

# Decomposing Farkas Interpolants

Martin Blicha<sup>1,2</sup>, Antti E. J. Hyvärinen<sup>1</sup>, Jan Kofroně<sup>2</sup>, and Natasha Sharygina<sup>1</sup>

<sup>1</sup> Università della Svizzera italiana (USI), Switzerland

<sup>2</sup> Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic

**Abstract.** Modern verification commonly models software with Boolean logic and a system of linear inequalities over reals and over-approximates the reachable states of the model with Craig interpolation to obtain, for example, candidates for inductive invariants. Interpolants for the linear system can be efficiently constructed from a Simplex refutation by applying the Farkas’ lemma. However, this approach results in the worst case in incompleteness, since Farkas interpolants do not always suit the verification task. This work introduces the decomposed interpolants, a fundamental extension of the Farkas interpolants obtained by identifying and separating independent components from the interpolant structure using methods from linear algebra. We integrate our approach to the model checker Sally and show experimentally that a portfolio of decomposed interpolants results in immediate convergence on instances where state-of-the-art approaches diverge. Being based on the efficient Simplex method, the approach is very competitive also outside these diverging cases.

**Keywords:** Model checking · Satisfiability modulo theory · Linear arithmetic · Craig interpolation

## 1 Introduction

A central task in model checking systems with respect to safety properties [25] consists of proving facts and attempting to generalize the obtained proofs. The generalizations serve as a basis for inductive invariants needed for guiding the search for a correctness proof in approaches such as IC3 [7] and  $k$ -induction [27], both known to scale to the verification of highly complex systems.

Finding good proofs and generalizing them is hard. A widely used approach, Satisfiability Modulo Theories (SMT) [12,6], models a system with propositional logic and a range of first-order logics. Solvers for SMT combine a resolution-based variant of the DPLL-algorithm [11,10] for propositional logic with decision procedures for first-order logics. A vast range of first-order logics is maintained as part of the SMT-LIB Initiative [5]. What is common to these logics is that their solving requires typically only a handful of algorithms. Arguably, the two most important algorithms are a congruence closure algorithm for deciding quantifier-free equality logic with uninterpreted functions [29], and a Simplex-based procedure for linear arithmetic over real or rational numbers [15].

Generalizing proofs to inductive invariants is commonly done by Craig interpolation [9]. Here, the model is split into two parts, say,  $A$  and  $B$ , resulting in an *interpolation problem*  $(A, B)$ . The proof of unsatisfiability for  $A \wedge B$  is used to extract an *interpolant*  $I$ , a formula that is defined over the common symbols of  $A$  and  $B$ , is implied by  $A$ , and is unsatisfiable with  $B$ . Several interpolants can be computed for a given interpolation problem, and not all of them are useful for proving safety. Typically, this is a phenomenon used to construct a *portfolio* [18] of interpolation algorithms that is then applied in the hopes of aiding to find the safety proof.

The approaches to interpolation based on Farkas’ lemma construct an LRA interpolant by summing all inequalities appearing in  $A$  into a single inequality. We call the resulting interpolant the *Farkas interpolant*. While a single inequality is desirable in some cases, it prevents IC3-style algorithms from converging in other ones [34]. We present how methods from linear algebra can be applied on a Farkas interpolant to obtain *decomposed interpolants* that do not consist of a single inequality and guarantee the convergence of the model-checking algorithm for some of the cases where Farkas interpolants do not converge. A major advantage of decomposed interpolants is that they can be computed using Simplex-based decision procedures as a black box, allowing us to make use of the highly tuned implementations present in many state-of-the-art SMT solvers.

Intuitively, while computing the decomposed interpolants we do not directly sum the inequalities in  $A$ , but, instead, we split the sum into sub-sums. The result is an interpolant that is a conjunction of often more than one component of the Farkas interpolant. This allows us not only to solve the convergence problem observed in model checking examples, but also to gain more control over the strength of LRA interpolants. In summary, the main contributions of this paper are

1. a new Farkas-lemma-based interpolation algorithm for LRA that is able to deal with convergence problems in model-checking benchmarks while still relying on a highly efficient Simplex-based decision procedure,
2. establishing properties regarding logical strength of interpolants produced by our interpolation algorithm with respect to the original Farkas interpolants,
3. implementation of our new interpolation algorithm in OPENSMT, our SMT solver, and integration of our approach with the model checker SALLY
4. experiments showing that the new approach is efficient in model checking, in particular in showing systems unsafe.

While the underlying intuition is simple, we quote here Jean D’Alembert (1717 – 1783) in saying that *Algebra is generous; she often gives more than is asked of her*: Our detailed analysis in Sec. 4 and Sec. 5 shows that the structure of the problem is surprisingly rich. Our experiments in Sec. 6 verify that the phenomena are practically relevant. Overall a portfolio constructed from our interpolation algorithm is significantly better than a portfolio based purely on Farkas interpolants. We furthermore show for individual instances that the effect is consistent instead of arising from random effects.

*Related work.* The work on interpolation in LRA dates back to [30]. A compact set of rules for deriving LRA interpolants from the proof of unsatisfiability in an inference system was presented in [22]. The interpolants in these works were the Farkas interpolants. Current methods usually compute Farkas interpolants from explanations of unsatisfiability extracted directly from the Simplex-based decision procedure inside the SMT solver [15]. Recently in [2], we presented a way of computing an infinite family of interpolants between a primal and a dual interpolant with variable strength. However, those interpolants are still restricted to single inequalities.

The work most closely related to ours is [34] where the authors independently recognized the weakness of interpolation based on Farkas coefficients. They introduce a new interpolation procedure that gives guarantees of convergence of a special sequence of interpolation problems often occurring in model checking problems. However, this interpolation algorithm is based on a different decision procedure, called conflict resolution [24], which, based on the results reported in [34], is not as efficient as the Simplex-based decision procedure. In contrast, we show how the original approach based on the Simplex-based decision procedure and Farkas coefficients can be modified to produce interpolants not restricted to the single-inequality form, while additionally obtaining strength guarantees with respect to the original Farkas interpolants.

Other work on LRA interpolants include e.g. [1,35,33]. Both [1] and [35] focus on producing simple overall interpolants by attempting to reuse (partial) interpolants from pure LRA conflicts. Our focus is not on the overall interpolant, but on a single LRA conflict. However, in the context of interpolants from proofs produced by SMT solvers, our approach also has a potential for re-using components of interpolants for LRA conflicts across the whole proof. Beside algorithms for interpolants for LRA conflicts, there exist a large body of work on propositional interpolation. See, e.g. [13].

The structure of the paper is as follows. In Sec. 2 we provide a concrete example model-checking problem where our approach guarantees immediate convergence but Farkas interpolation deviates. In Sec. 3 we define the notation used in the paper, and in Sec. 4 and 5 detail our main theoretical contribution. We provide experimental results in Sec. 6, and finally conclude in Sec. 7.

## 2 Motivation

Consider the transition system  $S = (I, T, Err)$ , where  $I$  and  $Err$  are, respectively, predicates that capture the initial and error states, and  $T$  is the transition function. The symbols  $x, y$  are real variables, and  $x', y'$  are their next-state versions.<sup>3</sup>

$$S = \begin{cases} I \equiv (x = 0) \wedge (y = 0), \\ T \equiv (x' = x + y) \wedge (y' = y + 1), \\ Err \equiv (x < 0) \end{cases} \quad (1)$$

---

<sup>3</sup> This example was first brought to our attention by Prof. Arie Gurfinkel. A similar example appears in [34].

The system is one variant from a family of similar transition systems that are known to not converge in straightforward implementations of IC3-based algorithms using LRA interpolation. For example, both SPACER [23] (using interpolation algorithm of Z3 [28]) and SALLY [21] (using interpolation algorithm of MATHSAT [8]) fail to compute a safe inductive invariant for this transition system. However, SALLY with our interpolation algorithm succeeds in computing the safe inductive invariant.<sup>4</sup> Closer examination of SALLY and SPACER reveals that the tools in their default configurations produce a divergent series of candidate invariants of the form  $0 \leq kx + y$  for  $k = 1, 2, 3, \dots$ . The reason for producing such a series is that both tools rely on Farkas interpolants that always consist of a single inequality. Instead of generalizing the Farkas interpolants, an approach advocated in this work, interpolation based on a different decision procedure was proposed for SALLY in [34], whereas SEAHORN [17] with SPACER as its underlying reasoning engine solves this issue with abstract interpretation.

In Sec. 4 we show how to modify the interpolation algorithm to produce in the general case a *conjunction* of multiple inequalities, leading, in this case, to the discovery of an inductive safe invariant  $x \geq 0 \wedge y \geq 0$ . To avoid here a lengthy discussion on internals of IC3 but nevertheless provide a concrete example of the power of decomposed interpolants, we apply decomposed interpolants in a simple, interpolation-based procedure for computing inductive invariants for transition systems. This approach is a simplified version of  $k$ -induction (see, e.g., [26]). When applied to the system in Eq. (1), we show that computing the Farkas interpolant fails and decomposed interpolant succeeds in producing a safe inductive invariant. A safe, inductive invariant for  $(I, T, Err)$  is a predicate  $R$  that satisfies (1)  $I(X) \rightarrow R(X)$ , (2)  $R(X) \wedge T(X, X') \rightarrow R(X)$ , and (3)  $R(X) \wedge Err(X) \rightarrow \perp$ . We may opportunistically try to synthesise  $R$  by interpolating over the interpolation problem  $(I(X), T(X, X') \wedge Err(X'))$ . Using the system  $S$  of Eq. (1), we obtain  $(x \geq 0 \wedge y \geq 0, x' = x + y \wedge y' = y + 1 \wedge x' < 0)$ . A Farkas interpolant, the sum of the components from the  $A$ -part, is  $x + y \geq 0$ , which is neither safe nor inductive for  $S$ . However, the *decomposed interpolant* (see Sec. 4)  $x \geq 0 \wedge y \geq 0$  is an inductive invariant.

### 3 Preliminaries

We work in the domain of *Satisfiability Modulo Theories* (SMT) [12,6], where satisfiability of formulas is determined with respect to some background theory. In particular, we are concerned with the *lazy* approach to SMT, that combines SAT solver dealing with the propositional structure of a formula and *theory* solver for checking consistency of a conjunction of theory literals. The proof of unsatisfiability in this approach is basically a propositional proof that incorporates *theory lemmas* learnt by the theory solver and propagated to the SAT solver.

---

<sup>4</sup> Current implementation of SPACER does not support conjunctions of inequalities as interpolants, and therefore we are at the moment unable to try our approach on SPACER.

The proof-based interpolation algorithm then combines any propositional-proof-based interpolation algorithm with *theory interpolator*. Theory interpolator provides an interpolant for each theory conflict—an unsatisfiable conjunction of theory literals.

*Linear arithmetic and linear algebra.* We use the letters  $x, y, z$  to denote variables and  $c, k$  to denote constants. Vector of  $n$  variables is denoted by  $\mathbf{x} = (x_1, \dots, x_n)^\top$  where  $n$  is usually known from context.  $\mathbf{x}[i]$  denotes the element of  $\mathbf{x}$  at position  $i$ , i.e.  $\mathbf{x}[i] = x_i$ . The vector of all zeroes is denoted as  $\mathbf{0}$  and  $\mathbf{e}_i$  denotes the unit vector with  $\mathbf{e}_i[i] = 1$  and  $\mathbf{e}_i[j] = 0$  for  $j \neq i$ .  $\mathbb{Q}$  denotes the set of rational numbers,  $\mathbb{Q}^n$  the  $n$ -dimensional vector space of rational numbers and  $\mathbb{Q}^{m \times n}$  the set of rational matrices with  $m$  rows and  $n$  columns. A transpose of matrix  $M$  is denoted as  $M^\top$ . A kernel (also nullspace) of a matrix  $M$  is the vector space  $\ker(M) = \{\mathbf{x} \mid M\mathbf{x} = \mathbf{0}\}$

We adopt the notation of matrix product for linear arithmetic. For a linear term  $l = c_1x_1 + \dots + c_nx_n$ , we write  $\mathbf{c}^\top\mathbf{x}$  to denote  $l$ . Without loss of generality we assume that all linear inequalities are of the form  $\mathbf{c}^\top\mathbf{x} \bowtie c$  with  $\bowtie \in \{\leq, <\}$ . By linear system over variables  $\mathbf{x}$  we mean a finite set of linear inequalities  $S = \{C_i \mid i = 1, \dots, m\}$ , where each  $C_i$  is a linear inequality over  $\mathbf{x}$ . Note that from the logical perspective, each  $C_i$  is an atom in the language of the theory of linear arithmetic, thus system  $S$  can be expressed as a formula  $\bigwedge_{i=1}^m C_i$  and we use these representations interchangeably. A linear system is satisfiable if there exists an evaluation of variables that satisfies all inequalities; otherwise, it is unsatisfiable. This is the same as the (un)satisfiability of the formula representing the system.

We extend the matrix notation also to the linear system and we write  $C\mathbf{x} \leq \mathbf{c}$  to denote the linear system  $S$  where  $C$  denotes the matrix of all coefficients of the system,  $\mathbf{x}$  are the variables and  $\mathbf{c}$  is the vector of the right sides of the inequalities. With the matrix notation, we can easily express the sum of (multiples) of inequalities. Given a system of inequalities  $C\mathbf{x} \leq \mathbf{c}$  and a vector of “weights” (multiples) of the inequalities  $\mathbf{k} \geq \mathbf{0}$ , the inequality that is the (weighted) sum of the system can be expressed as  $\mathbf{k}^\top C\mathbf{x} \leq \mathbf{k}^\top \mathbf{c}$ . Note that if at least one inequality in the system is strict, then its sum is a strict inequality. For the sake of simplicity, we use  $\leq$  notation, as the theory and the algorithm work the same way also in the presence of strict inequalities.

*Craig interpolation.* Given two formulas  $A(\mathbf{x}, \mathbf{y})$  and  $B(\mathbf{y}, \mathbf{z})$  such that  $A \wedge B$  is unsatisfiable, a *Craig interpolant* [9] is a formula  $I(\mathbf{y})$  such that  $A \implies I$  and  $I \implies \neg B$ .

The pair of formulas  $(A, B)$  is also referred to as an *interpolation problem*. In linear arithmetic, the interpolation problem is a linear system  $S$  partitioned into two parts:  $A$  and  $B$ .

One way to compute a solution to an interpolation problem in linear arithmetic, used in many modern SMT solvers, is based on Farkas’ lemma [16,36]. Farkas’ lemma states that for an unsatisfiable system of linear inequalities  $S \equiv C\mathbf{x} \leq \mathbf{c}$  there exist *Farkas* coefficients  $\mathbf{k} \geq \mathbf{0}$  such that  $\mathbf{k}^\top C\mathbf{x} \leq \mathbf{k}^\top \mathbf{c} \equiv 0 \leq -1$ .

In other words, the weighted sum of the system given by the Farkas coefficients is a contradictory inequality. If  $S$  contains strict inequality, the witness weighted sum might also be  $0 < 0$ .

The idea behind the interpolation algorithm based on Farkas coefficients is simple. Intuitively, given the partition of the linear system into  $A$  and  $B$ , we compute only the weighted sum of  $A$ . It is not hard to see that this sum is an interpolant. It follows from  $A$  because a weighted sum of a linear system with non-negative weights is always implied by the system. It is inconsistent with  $B$  because its sum with the weighted sum of  $B$  (using Farkas coefficients) is a contradictory inequality by Farkas lemma. Finally, it cannot contain any  $A$ -local variables, because in the weighted sum of the whole system all variables are eliminated,  $A$ -local variables are not present in  $B$ , so they must be eliminated already in the weighted sum of  $A$ .

More formally, for an unsatisfiable linear system  $S \equiv C\mathbf{x} \leq \mathbf{c}$  over  $n$  variables, where  $C \in \mathbb{Q}^{m \times n}$ ,  $\mathbf{c} \in \mathbb{Q}^m$ , and its partition to  $A \equiv C_A\mathbf{x} \leq \mathbf{c}_A$  and  $B \equiv C_B\mathbf{x} \leq \mathbf{c}_B$ , where  $C_A \in \mathbb{Q}^{k \times n}$ ,  $C_B \in \mathbb{Q}^{l \times n}$ ,  $\mathbf{c}_A \in \mathbb{Q}^k$ ,  $\mathbf{c}_B \in \mathbb{Q}^l$  and  $k + l = m$ , there exist Farkas coefficients  $\mathbf{k}^\top = (\mathbf{k}_A^\top \mathbf{k}_B^\top)$  such that

$$(\mathbf{k}_A^\top \mathbf{k}_B^\top) \begin{pmatrix} C_A \\ C_B \end{pmatrix} = 0, (\mathbf{k}_A^\top \mathbf{k}_B^\top) \begin{pmatrix} \mathbf{c}_A \\ \mathbf{c}_B \end{pmatrix} = -1,$$

and the *Farkas interpolant* for  $(A, B)$  is the inequality

$$I^f \equiv \mathbf{k}_A^\top C_A \mathbf{x} \leq \mathbf{k}_A^\top \mathbf{c}_A \tag{2}$$

## 4 Decomposed Interpolants

In this section, we present our new approach to computing interpolants in linear arithmetic based on Farkas coefficients.

Recall the definition of Farkas interpolant (2). This corresponds to the weighted sum of  $A$ -part of the unsatisfiable linear system.

This sum can be decomposed into  $j$  “smaller” sums by decomposing the vector  $\mathbf{k}_A$  into  $j$  vectors

$$\mathbf{k}_A = \sum_{i=1}^j \mathbf{k}_{A,i} \tag{3}$$

such that  $\mathbf{0} \leq \mathbf{k}_{A,i} \leq \mathbf{k}_A$  for all  $i$ , thus obtaining  $j$  inequalities

$$I_i \equiv \mathbf{k}_{A,i}^\top C_A \mathbf{x} \leq \mathbf{k}_{A,i}^\top \mathbf{c}_A \tag{4}$$

If  $\mathbf{k}_{A,i}$  are such that the left-hand side of the inequalities  $I_i$  contains only shared variables, the decomposition has an interesting application in interpolation, as illustrated below.

**Definition 1 (decomposed interpolants).** Given an interpolation instance  $(A, B)$ , if there exists a sum of the form Eq. (3) such that the left side of Eq. (4) contains only shared variables for all  $1 \leq i \leq j$ , then the set of inequalities  $S = \{I_1, \dots, I_j\}$  is a decomposition. In that case the formula  $\bigwedge_{i=1}^j I_i$  is a decomposed interpolant (DI) of size  $j$  for  $(A, B)$ .

The decomposed interpolants are proper interpolants, as stated in the following theorem.

**Theorem 1.** Let  $(A, B)$  be an interpolation problem in linear arithmetic. If  $S = \{I_1, \dots, I_k\}$  is a decomposition, then  $I^{DI} = I_1 \wedge \dots \wedge I_k$  is an interpolant for  $(A, B)$ .

*Proof.* Let  $I^{DI} = I_1 \wedge \dots \wedge I_k$ . First,  $A \implies I^{DI}$  holds since for all  $I_i$ ,  $A \implies I_i$ . This is immediate from the fact that  $A$  is a system of linear inequalities  $C_A \mathbf{x} \leq \mathbf{c}_A$ ,  $I_i \equiv \mathbf{k}_{A,i}^\top C_A \mathbf{x} \leq \mathbf{k}_{A,i}^\top \mathbf{c}_A$  and  $\mathbf{0} \leq \mathbf{k}_{A,i}$ . Second,  $I^{DI} \wedge B \implies \perp$  since  $I^{DI}$  implies Farkas interpolant  $I^f$ . This holds because  $\mathbf{k}_A = \sum_i \mathbf{k}_{A,i}$  and  $\mathbf{0} \leq \mathbf{k}_{A,i}$ . Third,  $I^{DI}$  contains only shared variables by the definition of decomposition (Def 1). Therefore,  $I^{DI}$  is an interpolant.  $\square$

Each interpolation instance has a DI of size one, a *trivial* decomposition, corresponding to the Farkas interpolant of Eq. (2). However, interpolation problems in general can admit bigger decompositions. In the following we give a concrete example of an instance with decomposition of size two.

*Example 1.* Let  $(A, B)$  be an interpolation problem in linear arithmetic with  $A = (x_1 + x_2 \leq 0) \wedge (x_1 + x_3 \leq 0) \wedge (-x_1 \leq 0)$  and  $B = (-x_2 - x_3 \leq -1)$ . The linear systems corresponding to  $A$  and  $B$  are

$$C_A = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ -1 & 0 & 0 \end{pmatrix}, \quad \mathbf{c}_A = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \quad \text{and} \quad C_B = (0 \ -1 \ -1), \quad \mathbf{c}_B = (-1).$$

Farkas coefficients are

$$\mathbf{k}_A^\top = (1 \ 1 \ 2) \quad \text{and} \quad \mathbf{k}_B^\top = (1),$$

while Farkas interpolant for  $(A, B)$  is the inequality  $I^f \equiv x_2 + x_3 \leq 0$ . However, if we decompose  $\mathbf{k}_A$  into

$$\mathbf{k}_{A,1}^\top = (1 \ 0 \ 1) \quad \text{and} \quad \mathbf{k}_{A,2}^\top = (0 \ 1 \ 1),$$

we obtain the decomposition  $\{x_2 \leq 0, x_3 \leq 0\}$  corresponding to the decomposed interpolant  $I^{DI} \equiv x_2 \leq 0 \wedge x_3 \leq 0$  of size two.

#### 4.1 Strength-Based Ordering of Decompositions

Here we define an order on decompositions for a single interpolation problem.

**Definition 2.** Let  $D_1, D_2$  denote two decompositions of the same interpolation problem of size  $m, n$ , respectively. Let  $(\mathbf{q}_1, \dots, \mathbf{q}_m)$  denote the decomposition of Farkas coefficients corresponding to  $D_1$  and let  $(\mathbf{r}_1, \dots, \mathbf{r}_n)$  denote the decomposition of Farkas coefficients corresponding to  $D_2$ . We say that decomposition  $D_1$  is finer than  $D_2$  (or equivalently  $D_2$  is coarser than  $D_1$ ) and denote this as  $D_1 \prec D_2$  when there exists a partition  $P = \{p_1, \dots, p_n\}$  of the set  $\{\mathbf{q}_1, \dots, \mathbf{q}_m\}$  such that for each  $i$  with  $1 \leq i \leq n$ ,  $\mathbf{r}_i = \sum_{\mathbf{q} \in p_i} \mathbf{q}$ .

**Lemma 1.** Assume  $D_1, D_2$  are two decompositions of the same interpolation problem such that  $D_1 \prec D_2$ . Let  $I^{D_1}, I^{D_2}$  be the decomposed interpolants corresponding to  $D_1, D_2$ . Then  $I^{D_1}$  implies  $I^{D_2}$ .

*Proof.* Informally, the implication follows from the fact that each linear inequality of  $I^{D_2}$  is a sum of some inequalities in  $I^{D_1}$ .

Formally, let  $I_i$  denote the  $i$ -th inequality in  $I^{D_2}$ . Then  $I_i \equiv \mathbf{r}_i^\top C_A \mathbf{x} \leq \mathbf{r}_i^\top \mathbf{c}_A$ . Since  $D_1 \prec D_2$ , there is a set  $\{I_{i_1}, \dots, I_{i_j}\} \subseteq D_1$  such that for each  $k$  with  $1 \leq k \leq j$ ,  $I_{i_k} \equiv \mathbf{q}_{i_k}^\top C_A \mathbf{x} \leq \mathbf{q}_{i_k}^\top \mathbf{c}_A$  and  $\mathbf{r}_i = \sum_{k=1}^j \mathbf{q}_{i_k}$ .

Since  $\mathbf{q}_{i_k} \geq \mathbf{0}$ , it holds that  $I_{i_1} \wedge \dots \wedge I_{i_j} \implies I_i$ . This means that  $I^{D_1}$  implies every conjunct of  $I^{D_2}$ .  $\square$

Note that the trivial, single-element decomposition corresponding to Farkas interpolant is the greatest element of this decomposition ordering. Also, for any decomposition of size more than one, replacing any number of elements by their sum yields a coarser decomposition. A possible reason to use a coarser decomposition may be that summing up some of the elements of a decomposition may result in eliminating a shared variable from the decomposition.

## 4.2 Duality

Let  $Itp$  denote an interpolation procedure and let  $Itp(A, B)$  stand for the interpolant computed by  $Itp$  for an interpolation problem  $(A, B)$ . Then by  $Itp'$  we denote the *dual* interpolation procedure, which works as follows:  $Itp'(A, B) = \neg Itp(B, A)$ . The duality theorem for interpolation states that  $Itp'$  is correct interpolation procedure. This can be shown by verifying that the three interpolation conditions hold for  $Itp'(A, B)$ , given they hold for  $Itp(B, A)$ .

Let us denote the interpolation procedure based on Farkas' lemma as  $Itp_F$  and the interpolation procedure computing decomposed interpolants as  $Itp_{DI}$ . The relation between  $Itp_F$  and its dual  $Itp'_F$  has been established in [2], namely that  $Itp_F(A, B) \implies Itp'_F(A, B)$ . We have shown in Lemma 1 that decomposed interpolant always implies Farkas interpolant computed from the same Farkas coefficients. This means that  $Itp_{DI}(A, B) \implies Itp_F(A, B)$ .

We can use this result to establish similar result for the dual interpolation procedures. Since  $Itp_{DI}(B, A) \implies Itp_F(B, A)$ , it follows that  $\neg Itp_F(B, A) \implies \neg Itp_{DI}(B, A)$  and consequently  $Itp'_F(A, B) \implies Itp'_{DI}(A, B)$ .

Putting all the results on logical strength together, we obtain

$$Itp_{DI}(A, B) \implies Itp_F(A, B) \implies Itp'_F(A, B) \implies Itp'_{DI}(A, B).$$



Note that while both  $Itp_F$  and  $Itp'_F$  produce interpolants which are a single inequality and interpolants produced by  $Itp_{DI}$  are *conjunctions* of inequalities, interpolants produced by  $Itp'_{DI}$  are *disjunctions* of inequalities.

In the following section, we describe the details of the  $Itp_{DI}$  interpolation procedure.

## 5 Finding Decompositions

In this section we present our approach for finding decompositions for linear arithmetic interpolation problems given their Farkas coefficients.

We focus on the task of finding decomposition of  $\mathbf{k}_A^\top C_A \mathbf{x}$ . Recall that  $C_A \in \mathbb{Q}^{l \times n}$  and  $\mathbf{x}$  is a vector of variables of length  $n$ . Without loss of generality assume that there are no  $B$ -local variables since columns of  $C_A$  corresponding to  $B$ -local variables would contain all zeroes by definition in any case.

Furthermore, without loss of generality, assume the variables in the inequalities of  $A$  are ordered such that all  $A$ -local variables are before the shared ones. Then let us write

$$C_A = (L \ S), \quad \mathbf{x}^\top = (\mathbf{x}_L^\top \ \mathbf{x}_S^\top) \quad (5)$$

with  $\mathbf{x}_L$  the vector of  $A$ -local variables of size  $p$ ,  $\mathbf{x}_S$  the vector of shared variables of size  $q$ ,  $n = p + q$ ,  $L \in \mathbb{Q}^{l \times p}$  and  $S \in \mathbb{Q}^{l \times q}$ . We know that  $\mathbf{k}_A^\top L = \mathbf{0}$  and the goal is to find  $\mathbf{k}_{A,i}$  such that  $\sum_i \mathbf{k}_{A,i} = \mathbf{k}_A$  and for each  $i$   $\mathbf{0} \leq \mathbf{k}_{A,i} \leq \mathbf{k}_A$  and  $\mathbf{k}_{A,i}^\top L = \mathbf{0}$ .

In the following we will consider two cases for computing the decompositions. We first study a common special case where the system  $A$  contains rows with no local variables, and give a linear-time algorithm for computing the decompositions. We then move to the general case where the rows of  $A$  contain local variables, and provide a decomposition algorithm based on computing a vector basis for a null space of a matrix obtained from  $A$ .

### 5.1 Trivial Elements

First, consider a situation where there is a linear inequality with no local variables. This means there is a row  $j$  in  $C_A$  (denoted as  $C_{A_j}$ ) such that all entries in columns corresponding to local variables are 0, i.e.,  $L_j = \mathbf{0}^\top$ . Then  $\{I_1, I_2\}$  for  $\mathbf{k}_{A,1} = \mathbf{k}_A[j] \times \mathbf{e}_j$  and  $\mathbf{k}_{A,2} = \mathbf{k}_A - \mathbf{k}_{A,1}$  is a decomposition. Intuitively, any linear inequality that contains only shared variables can form a stand-alone element of a decomposition. When looking for finest decomposition, we do this iteratively for all inequalities with no local variables. In the next part we show how to look for a non-trivial decomposition when dealing with local variables.

### 5.2 Decomposing in the Presence of Local Variables

For this section, assume that  $L$  has no zero rows (we have shown how to deal with such rows in 5.1). We are going to search for a non-trivial decomposition starting with the following observation:

<p><b>input</b> : matrix <math>M</math>, vector <math>\mathbf{v}</math> such that <math>\mathbf{v} \in \ker(M)</math> and <math>\mathbf{v} &gt; \mathbf{0}</math></p> <p><b>output</b>: <math>(\mathbf{w}_1, \dots, \mathbf{w}_m)</math>, a decomposition of <math>\mathbf{v}</math>, such that <math>\mathbf{w}_i \in \ker(M)</math>, <math>\mathbf{w}_i \geq \mathbf{0}</math> and <math>\sum \mathbf{w}_i = \mathbf{v}</math></p> <p>1 <math>n \leftarrow \text{Nullity}(M)</math></p> <p>2 <b>if</b> <math>n \leq 1</math> <b>then return</b> <math>(\mathbf{v})</math></p> <p>3 <math>(\mathbf{b}_1, \dots, \mathbf{b}_n) \leftarrow \text{KernelBasis}(M)</math></p> <p>4 <math>(\alpha_1, \dots, \alpha_n) \leftarrow \text{Coordinates}(\mathbf{v}, (\mathbf{b}_1, \dots, \mathbf{b}_n))</math></p> <p>5 <math>(\mathbf{w}_1, \dots, \mathbf{w}_n) \leftarrow (\alpha_1 \mathbf{b}_1, \dots, \alpha_n \mathbf{b}_n)</math></p> <p>6 <b>if</b> <math>\mathbf{w}_i \geq \mathbf{0}</math> for each <math>i</math> <b>then return</b> <math>(\mathbf{w}_1, \dots, \mathbf{w}_n)</math></p> <p>7 <b>else return</b> <math>(\mathbf{v})</math></p>
---

**Algorithm 1:** Algorithm for decomposition of Farkas coefficients

**Observation.**  $\mathbf{k}_A^\top L = 0$ . Equivalently, there are no  $A$ -local variables in the Farkas interpolant. It follows that  $L^\top \mathbf{k}_A = 0$  and  $\mathbf{k}_A$  is in the kernel of matrix  $L^\top$ .

Let us denote by  $\mathbb{K} = \ker(L^\top)$  the kernel of  $L^\top$ .

**Theorem 2.** Let  $\mathbf{v}_1, \dots, \mathbf{v}_n$  be  $n$  linearly independent vectors from  $\mathbb{K}$  such that  $\exists \alpha_1, \dots, \alpha_n$  with  $\alpha_i \mathbf{v}_i \geq \mathbf{0}$  for all  $i$  and  $\alpha_1 \mathbf{v}_1 + \dots + \alpha_n \mathbf{v}_n = \mathbf{k}_A$ . Then  $\{\mathbf{w}_1, \dots, \mathbf{w}_n\}$  for  $\mathbf{w}_i = \alpha_i \mathbf{v}_i$  is a decomposition of  $\mathbf{k}_A$  and  $\{I_1, \dots, I_n\}$  for  $I_i \equiv \mathbf{w}_i^\top C_A \mathbf{x} \leq \mathbf{c}_A$  is a decomposition.

*Proof.* The theorem follows from the definition of decomposition (Def. 1). From the assumptions of the theorem we immediately obtain  $\mathbf{k}_A = \sum_{i=1}^n \mathbf{w}_i$  and  $\mathbf{w}_i \geq \mathbf{0}$ . Moreover,  $\mathbf{w}_i \in \mathbb{K}$ , since  $\mathbf{v}_i \in \mathbb{K}$  and  $\mathbf{w}_i = \alpha_i \mathbf{v}_i$ . As a consequence,  $L^\top \mathbf{w}_i = 0$  and it follows that there are no  $A$ -local variables in  $\mathbf{w}_i^\top C_A \mathbf{x}$ .  $\square$

**Corollary 1.** Let  $B$  denote a basis of  $\mathbb{K}$  such that  $\forall \mathbf{w} \in B$ ,  $\mathbf{w} \geq \mathbf{0}$  and  $\mathbf{k}_A$  belongs to the conical hull of  $B$ . Then  $B$  forms a decomposition. In addition, it is the finest decomposition in the sense that it is the minimal element of the decomposition order  $\prec$ .

Based on Corollary 1, we search for decomposition of Farkas coefficients  $\mathbf{k}_A$  by computing basis of the kernel of the matrix of  $A$ -local variables  $L$ , as in Algorithm 1.

Function `Nullity` returns the dimension of the kernel. This can be efficiently computed for example using *Rank-Nullity Theorem* by computing Row Echelon Form of  $M$  by Gaussian elimination. Only if nullity is at least 2, we can hope to find any non-trivial decomposition.

Function `KernelBasis` returns a basis of the kernel of given matrix while function `Coordinates` returns the coordinates of the given vector with respect to the given basis. If the components of the linear combination are all non-negative, it is a true decomposition, otherwise, it cannot be used and we fall back to the trivial decomposition leading to the original Farkas interpolant. An algorithm to compute a basis of the kernel of a matrix can be found in any good introductory book on Linear Algebra, see e.g. [4].

We noticed that sometimes the computed basis is not good, but it can be modified to get a suitable basis and we implemented a simple heuristic to do this.

## 6 Experiments

We have implemented our algorithm in our SMT solver OPENSMT [19], which had already provided a variety of interpolation algorithms for propositional logic [31,20], theory of uninterpreted functions [3] and theory of linear real arithmetic [2]. We implemented both primal and dual versions of decomposed interpolation algorithm, which return the finest decomposition they can find.

We evaluated the effect of decomposed interpolants in a model-checking scenario using the model checker SALLY, which relies on OPENSMT for interpolation.<sup>5</sup> The PDKIND engine of SALLY was used, relying on YICES [14] for satisfiability queries and OPENSMT for interpolation queries. We experimented with four LRA interpolation algorithms: the original interpolation algorithms based on Farkas’ lemma,  $Itp_F$  and  $Itp'_F$ , and the interpolation algorithm computing decomposed interpolants,  $Itp_{DI}$  and  $Itp'_{DI}$ . In each case, we used McMillan’s interpolation rules [26] for the Boolean part. For comparison, we ran also a version of SALLY using MATHSAT in its default settings as an interpolation engine instead of OPENSMT. Since OPENSMT does not support the combination of incremental functionality and interpolation, SALLY was also using MATHSAT in non-incremental mode in this setting. The results are summarised in Figures 1 and 2, and Table 1. The result of a portfolio is the virtual best of the results of individual algorithms in the portfolio. The original portfolio of OPENSMT consists of  $Itp_F$  and  $Itp'_F$ , while in the new portfolio  $Itp_{DI}$  and  $Itp'_{DI}$  are added.

We used the same benchmarks as in [34]. They consist of several problem sets related to fault-tolerant algorithms (**om**, **ttesynchro**, **ttastartup**, **unifapprox**, **azadmanesh**, **approxagree**, **hacms**, **misc**), benchmarks from software model checking (**cav12**, **ctigar**), benchmarks from the benchmark suite of KIND model checker (**lustre**), simple concurrent programs (**conc**), and problems modelling a lock-free hash table (**lfht**). Each benchmark is a transition system defined by formulas characterizing initial states and the transition relation. In addition, each benchmark defines a property that should hold for the system. SALLY can finish with two possible answers: *valid* means the property holds and a safe invariant implying the property has been found; *invalid* means the property does not hold and a counterexample leading to a state where the property does not hold has been found. In the plots, we denote the answers as + and o, respectively.

The benchmarks were run on Linux machines with Intel E5-2650 v3 processor (2.3 GHz) with 64GB of available memory. Each benchmark was restricted to 600 seconds of running time and to 4GB of memory.

<sup>5</sup> Detailed description of the set-up and specifications of the experiments, together with all the results, can be found at <http://verify.inf.usi.ch/content/decomposed-interpolants>

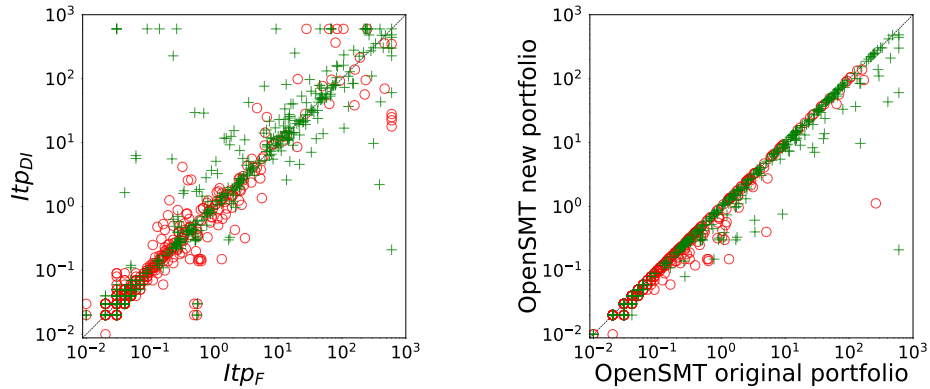


Fig. 1: Evaluation of the decomposed interpolants in model checking scenario. On the left, comparison of performance of SALLY using OPENSMT with different interpolation procedures,  $Itp_F$  and  $Itp_{DI}$ . On the right, the benefit of adding  $Itp_{DI}$  and  $Itp'_{DI}$  to the portfolio of interpolation procedures.

Figure 1 illustrates the benefit of adding  $Itp_{DI}$  and  $Itp'_{DI}$  to the portfolio of OPENSMT interpolation algorithms. The direct comparison of  $Itp_F$  and  $Itp_{DI}$  clearly shows that in many cases, the use of decomposed interpolants outperforms the use of original interpolation procedure, sometimes by an order of magnitude. The comparison of the capabilities of the old and the new portfolio shows that the importance of decomposition is still significant even after taking the capabilities of dual versions into account.

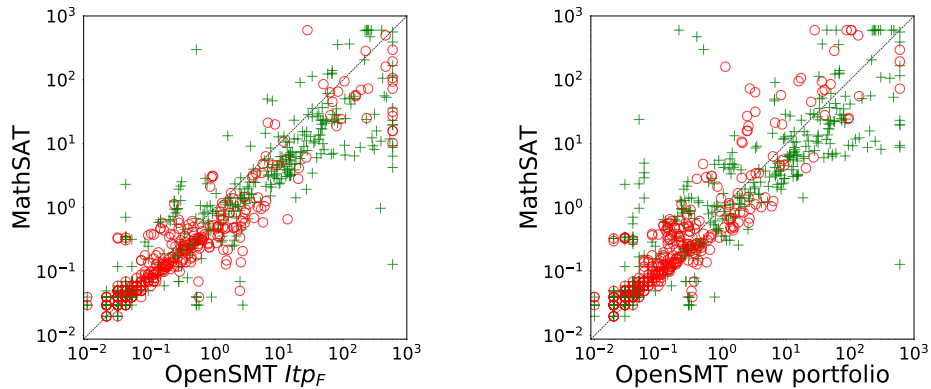


Fig. 2: Comparison of performance for the use of MATHSAT interpolation procedure and the use of OPENSMT interpolation procedure—original  $Itp_F$  and the whole portfolio, respectively.

Table 1: Performance of SALLY with old and new OPENSMT interpolation capabilities. Comparison with MATHSAT with its default interpolation included.

Problem set	OPENSMT portfolio $Itp_F, Itp'_F$		OPENSMT portfolio $Itp_F, Itp'_F, Itp_{DI}, Itp'_{DI}$		MATHSAT	
	solved (V/I)	$\Sigma$ time(s)	solved (V/I)	$\Sigma$ time(s)	solved (V/I)	$\Sigma$ time(s)
approxagree (9)	9 (8/1)	101	9 (8/1)	72	9 (8/1)	173
azadmanesh (20)	16 (13/3)	74	16 (13/3)	69	19 (16/3)	102
cav12 (99)	63 (46/17)	3,427	63 (46/17)	2,960	64 (47/17)	796
conc (6)	3 (3/0)	48	4 (4/0)	347	3 (3/0)	38
ctigar (110)	70 (51/19)	1,812	72 (53/19)	1,493	75 (55/20)	1,803
hacms (5)	1 (1/0)	147	1 (1/0)	84	1 (1/0)	55
lfht (27)	17 (17/0)	502	17 (17/0)	502	16 (16/0)	518
lustre (790)	757 (423/334)	5,122	759 (425/334)	4,903	752 (420/332)	5,610
misc (10)	7 (6/1)	80	7 (6/1)	80	7 (6/1)	36
om (9)	9 (7/2)	7	9 (7/2)	7	9 (7/2)	6
ttastartup (3)	1 (1/0)	2	1 (1/0)	2	1 (1/0)	13
ttesynchro (6)	6 (3/3)	11	6 (3/3)	10	6 (3/3)	6
unifapprox (11)	10 (7/3)	21	10 (7/3)	20	11 (8/3)	125
	969 (586/383)	11,354	974 (591/383)	10,549	973 (591/382)	9,281

Fig. 2 shows the benefit of the new portfolio by comparing the model checker performance to one using a different SMT solver. As far as we know, MATHSAT also computes interpolants from the proof of unsatisfiability and uses interpolation algorithm based on Farkas' lemma for LRA conflicts. Comparing to OPENSMT's  $Itp_F$ , we see that most probably due to the differences in the implementation of the SMT solver, the version of SALLY using MATHSAT is superior. However, using the portfolio of interpolation procedures available in OPENSMT bridges the gap and allows SALLY to solve more benchmarks as can be seen in Table 1. This also shows a potential improvement for MATHSAT if it would offer the same portfolio of interpolation procedures as OPENSMT does.

Table 1 demonstrates the gain in the performance of the model checker from expanding the interpolation portfolio to include  $Itp_{DI}$  and  $Itp'_{DI}$ . The results are summarised *by category* with first column showing the name of the category and the number of benchmarks in this category, and the two columns per interpolation engine show how many benchmarks were successfully solved (validated/invalidated) within the limits and the total running time spent on *solved* benchmarks. With the new interpolation algorithm added to the portfolio not only is the model checker able to solve *more* instances, but it is also able to do so in *less* time.

Table 2 answers the question how often in these model-checking benchmarks the new interpolation procedure manages to decompose the Farkas coefficients, thus returning a different interpolant than the original procedure would.

The statistics differ for  $Itp_{DI}$  and  $Itp'_{DI}$  due to the special nature of the interpolation problems in the model checking algorithm from IC3 family, where the  $B$ -part of the interpolation problem *always* contains only shared symbols.

Table 2: Interpolation statistics – pwd stands for “Number of problems with at least one decomposition”. The numbers in parentheses denote “Decompositions with trivial and with non-trivial elements” (trivial/non-trivial).

Problem set	$Itp_{DI}$			$Itp'_{DI}$		
	pwd	#non-triv. LRA itps	#decomp. itps	pwd	#non-triv. LRA itps	#decomp. itps
approxagree (9)	1 (1/0)	7	7 (7/0)	1	18	18
azadmanesh (20)	4 (0/4)	4,831	266 (0/266)	4	4,353	4,353
cav12 (99)	31 (25/15)	1,368,187	7,399 (1,690/5,738)	45	204,036	57,127
conc (6)	3 (3/3)	424,145	215,376 (1,256/214,120)	3	13	13
ctigar (110)	73 (56/70)	2,982,559	826,621 (29,378/797,871)	77	152,613	152,612
hacms (5)	5 (5/5)	363,265	15,282 (532/14,750)	5	58,416	58,416
lftt (27)	13 (12/13)	838,094	12,785 (169/12,616)	14	111,060	111,060
lustre (790)	356 (356/192)	2,571,091	1,851,213 (855,958/1,054,516)	500	1,833,310	1,833,310
misc (10)	5 (4/5)	195,819	62,865 (8,700/55,042)	6	35,131	35,108
om (9)	4 (4/3)	1,150	236 (206/30)	3	168	168
ttastartup (3)	2 (2/2)	69,699	924 (16/908)	3	11,528	11,528
ttesynchro (6)	4 (4/0)	64	38 (38/0)	5	310	310
unifapprox (11)	0 (0/0)	0	0 (0/0)	2	25	25

Theoretically, this means there is no chance for  $Itp'_{DI}$  to discover any non-trivial elements of decomposition as there are no  $B$ -local variables. On the other hand, this means that decomposition to trivial elements is always possible, as all  $B$ -inequalities contain only shared variables. In our implementation, however, we consider the locality of a variable not from the global point of the whole interpolation problem, but from of the local point of current theory conflict. As a consequence, even if a variable is shared in the whole problem, it is sometimes considered local in the interpolation problem for the current theory conflict and the interpolant is not decomposed even if, from a global point of view, it could have been.

For  $Itp_{DI}$ , we show first in how many benchmarks there was at least a *single* decomposition (any vs. with at least one trivial element vs. with at least one non-trivial element). Then, we counted how many interpolation problems for theory conflict occurred during solving and removed those where there was not even theoretical possibility for decomposition. These include the problems where all inequalities were from one part of the problem (resulting in trivial interpolants, either  $\top$  or  $\perp$ ) and also the problems with a single inequality in the  $A$ -part (trivially yielding an interpolant equal to that inequality). After subtracting these, we got the numbers in the column “#non-triv. LRA itps”. Finally, the last column reports how many of these non-trivial interpolants were actually successfully decomposed (with at least one trivial element vs. with at least one non-trivial element). Note that it can happen that a successful decomposi-

tion contains both trivial and non-trivial elements. For  $Itp'_{DI}$  we leave out the statistics for decomposition with non-trivial elements, since they occur in these benchmarks very rarely only. We see that at least one decomposition was possible in only roughly half of all the benchmarks. This explains why there are many points on (or extremely close to) the diagonal in Fig. 1. On the other hand, it shows that the test for the *possibility* of decomposition is very cheap and does not present a significant overhead. Another conclusion we can draw is that when the structure of the benchmark allows decomposition, the decomposition can often be discovered in many of the interpolation problems for theory conflicts that appear during the solving.

During the evaluation, we noticed that a small change in the solver sometimes had a huge effect on the performance of the model checker for a particular benchmark. It made previously unsolved instance solvable in tens of seconds or the other way around. To confirm that on some benchmarks  $Itp_{DI}$  is really better than  $Itp_F$ , we ran the model checker 100 times on chosen benchmarks, each time with a different random seed set for the interpolating solver.

We found out that benchmark **dillig03.c.mcmt** from category **ctigar** is exactly of the type where model checker using  $Itp_F$  does *not* converge (in all 100 runs) while  $Itp_{DI}$  ensures convergence in 0.2 seconds (in all 100 runs). Other benchmarks that  $Itp_F$  was not able to solve at all are **fib\_bench\_safe\_v1.mcmt** from category **conc** and **large\_const.c.mcmt** from category **ctigar**, while  $Itp_{DI}$  was able to solve them in 42 runs on average in 377 seconds, and in 80 runs on average in 97 seconds, respectively. Finally, the benchmark **DRAGON\_13.mcmt** from category **lustre** was solved by  $Itp_F$  in 5 runs on average in 539 seconds, while it was solved by  $Itp_{DI}$  in 23 runs on average in 441 seconds.

## 7 Conclusion

In this paper, we have presented a new interpolation algorithm for linear real arithmetic that generalizes the interpolation algorithm based on Farkas' lemma used in modern SMT solvers. We showed that the algorithm is able to compute interpolants in the form of a *conjunction* of inequalities that are logically stronger than the single inequality returned by the original approach. This is useful in the IC3-style model-checking algorithms where Farkas interpolants have been shown to be a source of incompleteness. In our experiments, we have demonstrated that the opportunity to decompose Farkas interpolants occurs frequently in practice and that the decomposition often leads to (i) shortening of solving time and, in some cases, to (ii) solving a problem not solvable by the previous approach.

As the next steps, we want to investigate how to automatically determine what kind of interpolant would be more useful for the current interpolation query in IC3-style model-checking algorithms. We also plan to investigate other uses of interpolation in model checking where stronger (or weaker) interpolants are desirable [32].

*Acknowledgements.* We would like to thank Dejan Jovanović for providing the benchmarks and for the help with integrating OPENSMT into SALLY.

## References

1. Albarghouthi, A., McMillan, K.L.: Beautiful interpolants. In: Proc. CAV 2013. LNCS, vol. 8044, pp. 313–329. Springer (2013)
2. Alt, L., Hyvärinen, A.E.J., Sharygina, N.: LRA interpolants from no man’s land. In: Strichman, O., Tzoref-Brill, R. (eds.) Proc. HVC 2017. LNCS, vol. 10629, pp. 195–210. Springer (2017)
3. Alt, L., Hyvärinen, A.E.J., Asadi, S., Sharygina, N.: Duality-based interpolation for quantifier-free equalities and uninterpreted functions. In: Stewart, D., Weissenbacher, G. (eds.) Proc. FMCAD 2017. pp. 39–46. IEEE (2017)
4. Andriilli, S., Hecker, D.: Chapter 5 - linear transformations. In: Andriilli, S., Hecker, D. (eds.) Elementary Linear Algebra (Fifth Edition), pp. 319 – 412. Academic Press, Boston, fifth edition edn. (2016). <https://doi.org/https://doi.org/10.1016/B978-0-12-800853-9.00005-0>, <http://www.sciencedirect.com/science/article/pii/B9780128008539000050>
5. Barrett, C., de Moura, L.M., Ranise, S., Stump, A., Tinelli, C.: The SMT-LIB initiative and the rise of SMT - (HVC 2010 award talk). In: Barner, S., Harris, I.G., Kroening, D., Raz, O. (eds.) Proc. HVC 2010. LNCS, vol. 6504, p. 3. Springer (2011)
6. Barrett, C., Sebastiani, R., Seshia, S., Tinelli, C.: Satisfiability modulo theories, Frontiers in Artificial Intelligence and Applications, vol. 185, pp. 825–885. 1 edn. (2009). <https://doi.org/10.3233/978-1-58603-929-5-825>
7. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D.A. (eds.) Proc. VMCAI 2011. LNCS, vol. 6538, pp. 70–87. Springer (2011)
8. Cimatti, A., Griggio, A., Schaafsma, B., Sebastiani, R.: The MathSAT5 SMT Solver. In: Piterman, N., Smolka, S. (eds.) Proceedings of TACAS. LNCS, vol. 7795. Springer (2013)
9. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *The Journal of Symbolic Logic* **22**(3), 269–285 (1957)
10. Davis, M., Logemann, G., Loveland, D.W.: A machine program for theorem-proving. *Commun. ACM* **5**(7), 394–397 (1962)
11. Davis, M., Putnam, H.: A computing procedure for quantification theory. *J. ACM* **7**(3), 201–215 (1960)
12. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: A theorem prover for program checking. *J. ACM* **52**(3), 365–473 (2005)
13. D’Silva, V., Kroening, D., Purandare, M., Weissenbacher, G.: Interpolant strength. In: Proc. VMCAI 2010. LNCS, vol. 5944, pp. 129–145. Springer (2010)
14. Dutertre, B.: Yices 2.2. In: Biere, A., Bloem, R. (eds.) Computer-Aided Verification (CAV’2014). Lecture Notes in Computer Science, vol. 8559, pp. 737–744. Springer (July 2014)
15. Dutertre, B., de Moura, L.M.: A fast linear-arithmetic solver for DPLL(T). In: Ball, T., Jones, R.B. (eds.) Proc. CAV 2006. LNCS, vol. 4144, pp. 81–94. Springer (2006)
16. Farkas, G.: A Fourier-féle mechanikai elv alkalmazásai (Hungarian) [On the applications of the mechanical principle of Fourier] (1894)
17. Gurfinkel, A., Kahsai, T., Komuravelli, A., Navas, J.A.: The SeaHorn verification framework. In: Proc. CAV 2015. LNCS, vol. 9206, pp. 343–361. Springer (2015). <https://doi.org/10.1007/978-3-319-21690-4>



18. Huberman, B.A., Lukose, R.M., Hogg, T.: An economics approach to hard computational problems. *Science* **275**(5296), 51 – 54 (1997)
19. Hyvärinen, A.E.J., Marescotti, M., Alt, L., Sharygina, N.: OpenSMT2: An SMT solver for multi-core and cloud computing. In: *Proc. SAT 2016*. LNCS, vol. 9710, pp. 547–553. Springer (2016)
20. Jančík, P., Alt, L., Fedyukovich, G., Hyvärinen, A.E.J., Kofron, J., Sharygina, N.: Pvair: Partial variable assignment interpolator. In: *FASE2016* (2016), <http://verify.inf.usi.ch/sites/default/files/main-2.pdf>
21. Jovanović, D., Dutertre, B.: Property-directed k-induction. In: *2016 Formal Methods in Computer-Aided Design (FMCAD)*. pp. 85–92 (Oct 2016). <https://doi.org/10.1109/FMCAD.2016.7886665>
22. Kenneth L. McMillan: Applications of Craig Interpolation in Model Checking. In: *Tools and Alg. for the Const. and Anal. of Systems (TACAS '05)*. LNCS, vol. 3440, pp. 1–12 (2005)
23. Komuravelli, A., Gurfinkel, A., Chaki, S.: SMT-based model checking for recursive programs. In: Biere, A., Bloem, R. (eds.) *Computer Aided Verification*. pp. 17–34. Springer International Publishing, Cham (2014)
24. Korovin, K., Tsiskaridze, N., Voronkov, A.: Conflict resolution. In: Gent, I.P. (ed.) *Principles and Practice of Constraint Programming - CP 2009*. pp. 509–523. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
25. Manna, Z., Pnueli, A.: Springer (1995)
26. McMillan, K.L.: Interpolation and SAT-based model checking. In: *Proc. CAV 2003*. LNCS, vol. 2725, pp. 1–13. SV (2003)
27. McMillan, K.L.: Interpolation and SAT-Based Model Checking. In: *Computer Aided Verification (CAV '03)*. LNCS, vol. 2725, pp. 1–13 (2003)
28. de Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems*. pp. 337–340. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
29. Nieuwenhuis, R., Oliveras, A.: Proof-producing congruence closure. In: *Proc. RTA 2005*. LNCS, vol. 3467, pp. 453–468. Springer (2005)
30. Pudlák, P.: Lower bounds for resolution and cutting plane proofs and monotone computations. *Journal of Symbolic Logic* **62**(3), 981–998 (1997)
31. Rollini, S.F., Alt, L., Fedyukovich, G., Hyvärinen, A.E.J., Sharygina, N.: PeRIPLO: A framework for producing effective interpolants in sat-based software verification. In: *Proc. LPAR 2013*. LNCS, vol. 8312, pp. 683–693. Springer (2013)
32. Rollini, S.F., Sery, O., Sharygina, N.: Leveraging interpolant strength in model checking. In: *Proc. CAV 2012*. LNCS, vol. 7358, pp. 193–209. Springer (2012)
33. Rybalchenko, A., Sofronie-Stokkermans, V.: Constraint solving for interpolation. In: *Proc. VMCAI 2007*. LNCS, vol. 4349, pp. 346–362. Springer (2007)
34. Schindler, T., Jovanović, D.: Selfless interpolation for infinite-state model checking. In: Dillig, I., Palsberg, J. (eds.) *Verification, Model Checking, and Abstract Interpretation*. pp. 495–515. Springer International Publishing, Cham (2018)
35. Scholl, C., Pigorsch, F., Disch, S., Althaus, E.: Simple interpolants for linear arithmetic. In: *2014 Design, Automation Test in Europe Conference Exhibition (DATE)*. pp. 1–6 (March 2014). <https://doi.org/10.7873/DATE.2014.128>
36. Schrijver, A.: *Theory of Linear and Integer Programming* (1998)