

Verifying a Solver for Linear Mixed Integer Arithmetic in Isabelle/HOL[★]

Ralph Bottesch¹, Max W. Haslbeck¹, Alban Reynaud², and René Thiemann¹

¹ University of Innsbruck, Austria

² ENS Lyon, France

Abstract. We implement a decision procedure for linear mixed integer arithmetic and formally verify its soundness in Isabelle/HOL. We further integrate this procedure into one application, namely into **CeTA**, a formally verified certifier to check untrusted termination proofs. This checking involves assertions of unsatisfiability of linear integer inequalities; previously, only a sufficient criterion for such checks was supported. To verify the soundness of the decision procedure, we first formalize the proof that every satisfiable set of linear integer inequalities also has a small solution, and give explicit upper bounds. To this end we mechanize several important theorems on linear programming, including statements on integrality and bounds. The procedure itself is then implemented as a branch-and-bound algorithm, and is available in several languages via Isabelle’s code generator. It internally relies upon an adapted version of an existing verified incremental simplex algorithm.

Keywords: branch-and-bound, Isabelle/HOL, linear programming, polyhedra, simplex algorithm

1 Introduction

The computational problem of deciding whether a system of linear inequalities with integer coefficients has an integral solution arises in many practical situations. Since it is NP-complete, no currently known algorithm can in general avoid searches of exponential length. Furthermore, while satisfiable instances always have short solutions that can be efficiently checked, there need not be short certificates for the fact that an instance has no solutions, unless $\text{NP} = \text{co-NP}$. (Contrast this with the related problem of deciding whether a system of linear inequalities with integer coefficients has a *rational* solution – this problem is in P, and Farkas’ lemma provides a short and efficiently checkable certificate that an unsatisfiable instance indeed has no solution.) Thus, if a solver declares that a given instance is unsatisfiable over the integers, the length of any certificate

[★] This research was supported by the Austrian Science Fund (FWF) project Y757. The authors are listed in alphabetical order regardless of individual contributions or seniority.

that allows this fact to be checked may be exponential in the size of the input instance, in which case the computational effort required to check such a certificate would be exponential as well.

Instead of repeatedly performing certification tasks that require immense amounts of data and computational effort, it may be more fruitful to formally verify the soundness of a solver *once*, so that it can then be trusted without instance-by-instance certification of its output. The implementation of such a solver, together with a formal proof of its soundness, is the goal of the present work. Specifically, we use Isabelle/HOL [17] to implement and prove the correctness of a branch-and-bound algorithm [19, Chapter 24.1], and then use Isabelle’s code generator [10] to obtain verified executable code. Along the way, we also give the first formalized proofs for several important results on integer programming.

A concrete example of an application for our solver comes from *termination analysis*, where a program is given as input to a termination tool that tries to determine whether the given program terminates on all inputs. Since termination tools get patched and improved repeatedly, maintaining an up-to-date formal proof of soundness would be extremely difficult. Therefore, the approach that is typically used is to have the (unverified) termination tool output a certificate for its analysis, which can then be checked by a verified certificate checker. One such certificate checker is **CeTA** [5,21]. It has been verified in Isabelle/HOL, so that whenever it accepts a proof of termination for some program, the formal proof of **CeTA**’s soundness ensures that the program does indeed terminate.

As an example, consider a program to compute the binary logarithm.

```
int log2(int x) {
  int n := 0;
  while (x > 0) {
    x := x div 2;
    n := n + 1;
  }
  return n;
}
```

This program can be translated into an integer transition system and termination can be proved by showing that the value of x is decreased by at least 1 after every loop iteration. This property can be expressed in linear integer arithmetic (LIA): it is equivalent to the validity of formula (1), where x' and n' represent the new values of x and n , respectively, after an iteration of the loop.

$$x > 0 \wedge 2x' \leq x \wedge x \leq 2x' + 1 \wedge n' = n + 1 \longrightarrow x \geq x' + 1 \quad (1)$$

Validity of (1) is equivalent to unsatisfiability of the negated formula, which is simply a conjunction of linear inequalities:

$$x > 0 \wedge 2x' \leq x \wedge x \leq 2x' + 1 \wedge n' = n + 1 \wedge x < x' + 1 \quad (2)$$

A *sufficient* condition for the unsatisfiability of (2) over the integers (LIA) is the unsatisfiability of the same system over the rationals (LRA); the latter can be shown, for instance, via the simplex algorithm [8]. Indeed, a verified implementation [20] of the simplex algorithm is currently integrated into **CeTA** [5]. However, whereas (2) is unsatisfiable over the integers, it has a rational solution $x = n' = 1$, $x' = \frac{1}{2}$, $n = 0$. For such examples, considering the problem over the rationals may prohibit **CeTA** from detecting unsatisfiability over the integers.

Therefore, in this paper we develop a verified theory solver for LIA (in fact, for linear *mixed* integer arithmetic, where only a user-specified part of the solution is required to be integral). The verified solver takes a conjunction of strict and non-strict linear inequalities as input, and decides whether they are simultaneously solvable. We fully integrate the LIA solver into **CeTA**, so that the new version can handle all instances that are unsatisfiable over the integers and not only those that are unsatisfiable over the rationals as well. Of course, the LIA solver can also be used as a stand-alone theory solver, e.g., to perform verified SMT solving.

We verify our LIA solver in two major steps.

1. First, we show that for every set of LIA constraints it suffices to search for small solutions. To this end, we formally verify an a priori bound in the style of Papadimitriou [18]: If there is an integer solution to a set of LIA constraints, then there is also one that is bounded by $b := n(ma)^{2m+1}$. To be more precise, the small solution satisfies $|x| \leq b$ for each variable x , where n is the number of variables, m the number of inequalities, and a the largest absolute value of any number occurring in the inequalities.

Our verified upper bound matches the one given in a textbook [19, Thm. 17.1] (which is considerably lower than the one by Papadimitriou).³ Specifically, we establish a bound of $(n+1)!a^n$ (with no dependence on m). To prove this bound in Isabelle/HOL we mostly follow the textbook proofs and formalize several important results from linear programming, often with additional statements on bounds and integrality. These results include: the fundamental theorem of linear inequalities, the Farkas–Minkowski–Weyl theorem, Carathéodory’s theorem, and the decomposition theorem for polyhedra. Note that the bound on the size of solutions also implies the fact that the problem of deciding satisfiability for linear integer inequalities is in NP.

2. Using the upper bound, we can decide satisfiability via a finite search. For instance, for formula (2) we have $n = 4$, $a = 2$ and $m = 6$ (the equality counts as two inequalities), and we know that if (2) is satisfiable, then there is an integer solution with absolute values at most 1920.

To perform this search, we implement and verify a basic branch-and-bound algorithm. It is based on an incremental version of the simplex algorithm by Dutertre and de Moura [9], which is used to deliver candidate solutions and to prune the search tree by detecting unsatisfiability in LRA. Although the incremental simplex algorithm has recently been verified in Isabelle/HOL [3],

³ The textbook bound is somewhat more precise than ours, as it is phrased in terms of sub-determinants, whereas we use a generic bound on sub-determinants.

its integration into the branch-and-bound algorithm is not immediate: the branch-and-bound algorithm requires frequent updates of bounds on variables, and this operation is not supported by the existing verified incremental simplex algorithm.

Note that our verified LIA solver is missing several possible optimizations [6,7,13], some of which might be integrated in future work. Therefore, it clearly cannot compete with state-of-the-art (unverified) solvers. Still, our experimental results show that there are some examples from SMT-LIB where our solver is successful, but both CVC4 [2] and Z3 [15] fail.

Structure We give some preliminaries on linear (integer) programming and Isabelle in Section 2. Afterwards, we present our formalization of linear programming and the mentioned bound in Section 3. The branch-and-bound algorithm with the adaptation of the incremental simplex algorithm are covered in Section 4. We provide experimental results in Section 5 and conclude in Section 6.

The collection of theorems on polyhedra and small solutions is available as part of the archive of formal proofs (AFP) in the entry on linear inequalities [4], and the branch-and-bound algorithm is part of *IsaFoR/CeTA* [21]. All of the theorems of this paper are linked to the formalization on an accompanying website. It also provides details on the experiments.

<http://cl-informatik.uibk.ac.at/software/ceta/experiments/lia/>

Related work Allamigeon and Katz [1] have implemented the simplex algorithm in Coq and used it to give constructive proofs of a number of important theorems about convex polyhedra. The overlap between our work and [1] consists of formalizations of basic facts concerning cones and polyhedra, the fundamental theorem of linear inequalities, and Farkas' lemma. However, whereas in [1] a simplex algorithm for optimization problems is implemented in order to be used in constructive mathematical proofs, we formalize theorems concerning integer programming, including bounds on the size of solutions, and use these together with the previously Isabelle-verified simplex algorithm to obtain formally verified, yet efficient, software.

There is also a formalization of theorems about polyhedra in HOL Light, due to Harrison [11], but it contains neither a formalization of the simplex algorithm nor does it cover integer programming.

Cooper's algorithm has been formalized by Nipkow [16] in Isabelle/HOL. Although this algorithm also solves linear integer arithmetic, it internally works completely differently and its formalization requires different proofs; therefore, we do not see any overlap between the two works. We nevertheless consider the verified version of Cooper's algorithm in our experiments.

2 Preliminaries

We briefly review some linear programming and Isabelle background.

2.1 Linear programming

We assume familiarity with vector spaces. Although our Isabelle theorems use a more general type, here we present our results in the context of Euclidean spaces (\mathbb{R}^n). We denote the usual inner product in \mathbb{R}^n by ‘ \cdot ’.

A (*non-strict*) *linear inequality* is an inequality of the form $a \cdot x \leq b$, where $a, x \in \mathbb{R}^n$ (a a row vector, x a column vector) and $b \in \mathbb{R}$. A system of linear inequalities can therefore be written as $Ax \leq b$, with $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ a column vector. A system of linear inequalities is a *mixed integer system* if, for some $I \subseteq \{1, \dots, n\}$, it is required that $x_i \in \mathbb{Z}$ for all $i \in I$. We also define *strict* linear inequalities to be inequalities of the form $ax < b$, with a, x and b as before.

In this work we consider mixed integer systems of linear inequalities containing both non-strict and strict inequalities.

For reference we collect below the definitions of several important concepts from linear algebra that are needed in order to state the theorems that we formalize. These definitions can also be found in [19, Chapters 7.1–2 and 16.2].

Definition 1 (Half-spaces, hyperplanes, polyhedra). For $c \in \mathbb{R}^n \setminus \{0_n\}$ (a row vector) and $d \in \mathbb{R}$, we say that the set $H = \{x \mid c \cdot x \leq d\}$ is an affine half-space, and that c is its normal vector. If $d = 0$, then H is called a linear half-space (or just a half-space). The set $\{x \mid c \cdot x = 0\}$ is called a hyperplane (of which c is a normal vector).

A set $P \subseteq \mathbb{R}^n$ is called a (convex) polyhedron if $P = \{x \mid Ax \leq b\}$, for some matrix $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. In words, a polyhedron is the intersection of a finite collection of affine halfspaces.

Definition 2 (Cones). A non-empty set $C \subseteq \mathbb{R}^n$ is a cone if, for all $x, y \in C$ and $\lambda, \mu \geq 0$, we have $\lambda x + \mu y \in C$. A cone C is generated by the set of vectors X if $C = \{\lambda_1 v_1 + \dots + \lambda_m v_m \mid \lambda_1, \dots, \lambda_m \geq 0, \{v_1, \dots, v_m\} \subseteq X\}$, and C is finitely generated if it is generated by a finite set of vectors. A cone is polyhedral if it is the intersection of finitely many (linear) half-spaces.

Definition 3 (Convex hull, polytopes, integer hull). The convex hull of a vector set X is the set of all convex linear combinations of vectors from X . More precisely,

$$\text{conv.hull } X = \{\lambda_1 v_1 + \dots + \lambda_m v_m \mid \lambda_1, \dots, \lambda_m \geq 0, \sum \lambda_i = 1, \{v_1, \dots, v_m\} \subseteq X\}$$

The convex hull of a finite set of vectors is called a (convex) polytope.

Finally, if P is a polyhedron, then the integer hull of P , denoted P_I , is the convex hull of the set of integral vectors of P . (Integral vectors are vectors whose coordinates with respect to the standard basis are integers.)

2.2 Isabelle

For our formalization work we use the theorem prover Isabelle. Knowledge of Isabelle will be helpful, but is not necessary in order to read the paper, as we

have tried to make the formal source listings accessible even to a reader with a purely mathematical background.

Nevertheless, we briefly explain the meaning of some important notation here. First, we have $\text{carrier_vec } n = \mathbb{R}^n$, $\text{carrier_mat } m \ n = \mathbb{R}^{m \times n}$, and denote the zero-vector of dimension n by 0_n . Often, the statement that a vector or a matrix has a certain property will be expressed as membership in the set of all vectors or matrices with that property: $\text{Bounded_vec } bnd$ is the set of vectors (of finite dimension) with entries bounded in absolute value by bnd (similarly $\text{Bounded_mat } bnd$), $\text{indexed_Ints_vec } I$ is the set of vectors v with $v_i \in \mathbb{Z}$ for all $i \in I$, and, finally, \mathbb{Z}_v is the set of vectors (of finite dimension) with integer entries (similarly \mathbb{Z}_m , is a set of matrices). We also have notation for sets defined by some set of vectors or by a matrix: $\text{finite_cone } X$ denotes the cone generated by the finite set X ; other examples are $\text{cone } X$, $\text{polyhedral_cone } A$ and $\text{polyhedron } A \ b$, all with the obvious meanings.

3 Mixed-Integer Linear Problems

3.1 The main formalized theorems

We discuss our formalization of several results that are needed in order to formally prove the soundness of a branch-and-bound-based solver for mixed-integer linear systems of inequalities. The main theorem for this purpose states that if a mixed integer system of linear inequalities can be described using only integers, then it has a solution if and only if it also has a solution involving only numbers of bounded size.

```
lemma small_mixed_integer_solution:
  assumes  $A_1 \in \text{carrier\_mat } nr_1 \ n$  and ...
    and  $\{A_1, A_2\} \subseteq \mathbb{Z}_m \cap \text{Bounded\_mat } Bnd$ 
    and  $\{b_1, b_2\} \subseteq \mathbb{Z}_v \cap \text{Bounded\_vec } Bnd$ 
    and  $x \in \text{indexed\_Ints\_vec } I$ 
    and  $A_1 x \leq b_1$  and  $A_2 x < b_2$ 
  shows  $\exists x. x \in \text{carrier\_vec } n \wedge x \in \text{indexed\_Ints\_vec } I$ 
     $\wedge A_1 x \leq b_1 \wedge A_2 x < b_2 \wedge x \in \text{Bounded\_vec } (\text{fact } (n+1) * (\max 1 Bnd) ^ n)$ 
```

In order to derive this result, we require formalizations of several results from the theory of linear inequalities, beginning with the fundamental theorem of linear inequalities (FTLI). This theorem states that for any finite set of vectors A and vector b , either b is in the cone generated by a subset of A , or there exists a hyperplane $\{x \mid c \cdot x = 0\}$ separating b from A and containing some number of vectors of A .

```
lemma fundamental_theorem_of_linear_inequalities:4
```

⁴ The Isabelle statement given here matches the presentation of the theorem in [19]; in our formalization, the equivalence is written as an equality of sets.

```

assumes  $A \subseteq \text{carrier\_vec } n$  and  $\text{finite } A$  and  $b \in \text{carrier\_vec } n$ 
shows  $(\exists B. B \subseteq A \wedge b \in \text{finite\_cone } B \wedge \text{lin\_indpt } B) \longleftrightarrow$ 
 $\neg (\exists c. B. c \in \text{carrier\_vec } n \wedge B \subseteq A \wedge |B|+1 = \text{dim\_span } (A \cup \{b\}) \wedge$ 
 $\text{lin\_indpt } B \wedge (\forall a \in B. c \cdot a = 0) \wedge (\forall a \in A. c \cdot a \geq 0) \wedge c \cdot b < 0)$ 

```

To prove the theorem, one first considers an algorithm that iteratively applies a procedure that takes a subset of vectors from A and produces either the cone containing b from the theorem statement, or the separating hyperplane, or a new set of vectors from A . In case of the third outcome, the output set is used as the input for the next iteration. Thus, starting from some valid set of vectors, the above algorithm either never terminates (if the third outcome occurs in every iteration), or it produces an object satisfying the theorem statement. The proof is completed by showing that an infinite loop cannot occur.

The above argument could in principle be formalized in Isabelle by defining a function that incorporates the algorithm, and then proving that the function is well-defined (which implies the termination of the algorithm on all inputs). However, we are really only interested in the algorithm's termination; the fact that some input is mapped to a certain output is irrelevant for the proof of the theorem. Furthermore, we only need that the algorithm terminates when the set of input vectors is valid (i.e., of the right cardinality and linearly independent), but, due to the limitations of the Isabelle function-package [12], the domain of a function cannot be restricted in this manner. Consequently, we formalize the proof without modeling the algorithm as an Isabelle function. Instead, we define a relation on pairs of valid subsets of A : The pair (J', J) is in the relation if and only if, starting with J as input, one iteration of the algorithm produces output J' . In other words, the relation encodes all iterations of the algorithm where the third outcome occurs. Since A is finite, termination is equivalent to the fact that the above relation has no cycles. The latter fact is established by a proof by contradiction (here, our formalization closely follows the textbook proof [19, Chapter 7.1]).

We also need to formalize three corollaries of the FTLI. First, we have a theorem of Carathéodory that follows directly.

```

lemma Caratheodory_theorem: assumes  $A \subseteq \text{carrier\_vec } n$ 
shows  $\text{cone } A = \bigcup \{\text{finite\_cone } B \mid B. B \subseteq A \wedge \text{lin\_indpt } B\}$ 

```

Next, we have the Farkas-Minkowski-Weyl (FMW) theorem, which states that a cone is polyhedral if and only if it is finitely generated.

```

lemma farkas_minkowsky_weyl_theorem:
 $(\exists X. X \subseteq \text{carrier\_vec } n \wedge \text{finite } X \wedge P = \text{cone } X) \longleftrightarrow$ 
 $(\exists A \text{ nr. } A \in \text{carrier\_mat } \text{nr } n \wedge P = \text{polyhedral\_cone } A)$ 

```

The proofs of the FMW theorem and of the FTLI in [19] contain some simplifying assumptions that can be made without loss of generality. Of course, in Isabelle we must provide the full details of every proof, which often entails a

non-trivial amount of additional formalization work. For example, the textbook proof of the “ \longrightarrow ”-implication of the FMW theorem only covers the case where X spans \mathbb{R}^n . One way to recover this part of the theorem in full generality is to identify the span of X with \mathbb{R}^m for some $m < n$, apply the “ \longrightarrow ”-implication for dimension m , and then extend into \mathbb{R}^n the half-spaces (of span X) that define the polyhedral cone. In fact, this argument is essentially the justification for the *wlog* that is given in the book. Unfortunately, the Isabelle vector/matrix library we use does not support identifying an arbitrary proper subspace of \mathbb{R}^n with a Euclidean subspace of lower dimension: Even if we prove some statement for `carrier_vec m`, we cannot apply it to some arbitrary m -dimensional subspace of \mathbb{R}^n . Instead, our formalization of the general case involves adding suitable dummy vectors to X until the set spans all of \mathbb{R}^n , so that we can apply the full-dimension implication for `carrier_vec n`. This is one of several situations where filling in the “obvious” steps of a proof in a way that can be formally expressed in Isabelle requires some creativity.

The third corollary is the decomposition theorem for polyhedra, stating that every polyhedron can be written as the sum of a polytope and a polyhedral cone:

```
lemma decomposition_theorem_polyhedra: ( $\exists A \ b \ nr. A \in \text{carrier\_mat } nr \ n \wedge$ 
   $b \in \text{carrier\_vec } nr \wedge P = \text{polyhedron } A \ b$ )
 $\longleftrightarrow (\exists Q \ X. Q \cup X \subseteq \text{carrier\_vec } n \wedge \text{finite } (Q \cup X) \wedge$ 
   $P = \text{convex\_hull } Q + \text{cone } X)$ 
```

For both the FMW theorem and the decomposition theorem, the fact that we used a set-based matrix/vector library proved to be beneficial. To show the “ \longrightarrow ”-implication of the FMW theorem, one defines a matrix, the dimension of which is a function of X (and can therefore not be independently fixed just by the type of X). Constructing matrices of dimensions that depend on the value of some variable is easy when using `carrier_mat`, but would be very difficult with matrix libraries which utilize Harrison’s encoding of dimensions in types [11]. In the case of the decomposition theorem for polyhedra, the proof involves adding a new component to each vector from a set of n -dimensional vectors and then reasoning about the resulting set of $(n + 1)$ -dimensional vectors, while maintaining the correspondence between the two sets. Here, the use of `carrier_vec` makes it possible to easily switch between dimensions and reason about objects such as “the vector formed of the first n components of some $(n + 1)$ -dimensional vector”.

Since the set of (real) solutions of a system of linear inequalities is a polyhedron, the decomposition theorem for polyhedra allows us to write any solution vector x as $y + z$, with y an element of a polytope (and therefore bounded), and z an element of a finitely generated cone. This suggests the following approach to proving lemma `small_mixed_integer_solution`: If x is such that $x_i \in \mathbb{Z}$ for all $i \in I$, we may try to replace z with a vector z' of the same cone, with bounded entries, such that $(y + z')_i \in \mathbb{Z}$ for all $i \in I$ (thus, $y + z'$ would be the desired bounded solution). This approach does in fact work, but it clearly requires a more powerful version of the decomposition theorem, since the one

we have shown so far says nothing about bounds or integrality. The proof of the new decomposition theorem also requires a bounded integer version of the FMW theorem. This latter theorem in turn is based on a modified version of the FTLI theorem which describes more precisely how separating hyperplanes can be computed so that the normal vectors are integral and with components of bounded size.

```
lemma decomposition_theorem_polyhedra_1:
  and P = polyhedron A b
  shows  $\exists Q X. X \subseteq \text{carrier\_vec } n \wedge Q \subseteq \text{carrier\_vec } n \wedge \text{finite } X \cup Q$ 
     $\wedge P = \text{convex\_hull } Q + \text{cone } X \wedge$ 
     $(A \in \mathbb{Z}_m \cap \text{Bounded\_mat } \text{Bnd} \longrightarrow b \in \mathbb{Z}_v \cap \text{Bounded\_vec } \text{Bnd} \longrightarrow$ 
     $X \subseteq \mathbb{Z}_v \cap \text{Bounded\_vec } (...) \wedge Q \subseteq \text{Bounded\_vec } (...))$ 
```

The ‘ \longrightarrow ’-implication of this stronger version of the decomposition theorem for polyhedra states that if A and b have bounded integer entries, then the finite sets Q and X can be chosen such that they contain only bounded vectors and, furthermore, such that X contains only integral vectors. The integrality of the vectors in X is the crucial ingredient necessary for constructing the vector z' as required and completing the proof of the lemma.

In [19], only a weaker version of the *small_mixed_integer_solution* theorem is proved; it covers only the case of non-strict linear inequalities with integral solutions. Although our result trivially implies this weaker form, we have formalized the proof from the textbook as well, for the sake of completeness.

This proof relies on a decomposition theorem for the integer hull of a polyhedron, which also requires bounded integer versions of the FMW theorem and the decomposition theorem for polyhedra. Only a rough sketch is given in the book as to how the bounded integer versions of these theorems can be obtained. When formalizing this part, however, we encounter the following issue: In the course of a proof, it will be necessary to add new vectors to a set until it has a certain property, or to add half-spaces to a collection until its intersection coincides with some polyhedron. This suffices if we only wish to prove the *existence* of a set of vectors with some property, or of a specific representation of a polyhedron, but if we also need to prove bounds on the numbers needed to describe these objects, it becomes crucial *which* vectors or half-spaces are chosen, because some choices, while valid, will lead to results that do not respect the desired bounds.

For a concrete example, we return to the “ \longrightarrow ”-implication of the FMW theorem (this time in its bounded integer version):

```
lemma farkas_minkowsky_weyl_theorem_1:
  assumes  $X \subseteq \text{carrier\_vec } n$  and  $\text{finite } X$ 
  shows  $\exists nr A. A \in \text{carrier\_mat } nr \ n \wedge \text{cone } X = \text{polyhedral\_cone } A \wedge$ 
     $(X \subseteq \mathbb{Z}_v \cap \text{Bounded\_vec } \text{Bnd} \longrightarrow A \in \mathbb{Z}_m \cap \text{Bounded\_mat } (...))$ 
```

As mentioned earlier in this section, this implication is proved for the case where the span of X is \mathbb{R}^n , which is then used to prove the general implication,

but the switch from the special to the general case involves adding vectors to X until the set spans the entire space, and then applying the full-dimension statement to obtain the half-spaces that define the polyhedral cone. Now, the vectors that are added to X can affect the size of the entries of the resulting matrix A , and the fact that these vectors can also be chosen in such a way that the entries of A are bounded in terms of only Bnd and n , is not obvious, and in fact requires a careful construction. Whereas such matters are simply glossed over in the textbook, resolving the *wlogs* in the proofs of the FTLI and the FMW theorem resulted in Isabelle proofs of 176 lines and 110 lines, respectively.

In the end we achieve the following formalized version of the textbook theorem [19, Thm. 17.1].

```
lemma small_integer_solution_nonstrict_via_decomp:
  assumes  $A \in \text{carrier\_mat } nr \ n \cap \mathbb{Z}_m \cap \text{Bounded\_mat } Bnd$ 
    and  $b \in \text{carrier\_vec } nr \cap \mathbb{Z}_v \cap \text{Bounded\_vec } Bnd$ 
    and  $x \in \text{carrier\_vec } n \cap \mathbb{Z}_v$  and  $Ax \leq b$ 
  shows  $\exists y. y \in \text{carrier\_vec } n \cap \mathbb{Z}_v \wedge Ay \leq b$ 
     $\wedge y \in \text{Bounded\_vec } (\text{fact } (n+1) * (\max 1 Bnd)^n)$ 
```

3.2 Additional formalized theorems

In order to formalize the proofs of the main theorems, we collect a number of basic lemmas concerning cones, convex hulls, integer hulls, normal vectors and bases of vector spaces. On the one hand, these lemmas include very basic statements that would not normally require separate proofs, but were needed for the formalization, such as the fact that a set of vectors is a subset of the cone it generates, or that a convex combination of two vectors of a cone belongs to the cone. On the other hand, our supporting lemmas include statements that appear in standard mathematical texts, such as the fact that any linearly independent set of vectors can be extended to a basis of the vector space. We mention that we have proved all of these facts only for Euclidean vector spaces, making heavy use of the fact that the dimension is finite, because this case suffices for our application.

```
lemma expand_to_basis: assumes  $\text{lin\_indpt\_list } xs$ 
  shows  $\exists ys. \text{set } ys \subseteq \text{set } (\text{unit\_vecs } n) \wedge \text{lin\_indpt\_list } (xs @ ys)$ 
     $\wedge \text{length } (xs @ ys) = n$ 
```

We note that in the above lemma, $@$ is list concatenation and $\text{unit_vecs } n$ refers to the standard basis of \mathbb{R}^n . Of course, a linearly independent set can be extended in many other ways, but we use vectors from the standard basis because they allow us to obtain the same number bounds as for the original linearly independent set. Adding the standard basis vectors is also the reason for using $\max 1 Bnd$ instead of Bnd in the theorems that mention upper bounds. Indeed, the “ $\max 1$ ”-operation cannot be dropped from our main result *small_mixed_integer_solution*: For the single constraint $x > 0$ we would have $n =$

1 and $Bnd = 0$, but there is no solution within the bounds $fact(n+1)*Bnd \sim n = 2*0 \sim 1 = 0$, since $x = 0$ is not a solution.

Aside from the main theorems and supporting lemmas, we also formally prove two variants of Farkas' lemma. We do not need these for our work on the verified linear arithmetic solver, but obtaining them did not entail a prohibitively large additional effort, and they may be useful for other formalizations.

Although there already exists an entry for Farkas' lemma in the AFP, its proof there is based not on the FTLI, but on a separate formalization of the simplex algorithm (one that has been formalized solely for rational numbers). Since here we use the FTLI, we obtain a version of a lemma that allows for the use of a more general type than just the rationals. (In Isabelle, type annotation is denoted by $::$. Below, $'a$ is a type variable that stands for the type of the entries of a matrix/vector; it can be any type with the suitable algebraic properties.)

```
lemma Farkas_Lemma: fixes A :: 'a mat and b :: 'a vec
  assumes A ∈ carrier_mat n nr and b ∈ carrier_vec n
  shows (∃x. x ≥ 0nr ∧ Ax = b) ⟷
    (∀y. y ∈ carrier_vec n ∧ ATy ≥ 0nr ⟶ yb ≥ 0)

lemma Farkas_Lemma': fixes A :: 'a mat and b :: 'a vec
  assumes A ∈ carrier_mat nr nc and b ∈ carrier_vec nr
  shows (∃x. x ∈ carrier_vec nc ∧ Ax ≤ b)
    ⟷ (∀y. y ≥ 0nr ∧ ATy = 0nc ⟶ yb ≥ 0)
```

Finally, we remark that, while the first of the two variants of Farkas' lemma follows easily from the Fundamental theorem, the second variant (which, in [19], has a three-line proof that is based on the first variant) is somewhat more difficult to formalize. This is because its proof involves concatenating matrices and deducing inequalities involving the resulting matrix from facts about its components. Such operations require laborious low-level manipulations of vector inequalities, turning a three-line textbook proof into 102 lines of Isabelle code.

4 A Verified Branch-and-Bound Algorithm

4.1 The Branch-and-Bound Algorithm

Algorithm 1 shows the Isabelle/HOL function *bnb*, which is our implementation of a branch-and-bound algorithm for solving LIA problems. It takes as parameters a list of constraints *cs*, the list of variables *Is* that should get an integer assignment and (total) functions *lb* and *ub* that map the variables in *Is* to their lower and upper integer bounds. *bnb* returns either a satisfying assignment which maps variables to rational numbers and all variables in *Is* to integers or *None* if the mixed integer problem is unsatisfiable within the bounds *lb* and *ub*. *bnb* first uses the simplex algorithm to find a rational solution of the constraints within the bounds. If the constraints are already unsatisfiable in the rational numbers or if the solution is already integral for all values in *Is*, then *bnb* terminates

```

function bnb :: constraint list ⇒ var list ⇒ (var⇒int) ⇒ (var⇒int)
              ⇒ (var⇒rat) option where
  bnb cs Is lb ub =
    case simplex (bounds_to_constraints Is lb ub @ cs) of
      Unsat _ ⇒ None
    | Sat v ⇒ case find (λx. v x ∉ ℤ) Is of
      None ⇒ Some v
      | Some x ⇒ case bnb cs Is lb (ub(x := ⌊v x⌋)) of
        Some sol ⇒ Some sol
        | None ⇒ bnb cs Is (lb(x := ⌈v x⌉)) ub

```

Algorithm 1: A simple implementation of a branch-and-bound algorithm

accordingly. Otherwise, there exists an $x \in Is$ where $v\ x$ (the value assigned to x in the rational solution v) is not an integer. We update the bounds on x once in lb and once in ub and branch by running bnb with the new upper bound and then with the new lower bound.

To verify bnb in Isabelle/HOL we have to show that it always terminates. Note that in every recursive call, we either decrease one of the upper bounds ub or increase one of the lower bounds lb . This fact is used to show that in every recursive call, the range of possible values decreases for *some* x , and, hence, so does the search space. Thus, we use the following measure (of the size of the search space) to prove termination in Isabelle/HOL:

$$\max \left(0, \sum_{x_i \in Is} (ub(x_i) - lb(x_i)) \right)$$

We then prove two theorems about bnb : any detected solution is valid, and whenever bnb delivers *None*, no solution exists within the range that is specified by the lower- and upper-bounds.

```

lemma brach_and_bound_sat:
  assumes bnb cs Is lb ub = Some v
  shows v ⊨mcs (set cs, set Is)

```

```

lemma banch_and_bound_unsat:
  assumes bnb c Is lb ub = None
  and ∀ i ∈ set Is. lb i ≤ v i ∧ v i ≤ ub i
  shows v ⊭mcs (set cs, set Is)

```

At this point we connect the branch-and-bound algorithm with the bounds from Section 3 to obtain a decision procedure for linear (mixed) integer arithmetic:

```

definition branch_and_bound cs Is = (let B = compute_bound cs
                                     in bnb cs Is ( $\lambda\_.$   $-B$ ) ( $\lambda\_.$  B))

```

Here, `compute_bound` is an algorithm that extracts the relevant parameters (number of variables, maximal absolute value in constraints) and then calculates the upper bound as in Section 3. One complication comes from the fact that there are two different representations of constraints: the statements regarding bounds have been proved for constraints given in matrix-vector form, $Ax \leq b$ or $Ax < b$ with integral matrix A and integral vector b , whereas the input to the branch-and-bound algorithm is a set of constraints, where each constraint is represented via a (sparse) linear polynomial with rational entries, e.g., $x_5 + \frac{1}{10}x_{1041} \leq \frac{7}{3}$. Hence, `compute_bound` internally also normalizes the constraints, e.g., by canceling fractions, and by renaming the variables so that the indices form a contiguous block: x_0, \dots, x_{n-1} . The normalized constraints can then easily be translated into matrix-vector-form, which enables a lifting of lemma `small_mixed_integer_solution` to constraints that are represented via sparse polynomials.

```

lemma compute_bound:
  assumes  $v \models_{mcs} (\text{set } cs, Is)$ 
  shows  $\exists v. v \models_{mcs} (\text{set } cs, Is) \wedge (\forall i \in Is. |v\ i| \leq \text{compute\_bound } cs)$ 

```

At this point, it is easy to combine the results of `bnb` with `compute_bound` to finally show that `branch_and_bound` is a complete and sound decision procedure. Either it returns some assignment, which is then a solution to the mixed integer problem; or it returns `None`, and the mixed integer problem is unsatisfiable.

```

lemma branch_and_bound:
  branch_and_bound cs Is = Some v  $\implies v \models_{mcs} (\text{set } cs, \text{set } Is)$ 
  branch_and_bound cs Is = None  $\implies \nexists v. v \models_{mcs} (\text{set } cs, \text{set } Is)$ 

```

4.2 Using The Incremental Version of Simplex

One problem of the branch-and-bound algorithm from the previous section is in the way it invokes the simplex algorithm: although in every iteration only a single constraint changes, the simplex algorithm is always started from scratch.

Therefore, in this section we optimize the branch-and-bound algorithm to use an already existing verified *incremental* version of the simplex algorithm [3,14], which returns a state instead of only returning a satisfying assignment or stating unsatisfiability. The state contains for instance a tableau, i.e., a list of equations which is essential for the simplex algorithm. By reusing the state, expensive operations like creating the tableau can be avoided, making the incremental simplex very attractive to be used within the branch-and-bound algorithm.

A complication arises, since the verified incremental simplex algorithm was developed to be used in a DPLL(T)-solver, where all constraints are known beforehand and the constraints are not changed throughout one DPLL(T) run.

Therefore, the incremental interface does not allow for changing constraints or adding new ones. As a consequence, an integration of the incremental simplex into the branch-and-bound algorithm is not immediate, since there the bounds are changed in every iteration.

Our solution is a slight extension of the incremental simplex algorithm. To be more precise, we write a function which changes exactly one constraint in the state in a way that the relevant invariants of the incremental interface still hold. This extension allows us to reuse all the existing soundness properties and proofs of the incremental simplex algorithm without modifications. It is specifically tailored for running the branch-and-bound algorithm. We choose this approach instead of adding a feature to change arbitrary constraints in the incremental simplex interface, since such a feature would require a major rewrite.

Since the algorithmic structure and the soundness statement of the modified branch-and-bound algorithm is completely identical to the one of Section 4.1, we just refer to the formalization for further details.

5 Benchmarking

We tested two versions of our solver (based on incremental/non-incremental simplex) by comparing them with two well-established SMT-solvers, Z3 and CVC4. Testing was done on a subset of the non-incremental⁵ QF_LIA (quantifier-free linear integer arithmetic)⁶ benchmark set from SMT-LIB. For this experiment we had two goals in mind: 1. to see whether it is worthwhile to use the non-incremental version of simplex as a sub-routine in the branch-and-bound algorithm, and 2. to get an idea about the extent to which our verified, non-optimized solver can handle practical examples.

We did not go through the effort of making our solver compliant with the language of SMT-LIB, as we felt that for the above two goals, it would suffice to write a simple converter that could handle a reasonable portion of the QF_LIA benchmarks. Thus, we obtained a dataset of 1192 benchmarks, comprising 18% of the 6489 benchmarks in QF_LIA. (More specifically, the following sub-folders were fully converted to a format that is readable by our solver: 20180326-Bromberger, miplib2003, pb2010, pidgeons, prime-cone, and slacks.) All solvers were tested on this dataset, on the same hardware, with a 60s-timeout per benchmark. Z3 version 4.4.0pre-2, CVC4 version 1.5-4, and the 2019-05-09 release of SMT-LIB were used.

The only other verified LIA solver that we are aware of is an Isabelle formalization of Cooper’s algorithm in the AFP. This algorithm solves a more general problem than linear integer arithmetic (namely linear arithmetic with arbitrary quantifiers over integer variables). We obtained an implementation with minimal changes to make code generation possible (just as we produced executables for our own solver).

⁵ Here, “non-incremental” means that the tests are simply sets of constraints, as opposed to constraints together with an assert/check script that a solver must execute.

⁶ QF_LIRA (quantifier-free mixed integer real arithmetic) contains only 8 tests.

Table 1: Experimental results

	sat	unsat	total
non-incremental bnb	245	131	376
incremental bnb	314	131	445
CVC4	470	158	628
Z3	570	164	734
verified Cooper	2	0	2

Evaluation. Our branch-and-bound implementation solves 37% of the dataset with incremental simplex as a sub-routine, and 31% with non-incremental simplex. Since we have only implemented a naive branch-and-bound algorithm, without any additional heuristics for pruning the search space, it is unsurprising that its performance cannot match that of more mature solvers. Somewhat surprising is the fact that some benchmarks are solved by our solvers but not by Z3 or CVC4: of the benchmarks solved by incremental bnb, 28 are not solved by Z3, 29 are not solved by CVC4, and 8 are solved by neither Z3 nor CVC4.

Interestingly, the non-incremental simplex-based solver can handle a few instances that the incremental simplex-based solver does not. Although using an incremental simplex leads to better overall results, it appears that reusing valuations from previous simplex runs can sometimes lead the search astray in such a way that simple solutions are missed.

Cooper’s algorithm is known to have a very high asymptotic complexity, which means that its performance is not a matter of optimizing an implementation. As such, the outcome of our experiments with regards to Cooper’s algorithm is as expected, showing that this algorithm is not usable on medium-sized examples in practice.

6 Conclusion and Future Work

We have developed a verified solver for linear mixed integer arithmetic, and have formalized important results on linear integer programming that were needed in order to prove the soundness of the solver. To the extent of our knowledge, the main mathematical theorems of which we formalized proofs had not been previously verified in any formal system, and our solver is the first verified LIA solver that is also usable in practice. The two parts of our formalization amount to 9732 lines of Isabelle code and took roughly 10 person-months to implement.

Currently, our solver is essentially “proof of concept” software, and there are a number of known optimizations in which it could be improved, e.g., preprocessing of constraints, integration of cutting planes, unit-cube-tests, etc. [6,7,13]. We have also used run-time profiling in order to establish which sub-routines our solver spends most time on, and have identified parts of the incremental simplex algorithm that we could further modify in order to improve running times.

References

1. Allamigeon, X., Katz, R.D.: A formalization of convex polyhedra based on the simplex method. In: Interactive Theorem Proving, 8th International Conference, ITP 2017, Proceedings. LNCS, vol. 10499, pp. 28–45 (2017). https://doi.org/10.1007/978-3-319-66107-0_3
2. Barrett, C., Conway, C.L., Deters, M., Hadarean, L., Jovanović, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Computer Aided Verification, 23rd International Conference, CAV ’11, Proceedings. LNCS, vol. 6806, pp. 171–177 (2011)
3. Bottesch, R., Haslbeck, M.W., Thiemann, R.: Verifying an incremental theory solver for linear arithmetic in Isabelle/HOL. In: Frontiers of Combining Systems, 12th International Symposium, FroCoS ’19, Proceedings. LNCS, vol. 11715, pp. 223–239 (2019). https://doi.org/10.1007/978-3-030-29007-8_13
4. Bottesch, R., Reynaud, A., Thiemann, R.: Linear inequalities. Archive of Formal Proofs (Jun 2019), http://isa-afp.org/entries/Linear_Inequalities.html, Formal proof development
5. Brockschmidt, M., Joosten, S.J.C., Thiemann, R., Yamada, A.: Certifying safety and termination proofs for integer transition systems. In: Automated Deduction, 26th International Conference, CADE 26, Proceedings. LNCS, vol. 10395, pp. 454–471 (2017). https://doi.org/10.1007/978-3-319-63046-5_28
6. Bromberger, M.: A reduction from unbounded linear mixed arithmetic problems into bounded problems. In: Automated Reasoning, 9th International Joint Conference, IJCAR 2018, Proceedings. LNCS, vol. 10900, pp. 329–345 (2018). https://doi.org/10.1007/978-3-319-94205-6_22
7. Bromberger, M., Weidenbach, C.: New techniques for linear arithmetic: cubes and equalities. Formal Methods in System Design **51**(3), 433–461 (2017). <https://doi.org/10.1007/s10703-017-0278-7>
8. Dantzig, G.B.: Linear Programming and Extensions. Princeton University Press (1963)
9. Dutertre, B., de Moura, L.M.: A fast linear-arithmetic solver for DPLL(T). In: Computer Aided Verification, 18th International Conference, CAV 2006, Proceedings. LNCS, vol. 4144, pp. 81–94 (2006). https://doi.org/10.1007/11817963_11
10. Haftmann, F., Nipkow, T.: Code generation via higher-order rewrite systems. In: Functional and Logic Programming, 10th International Symposium, FLOPS ’10, Proceedings. LNCS, vol. 6009, pp. 103–117 (2010). https://doi.org/10.1007/978-3-642-12251-4_9
11. Harrison, J.: The HOL light theory of Euclidean space. J. Autom. Reasoning **50**, 173–190 (2013). <https://doi.org/10.1007/s10817-012-9250-9>
12. Krauss, A.: Partial and nested recursive function definitions in higher-order logic. J. Autom. Reasoning **44**(4), 303–336 (2010), <https://doi.org/10.1007/s10817-009-9157-2>
13. Marchand, H., Martin, A., Weismantel, R., Wolsey, L.A.: Cutting planes in integer and mixed integer programming. Discrete Applied Mathematics **123**(1-3), 397–446 (2002). [https://doi.org/10.1016/S0166-218X\(01\)00348-1](https://doi.org/10.1016/S0166-218X(01)00348-1)
14. Marić, F., Spasić, M., Thiemann, R.: An incremental simplex algorithm with unsatisfiable core generation. Archive of Formal Proofs (Aug 2018), <http://isa-afp.org/entries/Simplex.html>, Formal proof development
15. de Moura, L., Björner, N.: Z3: An Efficient SMT Solver. In: Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Proceedings. LNCS, vol. 4963, pp. 337–340. Springer (2008). https://doi.org/10.1007/978-3-540-78800-3_24

16. Nipkow, T.: Linear quantifier elimination. *J. Autom. Reasoning* **45**(2), 189–212 (2010). <https://doi.org/10.1007/s10817-010-9183-0>
17. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL – A Proof Assistant for Higher-Order Logic, LNCS, vol. 2283. Springer (2002). <https://doi.org/10.1007/3-540-45949-9>
18. Papadimitriou, C.H.: On the complexity of integer programming. *J. ACM* **28**(4), 765–768 (1981). <https://doi.org/10.1145/322276.322287>
19. Schrijver, A.: Theory of linear and integer programming. Wiley (1999)
20. Spasić, M., Marić, F.: Formalization of incremental simplex algorithm by step-wise refinement. In: Formal Methods, 18th International Symposium, FM 2012, Proceedings. LNCS, vol. 7436, pp. 434–449 (2012). https://doi.org/10.1007/978-3-642-32759-9_35
21. Thiemann, R., Sternagel, C.: Certification of termination proofs using **CeTA**. In: Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLs 2009, Proceedings. LNCS, vol. 5674, pp. 452–468 (2009). https://doi.org/10.1007/978-3-642-03359-9_31