4 - \frac{1}{6}s_3 - \frac{1}{3}s_1, \\ s_2 = -\frac{1}{3}s_4 - \frac{2}{3}s_3 - \frac{1}{3}s_1, \\ x_1 = -\frac{1}{4}s_4 - \frac{1}{2}s_3, \\ x_3 = -\frac{1}{2}s_4. \end{cases}$$

The final values of β are as follows:

$$\begin{aligned} \beta(x_1) &= \frac{7}{4}, & \beta(x_2) &= -\frac{1}{12}, & \beta(x_3) &= \frac{1}{2}, \\ \beta(s_1) &= -1, & \beta(s_2) &= -\frac{4}{3}, & \beta(s_3) &= 3, & \beta(s_4) &= -1. \end{aligned}$$

Therefore, the bound ($0 \leq s_2$) is violated. From the second row of T' , the set η and the conflict set η' are computed:

$$\begin{aligned} \eta &\stackrel{\text{def}}{=} \emptyset \cup \{(-1 \leq s_4), (3 \leq s_3), (-1 \leq s_1)\} \cup \{(0 \leq s_2)\}, \\ \eta' &\stackrel{\text{def}}{=} \emptyset \cup \{(0 \leq 1 - 2x_3), (0 \leq x_3 - 2x_1 - 3), (0 \leq x_1 - 3x_2 + 1)\} \cup \{(0 \leq x_1 + x_2)\}. \end{aligned}$$

The generated proof of unsatisfiability P is

$$\frac{\frac{1}{3} * (0 \leq 1 - 2x_3) \quad \frac{2}{3} * (0 \leq x_3 - 2x_1 - 3)}{1 * (0 \leq -\frac{4}{3}x_1 - \frac{5}{3})} \quad \frac{\frac{1}{3} * (0 \leq x_1 - 3x_2 + 1)}{1 * (0 \leq -x_1 - x_2 - \frac{4}{3})} \quad \frac{1 * (0 \leq x_1 + x_2)}{(0 \leq -\frac{4}{3})} .$$

After replacing the inequalities of B with ($0 \leq 0$) in P , the new proof P' is

$$\frac{\frac{1}{3} * (0 \leq 0) \quad \frac{2}{3} * (0 \leq 0)}{1 * (0 \leq 0)} \quad \frac{\frac{1}{3} * (0 \leq x_1 - 3x_2 + 1)}{1 * (0 \leq \frac{1}{3}x_1 - x_2 + \frac{1}{3})} \quad \frac{1 * (0 \leq x_1 + x_2)}{(0 \leq \frac{4}{3}x_1 + \frac{1}{3})} .$$

Thus the computed interpolant is ($0 \leq \frac{4}{3}x_1 + \frac{1}{3}$) (which is equivalent to that of Example 2.2).

3.2 Interpolation with Strict Inequalities and Disequalities

Another benefit of the Dutertre-de Moura algorithm is that it can handle strict inequalities directly. Its method is based on the following lemma.

LEMMA 3.3 (LEMMA 1 IN DUTERTRE AND DE MOURA [2006]). *A set of linear arithmetic atoms Γ containing strict inequalities $S = \{0 < t_1, \dots, 0 < t_n\}$ is satisfiable iff there exists a rational number $\varepsilon > 0$ such that $\Gamma_\varepsilon \stackrel{\text{def}}{=} (\Gamma \cup S_\varepsilon) \setminus S$ is satisfiable, where $S_\varepsilon \stackrel{\text{def}}{=} \{\varepsilon \leq t_1, \dots, \varepsilon \leq t_n\}$.*

The idea of Dutertre and de Moura [2006] is that of treating the *infinitesimal parameter* ε symbolically instead of explicitly computing its value. Strict bounds ($x < b$) are replaced with weak ones ($x \leq b - \varepsilon$), and the operations on bounds are adjusted to take ε into account.

We extend the same idea to the computation of interpolants. We transform every atom ($0 < t_i$) occurring in the proof of unsatisfiability into ($0 \leq t_i - \varepsilon$). Then we compute an interpolant I_ε in the usual way. As a consequence of the rules of McMillan [2005], I_ε is always a single atom. As shown by the following lemma, if I_ε contains ε , then it must be in the form ($0 \leq t - c\varepsilon$) with $c > 0$, and we can rewrite I_ε into ($0 < t$).

THEOREM 3.4 (INTERPOLATION WITH STRICT INEQUALITIES). *Let Γ , S , Γ_ε , and S_ε be defined as in Lemma 3.3. Let Γ be partitioned into A and B , and let A_ε and B_ε be obtained from A and B by replacing atoms in S with the corresponding ones in S_ε . Let I_ε be an interpolant for $(A_\varepsilon, B_\varepsilon)$. Then*

- if $\varepsilon \not\leq I_\varepsilon$, then I_ε is an interpolant for (A, B) .
- if $\varepsilon \leq I_\varepsilon$, then $I_\varepsilon \equiv (0 \leq t - c\varepsilon)$ for some $c > 0$, and $I \stackrel{\text{def}}{=} (0 < t)$ is an interpolant for (A, B) .

PROOF. Since the side condition of the COMB rule ensures that equations are combined only using positive coefficients, and since the atoms introduced in the proof either do not contain ε or contain it with a negative coefficient, if ε appears in I_ε , it must have a negative coefficient.

If ε does not appear in I_ε , then I_ε has been obtained from atoms appearing in A or B , so that I_ε is an interpolant for (A, B) .

If ε appears in I_ε , since its value has not been explicitly computed, it can be arbitrarily small, so thanks to Lemma 3.3 we have that $B_\varepsilon \wedge I_\varepsilon \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$ implies $B \wedge I \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$.

We can prove that $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I$ as follows. We consider some interpretation μ which is a model for A . Since ε does not occur in A , we can extend μ by setting $\mu(\varepsilon) = \delta$ for some $\delta > 0$ such that μ is a model also for A_ε . As $A_\varepsilon \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I_\varepsilon$, μ is also a model for I_ε , and hence μ is also a model for I . Thus, we have that $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I$. \square

Notice that Theorem 3.4 can be extended straightforwardly to the case in which the interpolant is a conjunction of inequalities.

Thus, in case of strict inequalities, Theorem 3.4 gives us a way for constructing interpolants with no need of expensive theory combination (as instead was the case in McMillan [2005]). Moreover, thanks to it we can handle also negated equalities ($0 \neq t$) directly. Suppose our set S of input atoms (partitioned into A and B) is the union of a set S' of equalities and inequalities (both weak and strict) and a set S^\neq of disequalities, and suppose that S' is consistent. (If not so, an interpolant can be computed from S' .) Since $\mathcal{L}\mathcal{A}(\mathbb{Q})$ is convex, S is

inconsistent iff exists $(0 \neq t) \in S^\neq$ such that $S' \cup \{(0 \neq t)\}$ is inconsistent, that is, such that both $S' \cup \{(0 < t)\}$ and $S' \cup \{(0 > t)\}$ are inconsistent.

Therefore, we pick one element $(0 \neq t)$ of S^\neq at a time, and check the satisfiability of $S' \cup \{(0 < t)\}$ and $S' \cup \{(0 > t)\}$. If both are inconsistent, from the two proofs we can generate two interpolants I^- and I^+ . We combine I^+ and I^- to obtain an interpolant I for (A, B) : if $(0 \neq t) \in A$, then I is $I^+ \vee I^-$; if $(0 \neq t) \in B$, then I is $I^+ \wedge I^-$, as shown by the following theorem.

THEOREM 3.5 (INTERPOLATION FOR NEGATED EQUALITIES). *Let A and B two conjunctions of $\mathcal{L}\mathcal{A}(\mathbb{Q})$ atoms, and let $n \stackrel{\text{def}}{=} (0 \neq t)$ be one such atom. Let $g \stackrel{\text{def}}{=} (0 < t)$ and $l \stackrel{\text{def}}{=} (0 > t)$.*

- If $n \in A$, then let $A^+ \stackrel{\text{def}}{=} A \setminus \{n\} \cup \{g\}$, $A^- \stackrel{\text{def}}{=} A \setminus \{n\} \cup \{l\}$, and $B^+ \stackrel{\text{def}}{=} B^- \stackrel{\text{def}}{=} B$.
- If $n \in B$, then let $A^+ \stackrel{\text{def}}{=} A^- \stackrel{\text{def}}{=} A$, $B^+ \stackrel{\text{def}}{=} B \setminus \{n\} \cup \{g\}$, and $B^- \stackrel{\text{def}}{=} B \setminus \{n\} \cup \{l\}$.

Assume that $A^+ \wedge B^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$ and that $A^- \wedge B^- \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$, and let I^+ and I^- be two interpolants for (A^+, B^+) and (A^-, B^-) , respectively, and let

$$I \stackrel{\text{def}}{=} \begin{cases} I^+ \vee I^-, & \text{if } n \in A, \\ I^+ \wedge I^-, & \text{if } n \in B. \end{cases}$$

Then I is an interpolant for (A, B) .

PROOF. We have to prove that

- (i) $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I$;
- (ii) $B \wedge I \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$;
- (iii) $I \leq A$ and $I \leq B$.

- (i) If $n \in A$, then $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} g \vee l$. By hypothesis, we know that $A^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^+$ and $A^- \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^-$. Then trivially $A \cup \{g\} \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^+$ and $A \cup \{l\} \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^-$. Therefore $A \cup \{g\} \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^+ \vee I^-$ and $A \cup \{l\} \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^- \vee I^+$, so that $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I$.

If $n \in B$, then $A^+ \equiv A^- \equiv A$. By hypothesis $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^+$ and $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I^-$, so that $A \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} I$.

- (ii) If $n \in A$, then $B^+ \equiv B^- \equiv B$. By hypothesis $B \wedge I^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$ and $B \wedge I^- \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$, so that $B \wedge I \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$.

If $n \in B$, then $B \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} g \vee l$, so that either $B \rightarrow g$ or $B \rightarrow l$ must hold. By hypothesis we have $B^+ \wedge I^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$, so that $B \cup \{g\} \wedge I^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$. If $B \rightarrow g$ holds, then $B \wedge I^+ \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$, and hence $B \wedge I \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$. Similarly, if $B \rightarrow l$ holds, then $B \wedge I^- \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$, and so again $B \wedge I \models_{\mathcal{L}\mathcal{A}(\mathbb{Q})} \perp$.

- (iii) By the hypothesis, both I^+ and I^- contain only symbols common to A and B , so that $I \leq A$ and $I \leq B$. \square

Example 3.2. Consider the following sets of $\mathcal{L}\mathcal{A}(\mathbb{Q})$ atoms:

$$A \stackrel{\text{def}}{=} \{(0 \neq x_1 - 3x_2 + 1), (0 = x_1 + x_2)\},$$

$$B \stackrel{\text{def}}{=} \{(0 = x_3 - 2x_1 - 1), (0 = 1 - 2x_3)\}.$$

To compute an interpolant for (A, B) , we first split $n \stackrel{\text{def}}{=} (0 \neq x_1 - 3x_2 + 1)$ into $g \stackrel{\text{def}}{=} (0 < x_1 - 3x_2 + 1)$ and $l \stackrel{\text{def}}{=} (0 < -x_1 + 3x_2 - 1)$, thus obtaining A^+ and A^- defined

as in Theorem 3.5. We then generate two $\mathcal{LA}(\mathbb{Q})$ proofs of unsatisfiability P^+ for $A^+ \wedge B$ and P^- for $A^- \wedge B$, and replace g in P^+ with $g_\varepsilon \stackrel{\text{def}}{=} (0 \leq x_1 - 3x_2 + 1 - \varepsilon)$ and l in P^- with $l_\varepsilon \stackrel{\text{def}}{=} (0 \leq -x_1 + 3x_2 - 1 - \varepsilon)$, obtaining P_ε^+ and P_ε^- (we omit the names of the inference rules):

$$\begin{array}{c}
 P_\varepsilon^+ \stackrel{\text{def}}{=} \frac{\frac{(0 = x_1 + x_2) \quad (0 = x_3 - 2x_1 - 1) \quad (0 = 1 - 2x_3)}{(0 \leq x_1 - 3x_2 + 1 - \varepsilon) \quad (0 \leq x_1 + x_2) \quad (0 \leq x_3 - 2x_1 - 1) \quad (0 \leq 1 - 2x_3)}}{(0 \leq 4x_1 + 1 - \varepsilon)} \quad \frac{(0 \leq x_3 - 2x_1 - 1) \quad (0 = 1 - 2x_3)}{(0 \leq x_3 - 2x_1 - 1) \quad (0 \leq 1 - 2x_3)}}{(0 \leq -4x_1 - 1),}{(0 \leq -\varepsilon)} \\
 \\
 P_\varepsilon^- \stackrel{\text{def}}{=} \frac{\frac{(0 = x_1 + x_2) \quad (0 = x_3 - 2x_1 - 1) \quad (0 = 1 - 2x_3)}{(0 \leq -x_1 + 3x_2 - 1 - \varepsilon) \quad (0 \leq -x_1 - x_2) \quad (0 \leq -x_3 + 2x_1 + 1) \quad (0 \leq -1 + 2x_3)}}{(0 \leq -4x_1 - 1 - \varepsilon)} \quad \frac{(0 = x_3 - 2x_1 - 1) \quad (0 = 1 - 2x_3)}{(0 \leq -x_3 + 2x_1 + 1) \quad (0 \leq -1 + 2x_3)}}{(0 \leq +4x_1 + 1).}{(0 \leq -\varepsilon)}
 \end{array}$$

We then compute the two interpolants I_ε^+ from P_ε^+ and I_ε^- from P_ε^- :

$$I_\varepsilon^+ \stackrel{\text{def}}{=} (0 \leq 4x_1 + 1 - \varepsilon), \quad I_\varepsilon^- \stackrel{\text{def}}{=} (0 \leq -4x_1 - 1 - \varepsilon).$$

Therefore, according to Theorem 3.4 the two interpolants I^+ for (A^+, B) and I^- for (A^-, B) are

$$I^+ \stackrel{\text{def}}{=} (0 < 4x_1 + 1), \quad I^- \stackrel{\text{def}}{=} (0 < -4x_1 - 1).$$

Finally, since $n \in B$, according to Theorem 3.5, the interpolant I for (A, B) is

$$I \stackrel{\text{def}}{=} I^+ \vee I^- \equiv (0 < 4x_1 + 1) \vee (0 < -4x_1 - 1).$$

3.3 Obtaining Stronger Interpolants

We conclude this section by illustrating a simple technique for improving the strength of interpolants in $\mathcal{LA}(\mathbb{Q})$. The technique is orthogonal to our proof-generation algorithm described in Section 3.1.2, and it is therefore of independent interest. It is an improvement of the general algorithm of McMillan [2005] (and outlined in Section 2.3.1) for generating interpolants from $\mathcal{LA}(\mathbb{Q})$ proofs of unsatisfiability.

Definition 3.6. Given two interpolants I_1 and I_2 for the same pair (A, B) of conjunctions of $\mathcal{LA}(\mathbb{Q})$ literals, we say that I_1 is *stronger* than I_2 if and only if $I_1 \models_{\mathcal{LA}(\mathbb{Q})} I_2$ but $I_2 \not\models_{\mathcal{LA}(\mathbb{Q})} I_1$.

A simple way of obtaining strong interpolants in $\mathcal{LA}(\mathbb{Q})$ is to exploit *quantifier elimination* procedures. In this context, a quantifier elimination procedure (denoted with `QuantElim` in the following) is a procedure that takes as input an existentially quantified formula in prenex normal form⁹ and produces an equivalent quantifier-free formula. Given an $\mathcal{LA}(\mathbb{Q})$ -inconsistent pair (A, B) of conjunctions of constraints, an interpolant can be computed by existentially quantifying all the A -local variables in A , and then performing quantifier elimination:

$$I \stackrel{\text{def}}{=} \text{QuantElim}(\exists(\{x_i \mid x_i \notin B\}).A). \quad (11)$$

⁹We recall that a formula φ is in *prenex normal form* if and only if it is written as a string of quantifiers followed by a quantifier-free part.

The interpolant (11) obtained via quantifier elimination is the strongest possible for (A, B) . Moreover, such a technique can be applied also when the variables are interpreted over the integers rather than over the rationals. In general, however, quantifier elimination for $\mathcal{LA}(\mathbb{Q})$ is very expensive (and even more so for $\mathcal{LA}(\mathbb{Z})$), and for large problems it can be impractical: even in the restricted case of conjunctions of constraints only, in fact, it can cause an exponential blow-up in the size of the formula in the worst case (see, e.g., Kroening and Strichman [2008]).

We present instead an alternative algorithm which, although it does not guarantee the generation of the strongest interpolant, can be implemented very efficiently on top of proof-generating $\mathcal{LA}(\mathbb{Q})$ -decision procedures. Our technique is based on the simple observation that the only purpose of the summations performed during the traversal of proof trees for computing the interpolant (as described in Section 2.3.1) is that of eliminating A -local variables. In fact, it is easy to see that the conjunction of the constraints of A occurring as leaves in an $\mathcal{LA}(\mathbb{Q})$ proof of unsatisfiability satisfies the first two points of the definition of interpolant (Definition 2.2): if such constraints do not contain A -local variables, therefore, their conjunction is already an interpolant; if not, it suffices to perform only the summations of constraints of A that are necessary to eliminate A -local variables. Moreover, such an interpolant is stronger than that obtained by performing the summations with the coefficients found in the proof tree, since for any set of constraints $\{s_1, \dots, s_n\}$ and any set of positive coefficients $\{c_1, \dots, c_n\}$, $s_1 \wedge \dots \wedge s_n \models_{\mathcal{LA}(\mathbb{Q})} \sum_{i=1}^n c_i * s_i$ holds.

According to this observation, our proposal can be described as: *perform only those summations which are necessary for eliminating A -local variables.*

Example 3.3. Consider the following sets of $\mathcal{LA}(\mathbb{Q})$ atoms:

$$A \stackrel{\text{def}}{=} \left\{ (0 \leq 3x_1 - 4x_2 + 1), (0 \leq x_2 - 2x_1), \left(0 \leq x_2 - \frac{1}{3}x_3\right), \left(0 \leq x_4 - \frac{3}{2}x_5 - 1\right) \right\},$$

$$B \stackrel{\text{def}}{=} \{(0 \leq 3x_5 - x_1), (0 \leq x_3 - 2x_4)\},$$

and the following $\mathcal{LA}(\mathbb{Q})$ proof of unsatisfiability of $A \wedge B$:

$$\frac{\frac{\frac{0 \leq 3x_1 - 4x_2 + 1}{(0 \leq x_1 - 3x_2 + 1)} \quad (0 \leq x_2 - 2x_1)}{3 * (0 \leq x_2 - \frac{1}{3}x_3)}}{\frac{(0 \leq x_1 - x_3 + 1) \quad 2 * (0 \leq x_4 - \frac{3}{2}x_5 - 1)}{(0 \leq x_1 - x_3 + 2x_4 - 3x_5 - 1)} \quad (0 \leq 3x_5 - x_1)}{(0 \leq -x_3 + 2x_4 - 1) \quad (0 \leq x_3 - 2x_4)}.$$

$$(0 \leq -1)$$

Here, the variable x_2 is A -local, whereas all the others are AB -common. The interpolant computed with the algorithm of Section 2.3.1 is

$$(0 \leq x_1 - x_3 + 2x_4 - 3x_5 - 1),$$

which is the result of the linear combination of *all* the atoms of A in the proof. However, in order to eliminate the A -local variable x_2 , it is enough to combine $(0 \leq x_1 - 3x_2 + 1)$ (with coefficient 1) and $(0 \leq x_2 - \frac{1}{3}x_3)$ (with coefficient 3),

obtaining $(0 \leq x_1 - x_3 + 1)$. Therefore, a stronger interpolant is

$$(0 \leq x_1 - x_3 + 1) \wedge \left(0 \leq x_4 - \frac{3}{2}x_5 - 1\right).$$

Notice that in this case the interpolant is weaker than that computed with quantifier elimination (11), which is

$$\left(0 \leq -\frac{5}{4}x_1 + \frac{1}{4}\right) \wedge \left(0 \leq \frac{3}{4}x_1 - \frac{1}{3}x_3 + \frac{1}{4}\right) \wedge \left(0 \leq x_4 - \frac{3}{2}x_5 - 1\right).$$

The technique can be implemented with a small modification of the proof-based algorithm described in Section 2.3.1. We associate with each node in the proof P' (which is obtained from the original proof P by replacing inequalities from B with $(0 \leq 0)$) a list of pairs $\langle \text{coefficient}, \text{inequality} \rangle$. For a leaf, this list is a singleton in which the coefficient is 1 and the inequality is the atom in the leaf itself. For an inner node (which corresponds to an application of the COMB rule), the list l is generated from the two lists l_1 and l_2 of the premises as follows:

- (1) set l as the concatenation of l_1 and l_2 ;
- (2) let c_1 and c_2 be the coefficients used in the COMB rule; multiply each coefficient c'_i occurring in a pair $\langle c'_i, 0 \leq t_i \rangle$ of l by c_1 if the pair comes from l_1 , and by c_2 otherwise;
- (3) while there is an A -local variable x occurring in more than one pair $\langle c', 0 \leq t \rangle$ of l ¹⁰:
 - (a) collect all the pairs $\langle c'_i, 0 \leq t_i \rangle$ in which x occurs;
 - (b) generate a new pair $p \stackrel{\text{def}}{=} \langle 1, 0 \leq \sum_i c'_i * t_i \rangle$;
 - (c) add p to l , and remove all the pairs $\langle c'_i, 0 \leq t_i \rangle$.

After having applied the above algorithm, we can take the conjunction of the inequalities in the list associated with the root of P' as an interpolant.

THEOREM 3.7. *Let P be a $\mathcal{LA}(\mathbb{Q})$ proof of unsatisfiability for a conjunction $A \wedge B$ of inequalities, and P' be obtained from P by replacing each inequality of B with $(0 \leq 0)$. Let $l \stackrel{\text{def}}{=} \langle c_1, 0 \leq t_1 \rangle, \dots, \langle c_n, 0 \leq t_n \rangle$ be the list associated with the root of P' , computed as described above. Then $I \stackrel{\text{def}}{=} \bigwedge_{i=1}^n (0 \leq t_i)$ is an interpolant for (A, B) . Moreover, I is always stronger than or equal to the interpolant obtained with the algorithm of Section 2.3.1 for the same proof P' .*

PROOF. By induction on the structure of P' , it is easy to prove that, for each constraint $(0 \leq t)$ in P' with its associated list $l \stackrel{\text{def}}{=} \langle c_1, 0 \leq t_1 \rangle, \dots, \langle c_n, 0 \leq t_n \rangle$:

- (1) $A \models \bigwedge_{i=1}^n (0 \leq t_i)$; and
- (2) $(0 \leq t) \equiv \sum_{i=1}^n c_i \cdot (0 \leq t_i)$

Since the root of P' is an interpolant for (A, B) , this immediately proves the theorem. \square

¹⁰That is, x occurs in t .

4. FROM SMT(\mathcal{DL}) SOLVING TO SMT(\mathcal{DL}) INTERPOLATION

Several interesting verification problems can be encoded using only a subset of \mathcal{LA} , the theory of difference logic (\mathcal{DL}), either over the rationals ($\mathcal{DL}(\mathbb{Q})$) or over the integers ($\mathcal{DL}(\mathbb{Z})$). \mathcal{DL} is much simpler than \mathcal{LA} , since in \mathcal{DL} all atoms are inequalities of the form $(0 \leq y - x + c)$, where x and y are variables and c is an integer constant.¹¹ Equalities can be handled as conjunctions of inequalities. Here we do not consider the case when we also have strict inequalities ($0 < y - x + c$) and disequalities ($0 \neq y - x + c$), because in $\mathcal{DL}(\mathbb{Q})$ they can be handled in a way which is similar to that described in Section 3.2 for $\mathcal{LA}(\mathbb{Q})$, while in $\mathcal{DL}(\mathbb{Z})$ a strict inequality ($0 < y - x + c$) can be rewritten a priori into a weak one ($0 \leq y - x + c - 1$), and a disequality can be replaced by a disjunction of strict inequalities.

Very efficient solving algorithms have been conceived for \mathcal{DL} [Cotton and Maler 2006; Nieuwenhuis and Oliveras 2005]. In this section we present a specialized technique for computing interpolants in \mathcal{DL} which exploits such state-of-the-art decision procedures. Since a set of weak inequalities in \mathcal{DL} is consistent over the rationals if and only if it is consistent over the integers, our algorithm is applicable without any modifications to both $\mathcal{DL}(\mathbb{Q})$ and $\mathcal{DL}(\mathbb{Z})$ (see, e.g., Nieuwenhuis and Oliveras [2005]).

Many SMT solvers use dedicated, graph-based algorithms for checking the consistency of a set of $\mathcal{DL}(\mathbb{Q})$ atoms [Cotton and Maler 2006; Nieuwenhuis and Oliveras 2005]. Intuitively, a set S of $\mathcal{DL}(\mathbb{Q})$ atoms induces a graph whose vertexes are the variables of the atoms, and there exists an edge $x \xrightarrow{c} y$ for every $(0 \leq y - x + c) \in S$. S is inconsistent if and only if the induced graph has a cycle of negative weight.

We now extend the graph-based approach to generate interpolants. Consider the interpolation problem (A, B) where A and B are sets of inequalities as above, and let C be (the set of atoms in) a negative cycle in the graph corresponding to $A \cup B$.

If $C \subseteq A$, then A is inconsistent, in which case the interpolant is \perp . Similarly, when $C \subseteq B$, the interpolant is \top . If neither of these occurs, then the edges in the cycle can be partitioned into subsets of A and B . We call the maximal A path of C a path $x_1 \xrightarrow{c_1} \dots \xrightarrow{c_{n-1}} x_n$ such that (i) $x_i \xrightarrow{c_i} x_{i+1} \in A$ for $i \in [1, n-1]$, and (ii) C contains $x' \xrightarrow{c'} x_1$ and $x_n \xrightarrow{c''} x''$ that are in B . Clearly, the end-point variables x_1, x_n of the maximal A path are such that $x_1, x_n \leq A$ and $x_1, x_n \leq B$. Let the *summary constraint* of a maximal A path $x_1 \xrightarrow{c_1} \dots \xrightarrow{c_{n-1}} x_n$ be the inequality $0 \leq x_n - x_1 + \sum_{i=1}^{n-1} c_i$.

THEOREM 4.1. *The conjunction of summary constraints of the A paths of C is an interpolant for (A, B) .*

PROOF. Using the rules for $\mathcal{LA}(\mathbb{Q})$ of Figure 3, we build a deduction of the summary constraint of an maximal A path from the conjunction of its

¹¹Notice that we can assume without loss of generality that all constants are in \mathbb{Z} because, if this is not so, then we can rewrite the whole formula into an equivalently-satisfiable one by multiplying all constant symbols occurring in the formula by their greatest common denominator.

corresponding set of constraints $\bigwedge_{i=1}^{n-1} (0 \leq x_{i+1} - x_i + c_i)$:

$$\frac{\frac{(0 \leq x_2 - x_1 + c_1) \quad (0 \leq x_3 - x_2 + c_2)}{(0 \leq x_3 - x_1 + c_1 + c_2)} \quad (0 \leq x_4 - x_3 + c_3)}{\dots} \quad \dots \quad (0 \leq x_n - x_{n-1} + c_{n-1})}{(0 \leq x_n - x_1 + \sum_{i=1}^{n-1} c_i)} .$$

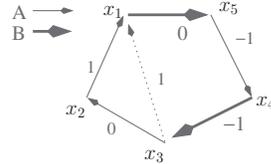
Hence, A entails the conjunction of the summary constraints of all maximal A paths. Then, we notice that the conjunction of the summary constraints is inconsistent with B . In fact, the weight of a maximal A path and the weight of its summary constraint are the same. Thus the cycle obtained from C by replacing each maximal A path with the corresponding summary constraint is also a negative cycle. Finally, we notice that every variable x occurring in the conjunction of the summary constraints is an end-point variable, and thus $x \preceq A$ and $x \preceq B$. \square

A final remark is in order. In principle, in order to generate a proof of unsatisfiability for a conjunction of $\mathcal{DL}(\mathbb{Q})$ atoms $A \wedge B$, the same rules used for $\mathcal{LA}(\mathbb{Q})$ [McMillan 2005] could be used. For instance, it is easy to build a proof which repeatedly applies the COMB rule with $c_1 = c_2 = 1$. In general, however, the interpolants generated from such proofs are not $\mathcal{DL}(\mathbb{Q})$ formulas anymore and, if computed starting from the same inconsistent set C , they are either identical or weaker than those generated with our method. In fact, it is easy to see that, unless our technique of Section 3.3 is adopted, such interpolants are in the form $(0 \leq \sum_i t_i)$ such that $\bigwedge_i (0 \leq t_i)$ is the corresponding interpolant generated with our graph-based method.

Example 4.1. Consider the following sets of $\mathcal{DL}(\mathbb{Q})$ atoms:

$$A \stackrel{\text{def}}{=} \{(0 \leq x_1 - x_2 + 1), (0 \leq x_2 - x_3), (0 \leq x_4 - x_5 - 1)\},$$

$$B \stackrel{\text{def}}{=} \{(0 \leq x_5 - x_1), (0 \leq x_3 - x_4 - 1)\},$$



corresponding to the negative cycle on the right. It is straightforward to see from the graph that the resulting interpolant is $(0 \leq x_1 - x_3 + 1) \wedge (0 \leq x_4 - x_5 - 1)$, because the first conjunct is the summary constraint of the first two conjuncts in A .

Applying instead the rules of Figure 3 with coefficients 1, the proof of unsatisfiability is

$$\frac{\frac{(0 \leq x_1 - x_2 + 1) \quad (0 \leq x_2 - x_3)}{(0 \leq x_1 - x_3 + 1)} \quad (0 \leq x_4 - x_5 - 1)}{(0 \leq x_1 - x_3 + x_4 - x_5)} \quad (0 \leq x_5 - x_1)}{(0 \leq -x_3 + x_4)} \quad (0 \leq x_3 - x_4 - 1)}{(0 \leq -1)}$$

By using the interpolation rules for $\mathcal{LA}(\mathbb{Q})$, the interpolant we obtain is $(0 \leq x_1 - x_3 + x_4 - x_5)$, which is not in $\mathcal{DL}(\mathbb{Q})$, and is weaker than that computed

above:

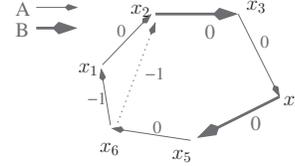
$$\frac{\frac{(0 \leq x_1 - x_2 + 1) \quad (0 \leq x_2 - x_3)}{(0 \leq x_1 - x_3 + 1)} \quad (0 \leq x_4 - x_5 - 1)}{(0 \leq x_1 - x_3 + x_4 - x_5)} \quad (0 \leq 0)}{\frac{(0 \leq x_1 - x_3 + x_4 - x_5)}{(0 \leq x_1 - x_3 + x_4 - x_5)} \quad (0 \leq 0)}{(0 \leq 0)}.$$

Finally, we observe that our technique of Section 3.3 can be easily adapted to the graph-based approach used for $\mathcal{DL}(\mathbb{Q})$. In such setting, the technique amounts to *avoiding unnecessary summarizations of A paths*: instead of building summary constraints for *maximal A* paths, we build summaries for all the *A* paths whose end-point variables x_1, x_n are shared between *A* and *B* (i.e., they are such that $x_1, x_n \leq A$ and $x_1, x_n \leq B$). Unlike with $\mathcal{LA}(\mathbb{Q})$ (see Example 3.3), in $\mathcal{DL}(\mathbb{Q})$ this procedure yields the same result as performing an existential quantification followed by quantifier elimination of all the *A*-local variables in the subset of constraints of *A* occurring in the negative cycle. (Proof sketch: each variable in a negative cycle occurs exactly in two inequalities, once positively and once negatively. Therefore, its elimination using the Fourier-Motzkin algorithm can be performed by simply taking the linear combination of such two constraints, which corresponds to a summarization step. Since we summarize only the *A* paths whose non-end-point variables are *A*-local, we do not remove any *AB*-common variable.)

Example 4.2. Consider the following sets of $\mathcal{DL}(\mathbb{Q})$ atoms:

$$A \stackrel{\text{def}}{=} \{(0 \leq x_1 - x_2), (0 \leq x_3 - x_4), (0 \leq x_5 - x_6), \\ (0 \leq x_6 - x_1 - 1)\},$$

$$B \stackrel{\text{def}}{=} \{(0 \leq x_2 - x_3), (0 \leq x_4 - x_5), (0 \leq x_4 - x_6 + 1)\},$$



$A \wedge B$ is $\mathcal{DL}(\mathbb{Q})$ -inconsistent, as demonstrated by the negative cycle on the right. By computing summaries for the maximal *A* paths in the cycle, the generated interpolant would be $(0 \leq x_5 - x_2 - 1) \wedge (0 \leq x_3 - x_4)$. However, there is no need of summarizing the whole *A* path $x_5 \xrightarrow{0} x_6 \xrightarrow{-1} x_1 \xrightarrow{0} x_2$ for eliminating the only *A*-local variable, which is x_1 . For this, it is enough to summarize the nonmaximal *A* path $x_6 \xrightarrow{-1} x_1 \xrightarrow{x_2}$, thus obtaining the stronger interpolant $(0 \leq x_5 - x_6) \wedge (0 \leq x_6 - x_2 - 1) \wedge (0 \leq x_3 - x_4)$.

5. FROM SMT(*UTVPI*) SOLVING TO SMT(*UTVPI*) INTERPOLATION

The Unit-Two-Variables-Per-Inequality (*UTVPI*) theory is a subtheory of linear arithmetic, in which all constraints are in the form $(0 \leq ax_1 + bx_2 + k)$, where k is a numerical constant, $a, b \in \{-1, 0, 1\}$, and variables x_i, x_2 range either over the rationals (for $UTVPI(\mathbb{Q})$) or over the integers (for $UTVPI(\mathbb{Z})$). Consequently, $\mathcal{DL}(\mathbb{Q})$ is a subtheory of $UTVPI(\mathbb{Q})$, which is itself a subtheory of $\mathcal{LA}(\mathbb{Q})$, and $\mathcal{DL}(\mathbb{Z})$ is a subtheory of $UTVPI(\mathbb{Z})$, which is itself a subtheory of $\mathcal{LA}(\mathbb{Z})$.

Table I. The Conversion Map from $UTVPI(\mathbb{Q})$ to $\mathcal{DL}(\mathbb{Q})$.

$UTVPI(\mathbb{Q})$ Constraints	$\mathcal{DL}(\mathbb{Q})$ Constraints
$(0 \leq x_1 - x_2 + k)$	$(0 \leq x_1^+ - x_2^+ + k), (0 \leq x_2^- - x_1^- + k)$
$(0 \leq -x_1 - x_2 + k)$	$(0 \leq x_1^- - x_2^+ + k), (0 \leq x_2^- - x_1^+ + k)$
$(0 \leq x_1 + x_2 + k)$	$(0 \leq x_1^+ - x_2^- + k), (0 \leq x_2^+ - x_1^- + k)$
$(0 \leq -x_1 + k)$	$(0 \leq x_1^- - x_1^+ + 2 \cdot k)$
$(0 \leq x_1 + k)$	$(0 \leq x_1^+ - x_1^- + 2 \cdot k)$

As for \mathcal{DL} , $UTVPI$ can be treated more efficiently than the full \mathcal{LA} , and several specialized algorithms for $UTVPI$ have been proposed in the literature. Traditional techniques are based on the iterative computation of the transitive closure of the constraints [Harvey and Stuckey 1997; Jaffar et al. 1994]; more recently [Lahiri and Musuvathi 2005] proposed a novel technique based on a reduction to \mathcal{DL} , so that graph-based techniques can be exploited, resulting into an asymptotically faster algorithm. We adopt the latter approach and show how the graph-based interpolation technique of Section 4 can be extended to $UTVPI$, for both the rationals (Section 5.1) and the integers (Section 5.2).

5.1 Graph-Based Interpolation for $UTVPI$ on the Rationals

We analyze first the simpler case of $UTVPI(\mathbb{Q})$. Miné [2001] showed that it is possible to encode a set of $UTVPI(\mathbb{Q})$ constraints into a $\mathcal{DL}(\mathbb{Q})$ one in a satisfiability-preserving way. The encoding works as follows. We use x_i to denote variables in the $UTVPI(\mathbb{Q})$ domain and u, v for variables in the $\mathcal{DL}(\mathbb{Q})$ domain. For every variable x_i in $UTVPI(\mathbb{Q})$, we introduce two distinct variables x_i^+ and x_i^- in $\mathcal{DL}(\mathbb{Q})$. We introduce a mapping Υ from $\mathcal{DL}(\mathbb{Q})$ variables to $UTVPI(\mathbb{Q})$ signed variables, such that $\Upsilon(x_i^+) = x_i$ and $\Upsilon(x_i^-) = -x_i$. Υ extends to (sets of) constraints in the natural way: $\Upsilon(0 \leq ax_1 + bx_2 + k) \stackrel{\text{def}}{=} (0 \leq a\Upsilon(x_1) + b\Upsilon(x_2) + c)$, and $\Upsilon(\{c_i\}_i) \stackrel{\text{def}}{=} \{\Upsilon(c_i)\}_i$. We say that $(x_i^+)^- = x_i^-$ and $(x_i^-)^- = x_i^+$. We say that the constraints $(0 \leq u - v)$ and $(0 \leq (v)^- - (u)^-)$ such that $u, v \in \{x_i^+, x_i^-\}_i$ are *dual*. We encode each $UTVPI$ constraint into the conjunction of two dual $\mathcal{DL}(\mathbb{Q})$ constraints, as represented in Table I. For each $\mathcal{DL}(\mathbb{Q})$ constraint $(0 \leq v - u + k)$, $(0 \leq \Upsilon(v) - \Upsilon(u) + k)$ is the corresponding $UTVPI(\mathbb{Q})$ constraint. Notice that the two dual $\mathcal{DL}(\mathbb{Q})$ constraints in the right column of Table I are just different representations of the original $UTVPI(\mathbb{Q})$ constraint. (The two dual constraints encoding a single-variable constraint are identical, so that their conjunction is collapsed into one constraint only.) The resulting set of constraints is satisfiable in $\mathcal{DL}(\mathbb{Q})$ if and only if the original one is satisfiable in $UTVPI(\mathbb{Q})$ [Miné 2001; Lahiri and Musuvathi 2005].

Consider the pair (A, B) where A and B are sets of $UTVPI(\mathbb{Q})$ constraints. We apply the map of Table I and we encode (A, B) into a $\mathcal{DL}(\mathbb{Q})$ pair (A', B') , and build the constraint graph $G(A' \wedge B')$. If $G(A' \wedge B')$ has no negative cycle, we can conclude that $A' \wedge B'$ is $\mathcal{DL}(\mathbb{Q})$ -consistent, and hence that $A \wedge B$ is $UTVPI(\mathbb{Q})$ -consistent; otherwise, $A' \wedge B'$ is $\mathcal{DL}(\mathbb{Q})$ -inconsistent, and hence $A \wedge B$ is $UTVPI(\mathbb{Q})$ -inconsistent [Miné 2001; Lahiri and Musuvathi 2005]. In fact, it is straightforward to observe that, for any set of $\mathcal{DL}(\mathbb{Q})$ constraints $\{C_1, \dots, C_n, C\}$ resulting from the encoding of some $UTVPI(\mathbb{Q})$ constraints, if $\bigwedge_{i=1}^n C_i \models_{\mathcal{DL}(\mathbb{Q})} C$ then $\bigwedge_{i=1}^n \Upsilon(C_i) \models_{UTVPI(\mathbb{Q})} \Upsilon(C)$.

When $A \wedge B$ is inconsistent, we can generate an $UTVPI(\mathbb{Q})$ interpolant by extending the graph-based approach used for $\mathcal{DL}(\mathbb{Q})$.

THEOREM 5.1. *Let $A \wedge B$ be an inconsistent conjunction of $UTVPI(\mathbb{Q})$ constraints, and let $G(A \wedge B)$ be the corresponding graph of $\mathcal{DL}(\mathbb{Q})$ constraints. Let I' be a $\mathcal{DL}(\mathbb{Q})$ interpolant built from $G(A \wedge B)$ with the technique described in Section 4. Then $I \stackrel{\text{def}}{=} \Upsilon(I')$ is an interpolant for (A, B) .*

PROOF. (i) I' is a conjunction of summary constraints, so it is in the form $\bigwedge_i C_i$. Therefore $A' \models_{\mathcal{DL}(\mathbb{Q})} C_i$ for all i , and so by the observation above $A \models_{UTVPI(\mathbb{Q})} \Upsilon(C_i)$. Hence, $A \models_{UTVPI(\mathbb{Q})} I$. (ii) From the $\mathcal{DL}(\mathbb{Q})$ inconsistency of $I' \wedge B'$ we immediately derive that $I \wedge B$ is $UTVPI(\mathbb{Q})$ -inconsistent. (iii) $I \preceq A$ and $I \preceq B$ derive from $I' \preceq A'$ and $I' \preceq B'$ by the definitions of Υ and the map of Table I. \square

As with the $\mathcal{DL}(\mathbb{Q})$ case, in principle, it is possible to generate a proof of unsatisfiability for a conjunction of $UTVPI(\mathbb{Q})$ atoms $A \wedge B$ by repeatedly applying the COMB rule for $\mathcal{LA}(\mathbb{Q})$ [McMillan 2005] with $c_1 = c_2 = 1$. As with $\mathcal{DL}(\mathbb{Q})$, however, the interpolants generated from such proofs may not be $UTVPI(\mathbb{Q})$ formulas anymore. Moreover, if computed starting from the same inconsistent set C and unless our technique of Section 3.3 is adopted, they are either identical or weaker than those generated with our graph-based method, since they are in the form $(0 \leq \sum_i t_i)$ such that $\bigwedge_i (0 \leq t_i)$ is the interpolant generated with our method.

Example 5.1. Consider the following sets of $UTVPI(\mathbb{Q})$ constraints:

$$\begin{aligned} A &= \{(0 \leq -x_2 - x_1 + 3), (0 \leq x_1 + x_3 + 1) \\ &\quad (0 \leq -x_3 - x_4 - 6), (0 \leq x_5 + x_4 + 1)\}, \\ B &= \{(0 \leq x_2 + x_3 + 3)(0 \leq x_6 - x_5 - 1), (0 \leq x_4 - x_6 + 4)\}. \end{aligned}$$

By the map of Table I, they are converted into the following sets of $\mathcal{DL}(\mathbb{Q})$ constraints:

$$\begin{aligned} A' &= \{(0 \leq x_1^- - x_2^+ + 3), (0 \leq x_2^- - x_1^+ + 3), \\ &\quad (0 \leq x_3^+ - x_1^- + 1), (0 \leq x_1^+ - x_3^- + 1), \\ &\quad (0 \leq x_4^- - x_3^+ - 6), (0 \leq x_3^- - x_4^+ - 6), \\ &\quad (0 \leq x_4^+ - x_5^- + 1), (0 \leq x_5^+ - x_4^- + 1)\}, \\ B' &= \{(0 \leq x_3^+ - x_2^- + 3), (0 \leq x_2^+ - x_3^- + 3), \\ &\quad (0 \leq x_6^+ - x_5^+ - 1), (0 \leq x_5^- - x_6^- - 1), \\ &\quad (0 \leq x_4^+ - x_6^+ + 4), (0 \leq x_6^- - x_4^- + 4)\}, \end{aligned}$$

whose conjunction corresponds to the constraint graph of Figure 4. This graph has a negative cycle

$$C' \stackrel{\text{def}}{=} x_2^+ \xrightarrow{3} x_1^- \xrightarrow{1} x_3^+ \xrightarrow{-6} x_4^- \xrightarrow{4} x_6^- \xrightarrow{-1} x_5^- \xrightarrow{1} x_4^+ \xrightarrow{-6} x_3^- \xrightarrow{3} x_2^+.$$

Thus, $A \wedge B$ is inconsistent in $UTVPI(\mathbb{Q})$. From the negative cycle C' we can extract the set of A' -paths $\{x_2^+ \xrightarrow{-2} x_4^-, x_5^- \xrightarrow{-5} x_3^-\}$, corresponding to the formula

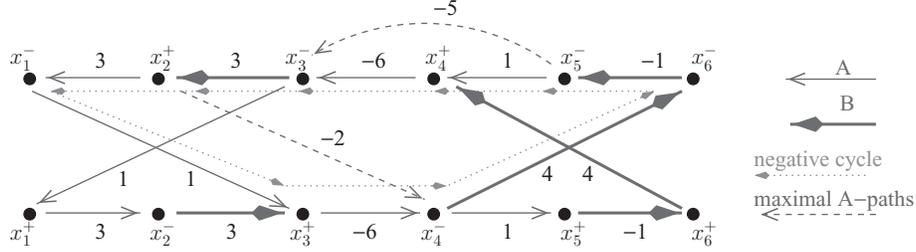


Fig. 4. The constraint graph of Example 5.1. (We represent only one negative cycle with its corresponding A-paths, because the other is dual.)

$I \stackrel{\text{def}}{=} (0 \leq x_4^- - x_2^+ - 2) \wedge (0 \leq x_3^- - x_5^- - 5)$, which is an interpolant for (A', B') . I' is thus mapped back into $I \stackrel{\text{def}}{=} \Upsilon(I) \stackrel{\text{def}}{=} (0 \leq -x_2 - x_4 - 2) \wedge (0 \leq x_5 - x_3 - 5)$, which is an interpolant for (A, B) .

Applying instead the $\mathcal{LA}(\mathbb{Q})$ interpolation technique of McMillan [2005], we find the interpolant $(0 \leq -x_2 - x_4 + x_5 - x_3 - 7)$, which is not in $UTVPI(\mathbb{Q})$ and is strictly weaker than that computed with our method.

5.2 Graph-Based Interpolation for $UTVPI$ on the Integers

In order to deal with the more complex case of $UTVPI(\mathbb{Z})$, we adopt a layered approach [Sebastiani 2007]. First, we check the consistency in $UTVPI(\mathbb{Q})$ using the technique of Miné [2001]. If this results in an inconsistency, we compute an $UTVPI(\mathbb{Q})$ interpolant as described in Section 5.1. If the $UTVPI(\mathbb{Q})$ procedure does not detect an inconsistency, we check the consistency in $UTVPI(\mathbb{Z})$ using the algorithm proposed by Lahiri and Musuvathi [2005], which extends the ideas of Miné [2001] to the integer domain. In particular, it gives necessary and sufficient conditions to decide unsatisfiability by detecting particular kinds of zero-weight cycles in the induced \mathcal{DL} constraint graph. This procedure works in $O(n \cdot m)$ time and $O(n + m)$ space, m and n being the number of constraints and variables, respectively, which improves the previous $O(n^2 \cdot m)$ time and $O(n^2)$ space complexity of the previous procedure of Jaffar et al. [1994].

We build on top of this algorithm and we extend the graph-based approach of Section 5.1 for producing interpolants also in $UTVPI(\mathbb{Z})$. In particular, we use the following reformulation of a result of Lahiri and Musuvathi [2005].

THEOREM 5.2. *Let ϕ be a conjunction of $UTVPI(\mathbb{Z})$ constraints such that ϕ is satisfiable in $UTVPI(\mathbb{Q})$. Then ϕ is unsatisfiable in $UTVPI(\mathbb{Z})$ iff the constraint graph $G(\phi)$ generated from ϕ has a cycle C of weight zero containing two vertices x_i^+ and x_i^- such that the weight of the path $x_i^- \rightsquigarrow x_i^+$ along C is odd.*

PROOF. The “only if” part is a corollary of lemmas 1, 2, and 4 in Lahiri and Musuvathi [2005]. The “if” comes straightforwardly from the analysis done in Lahiri and Musuvathi [2005], whose main intuitions we recall in what follows. Assume the constraint graph $G(\phi)$ generated from ϕ has one cycle C of weight zero containing two vertices x_i^+ and x_i^- such that the weight of the

path $x_i^- \rightsquigarrow x_i^+$ along C is $2k + 1$ for some integer value k . (Since C has weight zero, the weight of the other path $x_i^+ \rightsquigarrow x_i^-$ along C is $-2k - 1$.) Then the paths $x_i^- \rightsquigarrow x_i^+$ and $x_i^+ \rightsquigarrow x_i^-$ contain at least two constraints, because otherwise their weights would be even (see the last two lines of Table I). Then $x_i^- \rightsquigarrow x_i^+$ is in the form $x_i^- \rightsquigarrow v \xrightarrow{n} x_i^+$, for some v and n . From $x_i^- \rightsquigarrow v$, we can derive the summary constraint $(0 \leq v - x_i^- + (2k + 1 - n))$, which corresponds to the $UTVPPI(\mathbb{Z})$ constraint $(0 \leq \Upsilon(v) + x_i + (2k + 1 - n))$. (This corresponds to $l - 2$ applications of the TRANSITIVE rule of Lahiri and Musuvathi [2005], l being the number of constraints in $x_i^- \rightsquigarrow x_i^+$.) Then, by observing that the $UTVPPI(\mathbb{Z})$ constraint corresponding to $v \xrightarrow{n} x_i^+$ is $(0 \leq x_i - \Upsilon(v) + n)$, we can apply the TIGHTENING rule of Lahiri and Musuvathi [2005] to obtain $(0 \leq x_i + \lfloor (2k + 1 - n + n)/2 \rfloor)$, which is equivalent to $(0 \leq x_i + k)$. Similarly, from $x_i^+ \rightsquigarrow x_i^-$ we can obtain $(0 \leq -x_i - k - 1)$, and thus an inconsistency using the CONTRADICTION rule of Lahiri and Musuvathi [2005]. \square

Consider a pair (A, B) of $UTVPPI(\mathbb{Z})$ constraints such that $A \wedge B$ is consistent in $UTVPPI(\mathbb{Q})$ but inconsistent in $UTVPPI(\mathbb{Z})$. By Theorem 1, the constraint graph $G(A' \wedge B')$ has a cycle C of weight zero containing two vertices x_i^+ and x_i^- such that the weight of the paths $x_i^- \rightsquigarrow x_i^+$ and $x_i^+ \rightsquigarrow x_i^-$ along C are $2k + 1$ and $-2k - 1$, respectively, for some value $k \in \mathbb{Z}$. Our algorithm computes an interpolant for (A, B) from the cycle C . Let C_A and C_B be the subsets of the edges in C corresponding to constraints in A' and B' , respectively. We have to distinguish four distinct subcases.

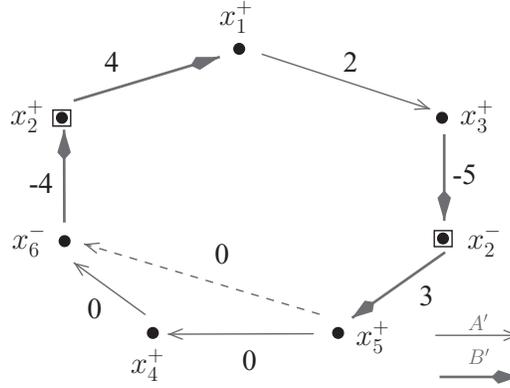
Case 1. x_i occurs in B but not in A . Consequently, x_i^+ and x_i^- occur in B' but not in A' , and hence they occur in C_B but not in C_A . Let I' be the conjunction of the summary constraints of the maximal C_A paths, and let I be the conjunction of the corresponding $UTVPPI(\mathbb{Z})$ constraints. The following theorem shows that I is an interpolant for (A, B) .

THEOREM 5.3. *Let (A, B) be a pair of sets of $UTVPPI(\mathbb{Z})$ constraints such that $A \wedge B \models_{UTVPPI(\mathbb{Z})} \perp$, let x_i be a variable that occurs in B but not in A , let $G(A' \wedge B')$ be the graph of the \mathcal{DL} encoding of $A \wedge B$, and let C be a zero-weight cycle in the graph such that the weight of the paths $x_i^- \rightsquigarrow x_i^+$ and $x_i^+ \rightsquigarrow x_i^-$ along it are odd. Let I' be the conjunction of the summary constraints of the maximal C_A paths, and let I be the conjunction of the corresponding $UTVPPI(\mathbb{Z})$ constraints. Then I is an interpolant for (A, B) .*

PROOF. (i) By construction, $A \models_{UTVPPI(\mathbb{Z})} I$, as in Section 5.1. (ii) The constraints in I' and C_B form a cycle matching the hypotheses of Theorem 5.2, from which $I \wedge B$ is $UTVPPI(\mathbb{Z})$ -inconsistent. (iii) We notice that every variable x_j^+, x_j^- occurring in the conjunction of the summary constraints is an end-point variable, so that $I' \leq C_A$ and $I' \leq C_B$, and thus $I \leq A$ and $I \leq B$. \square

Example 5.2. Consider the following set of constraints:

$$S = \{(0 \leq x_1 - x_2 + 4), (0 \leq -x_2 - x_3 - 5), (0 \leq x_2 + x_6 - 4), (0 \leq x_5 + x_2 + 3), \\ (0 \leq -x_1 + x_3 + 2), (0 \leq -x_6 - x_4), (0 \leq x_4 - x_5)\},$$


 Fig. 5. $UTVPI(\mathbb{Z})$ interpolation, Case 1.

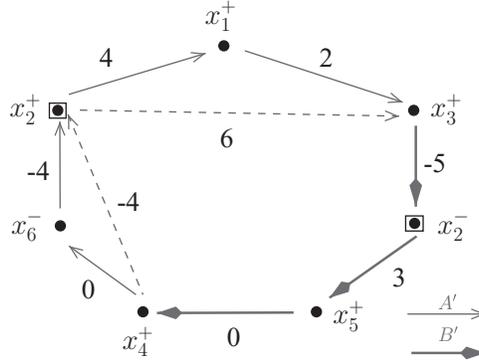
partitioned into A and B as follows:

$$A \begin{cases} (0 \leq x_3 - x_1 + 2) \\ (0 \leq -x_6 - x_4) \\ (0 \leq x_4 - x_5), \end{cases} \quad B \begin{cases} (0 \leq x_1 - x_2 + 4) \\ (0 \leq -x_2 - x_3 - 5) \\ (0 \leq x_2 + x_6 - 4) \\ (0 \leq x_5 + x_2 + 3). \end{cases}$$

Figure 5 shows a zero-weight cycle C in $G(A' \wedge B')$ such that the paths $x_2^- \rightsquigarrow x_2^+$ and $x_2^+ \rightsquigarrow x_2^-$ have an odd weight (-1 and 1 , respectively). Therefore, by Theorem 5.2 $A \wedge B$ is $UTVPI(\mathbb{Z})$ -inconsistent. The two summary constraints of the maximal C_A paths are $(0 \leq x_6^- - x_5^+)$ and $(0 \leq x_3^+ - x_1^+ + 2)$. It is easy to see that $I = (0 \leq -x_6 - x_5) \wedge (0 \leq x_3 - x_1 + 2)$ is an $UTVPI(\mathbb{Z})$ -interpolant for (A, B) .

Case 2. x_i occurs in both A and B . Consequently, x_i^+ and x_i^- occur in both A' and B' . If neither x_i^+ nor x_i^- is such that both the incoming and outgoing edges belong to C_A , then the cycle obtained by replacing each maximal C_A path with its summary constraint still contains both x_i^+ and x_i^- , so we can apply the same process of Case 1. Otherwise, if both the incoming and outgoing edges of x_i^+ belong to C_A , then we split the maximal C_A path $u_1 \xrightarrow{c_1} \dots \xrightarrow{c_k} x_i^+ \xrightarrow{c_{k+1}} \dots \xrightarrow{c_n} u_n$ containing x_i^+ into the two parts which are separated by x_i^+ : $u_1 \xrightarrow{c_1} \dots \xrightarrow{c_k} x_i^+$ and $x_i^+ \xrightarrow{c_{k+1}} \dots \xrightarrow{c_n} u_n$. We do the same for x_i^- . Let I' be the conjunction of the resulting summary constraints, and let I be corresponding set of $UTVPI(\mathbb{Z})$ constraints. The following theorem shows that I is an interpolant for (A, B) .

THEOREM 5.4. *Let (A, B) be a pair of sets of $UTVPI(\mathbb{Z})$ constraints such that $A \wedge B \models_{UTVPI(\mathbb{Z})} \perp$, let x_i be a variable that occurs in both A and B , let $G(A' \wedge B')$ be the graph of the \mathcal{DL} encoding of $A \wedge B$, and let C be a zero-weight cycle in the graph such that the weight of the paths $x_i^- \rightsquigarrow x_i^+$ and $x_i^+ \rightsquigarrow x_i^-$ along it are odd, and such that both the incoming and outgoing edges of x_i^+ in C belong to C_A . Let $u_1 \xrightarrow{c_1} \dots \xrightarrow{c_k} x_i^+ \xrightarrow{c_{k+1}} \dots \xrightarrow{c_n} u_n$ be the maximal C_A path containing x_i^+ , and let s_1 and s_2 be, respectively, the summary constraints of the two parts of the*

Fig. 6. $UTVPPI(\mathbb{Z})$ interpolation, Case 2.

above path which are separated by x_i^+ . Let I' be the conjunction of s_1 , s_2 and the summary constraints of the other maximal C_A paths, and let I be the conjunction of the corresponding $UTVPPI(\mathbb{Z})$ constraints. Then I is an interpolant for (A, B) .

PROOF. (i) As with Case 1, again, $A \models_{UTVPPI(\mathbb{Z})} I$. (ii) Since we split the maximal C_A paths as described above, the constraints in I' and C_B form a cycle matching the hypotheses of Theorem 5.2, from which $I \wedge B$ is $UTVPPI(\mathbb{Z})$ -inconsistent. (iii) x_i^+, x_i^- occur in both A' and B' by hypothesis, and every other variable x_j^+, x_j^- occurring in the conjunction of the summary constraints is an end-point variable, so that $I' \leq C_A$ and $I' \leq C_B$, and thus $I \leq A$ and $I \leq B$. \square

Example 5.3. Consider again the set of constraints S of Example 5.2, partitioned into A and B as follows:

$$A \begin{cases} (0 \leq x_3 - x_1 + 2) \\ (0 \leq -x_6 - x_4) \\ (0 \leq x_2 + x_6 - 4) \\ (0 \leq x_1 - x_2 + 4), \end{cases} \quad B \begin{cases} (0 \leq -x_2 - x_3 - 5) \\ (0 \leq x_5 + x_2 + 3) \\ (0 \leq x_4 - x_5). \end{cases}$$

and the zero-weight cycle C of $G(A' \wedge B')$ shown in Figure 6. As in the previous example, there is a path $x_2^- \rightsquigarrow x_2^+$ of weight -1 and a path $x_2^+ \rightsquigarrow x_2^-$ of weight 1 . In this case there is only one maximal C_A path, namely, $x_4^+ \rightsquigarrow x_3^+$. Since the cycle obtained by replacing it with its summary constraint ($0 \leq x_3^+ - x_4^+ + 2$) does not contain x_2^+ , we split $x_4^+ \rightsquigarrow x_3^+$ into two paths, $x_4^+ \rightsquigarrow x_2^+$ and $x_2^+ \rightsquigarrow x_3^+$, whose summary constraints are $(0 \leq x_2^+ - x_4^+ - 4)$ and $(0 \leq x_3^+ - x_2^+ + 6)$, respectively. By replacing the two paths above with the two summary constraints, we get a zero-weight cycle which still contains the two odd paths $x_2^- \rightsquigarrow x_2^+$ and $x_2^+ \rightsquigarrow x_2^-$. Therefore, $I \stackrel{\text{def}}{=} (0 \leq x_2 - x_4 - 4) \wedge (0 \leq x_3 - x_2 + 6)$ is an interpolant for (A, B) .

Notice that the $UTVPPI(\mathbb{Z})$ formula $J \stackrel{\text{def}}{=} (0 \leq x_3 - x_4 + 2)$ corresponding to the summary constraint of the maximal C_A path $x_4^+ \rightsquigarrow x_3^+$ is not an interpolant, since $J \wedge B$ is not $UTVPPI(\mathbb{Z})$ -inconsistent. In fact, if we replace the maximal C_A path $x_4^+ \rightsquigarrow x_3^+$ with the summary constraint $x_4^+ \xrightarrow{2} x_3^+$, the cycle we obtain has still weight zero, but it contains no odd path between two variables x_i^+ and x_i^- .

Case 3. x_i occurs in A but not in B , and one of the paths $x_i^+ \rightsquigarrow x_i^-$ or $x_i^- \rightsquigarrow x_i^+$ in C contains only constraints of C_A . In this case, x_i^+ and x_i^- occur in A but not in B . Suppose that $x_i^- \rightsquigarrow x_i^+$ consists only of constraints of C_A (the case $x_i^+ \rightsquigarrow x_i^-$ is analogous).

Let $2k + 1$ be the weight of the path $x_i^- \rightsquigarrow x_i^+$ (which is odd by hypothesis), and let \bar{C} be the cycle obtained by replacing such path with the edge $x_i^- \xrightarrow{2k} x_i^+$ in C . In the following, we call such a replacement *tightening summarization*. Since C has weight zero, \bar{C} has negative weight. Let C^P be the set of \mathcal{DL} constraints in the path $x_i^- \rightsquigarrow x_i^+$. Let I' be the \mathcal{DL} interpolant computed from \bar{C} for $(C_A \setminus C^P \cup \{(0 \leq x_i^+ - x_i^- + 2k)\}, C_B)$, and let I be the corresponding $UTVPI(\mathbb{Z})$ formula. The following theorem shows that I is an interpolant for (A, B) .

THEOREM 5.5. *Let (A, B) be a pair of sets of $UTVPI(\mathbb{Z})$ constraints such that $A \wedge B \models_{UTVPI(\mathbb{Z})} \perp$, let x_i be a variable that occurs in A but not in B , let $G(A \wedge B)$ be the graph of the \mathcal{DL} encoding of $A \wedge B$, and let C be a zero-weight cycle in the graph such that the weight of the paths $x_i^- \rightsquigarrow x_i^+$ and $x_i^+ \rightsquigarrow x_i^-$ along it are odd, and such that all the constraints in the path $x_i^- \rightsquigarrow x_i^+$ belong to C_A . Let \bar{C} be the cycle obtained by replacing the path $x_i^- \rightsquigarrow x_i^+$ with the edge $x_i^- \xrightarrow{2k} x_i^+$ in C , where $2k + 1$ is the weight of the replaced path. Let C^P be the set of \mathcal{DL} constraints in the path $x_i^- \rightsquigarrow x_i^+$. Let I' be the \mathcal{DL} interpolant computed from \bar{C} for $(C_A \setminus C^P \cup \{(0 \leq x_i^+ - x_i^- + 2k)\}, C_B)$, and let I be the corresponding $UTVPI(\mathbb{Z})$ formula. Then I is an interpolant for (A, B) .*

PROOF. (i) Let P be the set of $UTVPI(\mathbb{Z})$ constraints in the path $x_i^- \rightsquigarrow x_i^+$. Since the weight $2k + 1$ of such path is odd, we have that $P \models_{UTVPI(\mathbb{Z})} (0 \leq x_i + k)$ (cf. the proof of Theorem 5.2). Since $P \subseteq A$, therefore, $A \models_{UTVPI(\mathbb{Z})} (0 \leq x_i + k)$. By observing that $(0 \leq x_i^+ - x_i^- + 2k)$ is the \mathcal{DL} constraint corresponding to $(0 \leq x_i + k)$ we conclude that $C_A \setminus C^P \cup (0 \leq x_i^+ - x_i^- + 2k) \models_{\mathcal{DL}} I'$ implies that $A \setminus P \cup (0 \leq x_i + k) \models_{UTVPI(\mathbb{Z})} I$, and so that $A \models_{UTVPI(\mathbb{Z})} I$.

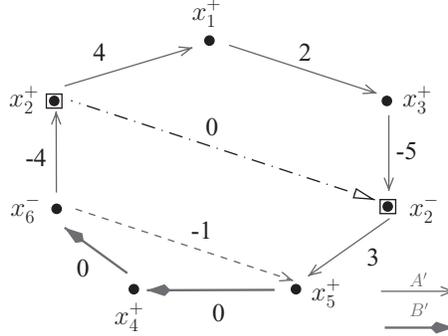
(ii) Since all the constraints in C_B occur in \bar{C} , we have that $B \wedge I$ is $UTVPI(\mathbb{Z})$ -inconsistent.

(iii) Since by hypothesis all the constraints in the path $x_i^- \rightsquigarrow x_i^+$ occur in C_A , from $I' \leq (C_A \setminus C^P \cup \{(0 \leq x_i^+ - x_i^- + 2k)\})$ we have that $I \leq A$. Finally, since all the constraints in C_B occur in \bar{C} , we have that $I \leq B$. \square

Example 5.4. Consider again the set S of constraints of Example 5.2, this time partitioned into A and B as follows:

$$A \begin{cases} (0 \leq x_1 - x_2 + 4), \\ (0 \leq x_3 - x_1 + 2), \\ (0 \leq -x_2 - x_3 - 5), \\ (0 \leq x_2 + x_6 - 4), \\ (0 \leq x_5 + x_2 + 3), \end{cases} \quad B \begin{cases} (0 \leq -x_6 - x_4), \\ (0 \leq x_4 - x_5). \end{cases}$$

Figure 7 shows a zero-weight cycle C of $G(A \wedge B)$. The only maximal C_A path is $x_6^- \rightsquigarrow x_5^+$. Since the path $x_2^+ \rightsquigarrow x_2^-$ has weight 1, we can add the tightening edge

Fig. 7. $UTVPPI(\mathbb{Z})$ interpolation, Case 3.

$x_2^+ \xrightarrow{1-1} x_2^-$ to $G(A' \wedge B')$ (shown in dots and dashes in Figure 7), corresponding to the constraint $(0 \leq x_2^- - x_2^+)$. Since all constraints in the path $x_2^+ \rightsquigarrow x_2^-$ belong to A' , $A' \models (0 \leq x_2^- - x_2^+)$. Moreover, the cycle obtained by replacing the path $x_2^+ \rightsquigarrow x_2^-$ with the tightening edge $x_2^+ \xrightarrow{0} x_2^-$ has a negative weight (-1). Therefore, we can generate a \mathcal{DL} interpolant $I' \stackrel{\text{def}}{=} (0 \leq x_2^- - x_6^- - 4)$ from such cycle, which corresponds to the $UTVPPI(\mathbb{Z})$ interpolant $I \stackrel{\text{def}}{=} (0 \leq -x_2 + x_6 - 4)$.

Notice that, similarly to Example 5.3, also in this case we cannot obtain an interpolant from the summary constraint $(0 \leq x_5^+ - x_6^-)$ of the maximal C_A path $x_6^- \rightsquigarrow x_5^+$, as $(0 \leq x_5 + x_6) \wedge B$ is not $UTVPPI(\mathbb{Z})$ -inconsistent.

Case 4. x_i occurs in A but not in B , and neither the path $x_i^+ \rightsquigarrow x_i^-$ nor the path $x_i^- \rightsquigarrow x_i^+$ in C consists only of constraints of C_A . As in the previous case, x_i^+ and x_i^- occur in A' but not in B' , and hence they occur in C_A but not in C_B . In this case, however, we can apply a tightening summarization neither to $x_i^+ \rightsquigarrow x_i^-$ nor to $x_i^- \rightsquigarrow x_i^+$, since none of the two paths consists only of constraints of C_A . We can, however, perform a *conditional tightening summarization* as follows. Let C_A^P and C_B^P be the sets of constraints of C_A and C_B , respectively, occurring in the path $x_i^- \rightsquigarrow x_i^+$, and let \overline{C}_A^P and \overline{C}_B^P be the sets of summary constraints of maximal paths in C_A^P and C_B^P . From $\overline{C}_A^P \cup \overline{C}_B^P$, we can derive $x_i^- \xrightarrow{2k} x_i^+$ (cf. Case 3), where $2k+1$ is the weight of the path $x_i^- \rightsquigarrow x_i^+$. Therefore, $\overline{C}_A^P \cup \overline{C}_B^P \models (0 \leq x_i^+ - x_i^- + 2k)$, and thus $\overline{C}_A^P \models \overline{C}_B^P \rightarrow (0 \leq x_i^+ - x_i^- + 2k)$. We say that $(0 \leq x_i^+ - x_i^- + 2k)$ is the summary constraint for $x_i^- \rightsquigarrow x_i^+$ *conditioned to* \overline{C}_B^P .

Using conditional tightening summarization, we generate an interpolant as follows. By replacing the path $x_i^- \rightsquigarrow x_i^+$ with $x_i^- \xrightarrow{2k} x_i^+$, we obtain a negative-weight cycle \overline{C} , as in Case 3. Let I' be the \mathcal{DL} interpolant computed from \overline{C} for $(C_A \setminus C_A^P \cup \{(0 \leq x_i^+ - x_i^- + 2k)\}, C_B \setminus C_B^P)$, and let I be the corresponding $UTVPPI(\mathbb{Z})$ formula. Finally, let \overline{P}_B be the conjunction of $UTVPPI(\mathbb{Z})$ constraints corresponding to \overline{C}_B^P . The following theorem shows that $(\overline{P}_B \rightarrow I)$ is an interpolant for (A, B) .

THEOREM 5.6. *Let (A, B) be a pair of sets of $UTVPPI(\mathbb{Z})$ constraints such that $A \wedge B \models_{UTVPPI(\mathbb{Z})} \perp$, let x_i be a variable that occurs in A but not in B , let $G(A' \wedge B')$*

Consider the zero-weight cycle C of $G(A \wedge B)$ shown in Figure 8. In this case, neither the path $x_2^+ \rightsquigarrow x_2^-$ nor the path $x_2^- \rightsquigarrow x_2^+$ consists only of constraints of A , and thus we cannot use any of the two tightening edges $x_2^+ \xrightarrow{1-1} x_2^-$ and $x_2^- \xrightarrow{-1-1} x_2^+$ directly for computing an interpolant. However, we can compute the summary $x_2^- \xrightarrow{-2} x_2^+$ for $x_2^- \rightsquigarrow x_2^+$ conditioned to $x_5^+ \xrightarrow{0} x_6^-$, which is the summary constraint of the B path $x_5^+ \rightsquigarrow x_6^-$, and whose corresponding $UTVPI(\mathbb{Z})$ constraint is $(0 \leq -x_6 - x_5)$. By replacing the path $x_2^- \rightsquigarrow x_2^+$ with such summary, we obtain a negative-weight cycle \bar{C} , from which we generate the \mathcal{DL} interpolant $(0 \leq x_1^+ - x_3^+ - 3)$, corresponding to the $UTVPI(\mathbb{Z})$ formula $(0 \leq x_1 - x_3 - 3)$. Therefore, the generated $UTVPI(\mathbb{Z})$ interpolant is $(0 \leq -x_6 - x_5) \rightarrow (0 \leq x_1 - x_3 - 3)$.

As in Example 5.4, notice that we cannot generate an interpolant from the conjunction of summary constraints of maximal C_A paths, since the formula we obtain (i.e., $(0 \leq x_1 + x_6) \wedge (0 \leq x_5 - x_3 - 2)$) is not inconsistent with B .

6. COMPUTING INTERPOLANTS FOR COMBINED THEORIES VIA DTC

In this Section, we consider the problem of generating interpolants for a pair of $\mathcal{T}_1 \cup \mathcal{T}_2$ formulas (A, B) , and propose a method based on the Delayed Theory Combination (DTC) approach [Bozzano et al. 2006]. First, in Section 6.1 we provide some background on Nelson-Oppen (NO) and DTC combination methods, and recall from Yorsh and Musuvathi [2005] the basics of interpolation for combined theories using NO; then, we present our novel technique for computing interpolants using DTC (Section 6.2); in Section 6.3 we discuss the advantages of the novel method; finally, in Section 6.4, we show how our novel technique can be used to generate multiple interpolants from the same proof.

6.1 Background

6.1.1 Resolution Proofs with NO Versus Resolution Proofs with DTC. One of the typical approaches to the SMT problem in combined theories, $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, is that of combining the solvers for \mathcal{T}_1 and for \mathcal{T}_2 with the Nelson-Oppen (NO) integration schema [Nelson and Oppen 1979]. The NO framework works for combinations of stably-infinite, signature-disjoint theories \mathcal{T}_i with equality. Moreover, it requires the input formula to be *pure* (i.e., such that all the atoms contain only symbols in one theory): if not, a *purification* step is performed, by recursively labeling terms t with fresh variables v_t , and by conjoining the definition atom $(v_t = t)$ to the formula. This process is linear in the size of the input formula.¹² For instance, the formula $(f(x + 3y) = g(2x - y))$ can be purified into $(f(v_{x+3y}) = g(v_{2x-y})) \wedge (v_{x+3y} = x + 3y) \wedge (v_{2x-y} = 2x - y)$.

In the NO setting, the two decision procedures for \mathcal{T}_1 and \mathcal{T}_2 cooperate by deducing and exchanging *interface equalities*,¹³ that is, equalities between variables appearing in atoms of different theories (*interface variables*).

¹²As shown in Barrett et al. [2002], the purification step is not strictly necessary. However, in the rest we shall assume that it is performed (as it is traditionally done in articles on combination of theories), since it makes the exposition easier.

¹³They deduce and exchange *disjunctions* of interface equalities if the theory is not convex.

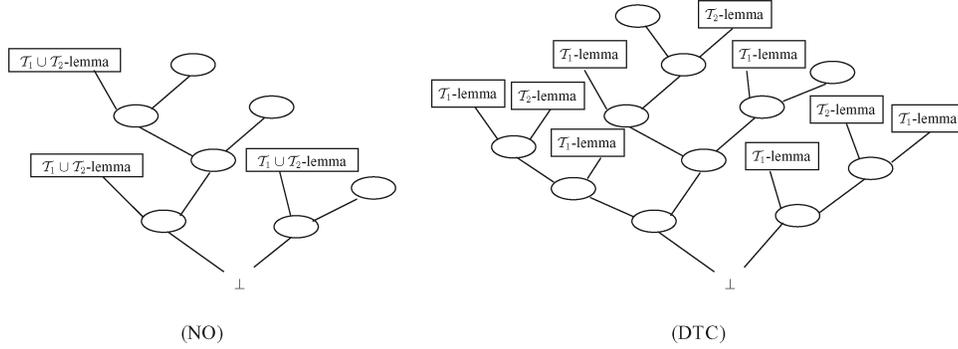


Fig. 9. Different structures of resolution proofs of unsatisfiability for $\mathcal{T}_1 \cup \mathcal{T}_2$ formulas, using NO (left) and DTC (right).

With an NO-based SMT solver, resolution proofs for formulas in a combination $\mathcal{T}_1 \cup \mathcal{T}_2$ of theories have the same structure as those for formulas in a single theory \mathcal{T} . The only difference is that theory lemmas in this case are the result of the NO-combination of \mathcal{T}_1 and \mathcal{T}_2 (i.e., they are $\mathcal{T}_1 \cup \mathcal{T}_2$ lemmas) (Figure 9, left). From the point of view of interpolation, the difference with respect to the case of a single theory \mathcal{T} is that the $\mathcal{T}_1 \cup \mathcal{T}_2$ interpolants for the negations of the $\mathcal{T}_1 \cup \mathcal{T}_2$ lemmas can be computed with the combination method of Yorsh and Musuvathi [2005] whenever it applies (see Section 6.1.2).

Recently, an alternative approach for combining theories in SMT has been proposed, called *Delayed Theory Combination* (DTC) [Bozzano et al. 2006]. With DTC, the solvers for \mathcal{T}_1 and \mathcal{T}_2 do not communicate directly. The integration is performed by the SAT solver, by augmenting the Boolean search space with up to all the possible interface equalities, so that each truth assignment on both original atoms and interface equalities is checked for consistency independently on both theories. DTC has several advantages with respect to NO, in terms of versatility, efficiency, and restrictions imposed to \mathcal{T} solvers [Bozzano et al. 2006; Bruttomesso et al. 2009], so that many current SMT tools implement variants and evolutions of DTC.

With DTC, resolution proofs are quite different from those obtained with NO. There is no $\mathcal{T}_1 \cup \mathcal{T}_2$ lemma anymore, because the two \mathcal{T}_i solvers don't communicate directly. Instead, the proofs contain both \mathcal{T}_1 lemmas and \mathcal{T}_2 lemmas (Figure 9, right), and—importantly—they contain also interface equalities. (Notice that \mathcal{T}_i lemmas derive either from \mathcal{T}_i conflicts and from \mathcal{T}_i propagation steps.) In this case, the combination of theories is encoded directly in the proofs (thanks to the presence of interface equalities), and not “hidden” in the $\mathcal{T}_1 \cup \mathcal{T}_2$ lemmas as with NO. This observation is at the heart of our DTC-based interpolant combination method.

Example 6.1. Consider the following formula ϕ :

$$\begin{aligned} \phi \stackrel{\text{def}}{=} & (a_1 = f(a_2)) \wedge (b_1 = f(b_2)) \wedge \\ & (y - a_2 = 1) \wedge (y - b_2 = 1) \wedge (a_1 + y = 0) \wedge (b_1 + y = 1). \end{aligned}$$

ϕ is expressed over the combined theory $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$: the first two atoms belong to \mathcal{EUF} , while the last four belong to $\mathcal{LA}(\mathbb{Q})$.

Using the NO combination method, ϕ can be proved unsatisfiable as follows:

- (1) from the conjunction $(y - a_2 = 1) \wedge (y - b_2 = 1)$, the $\mathcal{LA}(\mathbb{Q})$ solver deduces the interface equality $(a_2 = b_2)$, which is sent to the \mathcal{EUF} solver;
- (2) from $(a_2 = b_2)$ and the conjunction $(a_1 = f(a_2)) \wedge (b_1 = f(b_2))$ the \mathcal{EUF} solver deduces the interface equality $(a_1 = b_1)$, which is sent to the $\mathcal{LA}(\mathbb{Q})$ solver;
- (3) Together with the conjunction $(a_1 + y = 0) \wedge (b_1 + y = 1)$, $(a_1 = b_1)$ causes an inconsistency in the $\mathcal{LA}(\mathbb{Q})$ solver;
- (4) The $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ conflict set generated is $\{(y - a_2 = 1), (y - b_2 = 1), (a_1 = f(a_2)), (b_1 = f(b_2)), (a_1 + y = 0), (b_1 + y = 1)\}$, corresponding to the $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ Lemma $C \stackrel{\text{def}}{=} \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1) \vee \neg(a_1 = f(a_2)) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1)$.

The corresponding NO proof of unsatisfiability for ϕ is thus

$$\frac{\frac{\frac{\frac{\frac{\frac{C \quad (b_1 + y = 1)}{\dots} \quad (a_1 + y = 0)}{\dots} \quad (y - b_2 = 1)}{\dots} \quad (y - a_2 = 1)}{\dots} \quad (b_1 = f(b_2))}{\dots} \quad (a_1 = f(a_2))}{\perp}}{.}}$$

With DTC, the Boolean search space is augmented with the set of all possible interface equalities $Eg \stackrel{\text{def}}{=} \{(a_1 = a_2), (a_1 = b_1), (a_1 = b_2), (a_2 = b_1), (a_2 = b_2), (b_1 = b_2)\}$, so that the DPLL engine can branch on them. If we suppose that the negative branch is explored first (and we assume for simplicity that the \mathcal{T} solvers do not perform deductions), using the DTC combination method ϕ can be proved unsatisfiable as follows:

- (1) assigning $(a_2 = b_2)$ to false causes an inconsistency in the $\mathcal{LA}(\mathbb{Q})$ solver, which generates the $\mathcal{LA}(\mathbb{Q})$ lemma $C_1 \stackrel{\text{def}}{=} \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1) \vee (a_2 = b_2)$. C_1 is used by the DPLL engine to backjump and unit-propagate $(a_2 = b_2)$;
- (2) after such propagation, assigning $(a_1 = b_1)$ to false causes an inconsistency in the \mathcal{EUF} solver, which generates the \mathcal{EUF} lemma $C_2 \stackrel{\text{def}}{=} \neg(a_1 = f(a_2)) \vee \neg(b_1 = f(b_2)) \vee \neg(a_2 = b_2) \vee (a_1 = b_1)$. C_2 is used by the DPLL engine to backjump and unit-propagate $(a_1 = b_1)$;
- (3) this propagation causes an inconsistency in the $\mathcal{LA}(\mathbb{Q})$ solver, which generates the $\mathcal{LA}(\mathbb{Q})$ lemma $C_3 \stackrel{\text{def}}{=} \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1) \vee \neg(a_1 = b_1)$;
- (4) after learning C_3 , the DPLL engine detects the unsatisfiability of ϕ .

The corresponding DTC proof of unsatisfiability for ϕ is thus

$$\begin{array}{c}
 C_1 \quad (y - a_2 = 1) \\
 \hline
 \dots \quad (y - b_2 = 1) \\
 \hline
 \dots \quad C_2 \\
 \hline
 \dots \quad (b_1 = f(b_2)) \\
 \hline
 \dots \quad C_3 \\
 \hline
 \dots \quad (b_1 + y = 1) \\
 \hline
 \dots \quad (a_1 + y = 0) \\
 \hline
 \dots \quad (a_1 = f(a_2)) \\
 \hline
 \perp
 \end{array}$$

An important remark is in order. It is relatively easy to implement DTC in such a way that, if both \mathcal{T}_1 and \mathcal{T}_2 are convex, then all \mathcal{T} lemmas generated contain at most one positive interface equality. This is due to the fact that for convex theories \mathcal{T} it is possible to implement efficient \mathcal{T} solvers which generates conflict sets containing at most one negated equality between variables [Bozzano et al. 2005]¹⁴ (e.g., this is true for all the \mathcal{T}_i solvers on convex theories implemented in MATHSAT.) Thus, since we restrict to convex theories, in the rest of this article we can assume without loss of generality that every \mathcal{T} lemma occurring as a leaf in a resolution proof Π of unsatisfiability deriving from DTC contains at most one positive interface equality.

6.1.2 Interpolation with Nelson-Oppen. Yorsh and Musuvathi [2005] gave a method for generating an interpolant for a pair (A, B) of $\mathcal{T}_1 \cup \mathcal{T}_2$ formulas such that $A \wedge B \models_{\mathcal{T}_1 \cup \mathcal{T}_2} \perp$ by means of the NO schema. As in Yorsh and Musuvathi [2005], we assume that A and B have been purified using disjoint sets of auxiliary variables. We recall from Yorsh and Musuvathi [2005] a couple of definitions.

Definition 6.1 (AB-Mixed Equality). An equality between variables $(a = b)$ is an *AB-mixed equality* iff $a \not\leq B$ and $b \not\leq A$ (or vice versa).

Definition 6.2 (Equality-Interpolating Theory). A theory \mathcal{T} is said to be *equality-interpolating* iff, for all A and B in \mathcal{T} such that $A \wedge B \models_{\mathcal{T}} (a = b)$ and for all *AB mixed equalities* $(a = b)$, there exists a term t such that $A \wedge B \models_{\mathcal{T}} (a = t) \wedge (t = b)$ and $t \leq A$ and $t \leq B$.

Yorsh and Musuvathi [2005] described procedures for computing the term t from an *AB-mixed interface equality* $(a = b)$ for some convex theories of interest, including \mathcal{EUF} , $\mathcal{LA}(\mathbb{Q})$, and the theory of lists.

Notationally, with the letters x, x_i, y, y_i, z we denote generic variables, while with the letters a, a_i , and b, b_i we denote variables such that $a_i \not\leq B$ and $b_i \not\leq A$; hence, with the letters e_i we denote generic *AB-mixed interface equalities* in the form $(a_i = b_i)$; with the letters η, η_i we denote conjunctions of literals where no *AB-mixed interface equality* occurs, and with the letters μ, μ_i we denote conjunctions of literals where *AB-mixed interface equalities* may occur. If μ_i

¹⁴We recall that, if \mathcal{T} is convex, then $\mu \wedge \bigwedge_i \neg l_i \models_{\mathcal{T}} \perp$ iff $\mu \wedge \neg l_i \models_{\mathcal{T}} \perp$ for some i , where the l_i 's are positive literals.

(respectively, η_i) is $\bigwedge_i l_i$, we write $\neg\mu_i$ (respectively, $\neg\eta_i$) for the clause $\bigvee_i \neg l_i$.

Let $A \wedge B$ be a $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistent conjunction of $\mathcal{T}_1 \cup \mathcal{T}_2$ literals, such that $A \stackrel{\text{def}}{=} A_1 \wedge A_2$ and $B \stackrel{\text{def}}{=} B_1 \wedge B_2$ where each A_i and B_i is \mathcal{T}_i pure. The NO-based method of Yorsh and Musuvathi [2005] computes an interpolant for (A, B) by combining \mathcal{T}_i specific interpolants for subsets of A, B and the set of entailed interface equalities $\{e_j\}_j$ that are exchanged between the \mathcal{T}_i solvers for deciding the unsatisfiability of $A \wedge B$. In particular, let $Eq \stackrel{\text{def}}{=} \{e_j\}_j$ be the set of entailed interface equalities. Due to the fact that both \mathcal{T}_1 and \mathcal{T}_2 are equality-interpolating, it is possible to assume without loss of generality that Eq does not contain AB mixed equalities, because, instead of deducing an AB mixed interface equality $(a = b)$, a \mathcal{T} solver can always deduce the two corresponding equalities $(a = t) \wedge (t = b)$. (Notice that the other \mathcal{T} solver treats the term t as if it were a variable [Yorsh and Musuvathi 2005].) Let $A' \stackrel{\text{def}}{=} A \cup (Eq \downarrow A)$ and $B' \stackrel{\text{def}}{=} B \cup (Eq \downarrow B)$. Then, \mathcal{T}_i -specific partial interpolants are combined according to the following inductive definition:

$$I_{A,B}(e) \stackrel{\text{def}}{=} \begin{cases} \perp, & \text{if } e \in A, \\ \top, & \text{if } e \in B, \\ (I_{A',B'}^i(e) \vee \bigvee_{e_a \in A'} I_{A,B}(e_a)) \wedge \bigwedge_{e_b \in B'} I_{A,B}(e_b), & \text{otherwise,} \end{cases} \quad (12)$$

where e is either an entailed interface equality or \perp , and $I_{A',B'}^i(e)$ is a \mathcal{T}_i interpolant for $(A' \cup \neg e, B')$ if $e \leq A$, and for $(A', B' \cup \neg e)$ otherwise (if $e \leq B$). The computed interpolant for (A, B) is then $I_{A,B}(\perp)$. We refer the reader to Yorsh and Musuvathi [2005] for more details.

6.2 From DTC Solving to DTC Interpolation

We now discuss how to extend the DTC method to interpolation. As with Yorsh and Musuvathi [2005], we can handle the case that \mathcal{T}_1 and \mathcal{T}_2 are convex and equality-interpolating. The approach to generating interpolants for combined theories starts from the proof generated by DTC. Let Eq be the set of all interface equalities occurring in a DTC refutation proof for a $\mathcal{T}_1 \cup \mathcal{T}_2$ -unsatisfiable formula $\phi \stackrel{\text{def}}{=} A \wedge B$.

In the case Eq does not contain AB -mixed equalities, that is, Eq can be partitioned into two sets $(Eq \setminus B) \stackrel{\text{def}}{=} \{(x = y) \mid (x = y) \leq A \text{ and } (x = y) \not\leq B\}$ and $(Eq \downarrow B) \stackrel{\text{def}}{=} \{(x = y) \mid (x = y) \leq B\}$, *no interpolant-combination method is needed*: the combination is already encoded in the proof of unsatisfiability, and a direct application of Algorithm 1 to such proof yields an interpolant for the combined theory $\mathcal{T}_1 \cup \mathcal{T}_2$. Notice that this fact holds despite the fact that the interface equalities in Eq occur neither in A nor in B , but might be introduced in the resolution proof Π by \mathcal{T} lemmas. In fact, as observed in McMillan [2005], as long as for an atom p either $p \leq A$ or $p \leq B$ holds, it is possible to consider it part of A (respectively of B) simply by assuming the tautology clause $p \vee \neg p$ to be part of A (respectively of B). Therefore, we can treat the interface equalities in $(Eq \setminus B)$ as if they appeared in A , and those in $(Eq \downarrow B)$ as if they appeared in B .

When Eq contains AB -mixed equalities, instead, a proof-rewriting step is performed in order to obtain a proof that is free from AB -mixed equalities and

that is amenable for interpolation as described above. The idea is similar to that used in Yorsh and Musuvathi [2005] in the case of NO: using the fact that \mathcal{T}_1 and \mathcal{T}_2 are equality-interpolating, we reduce this case to the previous one by “splitting” every AB -mixed interface equality $(a_i = b_i)$ into the conjunction of two parts $(a_i = t_i) \wedge (t_i = b_i)$, such that $(a_i = t_i) \leq A$ and $(t_i = b_i) \leq B$. The main difference is that we do this a posteriori, after the construction of the resolution proof of unsatisfiability Π . In order to do this, we traverse Π and split each AB -mixed equality, performing also the necessary manipulations to ensure that the result is still a resolution proof of unsatisfiability.

We describe this process in two steps. In Section 6.2.1 we introduce a particular kind of resolution proofs of unsatisfiability, called *ie-local*, and show how to eliminate AB -mixed interface equalities from *ie-local* proofs; in Section 6.2.2 we show how to implement a variant of DTC so that to generate *ie-local* proofs.

6.2.1 Eliminating AB -Mixed Equalities by Exploiting *ie-Locality*.

Definition 6.3 (ie Local Proof). A resolution proof of unsatisfiability Π is local with respect to interface equalities (*ie-local*) iff the interface equalities occur only in subproofs Π_i^{ie} of Π , such that within each Π_i^{ie} :

- (i) all leaves are also \mathcal{T} lemma leaves of Π ;
- (ii) all the pivots are interface equalities;
- (iii) the root contains no interface equality;
- (iv) every right premise of an inner node is a leaf \mathcal{T} lemma containing exactly one positive interface equality.¹⁵

As a consequence of this definition, we also have that, within each Π_i^{ie} in Π :

- (v) all nodes are $\mathcal{T}_1 \cup \mathcal{T}_2$ -valid; (*Proof sketch:* they result from Boolean resolution steps from \mathcal{T}_1 -valid and \mathcal{T}_2 -valid clauses, and hence they are $\mathcal{T}_1 \cup \mathcal{T}_2$ -valid.)
- (vi) the only leaf \mathcal{T} lemma which is a left premise contains no positive interface equality. (*Proof sketch:* we notice that, in a resolution step $\frac{C_1 C_2}{C_3}$, if C_3 contains no positive interface equality, at least one between C_1 and C_2 contains no positive interface equality; since by (iv) the right premise contains one positive interface equality, only the left premise contains no positive interface equality. Thus the leftmost leaf \mathcal{T} -lemma of Π_i^{ie} contains no positive interface equality.)
- (vii) if an interface equality e_j occurs negatively in some \mathcal{T} lemma C_j , then e_j occurs positively in a leaf \mathcal{T} lemma C_k which is the right premise of a resolution step whose left premise derives from C_j and other \mathcal{T} lemmas. (*Proof sketch:* suppose that $\neg e_j$ occurs in C_j but e_j does not occur in any such C_k . Then e_j can not be a pivot; hence $\neg e_j$ occurs in the root of Π_i^{ie} , thus violating (iii).)

¹⁵We have adopted the graphical convention that at each resolution step in a Π_i^{ie} subproof, if $(a_i = b_i)$ is the pivot, then the premises containing $\neg(a_i = b_i)$ and $(a_i = b_i)$ are the left and right premises, respectively.

Intuitively, in ie-local proofs of unsatisfiability all the reasoning on interface equalities is circumscribed within Π_i^{ie} subproofs, which are linear subproofs involving only \mathcal{T} lemmas as leaves, starting from the one containing no positive interface equality, each time eliminating one negative interface equality by resolving it against the only positive one occurring in another leaf \mathcal{T} lemma.

Example 6.2. Consider the $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ formula ϕ of Example 6.1, and the \mathcal{T} lemmas C_1, C_2 , and C_3 introduced by DTC to prove its unsatisfiability. The proof Π of Example 6.1 is not ie-local, because resolution steps involving interface equalities are interleaved with resolution steps involving other atoms. The following proof Π' , instead, is ie-local: all the interface equalities are used as pivots in the Π^{ie} subproof:

$$\frac{\boxed{\begin{array}{c} C_3 \quad C_2 \\ \dots \\ \dots \end{array}} \text{ [pivot } (a_1 = b_1)] \quad C_1 \text{ [pivot } (a_2 = b_2)]}{\dots} \Pi^{\text{ie}} \quad \frac{\dots}{(a_2 + z = 1)} \quad \frac{\dots}{(a_1 + z = 0)} \quad \frac{\dots}{(z - x_2 = 1)} \quad \frac{\dots}{(a_1 = f(x_1))} \quad \frac{\dots}{(a_2 = f(x_2))} \quad \frac{\dots}{(z - x_1 = 1)} \quad \perp$$

$$C_1 \stackrel{\text{def}}{=} (a_2 = b_2) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1).$$

$$C_2 \stackrel{\text{def}}{=} (a_1 = b_1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(a_2 = b_2),$$

$$C_3 \stackrel{\text{def}}{=} \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(a_1 = b_1).$$

If Π is an ie-local proof containing AB -mixed interface equalities, then it is possible to eliminate all of them from Π by applying Algorithm 2 to every Π_i^{ie} subproof of Π . In a nutshell, each Π_i^{ie} subproof is explored bottom-up, starting from the right premise of the root, each time expanding the rightmost side \mathcal{T} lemma in the form $C_i \stackrel{\text{def}}{=} (a_i = b_i) \vee \neg\eta_i$ such that $(a_i = b_i)$ is AB -mixed into the (implicit) conjunction of two novel \mathcal{T} lemmas $C'_i \stackrel{\text{def}}{=} (a_i = t_i) \vee \neg\eta_i$ and $C''_i \stackrel{\text{def}}{=} (t_i = b_i) \vee \neg\eta_i$ (step (4)), where t_i is the AB -pure term computed from C_i as described in Section 6.1.2. Then the resolution step against C_i is substituted with the concatenation of two resolution steps against C'_i and C''_i (step (5)) and then the substitution $\neg(a_i = b_i) \mapsto \neg(a_i = t_i) \vee \neg(t_i = b_i)$ is propagated bottom-up along the left subproof Π . Notice that C'_i and C''_i are still \mathcal{T}_i -valid because \mathcal{T}_i is equality-interpolating and η_i does not contain other AB -mixed interfaced equalities.

Example 6.3. Consider the formula ϕ of Example 6.1 and its ie-local proof of unsatisfiability of Example 6.2. Suppose that ϕ is partitioned as follows:

$$\phi \stackrel{\text{def}}{=} A \wedge B,$$

$$A \stackrel{\text{def}}{=} (a_1 = f(a_2)) \wedge (y - a_2 = 1) \wedge (a_1 + y = 0),$$

$$B \stackrel{\text{def}}{=} (b_1 = f(b_2)) \wedge (y - b_2 = 1) \wedge (b_1 + y = 1).$$

Algorithm 2. Rewriting of Π^{ie} subproofs

- (1) Let σ be a mapping from negative *AB*-mixed interface equalities to a disjunction of two negative interface equalities, such that $\sigma[\neg(a_i = b_i)] \mapsto \neg(a_i = t_i) \vee \neg(t_i = b_i)$ and t_i is an *AB*-pure term as described in Section 6.1.2. Initially, σ is empty.
- (2) Let $C_i \stackrel{\text{def}}{=} (a_i = b_i) \vee \neg\mu_k$ be the right premise \mathcal{T} lemma of the root of the Π^{ie} subproof.
- (3) Replace each $\neg(a_j = b_j)$ in C_i with $\sigma[\neg(a_j = b_j)]$, to obtain $C_i^* \stackrel{\text{def}}{=} (a_i = b_i) \vee \neg\eta_i$. If $(a_i = b_i)$ is not *AB*-mixed, then let Π be the subproof rooted in the left premise, and go to step (7).
- (4) Split C_i^* into $C_i^{\text{def}}(a_i = t_i) \vee \neg\eta_i$ and $C_i^{\prime\text{def}}(t_i = b_i) \vee \neg\eta_i$.
- (5) Rewrite the subproof

$$\frac{\frac{\vdots}{\neg(a_i = b_i) \vee \neg\mu_k} C_i}{\neg\mu_k \vee \neg\mu_i} \text{ into } \frac{\frac{\frac{\vdots}{\neg(a_i = t_i) \vee \neg(t_i = b_i) \vee \neg\mu_k} C_i^{\prime}}{\neg(t_i = b_i) \vee \neg\eta_k \vee \neg\eta_i} C_i^{\prime\prime}}{\neg\eta_k \vee \neg\eta_i} C_i^{\prime\prime},$$

where $\neg\eta_k$ is obtained by $\neg\mu_k$ by substituting each negative *AB*-mixed interface equality $\neg(a_j = b_j)$ with $\sigma[\neg(a_j = b_j)]$.

- (6) Update σ by setting $\sigma[\neg(a_i = b_i)]$ to $\neg(a_i = t_i) \vee \neg(t_i = b_i)$.
- (7) If Π is of the form $\frac{\vdots}{\dots} C_j$, set C_i to C_j and go to step (3).
- (8) Otherwise, Π is the leaf $\neg(a_i = t_i) \vee \neg(t_i = b_i) \vee \neg\mu_k$. In this case, replace each $\neg(a_j = b_j)$ in $\neg\mu_k$ with $\sigma[\neg(a_j = b_j)]$, and then exit.

In this case, both interface equalities $(a_1 = b_1)$ and $(a_2 = b_2)$ are *AB*-mixed. Consider the Π^{ie} subproof of Example 6.2:

$$\begin{array}{l} C_1 \stackrel{\text{def}}{=} (a_2 = b_2) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1), \\ C_2 \stackrel{\text{def}}{=} (a_1 = b_1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(a_2 = b_2), \\ C_3 \stackrel{\text{def}}{=} \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(a_1 = b_1), \\ \Theta_1 \stackrel{\text{def}}{=} \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(a_2 = b_2), \\ \Theta_2 \stackrel{\text{def}}{=} \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1). \end{array}$$

$$\frac{\frac{C_3 \quad C_2}{\Theta_1} \quad C_1}{\Theta_2} \Pi^{\text{ie}}$$

The first \mathcal{T} lemma processed by Algorithm 2 is C_1 . Using the technique of Yorsh and Musuvathi [2005], $(a_2 = b_2)$ is split into $(a_2 = y - 1) \wedge (y - 1 = b_2)$ (step (4)), thus obtaining C'_1, C''_1 , and the new proof (in step (5)):

$$\begin{array}{l} C'_1 \stackrel{\text{def}}{=} (a_2 = y - 1) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1), \\ C''_1 \stackrel{\text{def}}{=} (y - 1 = b_2) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1), \\ \Theta'_2 \stackrel{\text{def}}{=} \neg(y - 1 = b_2) \vee \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(b_1 = f(b_2)) \vee \\ \quad \neg(a_1 = f(a_2)) \vee \neg(y - a_2 = 1) \vee \neg(y - b_2 = 1). \end{array}$$

$$\frac{\frac{C_3 \quad C_2}{\Theta_1} \quad C'_1}{\Theta'_2} \quad C''_1$$

Then, $\sigma[\neg(a_2 = b_2)]$ is set to $\neg(a_2 = y - 1) \vee \neg(y - 1 = b_2)$ (step (6)), and a new iteration of the loop (3)–(7) is performed, this time processing C_2 . First, $\neg(a_2 = b_2)$ is replaced by $\neg(a_2 = y - 1) \vee \neg(y - 1 = b_2)$ (step (3)). Then, $(a_1 = b_1)$

can be split into $(a_1 = f(y - 1)) \wedge (f(y - 1) = b_1)$ (step (4)). After the rewriting of step (5), the proof is

$$\begin{aligned}
C_2' &\stackrel{\text{def}}{=} (a_1 = f(y - 1)) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(a_2 = y - 1), \vee \\
&\quad \neg(y - 1 = b_2), \\
C_2'' &\stackrel{\text{def}}{=} (f(y - 1) = b_1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)) \vee \neg(a_2 = y - 1), \vee \\
&\quad \neg(y - 1 = b_2), \\
\Theta_1' &\stackrel{\text{def}}{=} \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(b_1 = f(b_2)) \vee \neg(a_1 = f(a_2)), \vee \\
&\quad \neg(a_2 = y - 1) \vee \neg(y - 1 = b_2), \\
\Theta_1'' &\stackrel{\text{def}}{=} \neg(a_1 = f(y - 1)) \vee \neg(a_1 + y = 0) \vee \neg(b_1 + y = 1) \vee \neg(b_1 = f(b_2)), \vee \\
&\quad \neg(a_1 = f(a_2)) \vee \neg(a_2 = b_2).
\end{aligned}$$

$$\boxed{
\begin{array}{c}
\frac{C_3 \quad C_2'}{\Theta_1'' \quad C_2''} \\
\frac{\Theta_1' \quad C_1'}{\Theta_2 \quad C_1''}
\end{array}$$

Finally, C_3 is processed in step (8), $\neg(a_1 = b_1)$ gets replaced with $\neg(a_1 = f(y - 1)) \vee \neg(f(y - 1) = b_1)$, and the following final proof Π^{ie} is generated:

$$\frac{\frac{C_3' \quad C_2'}{\Theta_1'' \quad C_2''} \quad \frac{\Theta_1' \quad C_1'}{\Theta_2 \quad C_1''}}{\Theta_2}$$

such that $C_3' \stackrel{\text{def}}{=} C_3[\neg(a_1 = b_1) \mapsto \neg(a_1 = f(y - 1)) \vee \neg(f(y - 1) = b_1)]$.

The following theorem states that Algorithm 2 is correct.

THEOREM 6.4. *Let Π be a Π^{ie} subproof, and let Π' be the result of applying Algorithm 2 to Π . Then*

- (a) Π' does not contain any *AB*-mixed interface equality; and
- (b) Π' is a valid subproof with the same root as Π .

PROOF.

- (a) Consider the \mathcal{T} lemma C_i of step (3). By item (vii) of Definition 6.3, all negative interface equalities occurring in C_i occur positively in leaf \mathcal{T} lemmas that are closer to the root of Π . For the same reason, the first \mathcal{T} lemma C_i analyzed in step (2) contains no negative *AB*-mixed interface equalities. Therefore, it follows by induction that all negative *AB*-mixed interface equalities in C_i must have been split in step (4) of a previous iteration of the loop (3)–(7) of Algorithm 2, and thus they occur in σ . The same argument can be used to show also that at steps (5) and (8) every negative *AB*-mixed interface equality in $\neg\mu_k$ occurs in σ .

- (b) We show that

- (i) every substep $\frac{\Theta' \quad \Theta''}{\Theta'''}$ of Π' is a valid resolution step;
- (ii) every leaf of Π' is a \mathcal{T} lemma; and
- (iii) the root of Π' is the same as that of Π .

- (i) The only problematic case is the resolution step

$$\frac{\neg(a_i = t_i) \vee \neg(t_i = b_i) \vee \neg\mu_k \quad C_i'}{\neg(t_i = b_i) \vee \neg\eta_k \vee \neg\eta_i}$$

introduced in step (5) of Algorithm 2. In this case, we have to show that, at the end of the algorithm, all the negative *AB*-mixed interface equalities in $\neg\mu_k$ have been replaced such that the result is identical to $\neg\eta_k$. We already know that all negative *AB*-mixed equalities in $\neg\mu_k$ occur in σ ; thus we only have to show that $\sigma[\neg e_j]$ cannot change between the time when $\neg e_j$ was rewritten to obtain $\neg\eta_k$ and the time in which it is rewritten in $\neg\mu_k$. The negative equality $\neg e_j$ is replaced in $\neg\mu_k$ at the next iteration of the algorithm (in step (5) for inner nodes, and in step (8) for the final leaf). In the meantime, the only update to σ is performed in step (6), but it involves the negative equality $\neg(a_i = b_i)$, which does not occur in $\neg\mu_k$.

- (ii) Let C_i be a \mathcal{T} lemma in Π . First, we observe that if $C_i \equiv \neg(a_i = b_i) \vee \neg\mu_i$, then for any t_i also the clause $C_i^* \stackrel{\text{def}}{=} \neg(a_i = t_i) \vee \neg(t_i = b_i) \vee \neg\mu_i$ is a \mathcal{T} lemma, since $(a_i = t_i) \wedge (t_i = b_i) \models_{\mathcal{T}} (a_i = b_i)$ by transitivity. Therefore, it follows by induction on the number of substitutions that the clauses obtained in steps (3) and (8) of Algorithm 2 are still \mathcal{T} lemmas. Finally, since we are considering equality-interpolating theories, after step (4) of Algorithm 2 both C'_i and C''_i are \mathcal{T} lemmas.
- (iii) Since the root of Π does not contain any interface equality (item (iii) of Definition 6.3), in step (5) $\neg\eta_i \equiv \neg\mu_i$ and $\neg\eta_k \equiv \neg\mu_k$, and therefore the root does not change. \square

Clearly, Algorithm 2 operates in linear time on the number of \mathcal{T} lemmas, and thus of *AB*-mixed interface equalities. Moreover, every time an interface equality is split, only two new nodes are added to the proof (a right leaf and an inner node), and therefore the size of Π' is linear in that of Π .

The advantage of having *ie*-local proofs is that they ease significantly the process of eliminating *AB*-mixed interface equalities. First, since all the reasoning involving interface equalities is confined in Π^{ie} subproofs, only such subproofs—which typically constitute only a small fraction of the whole proof—need to be traversed and manipulated. Second, the simple structure of Π^{ie} subproofs allows for an efficient application of the rewriting process of steps (5) and (3), preventing any explosion in size of the proof. In fact, for example, if in step (5) the right premise of the last step were instead the root of some subproof Π_i with C_i as a leaf, then two copies of Π'_i and Π''_i would be produced, in which each instance of $(a_i = b_i)$ must be replaced with $(a_i = t_i)$ and $(t_i = b_i)$, respectively.

6.2.2 Generating *ie*-Local Proofs in DTC. In this section we show how to implement a variant of DTC so that to generate *ie*-local proofs of unsatisfiability. For the sake of simplicity, we describe first a simplified algorithm which makes use of two distinct DPLL engines. We then describe how to avoid the need of a second DPLL engine with the use of a particular search strategy for DTC.

The simplified algorithm uses two distinct DPLL engines, a *main* one and an *auxiliary* one, which we shall call *DPL-1* and *DPL-2*, respectively. Consider Figure 10, left. DPLL-1 receives in input the clauses of the input problem ϕ (which we assume pure and $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistent), but no interface equality,

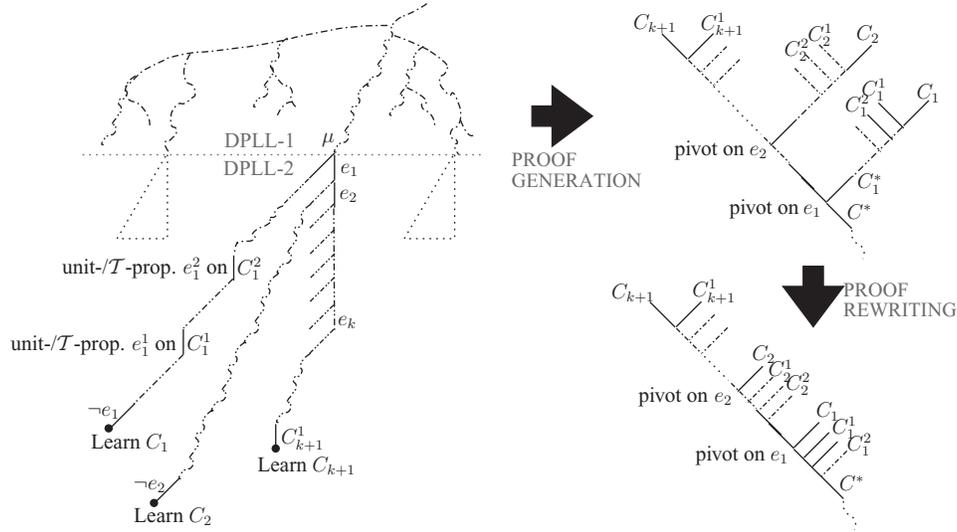


Fig. 10. Simple strategy for generating ie-local proofs. Left: DTC search; top-right: corresponding (sub)proof; bottom-right: Π^{ie} (sub)proof after rewriting.

which are instead given to DPLL-2. DPLL-1 enumerates total Boolean models μ of ϕ , and invokes the two \mathcal{T}_i solvers separately on the subsets $\mu_{\mathcal{T}_1}$ and $\mu_{\mathcal{T}_2}$ of μ . If one \mathcal{T}_i solver reports an inconsistency, then DPLL-1 backtracks. Otherwise, both $\mu_{\mathcal{T}_i}$ are \mathcal{T}_i -consistent, and DPLL-2 is invoked on the list of unit clauses composed of the $\mathcal{T}_1 \cup \mathcal{T}_2$ literals in μ , to check its $\mathcal{T}_1 \cup \mathcal{T}_2$ consistency.

DPLL-2 branches only on interface equalities, assigning them always to false first. Some interface equalities e_1^j , however, may be assigned to true by unit propagation on previously learned clauses in the form $C_1^j \stackrel{\text{def}}{=} \neg\mu_1^j \vee e_1^j$, or by \mathcal{T} propagation on deduction clauses C_1^j in the same form; we call C_1^j the *antecedent clause* of e_1^j .¹⁶ (As in Bruttomesso et al. [2009], we assume that when a \mathcal{T} propagation step $\mu_i^j \models_{\mathcal{T}} e_i^j$ occurs, μ_i^j being a subset of the current branch, the deduction clause $C_i^j \stackrel{\text{def}}{=} \neg\mu_i^j \vee e_i^j$ is learned, either temporarily or permanently; if so, we can see this step as a unit propagation on C_i^j .) When all the interface equalities have been assigned a truth value, the propositional model $\mu' \equiv \mu_{\mathcal{T}_1} \cup \mu_{\mathcal{T}_2} \cup \mu_{ie}$ is checked for $\mathcal{T}_1 \cup \mathcal{T}_2$ consistency by invoking each of the \mathcal{T}_i solvers on $\mu_{\mathcal{T}_i} \cup \mu_{ie}$.¹⁷ Since ϕ is inconsistent, one of the two \mathcal{T}_i solvers detects an inconsistency (if both do, we consider only the first). Therefore a \mathcal{T}_i lemma C_1 is generated. As stated at the end of Section 6.1.1, we can assume without loss of generality that C_1 contains at most one positive interface equality e_1 . (Notice also that all negative interface equalities $\neg e_1^j$ in C_1 , if any, have been assigned by unit propagation or \mathcal{T} propagation on some antecedent clause C_1^j .)

¹⁶Notationally, e_i^j denotes the j th most-recently unit-propagated interface equality in the branch in which C_i is learned, and $C_i^j \stackrel{\text{def}}{=} \neg\mu_i^j \vee e_i^j$ denotes the antecedent clause of e_i^j .

¹⁷In fact, it is not necessary to wait for all interface equalities to have a value before invoking the \mathcal{T}_i solvers. Rather, the standard *early pruning* optimization (see Section 2.2) can be applied.

DPLL-2 then learns C_1 and uses it as a conflicting clause to backjump: starting from C_1 , it eliminates from the clause every $\neg e_1^j$ by resolving the current clause against its antecedent clause C_1^j , until no negated equality occurs in the final clause C_1^* .¹⁸

If C_1 includes one positive interface equality e_1 , then also the final clause C_1^* includes it, so that DPLL-2 uses C_1^* as a conflict clause to jump up to μ and to unit propagate e_1 . Then DPLL-2 starts exploring a new branch. This process is repeated on several branches, learning a sequence of \mathcal{T} lemmas C_1, \dots, C_k each C_i containing only one positive interface equality e_i , until a branch causes the generation of a \mathcal{T} lemma C_{k+1} containing no positive interface equalities. Then C_{k+1} is resolved backward against the antecedent clauses of its negative interface equalities, generating a final conflict clause C^* which contains no interface equalities.

Overall, DPLL-2 has checked the $\mathcal{T}_1 \cup \mathcal{T}_2$ unsatisfiability of μ , building a resolution (sub)proof Π^* whose root is C^* . (Figure 10, top right.) Then the $\mathcal{T}_1 \cup \mathcal{T}_2$ lemma C^* is passed to DPLL-1, which uses it as a blocking clause for the assignment μ , it backtracks and continues the search. When the empty clause is obtained, it generates a proof of unsatisfiability in the usual way (see, e.g., van Gelder [2007]).

Since the main solver knows nothing about interface equalities, they can only appear inside the proofs of the blocking clauses generated by the auxiliary solver (like Π^*). Each Π^* is not yet a Π^{ie} subproof, since it complies only with items (i), (ii), and (iii) of Definition 6.3 but not with item (iv). The reason for the latter fact is that Π^* contains a set of right branches Π_{C_i} , one of each \mathcal{T} lemma C_i in $\{C_{k+1}, \dots, C_1\}$, representing the resolution steps to resolve away the interface equalities introduced by unit propagation/ \mathcal{T} propagation in each branch. Each such subbranch Π_{C_i} , however, can be reduced to length 1 by moving downward the resolution steps with the antecedent clauses C_i^1, C_i^2, \dots that C_i encounters in the branch. (Figure 10, bottom right.) This is done by recursively applying the following rewriting step to Π_{C_i} , until it reduces to the single clause C_i :

$$\frac{\begin{array}{c} \vdots \\ \hline \neg e_i \vee \neg \mu_i' \end{array} \quad \boxed{\begin{array}{c} \frac{\overbrace{C_i^j} \quad C_i^{j-1} \quad \vdots}{\neg \mu_i^j \vee e_i^j \quad \neg \mu_i'' \vee \neg e_i^j \vee e_i} \\ \hline \neg \mu_i^j \vee \neg \mu_i'' \vee e_i \end{array}}{\hline \neg \mu_i' \vee \neg \mu_i^j \vee \neg \mu_i''} \quad \Pi_{C_i}$$

¹⁸In order to determine the order in which to eliminate the interface equalities, the *implication graph* of the auxiliary DPLL engine can be used. This is a standard process in the conflict analysis in modern SAT and SMT solvers (see, e.g., van Gelder [2007]; Sebastiani [2007]).

⇒

$$\frac{\frac{\frac{\vdots}{\neg e_i \vee \neg \mu'_i} \quad \frac{\frac{C_i^1 \quad C_i}{C_i^{j-1} \quad \vdots}}{\neg \mu''_i \vee \neg e_i^j \vee e_i}}{\neg \mu'_i \vee \neg \mu''_i \vee \neg e_i^j}}{\neg \mu'_i \vee \neg \mu''_i \vee \neg \mu_i^j} \quad \underbrace{c_i^j}_{\neg \mu_i^j \vee e_i^j}}{\neg \mu'_i \vee \neg \mu_i^j \vee \neg \mu''_i} \quad . \quad (13)$$

As a result, each Π^* is transformed into a Π^{ie} subproof, so that the final proof is ie-local.

In an actual implementation, there is no need of having two distinct DPLL solvers for constructing ie-local proofs. In fact, we can obtain the same result by adopting a variant of the DTC Strategy 1 of Bruttomesso et al. [2009]. We never select an interface equality for case splitting if there is some other unassigned atom, and we always assign false to interface equalities first. Moreover, we “delay” \mathcal{T} propagation of interface equalities until all the original atoms have been assigned a truth value. Finally, when splitting on interface equalities, we restrict both the backjumping and the learning procedures of the DPLL engine as follows. Let d be the depth in the DPLL tree at which the first interface equality is selected for case splitting. If during the exploration of the current DPLL branch we have to backjump above d , then we generate by resolution a conflict clause that does not contain any interface equality, and “deactivate” all the \mathcal{T} lemmas containing some interface equality—that is, we do not use such \mathcal{T} lemmas for performing unit propagation—and we reactivate them only when we start splitting on interface equalities again. Using such a strategy, we obtain the same effect as in the simple algorithm using two DPLL engines: the search space is partitioned into two distinct subspaces, one of original atoms and the other of interface equalities, and the generated proof of unsatisfiability reflects this partition.

Finally, we remark that what described above is only *one* possible strategy for generating ie-local proofs, and not necessarily the most efficient one. Moreover, that of generating ie-local proofs is only a *sufficient* condition to obtain interpolants from DTC avoiding duplications of subproofs, and more general strategies may be conceived. The investigation of alternative strategies is part of ongoing and future work.

6.3 Discussion

Our new DTC-based combination method has several advantages over the traditional one of Yorsh and Musuvathi [2005] based on NO.

- (1) It inherits all the advantages of DTC over the traditional NO in terms of versatility, efficiency and restrictions imposed to \mathcal{T} solvers [Bozzano et al. 2006; Bruttomesso et al. 2009]. Moreover, it allows for using a more modern SMT solver, since many state-of-the-art solvers adopt variants or extensions of DTC instead of NO.

- (2) Instead of requiring an “ad hoc” method for performing the combination, it exploits the Boolean interpolation algorithm. In fact, thanks to the fact that interface equalities occur in the proof of unsatisfiability Π , once the AB -mixed terms in Π are split there is no need of any interpolant-combination method at all. In contrast, with the NO-based method of Yorsh and Musuvathi [2005] interpolants for $\mathcal{T}_1 \cup \mathcal{T}_2$ lemmas are generated by combining “theory-specific partial interpolants” for the two \mathcal{T}_i ’s with an algorithm that essentially duplicates the work that in our case is performed by the Boolean algorithm. This allows also for potentially exploiting optimization techniques for Boolean interpolation which are or will be made available from the literature.
- (3) By splitting AB -mixed terms only *after* the construction of the proof Π , it allows for computing several interpolants for several different partitions of the input problem into (A, B) *from the same proof* Π . This is particularly important for applications in abstraction refinement [Henzinger et al. 2004]. (This feature is discussed in Section 6.4.)

The work of Yorsh and Musuvathi [2005] can in principle deal with non-convex theories. Our approach is currently limited to the case of convex theories; however, we see no reason that would prevent from it being extensible at least theoretically to the case of nonconvex theories. Extending the approach to nonconvex theories is part of ongoing work. We also remark that implementing the algorithm of Yorsh and Musuvathi [2005] for nonconvex theories is a nontrivial task, and in fact we are not aware of any such implementation.

Another algorithm for computing interpolants in combined theories was given in Sofronie-Stokkermans [2006]. Rather than a combination of theories with disjoint signatures, that work considers the interpolation problem for extensions of a base (convex) theory with new function symbols, and it is therefore orthogonal to ours. The solution adopted is, however, similar to what we propose, in the sense that the algorithm of Sofronie-Stokkermans [2006] also works by splitting AB -mixed terms. The difference is that our algorithm is tightly integrated into an SMT context, as it is guided by the resolution proof generated by the DPLL engine.

6.4 Generating Multiple Interpolants

In Section 2.3 we remarked that a sufficient condition for generating multiple interpolants is that all the interpolants I_i ’s are computed from the same proof of unsatisfiability. When generating interpolants with our DTC-based algorithm, however, we generate a different proof of unsatisfiability Π_i for each partition of the input formula ϕ into A_i and B_i . In particular, every Π_i is obtained from the same “base” proof Π , by splitting all the $A_i B_i$ -mixed interface equalities with the algorithm described in Section 6.2. In this section, we show that (2) (at Section 2.3) holds also when each Π_i is obtained from the same ie-local proof Π by the rewriting of Algorithm 2 of Section 6.2.1. In order to do so, we need the following lemma.

LEMMA 6.5. *Let Θ be a $\mathcal{T}_1 \cup \mathcal{T}_2$ lemma, and let Π be a Π^{ie} proof for it which does not contain any AB-mixed term. Then the formula I_Θ associated to Θ in Algorithm 1 is an interpolant for $(\neg\Theta \setminus B, \neg\Theta \downarrow B)$.*

PROOF. By induction on the structure of Π , we have to prove that

- (1) $\neg\Theta \setminus B \models I_\Theta$;
- (2) $I_\Theta \wedge (\neg\Theta \downarrow B) \models \perp$;
- (3) I_Θ contains only common symbols.

The base case is when Π is just a single leaf. Then, the lemma trivially holds by definition of I_Θ in this case (see Algorithm 1).

For the inductive step, let $\Theta_1 \stackrel{\text{def}}{=} (x = y) \vee \phi_1$ and $\Theta_2 \stackrel{\text{def}}{=} \neg(x = y) \vee \phi_2$ be the antecedents of Θ in Π . (So $\Theta \stackrel{\text{def}}{=} \phi_1 \vee \phi_2$). Let I_{Θ_1} and I_{Θ_2} be the interpolants for Θ_1 and Θ_2 (by the inductive hypothesis).

If $(x = y) \not\leq B$, then $I_\Theta \stackrel{\text{def}}{=} I_{\Theta_1} \vee I_{\Theta_2}$.

- (1) By the inductive hypothesis, $(\neg\phi_1 \wedge \neg(x = y)) \setminus B \equiv (\neg\phi_1 \setminus B) \wedge \neg(x = y) \models I_{\Theta_1}$, and $(\neg\phi_2 \setminus B) \wedge (x = y) \models I_{\Theta_2}$. Then by resolution $(\neg\phi_1 \wedge \neg\phi_2) \setminus B \equiv \neg\Theta \setminus B \models I_\Theta$.
- (2) By the inductive hypothesis, $I_{\Theta_1} \models \phi_1 \downarrow B$ and $I_{\Theta_2} \models \phi_2 \downarrow B$, so $I_{\Theta_1} \vee I_{\Theta_2} \models (\phi_1 \vee \phi_2) \downarrow B$, that is, $I_\Theta \wedge (\neg\Theta \downarrow B) \models \perp$.
- (3) By the inductive hypothesis, both I_{Θ_1} and I_{Θ_2} contain only common symbols, and so also I_Θ does.

If $(x = y) \leq B$, then $I_\Theta \stackrel{\text{def}}{=} I_{\Theta_1} \wedge I_{\Theta_2}$.

- (1) By the inductive hypothesis, $\neg\phi_1 \setminus B \models I_{\Theta_1}$ and $\neg\phi_2 \setminus B \models I_{\Theta_2}$, so $(\neg\phi_1 \wedge \neg\phi_2) \setminus B \equiv \neg\Theta \setminus B \models I_\Theta$.
- (2) By the inductive hypothesis, we also have that $I_{\Theta_1} \models \phi_1 \downarrow B \vee (x = y)$ and $I_{\Theta_2} \models \phi_2 \downarrow B \vee \neg(x = y)$. Therefore, $I_{\Theta_1} \wedge I_{\Theta_2} \models (\phi_1 \vee \phi_2) \downarrow B$, that is, $I_\Theta \wedge (\neg\Theta \downarrow B) \models \perp$.
- (3) Finally, also in this case both I_{Θ_1} and I_{Θ_2} contain only common symbols, and so also I_Θ does. \square

We now formalize the sufficient condition of Henzinger et al. [2004] that (2) holds if the I_i 's are computed from the same Π . The proof of it will be useful for showing that (2) holds also if the I_i 's are computed from Π_i 's obtained from Π by splitting the $A_i B_i$ -mixed interface equalities.

THEOREM 6.6. *Let $\phi \stackrel{\text{def}}{=} \phi_1 \wedge \phi_2 \wedge \phi_3$, and let Π be a proof of unsatisfiability for it. Let $A \stackrel{\text{def}}{=} \phi_1$, $B \stackrel{\text{def}}{=} \phi_2 \wedge \phi_3$, $A' \stackrel{\text{def}}{=} \phi_1 \wedge \phi_2$, and $B' \stackrel{\text{def}}{=} \phi_3$, and let I' and I'' be two interpolants for (A, B) and (A', B') , respectively, both computed from Π . Then*

$$I' \wedge \phi_2 \models I''.$$

PROOF. Let Π_Θ be a proof whose root is the clause Θ . We will prove, by induction on the structure of Π_Θ , that

$$I'_\Theta \wedge \phi_2 \models I''_\Theta \vee (\Theta \setminus \phi_3),$$

where I_Θ is defined as in Algorithm 1. The validity of the theorem follows immediately, by observing that the root of Π is \perp .

We have to consider three cases:

- (1) The first is when Θ is an input clause. Then, we have three subcases:
 - (a) If $\Theta \in \phi_3$, then $I'_{\Theta} \stackrel{\text{def}}{=} \top$, $I''_{\Theta} \stackrel{\text{def}}{=} \top$ and $(\Theta \setminus \phi_3) \equiv \perp$, so the theorem holds.
 - (b) If $\Theta \in \phi_1$, then $I'_{\Theta} \stackrel{\text{def}}{=} (\Theta \downarrow (\phi_2 \cup \phi_3))$, $I''_{\Theta} \vee (\Theta \setminus \phi_3) \stackrel{\text{def}}{=} (\Theta \downarrow \phi_3) \vee (\Theta \setminus \phi_3) \equiv \Theta$, so the theorem holds also in this case.
 - (c) If $\Theta \in \phi_2$, then $I'_{\Theta} \wedge \phi_2 \equiv \phi_2$ and $I''_{\Theta} \vee (\Theta \setminus \phi_3) \equiv \Theta$, so again the implication holds.
- (2) The second is when Θ is a \mathcal{T} lemma. In this case, we have that I'_{Θ} is an interpolant for $(\neg\Theta \setminus (\phi_2 \cup \phi_3), \neg\Theta \downarrow (\phi_2 \cup \phi_3))$ and I''_{Θ} is an interpolant for $(\neg\Theta \setminus \phi_3, \neg\Theta \downarrow \phi_3)$. Therefore, by the definition of interpolant, $(\neg\Theta \setminus (\phi_2 \cup \phi_3)) \models I'_{\Theta}$ and $(\neg\Theta \setminus \phi_3) \models I''_{\Theta}$. Therefore, $I'_{\Theta} \vee (\Theta \setminus (\phi_2 \cup \phi_3))$ and $I''_{\Theta} \vee (\Theta \setminus \phi_3)$ are valid clauses, and so the implication trivially holds.
- (3) In this case Θ is obtained by resolution from $\Theta_1 \stackrel{\text{def}}{=} \phi \vee p$ and $\Theta_2 \stackrel{\text{def}}{=} \psi \vee \neg p$. If $p \in \phi_1$ or $p \in \phi_3$, then, by the inductive hypotheses that $I'_{\Theta_i} \wedge \phi_2 \models I''_{\Theta_i} \vee (\Theta_i \setminus \phi_3)$, we have that $I'_{\Theta} \wedge \phi_2 \models I''_{\Theta} \vee (\Theta \setminus \phi_3)$. If $p \in \phi_2$, then $I'_{\Theta} \stackrel{\text{def}}{=} I'_{\Theta_1} \wedge I'_{\Theta_2}$ and $I''_{\Theta} \stackrel{\text{def}}{=} I''_{\Theta_1} \vee I''_{\Theta_2}$. Again, by the inductive hypotheses $I'_{\Theta} \wedge \phi_2 \models I''_{\Theta} \vee (\Theta \setminus \phi_3)$ holds. \square

THEOREM 6.7. *Let $\phi \stackrel{\text{def}}{=} \phi_1 \wedge \phi_2 \wedge \phi_3$. Let $A \stackrel{\text{def}}{=} \phi_1$, $A'' \stackrel{\text{def}}{=} \phi_1 \wedge \phi_2$, $B \stackrel{\text{def}}{=} \phi_2 \wedge \phi_3$, and $B'' \stackrel{\text{def}}{=} \phi_3$. Let Π be a proof of unsatisfiability for ϕ , and let Π' and Π'' be obtained from Π by splitting all the $A'B$ -mixed and $A''B''$ -mixed interface equalities, respectively. Let I' be an interpolant for (A, B) computed from Π' , and I'' be an interpolant for (A'', B'') computed from Π'' . Then*

$$I' \wedge \phi_2 \models I''.$$

PROOF. We observe that Π' and Π'' are identical except for some Π^{ie} subproofs that contained some mixed interface equalities. Then we can proceed as in Theorem 6.6; we just need to consider one more case, namely, when Θ is a $\mathcal{T}_1 \cup \mathcal{T}_2$ lemma at the root of a Π^{ie} subproof. In this case, thanks to Lemma 6.5 we have the same situation as in the second case of the proof of Theorem 6.6, and so we can apply the same argument. \square

Thus, due to Theorem 6.7, we can use our DTC-based interpolation method in the context of abstraction refinement without any modification: it is enough to remember the original proof Π , and compute the interpolant I_i from the proof Π_i obtained by splitting the $A_i B_i$ -mixed terms in Π , for each partition of the input formula ϕ into A_i and B_i as in (1).

7. EXPERIMENTAL EVALUATION

The techniques presented in previous sections have been implemented within MATHSAT 4 [Bruttomesso et al. 2008]. MATHSAT is an SMT solver supporting a wide range of theories and their combinations. In the last SMT solvers competition (SMT-COMP'08), it proved to be competitive with the other state-of-the-art solvers. In this section, we experimentally evaluate our approach.

7.1 Description of the Benchmark Sets

We have performed our experiments on two different sets of benchmarks. The first was obtained by running the BLAST software model checker [Beyer et al. 2007] on some Windows device drivers; these are similar to those used in Rybalchenko and Sofronie-Stokkermans [2007]. This is one of the most important applications of interpolation in formal verification, namely, abstraction refinement in the context of CEGAR. The problem represents an abstract counterexample trace, and consists of a conjunction of atoms. In this setting, the interpolant generator is called very frequently, each time with a relatively simple input problem.

The second set of benchmarks originates from the SMT-LIB [Ranise and Tinelli 2006], and is composed of a subset of the unsatisfiable problems used in recent SMT solvers competitions.¹⁹ The instances have been converted to CNF and then split in two consistent parts of approximately the same size. The set consists of problems of varying difficulty and with a nontrivial Boolean structure.

The experiments were performed on a 3 GHz Intel Xeon machine with 4 GB of RAM running Linux. All the tools were run with a timeout of 600 s and a memory limit of 900 MB.

7.2 Comparison with the State-of-the-Art Tools Available

In this section, we compare with the other interpolant generators which are available: FOCI [McMillan 2005; Jhala and McMillan 2006], CLP-PROVER [Rybalchenko and Sofronie-Stokkermans 2007], and CSISAT [Beyer et al. 2008]. Other natural candidates for comparison would have been ZAP [Ball et al. 2005], and LIFTER [Kroening and Weissenbacher 2007]; however, it was not possible to obtain them from the authors. We also remark that no comparison with INT2 [Jain et al. 2009] is possible, since the domains of applications of MATHSAT and INT2 are disjoint: INT2 can handle $\mathcal{LA}(\mathbb{Z})$ equations/disequations and modular equations but only conjunctions of literals, whereas MATHSAT can handle formulas with arbitrary Boolean structure, but does not support $\mathcal{LA}(\mathbb{Z})$ except for its fragments $\mathcal{DL}(\mathbb{Z})$ and $UTVPI(\mathbb{Z})$.

The comparison had to be adapted to the limitations of FOCI, CLP-PROVER, and CSISAT. In fact, the current version of FOCI which is publically available does not handle the full $\mathcal{LA}(\mathbb{Q})$, but only the $\mathcal{DL}(\mathbb{Q})$ fragment.²⁰ We also notice that the interpolants it generates are not always $\mathcal{DL}(\mathbb{Q})$ formulas. (See, e.g., Example 4.2 of Section 4.) CLP-PROVER does handle the full $\mathcal{LA}(\mathbb{Q})$, but it accepts only conjunctions of atoms, rather than formulas with arbitrary Boolean structure. CSISAT, instead, can deal with $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ formulas with arbitrary Boolean structure, but it does not support Boolean variables. These limitations made it impossible to compare all the four tools on all the instances of our benchmark sets. Therefore, we performed the following comparisons.

¹⁹<http://www.smtcomp.org>.

²⁰For example, it fails to detect the $\mathcal{LA}(\mathbb{Q})$ unsatisfiability of the following problem: $(0 \leq y - x + w) \wedge (0 \leq x - z - w) \wedge (0 \leq z - y - 1)$.

Table II. Comparison of Execution Times of MATHSAT, FOCI, CLP-PROVER and CSISAT on Problems Generated by BLAST

Family	# of Problems	MATHSAT	FOCI	CLP-PROVER	CSISAT
kbfiltr.i	64	0.16	0.36	1.47	0.17
diskperf.i	119	0.33	0.78	3.08	0.39
floppy.i	235	0.73	1.64	5.91	0.86
cdaudio.i	130	0.35	1.07	2.98	0.47

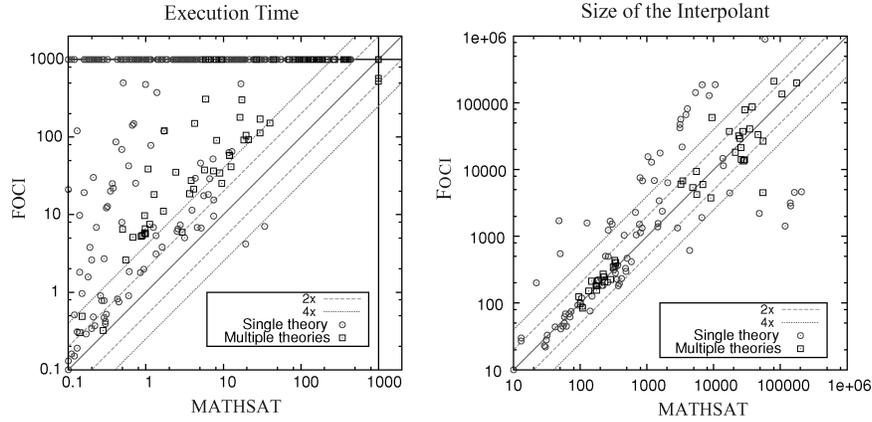


Fig. 11. Comparison of MATHSAT and FOCI on SMT-LIB instances: execution time (left), and size of the interpolant (right). In the left plot, points on the horizontal and vertical lines are timeouts/failures.

- We compared all the four solvers on the problems generated by BLAST.
- We compared MATHSAT with FOCI on SMT-LIB instances in the theories of \mathcal{EUF} , $\mathcal{DL}(\mathbb{Q})$ and their combination. In this case, we compared both the execution times and the sizes of the generated interpolants (in terms of number of nodes in the DAG representation of the formula). For computing interpolants in \mathcal{EUF} , we applied the algorithm of McMillan [2005], using an extension of the algorithm of Nieuwenhuis and Oliveras [2007] to generate \mathcal{EUF} proof trees. The combination $\mathcal{EUF} \cup \mathcal{DL}(\mathbb{Q})$ is handled with the technique described in Section 6;
- We compare MATHSAT, CLP-PROVER, and CSISAT on $\mathcal{LA}(\mathbb{Q})$ problems consisting of conjunctions of atoms. These problems are single branches of the search trees explored by MATHSAT for some $\mathcal{LA}(\mathbb{Q})$ instances in the SMT-LIB. We have collected several problems that took more than 0.1 s to MATHSAT to solve, and then randomly picked 50 of them. In this case, we did not compare the sizes of the interpolants as they are always atomic formulas.
- We compared MATHSAT and CSISAT on the subset (consisting of 78 instances of the about 400 collected) of the SMT-LIB instances without Boolean variables.

The results are collected in Table II and Figures 11 through 14. We can observe the following facts:

- Interpolation problems generated by BLAST are trivial for all the tools. In fact, we even had some difficulties in measuring the execution times reliably.

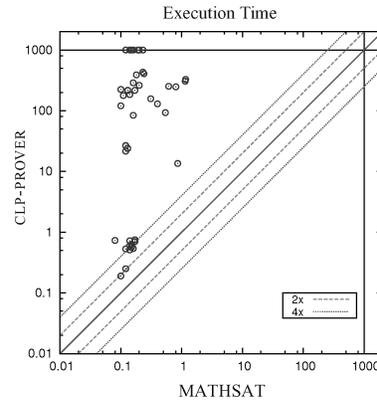


Fig. 12. Comparison of MATHSAT and CLP-PROVER on conjunctions of $\mathcal{L}\mathcal{A}(\mathbb{Q})$ atoms.

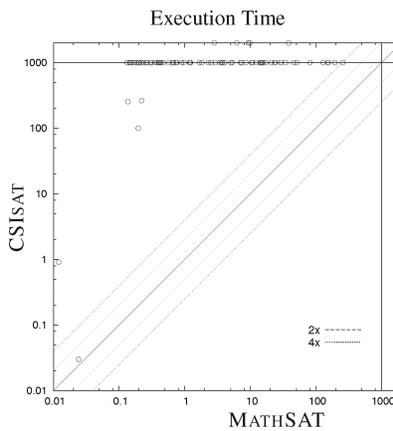


Fig. 13. Comparison of MATHSAT and CSISAT on SMT-LIB instances.

- Despite this, MATHSAT and CSISAT seem to be a little faster than the others.
- For problems with a nontrivial Boolean structure, MATHSAT outperformed FOCI in terms of execution time. This is true even for problems in the combined theory $\mathcal{EUF} \cup \mathcal{DL}(\mathbb{Q})$, despite the fact that the current implementation is still preliminary. As regards CSISAT, it could solve (within the time and memory limits) only 5 of the 78 instances it could potentially handle, and in all cases MATHSAT outperformed it.
 - In terms of size of the generated interpolants, the gap between MATHSAT and FOCI was smaller on average. However, the right plot of Figure 11 (which considers only instances for which both tools were able to generate an interpolant) shows that there were more cases in which MATHSAT produced a smaller interpolant.
 - On conjunctions of $\mathcal{L}\mathcal{A}(\mathbb{Q})$ atoms, MATHSAT outperformed CLP-PROVER, sometimes by more than two orders of magnitude. The performance of MATHSAT and CSISAT is comparable on such instances, with MATHSAT being slightly faster. However, there were several cases in which CSISAT computed a wrong

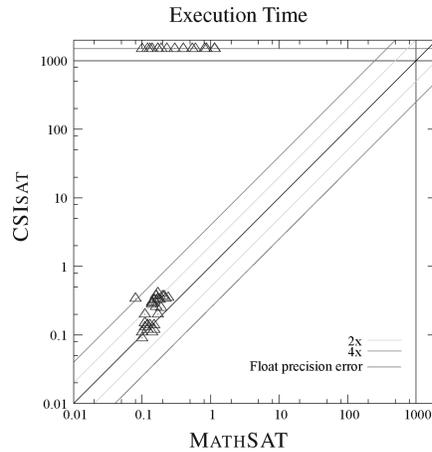


Fig. 14. Comparison of MATHSAT and CSISAT on conjunctions of $\mathcal{LA}(\mathbb{Q})$ atoms.

result, due to the use of floating-point arithmetic instead of infinite-precision arithmetic (which is used by MATHSAT).

8. CONCLUSIONS

In this article, we have shown how to efficiently build interpolants using state-of-the-art SMT solvers. Our methods encompass a wide range of theories (including \mathcal{EUF} , \mathcal{DL} , \mathcal{UTVPI} , and \mathcal{LA}), and their combination (based on the Delayed Theory Combination schema). A thorough experimental evaluation shows that the proposed methods retain the efficiency of the solvers, and are vastly superior to the state-of-the-art interpolants, both in terms of expressiveness, and in terms of efficiency.

REFERENCES

- BALL, T., LAHIRI, S. K., AND MUSUVATHI, M. 2005. Zap: Automated theorem proving for software analysis. In *Proceedings of LPAR*. G. Sutcliffe and A. Voronkov, Eds. Lecture Notes in Computer Science, vol. 3835. Springer, 2–22.
- BARRETT, C. W., DILL, D. L., AND STUMP, A. 2002. A Generalization of Shostak’s method for combining decision procedures. In *Proceedings of FroCos*. A. Armando, Ed. Lecture Notes in Computer Science, vol. 2309. Springer, Berlin, Germany, 132–146.
- BARRETT, C. W., SEBASTIANI, R., SESHIA, S. A., AND TINELLI, C. 2009. Satisfiability modulo theories. In *Handbook of Satisfiability*. A. Biere, M. Heule, H. van Maaren, and T. Walsh, Eds. Frontiers in Artificial Intelligence and Applications, vol. 185. IOS Press, Amsterdam, The Netherlands, 825–885.
- BEYER, D., HENZINGER, T. A., JHALA, R., AND MAJUMDAR, R. 2007. The Software Model Checker BLAST. *Int. J. Softw. Tools Tech. Trans.* 9, 5-6, 505–525.
- BEYER, D., ZUFFEREY, D., AND MAJUMDAR, R. 2008. CSISAT: Interpolation for LA+EUF. In *Proceedings of CAV*. A. Gupta and S. Malik, Eds. Lecture Notes in Computer Science, vol. 5123. Springer, Berlin, Germany, 304–308.
- BOZZANO, M., BRUTTOMESSO, R., CIMATTI, A., JUNTILA, T., ROSSUM, P., SCHULZ, S., AND SEBASTIANI, R. 2005. MathSAT: A tight integration of SAT and mathematical decision procedure. *J. Automat. Reason.* 35, 1-3.
- BOZZANO, M., BRUTTOMESSO, R., CIMATTI, A., JUNTILA, T., VAN ROSSUM, P., RANISE, S., AND SEBASTIANI, R. 2006. Efficient theory combination via Boolean search. *Inform. Comput.* 204, 10, 1411–1596.

- BRUTTOMESSO, R., CIMATTI, A., FRANZÉN, A., GRIGGIO, A., AND SEBASTIANI, R. 2008. The MATHSAT 4 SMT solver. In *Proceedings of CAV*. A. Gupta and S. Malik, Eds. Lecture Notes in Computer Science, vol. 5123. Springer, Berlin, Germany, 299–303.
- BRUTTOMESSO, R., CIMATTI, A., FRANZÉN, A., GRIGGIO, A., AND SEBASTIANI, R. 2009. Delayed theory combination vs. nelson-oppen for satisfiability modulo theories: A comparative analysis. Extended version. *Ann. Math. d. Artific. Intell.* 55, 1-2, 63–99.
- CABODI, G., MURCIANO, M., NOCCO, S., AND QUER, S. 2006. Stepping forward with interpolants in unbounded model checking. In *Proceedings of ICCAD*. S. Hassoun, Ed. ACM Press, New York, NY, 772–778.
- CIMATTI, A., GRIGGIO, A., AND SEBASTIANI, R. 2008. Efficient interpolant generation in satisfiability modulo theories. In *Proceedings of TACAS'08*. C. R. Ramakrishnan and J. Rehof, Eds. Lecture Notes in Computer Science, vol. 4963. Springer, Berlin, Germany, 397–412.
- CIMATTI, A., GRIGGIO, A., AND SEBASTIANI, R. 2009. Interpolant generation for UTVPI. In *Proceedings of CADE-22*. R. A. Schmidt, Ed. Lecture Notes in Computer Science, vol. 5663. Springer, Berlin, Germany, 167–182.
- COTTON, S. AND MALER, O. 2006. Fast and flexible difference constraint propagation for DPLL(T). In *Proceedings of SAT*. A. Biere and C. P. Gomes, Eds. Lecture Notes in Computer Science, vol. 4121. Springer, Berlin, Germany, 170–183.
- DAVIS, M., LOGEMANN, G., AND LOVELAND, D. W. 1962. A machine program for theorem-proving. *Comm. ACM* 5, 7, 394–397.
- DE MOURA, L. AND BJØRNER, N. 2008. Proofs and refutations, and Z3. In *Proceedings of the LPAR Workshops*. P. Rudnicki, G. Sutcliffe, B. Konev, R. A. Schmidt, and S. Schulz, Eds. CEUR Workshop Proceedings, vol. 418. CEUR-WS.org.
- DUTERTRE, B. AND DE MOURA, L. 2006. A fast linear-arithmetic solver for DPLL(T). In *Proceedings of CAV*. T. Ball and R. B. Jones, Eds. Lecture Notes in Computer Science, vol. 4144. Springer, Berlin, Germany, 81–94.
- ENDERTON, H. B. 2001. *A Mathematical Introduction to Logic*, 2nd ed. Academic Press, New York, NY.
- FUCHS, A., GOEL, A., GRUNDY, J., KRSTIC, S., AND TINELLI, C. 2009. Ground interpolation for the theory of equality. In *Proceedings of TACAS'09*. S. Kowalewski and A. Philippou, Eds. Lecture Notes in Computer Science, vol. 5505. Springer, Berlin, Germany, 413–427.
- GOEL, A., KRSTIC, S., AND TINELLI, C. 2009. Ground interpolation for combined theories. In *Proceedings of CADE-22*. R. A. Schmidt, Ed. Lecture Notes in Computer Science, vol. 5663. Springer, Berlin, Germany, 183–198.
- HARVEY, W. AND STUCKEY, P. 1997. A unit two variable per inequality integer constraint solver for constraint logic programming. In *Proceedings of the Australian Computer Science Conference*. 102–111.
- HENZINGER, T. A., JHALA, R., MAJUMDAR, R., AND McMILLAN, K. L. 2004. Abstractions from proofs. In *Proceedings of POPL*. N. D. Jones and X. Leroy, Eds. ACM Press, New York, NY, 232–244.
- JAFFAR, J., MAHER, M. J., STUCKEY, P. J., AND YAP, R. H. C. 1994. Beyond finite domains. In *Proceedings of PPCP*. Lecture Notes in Computer Science, vol. 874. Springer, Berlin, Germany, 86–94.
- JAIN, H., CLARKE, E. M., AND GRUMBERG, O. 2009. Efficient Craig interpolation for linear diophantine (dis)equations and linear modular equations. *Form. Meth. Syst. Des.* 35, 1, 6–39.
- JHALA, R. AND McMILLAN, K. 2005. Interpolant-based transition relation approximation. In *Proceedings of CAV*. K. Etessami and S. K. Rajamani, Eds. Lecture Notes in Computer Science, vol. 3576. Springer, Berlin, Germany, 39–51.
- JHALA, R. AND McMILLAN, K. L. 2006. A practical and complete approach to predicate refinement. In *Proceedings of TACAS*. H. Hermanns and J. Palsberg, Eds. Lecture Notes in Computer Science, vol. 3920. Springer, Berlin, Germany, 459–473.
- JHALA, R. AND McMILLAN, K. L. 2007. Array abstractions from proofs. In *Proceedings of CAV*. W. Damm and H. Hermanns, Eds. Lecture Notes in Computer Science, vol. 4590. Springer, Berlin, Germany, 193–206.

- KAPUR, D., MAJUMDAR, R., AND ZARBA, C. G. 2006. Interpolation for data structures. In *Proceedings of SIGSOFT FSE*. M. Young and P. T. Devanbu, Eds. ACM Press, New York, 105–116.
- KROENING, D. AND STRICHMAN, O. 2008. *Decision Procedures: An Algorithmic Point of View*. Springer, Berlin, Germany.
- KROENING, D. AND WEISSENBACHER, G. 2007. Lifting propositional interpolants to the word-level. In *Proceedings of FMCAD*. IEEE Computer Society Press, Los Alamitos, CA, 85–89.
- LAHIRI, S. K. AND MUSUVATHI, M. 2005. An efficient decision procedure for UTVPI constraints. In *Proceedings of FroCos*. B. Gramlich, Ed. Lecture Notes in Computer Science, vol. 3717. Springer, Berlin, Germany, 168–183.
- LI, B. AND SOMENZI, F. 2006. Efficient abstraction refinement in interpolation-based unbounded model checking. In *Proceedings of TACAS*. H. Hermanns and J. Palsberg, Eds. Lecture Notes in Computer Science, vol. 3920. Springer, Berlin, Germany, 227–241.
- LYNCH, C. AND TANG, Y. 2008. Interpolants for linear arithmetic in SMT. In *Proceedings of ATVA*. S. D. Cha, J.-Y. Choi, M. Kim, I. Lee, and M. Viswanathan, Eds. Lecture Notes in Computer Science, vol. 5311. Springer, Berlin, Germany, 156–170.
- MARQUES-SILVA, J. 2007. Interpolant learning and reuse in SAT-based model checking. *Electr. Notes Theor. Comput. Sci.* 174, 3, 31–43.
- McMILLAN, K. L. 2003. Interpolation and SAT-based model checking. In *Proceedings of CAV*. W. A. H. Jr. and F. Somenzi, Eds. Lecture Notes in Computer Science, vol. 2725. Springer, Berlin, Germany, 1–13.
- McMILLAN, K. L. 2005. An interpolating theorem prover. *Theor. Comput. Sci.* 345, 1, 101–121.
- McMILLAN, K. L. 2006. Lazy Abstraction with Interpolants. In *Proceedings of CAV*, T. Ball and R. B. Jones, Eds. Lecture Notes in Computer Science, vol. 4144. Springer, Berlin, Germany, 123–136.
- MINÉ, A. 2001. The octagon abstract domain. In *Proceedings of WCRE*. IEEE Computer Society Press, Los Alamitos, CA, 31–100.
- NELSON, G. AND OPPEN, D. 1979. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.* 1, 2, 245–257.
- NIEUWENHUIS, R. AND OLIVERAS, A. 2005. DPLL(T) with exhaustive theory propagation and its application to difference logic. In *Proceedings of CAV*. K. Etessami and S. K. Rajamani, Eds. Lecture Notes in Computer Science, vol. 3576. Springer, Berlin, Germany, 321–334.
- NIEUWENHUIS, R. AND OLIVERAS, A. 2007. Fast congruence closure and extensions. *Inf. Comput.* 2005, 4, 557–580.
- PUDLÁK, P. 1997. Lower bounds for resolution and cutting planes proofs and monotone computations. *J. Symb. Logic* 62, 3.
- RANISE, S. AND TINELLI, C. 2006. The satisfiability modulo theories library (SMT-LIB). www.SMT-LIB.org.
- RYBALCHENKO, A. AND SOFRONIE-STOKKERMANS, V. 2007. Constraint solving for interpolation. In *Proceedings of VMCAI*. B. Cook and A. Podelski, Eds. Lecture Notes in Computer Science, vol. 4349. Springer, Berlin, Germany, 346–362.
- SEBASTIANI, R. 2007. Lazy satisfiability modulo theories. *J. Satis. Boolean Model. Comput.* 3, 3-4, 141–224.
- SOFRONIE-STOKKERMANS, V. 2006. Interpolation in local theory extensions. In *Proceedings of IJCAR*, U. Furbach and N. Shankar, Eds. Lecture Notes in Computer Science, vol. 4130. Springer, Berlin, Germany, 235–250.
- TSEITIN, G. S. 1968. On the complexity of derivation in propositional calculus. In *Studies in Constructive Mathematics and Mathematical Logic, Part 2*, A. O. Slisenko, Ed. Nauka, Leningrad, Russia, 115–125.
- VAN GELDER, A. 2007. Verifying propositional unsatisfiability: Pitfalls to avoid. In *Proceedings of SAT*. J. Marques-Silva and K. A. Sakallah, Eds. Lecture Notes in Computer Science, vol. 4501. Springer, Berlin, Germany, 328–333.
- VANDERBEL, R. J. 2001. *Linear Programming: Foundations and Extensions*. Springer, Berlin, Germany.

YORSH, G. AND MUSUVATHI, M. 2005. A combination method for generating interpolants. In *Proceedings of CADE*. R. Nieuwenhuis, Ed. Lecture Notes in Computer Science, vol. 3632. Springer, Berlin, Germany, 353–368.

ZHANG, L. AND MALIK, S. 2002. The quest for efficient boolean satisfiability solvers. In *Proceedings of CADE*. A. Voronkov, Ed. Lecture Notes in Computer Science, vol. 2392. Springer, Berlin, Germany, 295–313.

Received June 2009; accepted December 2009