

Inferring linear invariants with parallelotopes[☆]

Gianluca Amato^a, Marco Rubino^a, Francesca Scozzari^a

^a *Università di Chieti-Pescara, Italy*

Abstract

We propose a new numerical abstract domain for inferring linear invariants based on parallelotopes. The domain may encode any linear constraint, as the polyhedra abstract domain, while maintaining the efficiency of weakly relational abstract domains, such as intervals and octagons. We provide the full set of abstract operators, define a reduced product with intervals and present an experimental comparison with polyhedra and octagons. According to these experiments, the reduced product we propose is much more precise than both polyhedra and octagons in inferring interval constraints.

Keywords: Static analysis, abstract interpretation, numerical abstract domain, linear invariant, parallelotopes.

1. Introduction

In the field of static analysis by abstract interpretation, much effort has been devoted in designing domains for inferring numerical properties such as “the value of the variable x in program point p is between 0 and 10”. In this paper we are mainly interested in linear invariants, which are typically expressed as linear inequalities such as $2x+3y \leq 42$. The polyhedra domain by Cousot and Halbwachs [1] is able to express any such invariant, but its computational complexity makes it difficult to use it in practice. Many other numerical domains have been proposed to gain efficiency. One of the simplest one is the Cousot and Cousot’s interval domain [2], a non-relational abstract domain which can represent only constraints involving a single variable such as $0 \leq x \leq 10$. Analyses using the interval domain are not very precise, since it cannot represent relationships between variables. Many weakly relational domains, more expressive than intervals but less expressive than polyhedra, have been proposed in the last years, such as octagons [3], octahedra [4], logahedra [5] and TVPI [6]. Weakly relational domains have proved to be quite efficient, but the invariants that can be inferred using these domains are seriously limited by syntactic restrictions,

[☆]©2017. This manuscript version is made available under the CC-BY-NC-ND 4.0 license: <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Email addresses: gamato@unich.it (Gianluca Amato), marco.rubino@unich.it (Marco Rubino), fscozzari@unich.it (Francesca Scozzari)

since expressivity is bartered for efficiency. For instance, the octagons abstract domain can only deal with inequalities involving two variables whose coefficients are $\{-1, 0, 1\}$, such as $x - y \leq 5$, while TVPI can only express inequalities with two variables, such as $5x + 2y \leq 8$.

In order to handle more expressive constraints, Sankaranarayanan et al. have proposed a different approach called template polyhedra [7]. For each program, they fix a constraint matrix A and consider all the polyhedra of the form $\{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{b}\}$, where \mathbf{x} is the vector of program variables. The *a priori* choice of the matrix differentiates template polyhedra from other domains, where the matrix is fixed for all the programs (such as intervals or octagons) or varies freely (such as polyhedra). Template polyhedra generalize most weakly relational domains, with the difference that its abstract operators are defined by means of linear programming. Along the same direction there are the proposals of generalized template polyhedra [8], which combine template polyhedra and bilinear forms, and template parallelotopes [9, 10], which are a special case of template polyhedra. A parallelotope is a polyhedron defined by at most n linearly independent constraints, where n is the number of variables. Any parallelotope can be described as the set of points $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{b}_1 \leq A\mathbf{x} \leq \mathbf{b}_2\}$, where \mathbf{x} is the vector of program variables and the constraint matrix A is square and invertible. In the special case of template parallelotopes, it is possible to derive efficient abstract operators, without resorting to linear programming tools. In general, when using templates, the result of the analyses greatly varies depending on the choice of the matrix A , and finding such matrices is a hard and still open problem. We are aware of only two proposals for generating templates, one in Sankaranarayanan et al. original paper [7] based on a syntactic inspection of the program, and another one based on the statistical analysis of partial execution traces [9, 10].

Here we propose a different approach which does not use templates and does not restrict the syntactic shape of constraints, still maintaining efficiency. We propose a domain whose abstract objects are parallelotopes, namely we only limit the number of constraints to n (the number of variables) and we require the constraints to be linearly independent.

While our abstract objects are parallelotopes as in [9, 10], there is the fundamental difference that the constraint matrix is not fixed *a priori* but may freely change during the analysis, as for the polyhedra domain. The domain of parallelotopes so defined can encode any linear constraints, does not require templates and can be equipped with very efficient abstract operations, whose complexity is comparable to that of octagons. The key to scalability is the restriction to n linearly independent constraints. This allows the use of algorithms adapted from linear equation solving which are simpler and more efficient than those used in the theory of polyhedra. For instance, minimizing a linear form on a parallelotope is cubic, since it essentially amounts to solving a linear equation system, while the corresponding operation on polyhedra and even on template polyhedra requires the use of the simplex algorithm.

Since many concrete operations are not closed with respect to parallelotopes, in most cases there is not a unique best parallelotope which approximates the

result of the concrete operation. The careful design of suitable abstract operators is then a key point in the parallelotope domain. In some cases, as for the abstract union operator, we propose different operators, and in general we resort to appropriate heuristics in order to ensure good precision in the overall analysis.

We also propose a combination of the domain of parallelotopes and intervals, and we experimentally show that it compares well with respect to polyhedra and octagons. Finally, we show that, if we are interested in inferring precise interval constraints, then the combination of parallelotopes and intervals is far more precise than both polyhedra and octagons.

Preliminary results on the domain of parallelotopes appeared in [11]. We provide here a new comprehensive presentation of the domain, containing several novelties such as new abstract operators on parallelotopes, the reduced product of parallelotopes and intervals, experimental evaluations, and proofs (which are in the Appendix).

2. Notation

2.1. Linear Algebra

We denote by \mathbb{R} the set of real numbers extended with $+\infty$ and $-\infty$. Addition and multiplication are partially extended, when possible, to \mathbb{R} in the obvious way. We avoid the use of indeterminate forms, with the exception of 0 times $\pm\infty$ which we define to be 0. We use boldface for elements \mathbf{v} of \mathbb{R}^n . Any vector $\mathbf{v} \in \mathbb{R}^n$ is intended as a column vector, and \mathbf{v}^T is the corresponding row vector. We denote by $\mathbf{u} \cdot \mathbf{v}$ the dot product of \mathbf{u} and \mathbf{v} and we use it indifferently for both row and column vectors. Given $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$, and a relation $\bowtie \in \{<, >, \leq, \geq, =\}$, we write $\mathbf{u} \bowtie \mathbf{v}$ if and only if $u_i \bowtie v_i$ for each $i \in \{1, \dots, n\}$. Given $\mathbf{u} \in \mathbb{R}^n$ and $i \in \{1, \dots, n\}$, we write $\mathbf{u}[i \mapsto x]$ to denote a vector \mathbf{v} such that $v_i = x$ and $v_j = u_j$ for $j \neq i$.

If $A = (a_{ij})$ is a matrix, we denote by A^T its *transpose*. If A is invertible, A^{-1} denotes its inverse, and $\text{GL}(n)$ is the group of $n \times n$ invertible matrices. The identity matrix in $\text{GL}(n)$ is denoted by I_n and the standard basis of \mathbb{R}^n is denoted by $\{\mathbf{e}^1, \dots, \mathbf{e}^n\}$. Clearly, any $1 \times n$ -matrix can be viewed as a vector: in particular, we denote by \mathbf{a}_{i*} the row vector given by the i -th row of any matrix A , and by \mathbf{a}_{*i} the column vector given by the i -th column of A .

More generally, if J is a set of indexes, then A_{J*} is the submatrix of A composed of the rows in J and A_{*J} the submatrix composed of the columns in J . We use A_{-J*} for the submatrix of A composed of the rows of A whose indexes are not in J .

In the rest of the paper, when talking about the computational complexity of matrix and vector operations, we always assume to use a dense representation, which is the one adopted in our implementation.

2.2. Convex sets

A set $C \subseteq \mathbb{R}^n$ is *convex* when, given $\mathbf{x}, \mathbf{y} \in C$, for each $\lambda \in [0, 1]$ we have that $\lambda\mathbf{x} + (1 - \lambda)\mathbf{y} \in C$. A *ray* in C is a non-null vector $\mathbf{v} \in \mathbb{R}^n$ such that for

each $\mathbf{x} \in C$ and $\alpha \geq 0$, $\mathbf{x} + \alpha \mathbf{v} \in C$. A convex set C is *bounded* iff it has no rays. A *line* in C is a vector $\mathbf{v} \in \mathbb{R}^n$ such that both \mathbf{v} and $-\mathbf{v}$ are rays. The lines in C form a vector space, called the *linearity space* of C , denote by $\text{lin}(C)$. Given a vector space V , we denote by V^\perp its orthogonal complement.

A set of vectors $S = \{v_1, \dots, v_m\}$ is *affinely dependent* if there are $\lambda_1, \dots, \lambda_m \in \mathbb{R}$ such that they are not all zero, $\lambda_1 + \dots + \lambda_m = 0$ and $\lambda_1 \mathbf{v}_1 + \dots + \lambda_m \mathbf{v}_m = 0$. The *affine hull* of S is the set $\text{aff.hull}(S) = \{\lambda_1 \mathbf{v}_1 + \dots + \lambda_m \mathbf{v}_m \mid \lambda_1 + \dots + \lambda_m = 1\}$.

A *flat* F is a set of points of the form $\mathbf{a} + H$ where H is a vector space. Given a flat, H is uniquely determined. Therefore, we call the dimension of F the dimension of the corresponding vector space H . By convention, the dimension of the empty set is -1 . The *dimension* of a set S is the dimension of the smallest flat containing S , and it is denoted by $\text{dim}(S)$.

2.3. Polyhedra and constraints

An (affine) *constraint* over a vector of n variables is a syntactic object $\mathbf{a} \cdot \mathbf{x} \leq b$, where $\mathbf{a} \in \mathbb{R}^n$ and $b \in \mathbb{R}$. We use $l \leq \mathbf{a} \cdot \mathbf{x} \leq u$ as a short form for the pair of constraints $\mathbf{a} \cdot \mathbf{x} \leq u$ and $-\mathbf{a} \cdot \mathbf{x} \leq -l$. If $C = \{\mathbf{a}_1 \cdot \mathbf{x} \leq b_1, \dots, \mathbf{a}_m \cdot \mathbf{x} \leq b_m\}$ is a finite set of constraints, its *solution set* is $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{a}_1 \cdot \mathbf{x} \leq b_1, \dots, \mathbf{a}_m \cdot \mathbf{x} \leq b_m\}$. Given a constraint $c \in C$, we say that c is *redundant* for C when the solution sets of C and $C \setminus \{c\}$ coincide. Equivalently, a constraint $c \equiv \mathbf{a} \cdot \mathbf{x} \leq b$ is redundant iff there is a family of constraints $\{c_i\}_{i \in I} \subseteq C \setminus \{c\}$ with $c_i \equiv \mathbf{a}_i \cdot \mathbf{x} \leq b_i$ and a family of positive numbers $\{\lambda_i\}_{i \in I}$ such that $\mathbf{a} = \sum_{i \in I} \lambda_i \mathbf{a}_i$ and $b \geq \sum_{i \in I} \lambda_i b_i$. Informally, this means that c is implied by a positive linear combinations of the c_i 's. This characterization of redundant constraints is known as Farkas' Lemma.

A set $P \subseteq \mathbb{R}^n$ is a *polyhedron* when there exists an $m \times n$ matrix A and a vector $\mathbf{u} \in \mathbb{R}^m$ such that

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid A\mathbf{x} \leq \mathbf{u}\} .$$

For each $i \in \{1, \dots, m\}$, $\mathbf{a}_{i*} \cdot \mathbf{x} \leq u_i$ is a constraint, therefore P may be viewed as the solution set of a finite set of constraints.

The Fourier-Motzkin elimination [12, Chapter 12.2] is a method to eliminate a variable from a set of affine constraints. Given two constraints $c_1 \equiv \mathbf{a} \cdot \mathbf{x} \leq b$ and $c_2 \equiv \mathbf{d} \cdot \mathbf{x} \leq e$ such that $a_1 < 0$ and $d_1 > 0$, we may combine c_1 and c_2 into a new constraint $c_3 \equiv (d_1 \mathbf{a} - a_1 \mathbf{d}) \cdot \mathbf{x} \leq d_1 b - a_1 e$ where the coefficient of the variable x_1 is zero. Given a set of constraints C , we partition C into C^+ , C^- and C^0 according to the sign of the coefficient of x_1 . Then, we denote by C^{+-} the set of all the constraints obtained combining a constraint in C^- and one in C^+ as shown above. The set $C^{+-} \cup C^0$ is a set of constraints over the variables x_2, \dots, x_n , and has the property that a point $(x_2, \dots, x_n) \in \mathbb{R}^{n-1}$ is in the solution set of $C^{+-} \cup C^0$ if and only if there exists $x_1 \in \mathbb{R}$ such that (x_1, \dots, x_n) is in the solution set of C .

2.4. Abstract interpretation

In this paper we adopt a framework for abstract interpretation which is weaker than the common one based on Galois connections (see [13, Section 7]).




Figure 1: Parallelotopes.

Given a partially-ordered set (C, \leq_C) — the *concrete domain* — and a preordered set (A, \leq_A) — the *abstract domain* — we establish an abstraction–concretization relationship between them with the use of a *concretization map*, which is a monotone function $\gamma : A \rightarrow C$. We say that $a \in A$ is a *correct approximation* of $c \in C$ when $c \leq_C \gamma(a)$. Given $c \in C$, there are many correct abstractions. The most interesting ones are those which are minimal w.r.t. the preorder \leq_A .

A function $f^\alpha : A \rightarrow A$ is a *correct abstraction* of $f : C \rightarrow C$ when it preserves correctness of approximation, i.e. when $c \leq_C \gamma(a)$ implies $f(c) \leq_C \gamma(f^\alpha(a))$. It is γ -*complete* when $\gamma \circ f^\alpha = f \circ \gamma$. It is *minimal* when, for any $a, a' \in A$, if $f(\gamma(a)) \leq_C \gamma(a') \leq_C \gamma(f^\alpha(a))$, then $\gamma(a') = \gamma(f^\alpha(a))$, i.e., when $f^\alpha(a)$ is a minimal correct approximation of $f(\gamma(a))$.

When \leq_A is a partial order and γ has a left adjoint, i.e., there exists a monotone function $\alpha : C \rightarrow A$ such that, $\forall a \in A \forall c \in C (\alpha(c) \leq a \iff c \leq \gamma(a))$, we say that α and γ form a *Galois connection*. The function α maps each $c \in C$ to its *best correct approximation* $\alpha(c)$, i.e., the least correct abstraction according to \leq_A . For each $f : C \rightarrow C$ it is possible to define its *best correct abstraction* as $f^\alpha = \alpha \circ f \circ \gamma$. When γ is injective, we say that α and γ form a *Galois insertion*.

3. Parallelotopes and their representation

A set $P \subseteq \mathbb{R}^n$ is a *parallelotope* when there is an invertible matrix $A \in \text{GL}(n)$ and vectors $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ such that

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}\} . \quad (1)$$

A parallelotope is a closed convex set. The matrix A is called the *constraint matrix*, while \mathbf{l} and \mathbf{u} are the *lower* and *upper bounds* respectively. While in the mathematical literature a parallelotope is generally considered to be bounded, we are also considering unbounded ones. We call *box* a parallelotope whose constraint matrix is the identity matrix. Examples of parallelotopes are given in Figure 1.

In this section we define the abstract domain of parallelotopes and its relationship with the concrete domain of subsets of \mathbb{R}^n .

3.1. Representation of parallelotopes

First of all, we need a finite representation for parallelotopes. Non-empty parallelotopes are represented with a constraint matrix and corresponding lower and upper bounds, while a special symbol ϵ is used for empty parallelotopes.

Definition 1 (Representations of parallelotopes). A representation \mathcal{P} of a parallelotope is either the symbol ϵ or a tuple $\langle A, \mathbf{l}, \mathbf{u} \rangle$ with $A \in \text{GL}(n)$, $\mathbf{l}, \mathbf{u} \in \bar{\mathbb{R}}^n$ subject to the following additional conditions:

- $\mathbf{l} \leq \mathbf{u}$;
- for all $i \in \{1, \dots, n\}$ $l_i \neq +\infty$ and $u_i \neq -\infty$.

We denote by Par_n the set of all the representations for parallelotopes in \mathbb{R}^n . We simply use Par when n is not relevant.

The last two conditions in Definition 1, together with the fact that A is invertible, ensure that $\langle A, \mathbf{l}, \mathbf{u} \rangle$ represents a non-empty parallelotope. We prefer to keep a different representation for the empty case, since this simplifies the definition of the abstract operators. The relationship between parallelotopes and their representations is formalized by the following concretization map.

Definition 2 (Concretization map). We define a concretization map $\gamma : \text{Par}_n \rightarrow \mathcal{P}(\mathbb{R}^n)$ from representations of parallelotopes to parallelotopes as follows:

$$\begin{aligned} \gamma(\epsilon) &= \emptyset, \\ \gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle) &= \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}\}. \end{aligned}$$

Theorem 1. For every parallelotope P there is a representation \mathcal{P} such that $\gamma(\mathcal{P}) = P$.

In the following, when this does not cause any ambiguities, we use a representation \mathcal{P} in place of the parallelotope $\gamma(\mathcal{P})$. Moreover, we refer to representations of parallelotopes simply as parallelotopes.

3.2. Some properties of parallelotopes

We want to study the connection between the properties of a parallelotope as a convex set (rays, boundedness, etc. . .) and its representation. Most of these results may be found in standard textbooks with slightly different notations. We provide the proofs in the appendix for clarity of presentation.

Definition 3 (Constrained, unconstrained and equality rows). A row \mathbf{a}_{i^*} of a representation $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ is said to be: (1) *constrained* when $l_i \in \mathbb{R}$ or $u_i \in \mathbb{R}$; (2) *unconstrained* when $l_i = -\infty$ and $u_i = +\infty$; (3) *equality* when $l_i = u_i$. We denote by $U(\mathcal{P})$ (resp. $E(\mathcal{P})$) the set of indexes of the unconstrained rows (resp. equality rows) in \mathcal{P} .

Theorem 2 (Parallelotopes and Representations). *If $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ is a representation of a parallelotope in \mathbb{R}^n , the following holds:*

1. $\dim(\mathcal{P})$ is equal to $n - |E(\mathcal{P})|$;¹
2. $\mathbf{v} \neq 0$ is a ray in \mathcal{P} iff for each $i \in \{1, \dots, n\}$, $l_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \geq 0$ and $u_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \leq 0$;
3. \mathcal{P} is bounded iff all the bounds are finite;
4. $\mathbf{v} \neq 0$ is a line in \mathcal{P} iff it is orthogonal to all the constrained rows.

Corollary 3. *The linearity space of a parallelotope \mathcal{P} is V^\perp where V is the linear space generated by all the constrained rows of \mathcal{P} . It is generated by the orthogonal projections of the unconstrained rows onto V^\perp .*

3.3. Minimization and maximization over a parallelotope

One of the basic operations on abstract domains is computing the minimum and maximum values of a linear form over the points in an abstract object. If \mathcal{P} is a parallelotope and $\mathbf{c} \in \mathbb{R}^n$, this means computing $\inf_{\mathbf{x} \in \mathcal{P}} \mathbf{c} \cdot \mathbf{x}$ and $\sup_{\mathbf{x} \in \mathcal{P}} \mathbf{c} \cdot \mathbf{x}$.

The minimization and maximization operator for parallelotopes may be obtained by viewing a parallelotope as a box over a non-canonical coordinate system. Given a box $\mathcal{B} = \langle I_n, \mathbf{l}, \mathbf{u} \rangle$ and a vector $\mathbf{c} \in \mathbb{R}^n$, it is well known that

$$\inf_{\mathbf{x} \in \mathcal{B}} \mathbf{c} \cdot \mathbf{x} = \inf_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c} \cdot \mathbf{x} = \mathbf{c} \cdot \mathbf{v} \quad \text{where} \quad \mathbf{v}_i = \begin{cases} l_i & \text{if } c_i \geq 0 \\ u_i & \text{otherwise.} \end{cases}$$

The computational complexity of this operation is $O(n)$.

Theorem 4. *Given the representation $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and a vector $\mathbf{c} \in \mathbb{R}^n$, we have that*

$$\inf_{\mathbf{x} \in \mathcal{P}} \mathbf{c} \cdot \mathbf{x} = \inf_{\mathbf{l} \leq \mathbf{y} \leq \mathbf{u}} \mathbf{c}^T A^{-1} \mathbf{y}.$$

The computational complexity of the minimization operation is $O(n^3)$.

The same properties hold for the maximization operator.

3.4. Parallelotopes over a fixed constraint matrix

We say that the parallelotope P is *definable over* $A \in \text{GL}(n)$ if it is empty or there is a representation for P with A as the constraint matrix. We denote by Par_n^A the subset of Par_n comprised of ϵ and the non-empty representations with A as the constraint matrix. Hence, P is definable over A if there is a representation for it in Par_n^A .

¹As an example of the problems which arise in considering representations with $\mathbf{l} \not\leq \mathbf{u}$, this property is false in that case, since \mathcal{P} is empty and $\dim(\mathcal{P}) = -1$ regardless of the number of equality rows.

The abstract domain of parallelotopes with a fixed constraint matrix has been thoroughly studied in [10]. Here we recall some results which are relevant to our paper, adapting the notations of [10] to the ones we are using here.

First of all, it is possible to define on Par_n^A a partial order as follows:

$$\langle A, \mathbf{l}, \mathbf{u} \rangle \leq \langle A, \mathbf{l}', \mathbf{u}' \rangle \text{ iff } \mathbf{l}' \leq \mathbf{l} \text{ and } \mathbf{u}' \geq \mathbf{u},$$

extended with the proviso that $\epsilon \leq \langle A, \mathbf{l}, \mathbf{u} \rangle$ for each $\mathbf{l}, \mathbf{u} \in \bar{\mathbb{R}}^n$. With respect to this ordering, Par_n^A is a complete lattice and γ is a complete join-morphism, hence it has left adjoint $\alpha_A : \mathcal{P}(\mathbb{R}^n) \rightarrow \text{Par}_n^A$ such that $\alpha_A(\emptyset) = \epsilon$ and, for $C \neq \emptyset$, $\alpha_A(C) = \langle A, \mathbf{l}, \mathbf{u} \rangle$ with $l_i = \inf_{\mathbf{x} \in C} \mathbf{a}_{i*} \mathbf{x}$ and $u_i = \sup_{\mathbf{x} \in C} \mathbf{a}_{i*} \mathbf{x}$. Maps α_A and γ form a Galois connection. Since γ is injective, it is actually a Galois insertion.

In particular, this means that for each $C \subseteq \mathcal{P}(\mathbb{R}^n)$, $\alpha_A(C)$ gives the least parallelotope which contains C and is definable over A . As a particular case, it is interesting, given a parallelotope P in \mathbb{R}^n and a matrix $A' \in \text{GL}(n)$, to find out the least parallelotope containing P and definable over A' .

Definition 4. Given a parallelotope $\langle A, \mathbf{l}, \mathbf{u} \rangle$ and a matrix $A' \in \text{GL}(n)$, the operator $\text{rot}_{A'}$ is defined as

$$\text{rot}_{A'}(\langle A, \mathbf{l}, \mathbf{u} \rangle) = \langle A', \mathbf{l}', \mathbf{u}' \rangle$$

where

$$B = A'A^{-1}, \quad l'_i = \inf_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{b}_{i*} \mathbf{x}, \quad u'_i = \sup_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{b}_{i*} \mathbf{x}.$$

This is extended to the empty parallelotope as $\text{rot}_{A'}(\epsilon) = \epsilon$.

Example 1. Given the parallelotope

$$P = \{\mathbf{x} \in \mathbb{R}^n \mid 0 \leq -x_1 + 3x_2 \leq 2, 3 \leq x_1 + 2x_2 \leq 8\}$$

depicted in gray in Figure 2, and given the matrix $A = \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$, $\text{rot}_A(P)$ is the dotted parallelotope.

Theorem 5. *Given a parallelotope \mathcal{P} and $A' \in \text{GL}(n)$, $\text{rot}_{A'}(\mathcal{P})$ is the least parallelotope definable over A' which contains \mathcal{P} , i.e., $\text{rot}_{A'}(\mathcal{P}) = \alpha_{A'}(\gamma(\mathcal{P}))$. The computational complexity of $\text{rot}_{A'}$ is $O(n^3)$.*

3.5. Parallelotopes over arbitrary constraint matrices

Given a set $C \subseteq \mathbb{R}^n$, we are interested in approximating C with a parallelotope $P \supseteq C$. In general, there is not a least parallelotope P which contains C , but there are several (possibly infinitely many) minimal parallelotopes with such a property. In particular, for each $A \in \text{GL}(n)$, $\gamma(\alpha_A(C))$ is a minimal parallelotope containing C .

It is possible to define a pre-order over parallelotopes in such a way that $\mathcal{P} \leq \mathcal{P}'$ iff $\gamma(\mathcal{P}) \subseteq \gamma(\mathcal{P}')$. The idea is that if $\gamma(\mathcal{P}) \subseteq \gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)$, then $\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)$ contains the least parallelotope definable over A which contains $\gamma(\mathcal{P})$.




Figure 2: The `rot` operator (see Example 1).

Definition 5. Given two representations of parallelotopes $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}' = \langle A', \mathbf{l}', \mathbf{u}' \rangle$, we say $\mathcal{P} \leq \mathcal{P}'$ iff $\mathbf{l}' \leq \mathbf{l}''$ and $\mathbf{u}'' \leq \mathbf{u}'$, where $\langle A', \mathbf{l}'', \mathbf{u}'' \rangle = \text{rot}_{A'}(\mathcal{P})$. Moreover, $\epsilon \leq \mathcal{P}$ for each representation \mathcal{P} .

Note that there is no ambiguity in using the same symbol \leq as for the ordering over Par_A^n . If $A = A'$, then $\text{rot}_{A'}(\langle A, \mathbf{l}, \mathbf{u} \rangle) = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and the two definitions coincide.

Theorem 6. Given two representations \mathcal{P} and \mathcal{P}' , we have that $\mathcal{P} \leq \mathcal{P}'$ iff $\gamma(\mathcal{P}) \subseteq \gamma(\mathcal{P}')$. Moreover, \leq is a pre-order. The computational complexity of deciding \leq is $O(n^3)$.

Theorem 6 allows us to state the precise abstract framework we use in this paper. The concrete domain is the powerset of \mathbb{R}^n , ordered by set inclusion, while the abstract domain is the set Par_n of all the representations of parallelotopes, according to Definition 1, ordered by \leq from Definition 5. The two domains are related by the monotonic concretization map γ of Definition 2.

3.6. On inverse matrices and factorizations

Most of the operations for parallelotopes require either computing the inverse of the constraint matrix A or solving a system of linear equations. In our implementation we use Gaussian elimination in both cases,

Solving systems of linear equations may be done more efficiently if we know a factorization (such as LU or QR) of the constraint matrix. Therefore, memoizing this factorization may be useful to improve efficiency. Moreover, it could be investigated whether it is possible to redefine the abstract operators in order to incrementally and automatically keep track of a factorization for the constraint matrix. This is left for future work.

4. Operators on parallelotopes

In this section we consider the operations on $\mathcal{P}(\mathbb{R}^n)$ commonly used when defining the collecting semantics of imperative programming languages, and

for each of them we introduce a correct approximation on parallelotopes. The main difficulty in defining the abstract operations on parallelotopes is that, in the general case, a subset of \mathbb{R}^n has not a best correct abstraction in Par_n . Different constraint matrices give rise to minimal but incomparable abstractions. Therefore, defining the abstract operations on parallelotopes essentially amounts to:

1. choosing a resulting constraint matrix A ;
2. computing bounds \mathbf{l} and \mathbf{u} to get a correct approximation of the result.

Once A is fixed, computing the tightest bounds is quite easy. In practice, the two steps of determining A and \mathbf{l}, \mathbf{u} will be carried out at the same time, for the sake of efficiency. However, there are generally several possible alternatives for A , which lead to results which are set-theoretically incomparable. The choice between different matrices may only be done under the basis of heuristic considerations, and validation requires extensive tests.

From a theoretical perspective, in evaluating the precision of an abstract operator we look for the following properties, in order of preference:

1. *γ -completeness*, if possible, i.e. when the result of the concrete operator is a parallelotope;
2. *minimality*, i.e. we compute one of the minimal parallelotopes which approximate the concrete result;
3. *relative optimality*, i.e. we fix a matrix A and compute the least parallelotope definable over A which approximates the concrete result.

It is easy to check that γ -completeness implies minimality which, in turn, implies relative optimality.

In general, our abstract operators will not be monotone. Given two inputs $\mathcal{P} \leq \mathcal{P}'$, different choices of the constraint matrix of the corresponding results may lead to non-monotonicity. However, abstract operators are monotone when they are γ -complete.

4.1. Invertible linear assignment

Linear assignment are used to analyze the behavior of the statement $x_i = c_1x_1 + \dots + c_nx_n + b$. The concrete linear assignment operation $\text{assign}(i, \mathbf{c}, b) : \mathcal{P}(\mathbb{R}^n) \rightarrow \mathcal{P}(\mathbb{R}^n)$ is defined as

$$\text{assign}(i, \mathbf{c}, b)(X) = \{\mathbf{x}[i \mapsto \mathbf{c} \cdot \mathbf{x} + b] \mid \mathbf{x} \in X\}.$$

If $c_i \neq 0$, then $\text{assign}(i, \mathbf{c}, b)$ is invertible and, most importantly, maps a parallelotope to a parallelotope. In this special case, it is possible to implement the abstract operator along the line of [1]. Intuitively, the operation $\text{assign}(i, \mathbf{c}, b)$ corresponds to the assignment $x'_i = \mathbf{c} \cdot \mathbf{x} + b$, where x'_i is the value of x_i after the assignment. From this equation, it is possible to recover x_i as a function of x'_i and the other elements of \mathbf{x} , namely $x_i = (x'_i - \sum_{j \neq i} c_j x_j - b) / c_i$. Replacing the variable x_i in $\mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}$ with the latter definition, we obtain a representation for the result of the assignment.

Definition 6. Given $\mathbf{c} \in \mathbb{R}^n$ such that $c_i \neq 0$ and $b \in \mathbb{R}$, we define the abstract linear assignment $\text{assign}^\alpha(i, \mathbf{c}, b) : \text{Par}_n \rightarrow \text{Par}_n$ as

$$\begin{aligned} \text{assign}^\alpha(i, \mathbf{c}, b)(\epsilon) &= \epsilon \\ \text{assign}^\alpha(i, \mathbf{c}, b)(\langle A, \mathbf{l}, \mathbf{u} \rangle) &= \langle A - \frac{1}{c_i} \mathbf{a}_{*i} (\mathbf{c} - \mathbf{e}^i)^T, \mathbf{l} + \frac{b}{c_i} \mathbf{a}_{*i}, \mathbf{u} + \frac{b}{c_i} \mathbf{a}_{*i} \rangle . \end{aligned}$$

Theorem 7. *The operation $\text{assign}^\alpha(i, \mathbf{c}, b)$ is correct and γ -complete. The computational complexity is $O(mn)$ where m is the number of non-zero components in \mathbf{c} .*

The case when $c_i = 0$ will be treated after the non-deterministic assignment.

4.2. Non-deterministic assignment

Consider the non-deterministic assignment operation $\text{forget}(i) : \mathcal{P}(\mathbb{R}^n) \rightarrow \mathcal{P}(\mathbb{R}^n)$, corresponding to the statement $x_i = ?$ and defined as

$$\text{forget}(i)(X) = \{\mathbf{x}[i \mapsto v] \mid \mathbf{x} \in X, v \in \mathbb{R}\} .$$

Note that, even if P is a parallelotope, $\text{forget}(i)(P)$ may not be parallelotope.

Example 2. Consider the parallelotope P given by the inequalities $-1 \leq x_1 + x_2 + x_3 \leq 1$, $-1 \leq x_1 + x_2 - x_3 \leq 1$ and $-1 \leq x_1 - x_2 + x_3 \leq 1$. The set $\text{forget}(1)(P)$ is a polyhedron whose constraints may be obtained by Fourier-Motzkin elimination:

$$\text{forget}(1)(P) = \{\mathbf{x} \in \mathbb{R}^3 \mid -1 \leq x_2 \leq 1, -1 \leq x_3 \leq 1, -1 \leq x_2 - x_3 \leq 1\} .$$

Although P is a parallelotope, $\text{forget}(1)(P)$ is not.

Given a matrix A , we can characterize in a syntactic way the parallelotopes definable over A which approximate the set $\text{forget}(i)(P)$. A parallelotope approximates $\text{forget}(i)(P)$ iff it approximates P and any constraint containing the variable x_i is unbounded.

Proposition 8. *Given two parallelotopes $\langle A, \mathbf{l}, \mathbf{u} \rangle \leq \langle A', \mathbf{l}', \mathbf{u}' \rangle$, we have that*

$$\text{forget}(i)(\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)) \subseteq \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle) \text{ iff } \forall j (a'_{ji} \neq 0 \rightarrow (l'_j = -\infty \wedge u'_j = +\infty)) .$$

The above proposition suggests that, in order to compute $\text{forget}^\alpha(i)(\mathcal{P})$, we need to find a parallelotope $\langle A', \mathbf{l}', \mathbf{u}' \rangle \geq \mathcal{P}$ such that each row of A' which contains a non-zero entry in the i -th position must be unconstrained. A naive definition of $\text{forget}^\alpha(i)(\langle A, \mathbf{l}, \mathbf{u} \rangle)$ would replace the bounds of the rows j such that $a_{ji} \neq 0$ with $-\infty$ and $+\infty$. However, this generally yields a gross over approximation.

Example 3. Consider the parallelotope given by the inequalities $0 \leq x_1 + x_2 \leq 0$ and $0 \leq x_1 - x_2 \leq 0$, which consists of a single point $\{(0, 0)\}$. After a non-deterministic assignment to x_2 , if we apply the naive procedure described above, we get $-\infty \leq x_1 + x_2 \leq +\infty$ and $-\infty \leq x_1 - x_2 \leq +\infty$ which is the entire space. However, by adding the two original inequalities, we get the new constraint $0 \leq 2x_1 \leq 0$, which does not contain x_2 and thus is preserved by the non-deterministic assignment.

Therefore, we need to make explicit the constraints hidden in \mathcal{P} which do not contain the variable x_i we want to forget. The problem is that, in general, there are more entailed constraints than we can represent with a parallelotope. We need a way to choose between competing constraints.

Example 4. Consider the parallelotope P in Example 2. When computing $\text{forget}(1)(P)$ we get the implicit constraints $-1 \leq x_2 \leq 1$, $-1 \leq x_3 \leq 1$ and $-1 \leq x_2 - x_3 \leq 1$. The problem is that the linear forms x_2 , x_3 and $x_2 - x_3$ are not linearly independent, hence we cannot keep all of them in the result. Note that although $x_2 - x_3$ is a linear combination of x_2 and x_3 , the constraint $-1 \leq x_2 - x_3 \leq 1$ is not implied by the other two: we lose precision when we omit one of them.

In order to deal with the above problems, we propose the following approach, which is detailed in Algorithm 1. First note that we may ignore the rows in \mathcal{P} which are unconstrained or whose i -th entry is zero: the first ones remain unconstrained, while the second ones are not affected by the assignment. Thus, in the following, we focus on the remaining rows only, whose indexes are in $J = \{j \mid a_{ji} \neq 0, l_j \neq -\infty \vee u_j \neq +\infty\}$. The idea is to transform the rows in J in such a way that they remain linearly independent and there is exactly one row whose i -th entry is not zero. Thus, we choose a row $r \in J$ and consider the linear combinations $R = \{a_{ji}\mathbf{a}_{r*} - a_{ri}\mathbf{a}_{j*} \mid j \in J \setminus \{r\}\} \cup \{\mathbf{a}_{r*}\}$. Vectors in R are linearly independent: any of its linear combination has the form

$$\sum_{j \in J \setminus \{r\}} \lambda_j (a_{ji}\mathbf{a}_{r*} - a_{ri}\mathbf{a}_{j*}) + \lambda_r \mathbf{a}_{r*} = \sum_{j \in J \setminus \{r\}} -a_{ri}\lambda_j \mathbf{a}_{j*} + \left(\lambda_r + \sum_{j \in J \setminus \{r\}} \lambda_j a_{ji} \right) \mathbf{a}_{r*} .$$

Since $a_{ri} \neq 0$, this linear combination is zero only if all λ_j with $j \in J$ are zero. Thus, R is a set of $|J|$ linearly independent vectors. Moreover, all the vectors in R are independent from the rows of A not in J . Combining them together, we get the resulting matrix A' . The last step is to ensure that $\mathcal{P}' \geq \mathcal{P}$, by computing the new bounds (with the exception of \mathbf{a}_{r*} which must become an unconstrained row).

The main question is how to choose the index r in J . Our intuition is that it is better to choose an index r such that both l_r and u_r are finite, possibly equal,

Algorithm 1 The $\text{forget}^\alpha(i)$ abstract operator

Require: $\langle A, \mathbf{l}, \mathbf{u} \rangle \in \text{Par}_n, i \in \{1, \dots, n\}$

- 1: $J = \{j \mid a_{ji} \neq 0, l_j \neq -\infty \vee u_j \neq +\infty\}$
- 2: **if** $J = \emptyset$ **then**
- 3: **return** $\langle A, \mathbf{l}, \mathbf{u} \rangle$
- 4: **end if**
- 5: $\langle A', \mathbf{l}', \mathbf{u}' \rangle \leftarrow \langle A, \mathbf{l}, \mathbf{u} \rangle$
- 6: $J_0 \leftarrow \{j \in J \mid l_j = u_j \in \mathbb{R}\}$
- 7: $J_1 \leftarrow \{j \in J \mid l_j, u_j \in \mathbb{R}\}$
- 8: **if** $J_0 \neq \emptyset$ **then**
- 9: $r \leftarrow$ an element in J_0
- 10: **else if** $J_1 \neq \emptyset$ **then**
- 11: $r \leftarrow$ an element in J_1
- 12: **else**
- 13: $r \leftarrow$ an element in J
- 14: **end if**
- 15: **for all** $j \in J \setminus \{r\}$ **do**
- 16: $\mathbf{a}'_{j*} \leftarrow a_{ji}\mathbf{a}_{r*} - a_{ri}\mathbf{a}_{j*}$
- 17: $(m_r, M_r) \leftarrow$ **if** $a_{ji} < 0$ **then** (u_r, l_r) **else** (l_r, u_r)
- 18: $(m_j, M_j) \leftarrow$ **if** $-a_{ri} < 0$ **then** (u_j, l_j) **else** (l_j, u_j)
- 19: $l'_j \leftarrow a_{ji}m_r - a_{ri}m_j$
- 20: $u'_j \leftarrow a_{ji}M_r - a_{ri}M_j$
- 21: **end for**
- 22: **return** $\langle A', \mathbf{l}'[r \mapsto -\infty], \mathbf{u}'[r \mapsto +\infty] \rangle$

since we get better bounds in \mathcal{P}' . In fact, we prove later that when we choose a row r whose lower and upper bounds coincide, then $\text{forget}^\alpha(i)$ is γ -complete.

Theorem 9. *The operator $\text{forget}^\alpha(i)$ described in Algorithm 1 is correct and relatively optimal. The computational complexity is $O(n^2)$.*

Given a parallelotope $P = \gamma(\mathcal{P})$, note that either $\text{forget}(i)(P)$ is equal to P or its linearity space is the sum of the linearity space of P plus \mathbf{e}^i . In both cases, it is easy to check that the linearity space of $\gamma(\text{forget}^\alpha(i)(\mathcal{P}))$ is the same as the linearity space of $\text{forget}(i)(P)$.

Although in the general case the result of $\text{forget}(i)(P)$ is not a parallelotope, in the particular case that $J_0 \neq \emptyset$ in Algorithm 1, $\text{forget}^\alpha(i)(\mathcal{P})$ is γ -complete. The same holds if \mathbf{e}^i is a line in P , i.e., $J = \emptyset$ in Algorithm 1.

Theorem 10. *If \mathbf{e}^i is a line in \mathcal{P} or $J_0 \neq \emptyset$, then it holds that $\gamma(\text{forget}^\alpha(i)(\mathcal{P})) = \text{forget}(i)(\gamma(\mathcal{P}))$.*

In the general case, we can prove that the $\text{forget}^\alpha(i)$ operator is minimal using the Fourier-Motzkin elimination method [12, Chapter 12.2].




Figure 3: Non-invertible assignment

Proposition 11. *Given a parallelotope $\langle A, \mathbf{l}, \mathbf{u} \rangle$ and a column index i , assume that $l_j < u_j$ for each constrained row j such that $a_{ji} \neq 0$. Consider the polyhedron obtained by Fourier-Motzkin elimination of the i -th variable. None of its constraints with finite bounds is redundant.*

Theorem 12. *The operator $\text{forget}^\alpha(i)$ is minimal.*

It is possible to think of other heuristics for selecting the implied constraints which appear in the result. For example, sparse constraints might be preferred to dense ones. This would require further investigations.

4.3. Non-invertible assignment

We consider the assignment operator $\text{assign}(i, \mathbf{c}, b)$ when $c_i = 0$. In this case, all the constraints involving the variable x_i need to be removed after the assignment, possibly replaced with other implied constraints. Note that if $c_i = 0$ we have $\text{assign}(i, \mathbf{c}, b) = \text{assign}(i, \mathbf{c}, b) \circ \text{forget}(i)$. This suggests to use the abstract forget operation to make implied constraints explicit.

Example 5. Given the parallelotope

$$\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n \mid 0 \leq -x_1 + x_2 \leq 2, 3 \leq x_1 + 2x_2 \leq 8\}$$

depicted in Figure 3(a), the parallelotope $\text{forget}^\alpha(x_1)(\mathcal{P})$ is depicted in Figure 3(b). The result of the non-invertible assignment $x_1 = 3$ to the parallelotope \mathcal{P} is in Figure 3(c).

The procedure is shown in Algorithm 2. Given a parallelotope \mathcal{P} , we first compute $\langle A', \mathbf{l}', \mathbf{u}' \rangle = \text{forget}^\alpha(i)(\mathcal{P})$ and choose a row j in A' with $a'_{ji} \neq 0$, which is certainly unconstrained. Lines 3–5 ensure that \mathbf{a}'_{j*} is the unique line with a non-zero i -th element. They do not change the parallelotope, since they operate on unconstrained rows. Then, we may replace \mathbf{a}'_{j*} with $\mathbf{e}^i - \mathbf{c}$, and both l_j, u_j with b . Since the j -th row of A' is unconstrained, we do not lose precision when we replace it. Since the steps 3–5 ensure that all the other rows in A' have a zero in the i -th column, then $\mathbf{e}^i - \mathbf{c}$ is linearly independent from them and the final matrix A' is invertible.

Algorithm 2 The non-invertible $\text{assign}^\alpha(i, \mathbf{c}, b)$ abstract operator

Require: $\langle A, \mathbf{l}, \mathbf{u} \rangle \in \text{Par}_n$, $i \in \{1, \dots, n\}$, $\mathbf{c} \in \mathbb{R}^n$, $b \in \mathbb{R}$, $c_i = 0$

- 1: $\langle A', \mathbf{l}', \mathbf{u}' \rangle \leftarrow \text{forget}^\alpha(i)(\langle A, \mathbf{l}, \mathbf{u} \rangle)$
 - 2: choose an index j with $a'_{ji} \neq 0$
 - 3: **for all** s such that $a'_{si} \neq 0$ and $s \neq j$ **do**
 - 4: $\mathbf{a}'_{s*} \leftarrow \mathbf{a}'_{s*} - a'_{si}/a'_{ji} \mathbf{a}'_{j*}$
 - 5: **end for**
 - 6: $\mathbf{a}'_{j*} \leftarrow \mathbf{e}^i - \mathbf{c}$
 - 7: $\mathbf{l}'_j \leftarrow b$
 - 8: $\mathbf{u}'_j \leftarrow b$
 - 9: **return** $\langle A', \mathbf{l}', \mathbf{u}' \rangle$
-

Theorem 13. *The operator $\text{assign}^\alpha(i, \mathbf{c}, b)$ described in Algorithm 2 is correct. The computational complexity is $O(n^2)$.*

Note that non-invertible assignment is γ -complete exactly in the same hypothesis under which non-deterministic assignment is γ -complete.

Theorem 14. *Consider $\mathbf{c} \in \mathbb{R}^n$, $b \in \mathbb{R}$ and $c_i = 0$. If \mathcal{P} is a parallelotope in \mathbb{R}^n and \mathbf{e}^i is a line in \mathcal{P} or $J_0 \neq \emptyset$ in Algorithm 1 then $\gamma(\text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P})) = \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P}))$.*

Theorem 15. *The operator $\text{assign}^\alpha(i, \mathbf{c}, b)$ with $c_i = 0$ is minimal.*

Even for non-invertible assignments, heuristics can be used to improve the result of the analysis. For example, it is possible that, in order to gain more benefits from the result of $\text{assign}^\alpha(i, \mathbf{c}, b)$, it is better to tweak $\text{forget}^\alpha(i)$ in order to get the best possible bounds for the row \mathbf{c} . However, we have not investigated this idea in details.

4.4. Refinement by linear inequality

Given $\mathbf{c} \in \mathbb{R}^n$ and $b \in \mathbb{R}$, consider the operation over $\mathcal{P}(\mathbb{R}^n)$ defined by

$$\text{refine}(\mathbf{c}, b)(X) = \{\mathbf{x} \mid \mathbf{x} \in X \wedge \mathbf{c} \cdot \mathbf{x} \leq b\} ,$$

which we call *linear refinement*, corresponding to the conditional statement “if $\mathbf{c} \cdot \mathbf{x} \leq b$ then”. In general, the linear refinement of a parallelotope is not a parallelotope.

Consider Algorithm 3. Given a parallelotope $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$, we first check if there exists an unconstrained row of A such that, if we replace that row with \mathbf{c} , the matrix is still invertible. This amounts to computing a vector $\mathbf{y} \in \mathbb{R}^n$ such that $A^T \mathbf{y} = \mathbf{c}$ and look for an index $j \in U(\mathcal{P})$ such that $y_j \neq 0$. When it does not exist, we simply compute the least parallelotope \mathcal{P}' containing $\text{refine}(\mathbf{c}, b)(\mathcal{P})$ and definable over A : we recall from [10] that $\mathcal{P}' = \langle A, \mathbf{l}', \mathbf{u}' \rangle$, where $\langle I_n, \mathbf{l}', \mathbf{u}' \rangle = \text{refine}(\mathbf{y}, b)\langle I_n, \mathbf{l}, \mathbf{u} \rangle$. Hence, we may use the known refine operator over boxes to define a refine operator over parallelotopes.

Algorithm 3 The $\text{refine}^\alpha(\mathbf{c}, b)$ abstract operator

Require: $\langle A, \mathbf{l}, \mathbf{u} \rangle \in \text{Par}_n$, $\mathbf{c} \in \mathbb{R}^n$, $b \in \mathbb{R}$

- 1: $\mathbf{y} \leftarrow$ solution of $A^T \mathbf{y} = \mathbf{c}$
- 2: **if** $\exists j. y_j \neq 0 \wedge l_j = -\infty \wedge u_j = +\infty$ **then**
- 3: $\mathbf{a}_{j*} \leftarrow \mathbf{c}^T$
- 4: $u_j \leftarrow b$
- 5: **return** $\langle A, \mathbf{l}, \mathbf{u} \rangle$
- 6: **else**
- 7: $\langle I_n, \mathbf{l}', \mathbf{u}' \rangle \leftarrow \text{refine}^\alpha(\mathbf{y}, b) \langle I_n, \mathbf{l}, \mathbf{u} \rangle$ {using operator on boxes}
- 8: **return** $\langle A, \mathbf{l}', \mathbf{u}' \rangle$
- 9: **end if**

Theorem 16. *The operator $\text{refine}^\alpha(\mathbf{c}, b)$ described in Algorithm 3 is correct and relatively optimal. If $\mathbf{c} \notin \text{lin}(\gamma(\mathcal{P}))^\perp$, then $\text{refine}^\alpha(\mathbf{c}, b)(\mathcal{P})$ is γ -complete. The computational complexity is $O(n^3)$.*

We could change the algorithm above in such a way that, when there is no $j \in U(\mathcal{P})$ with $y_j \neq 0$, we choose any j such that $y_j \neq 0$ and either l_j or u_j is infinite. Then, we proceed with lines 3–5 as if j were in $U(\mathcal{P})$. In this case, we replace a bounded constraint with a different bounded constraint. The result is generally incomparable with the original parallelotope \mathcal{P} . However, the rationale here is that the new constraint could be more useful in the remaining of the analysis than the old one which has been removed. We have not investigated in practice the impact of such a different approach to linear refinement.

4.5. Union

Let us come to the abstract union of parallelotopes. The result is obvious if one of the parallelotopes is empty. In the rest of this section and in the following ones we assume to be in the non-empty case. The same If we fix a matrix M , the least parallelotope definable over M which contains the parallelotopes $\mathcal{P}_A = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}_B = \langle B, \mathbf{j}, \mathbf{k} \rangle$ can be easily obtained by $\text{rot}_M(\mathcal{P}_A)$ and $\text{rot}_M(\mathcal{P}_B)$ simply selecting, for each row in M , the minimum of the two lower bounds and the maximum of the two upper bounds. By choosing M to be either A or B , we can use this method for a simple and fast implementation of abstract union. The biggest drawback of this choice is that it does not generate new constraints. The ability to generate new constraints precisely allows us to fully exploit the power of a non-template abstract domain, by changing the shape of the generated parallelotopes.

We now propose a more complex variant of abstract union, inspired by the *inversion join* operator [14]. The main idea of the algorithm is to generate a bunch of candidate linear forms. The corresponding constraints are obtained from the candidate linear forms by computing the lower and upper bounds on \mathcal{P}_A and \mathcal{P}_B . We then assign to each constraint a priority. In general, the candidate linear forms are not linearly independent. At the end, we select exactly n linearly independent constraints, according to their priorities.

Assigning the priorities to the candidate constraints is the key of the abstract union algorithm. As a first consideration, assigning the priorities must take into account whether the domain is used in combination with another abstract domain, thus allowing better optimizations. In particular, in the case of a reduced product of parallelotopes and intervals, it is immediate to see that constraints with a single non-null coefficient (interval constraints) are not useful, and thus we should avoid them. On the contrary, when parallelotopes are used alone, it is very convenient to choose such constraints. Thus, we have developed three slightly different algorithms for assigning priorities, according to whether we prefer to retain interval constraints or not. Choosing the algorithm which best fits depends on the analysis context.

We use a flag *favorAxes* to differentiates the behavior of the three algorithms: when the flag is set to 1, we assign to the interval constraints the highest priority, so that they are always preferred in the abstract union operation. When the flag is set to -1 , we assign to the interval constraints the lowest priority, so that the heuristic algorithm tries to avoid them, as much as reasonable. When the flag is set to 0, we assign priorities to intervals using the same algorithm as for non-interval constraints.

The priority of non-interval constraints is chosen according to the values of the bounds: equality constraints come first, followed by constraints which are saturated both in \mathcal{P}_A and \mathcal{P}_B . This order is mostly dictated by heuristic considerations. Algorithm 4 computes the bounds and assigns the priorities for a given linear form \mathbf{v} . The lines "*{other tests}*" in Algorithm 4 refers to a set of tests where we consider all the possible combinations of finite and infinite lower and upper bounds, in order to better assign priorities.

We now describe how to generate the candidate linear forms for the abstract union. Obvious candidates are the rows of the matrices A and B . Moreover, we also generate new linear forms using (a part of) the inversion join algorithm. Given two constraints in \mathcal{P}_A and/or \mathcal{P}_B , the inversion join computes a new linear form obtained as a linear combination of the two constraints, under the condition that they form an inversion.

Definition 7. The constraints $l_h \leq \mathbf{a}_{h*} \cdot \mathbf{x}$ taken from the parallelotope $\langle A, \mathbf{l}, \mathbf{u} \rangle$ and $j_i \leq \mathbf{b}_{i*} \cdot \mathbf{x}$ taken from $\langle B, \mathbf{j}, \mathbf{k} \rangle$ form an *inversion* when, provided $\langle A, \mathbf{j}', \mathbf{k}' \rangle = \text{rot}_A(\mathcal{P}_B)$, the following conditions hold:

1. all the bounds l_h, j'_h, l_i, j'_i are finite;
2. \mathbf{a}_{h*} and \mathbf{b}_{i*} are linearly independent;
3. $l_h < j'_h$ and $l_i > j'_i$ (or vice-versa, that is $l_h > j'_h$ and $l_i < j'_i$).

Whenever we find an inversion, we may generate the new constraint $\lambda_{hi} j'_h \leq (\mathbf{a}_{h*} + \lambda_{hi} \mathbf{b}_{i*}) \cdot \mathbf{x}$, where $\lambda_{hi} = \frac{l_h - j'_h}{j'_i - l_i}$. The same procedure can also be applied to upper bounds, just considering $\mathbf{a}_{h*} \cdot \mathbf{x} \leq u_h$ as equivalent to $-u_h \leq -\mathbf{a}_{h*} \cdot \mathbf{x}$. The complete procedure is illustrated in Algorithm 5.

Algorithm 4 Bounds and priorities for the linear form \mathbf{v} (Sketch)

Require: $\mathcal{P}_A, \mathcal{P}_B \in \text{Par}_n$, $\mathbf{v} \in \mathbb{R}^n$, $\text{favorAxes} \in \{-1, 0, 1\}$

- 1: $l_v \leftarrow \inf\{\mathbf{v} \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_A\}$
- 2: $u_v \leftarrow \sup\{\mathbf{v} \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_A\}$
- 3: $j_v \leftarrow \inf\{\mathbf{v} \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_B\}$
- 4: $k_v \leftarrow \sup\{\mathbf{v} \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_B\}$
- 5: **if** \mathbf{v} is an interval constraint and $\text{favorAxes} \neq 0$ **then**
- 6: **if** $\text{favorAxes} = 1$ **then**
- 7: $p \leftarrow -\infty$ // this is the highest priority
- 8: **else if** $\text{favorAxes} = -1$ **then**
- 9: $p \leftarrow \infty$ // this is the lowest priority
- 10: **end if**
- 11: **else**
- 12: **if** $l_v = u_v = j_v = k_v$ **then**
- 13: $p \leftarrow 0$
- 14: **else if** $l_v = j_v \in \mathbb{R}$ and $u_v = k_v \in \mathbb{R}$ **then**
- 15: $p \leftarrow 1$
- 16: **else if** ... **then**
- 17: : {other tests}
- 18: **else**
- 19: $p \leftarrow +\infty$
- 20: **end if**
- 21: **end if**
- 22: **return** $\langle \min(l_v, j_v), \max(u_v, k_v), p \rangle$

Example 6. Given the parallelotopes

$$\begin{aligned}\mathcal{P}_1 &= \{\mathbf{x} \in \mathbb{R}^n \mid 1 \leq x_1 \leq 2, 2 \leq x_2 \leq 4\} \\ \mathcal{P}_2 &= \{\mathbf{x} \in \mathbb{R}^n \mid 2 \leq x_1 \leq 3, 1 \leq x_2 \leq 3\}\end{aligned}$$

depicted in Figure 4 in grey and light-grey respectively, $\mathcal{P}_1 \cup^\alpha \mathcal{P}_2$ is the dotted parallelotope.

Theorem 17. *The abstract union operator described in Algorithm 5 is correct and relatively optimal. The computational complexity is $O(n^4)$.*

4.6. Widening

We use the standard definition of widening which appears in [15] since, although developed for monotonic abstract operators, it also works in the non-monotonic case [16].

A first widening for parallelotopes, denoted by ∇ , is essentially inspired by the standard widening on intervals. In order to compute $\mathcal{P}_A \nabla \mathcal{P}_B$, we first compute $\mathcal{P}' = \text{rot}_A(\mathcal{P}_B)$. Each bound of \mathcal{P}_A which has been enlarged in \mathcal{P}' is changed into $+\infty$ or $-\infty$.

Algorithm 5 The abstract union operator

Require: $\mathcal{P}_A = \langle A, \mathbf{l}, \mathbf{u} \rangle, \mathcal{P}_B = \langle B, \mathbf{j}, \mathbf{k} \rangle \in \text{Par}_n$

- 1: $Q \leftarrow \emptyset$ {a priority queue}
 - 2: **for all** $i \in \{1, \dots, n\}$ **do**
 - 3: $\langle c, d, p \rangle \leftarrow$ result of Algorithm 4 applied to \mathbf{a}_{i*}
 - 4: add $\langle \mathbf{a}_{i*}, c, d \rangle$ to Q with priority p
 - 5: **end for**
 - 6: same procedure of lines 2–5 applied to rows in B
 - 7: **for all** $\mathbf{v}_1, \mathbf{v}_2$ rows of A and B **do**
 - 8: {here we check if \mathbf{v}_1 and \mathbf{v}_2 form an inversion}
 - 9: $h_1 \leftarrow \inf\{\mathbf{v}_1 \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_A\}$
 - 10: $h_2 \leftarrow \inf\{\mathbf{v}_2 \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_A\}$
 - 11: $i_1 \leftarrow \inf\{\mathbf{v}_1 \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_B\}$
 - 12: $i_2 \leftarrow \inf\{\mathbf{v}_2 \mathbf{x} \mid \mathbf{x} \in \mathcal{P}_B\}$
 - 13: **if** $h_1, i_1, h_2, i_2 \in \mathbb{R}$ and $\mathbf{v}_1, \mathbf{v}_2$ are linearly independent
and $((h_1 < i_1 \wedge h_2 > i_2) \vee (h_1 > i_1 \wedge h_2 < i_2))$ **then**
 - 14: {we know that \mathbf{v}_1 and \mathbf{v}_2 form an inversion}
 - 15: $\mathbf{w} \leftarrow \mathbf{v}_1 + \frac{h_1 - i_1}{i_2 - h_2} \mathbf{v}_2$ { \mathbf{w} is the linear form obtained by inversion join}
 - 16: $\langle c, d, p \rangle \leftarrow$ result of Algorithm 4 applied to \mathbf{w}
 - 17: add $\langle \mathbf{w}, c, d \rangle$ to Q with priority p
 - 18: **end if**
 - 19: **end for**
 - 20: same procedure of lines 7–18 applied to upper bounds
 - 21: same proc. of lines 7–18 applied to lower bounds for \mathbf{v}_1 and upper bounds
for \mathbf{v}_2
 - 22: same proc. of lines 7–18 applied to upper bounds for \mathbf{v}_1 and lower bounds
for \mathbf{v}_2
 - 23: $\langle R, \mathbf{l}', \mathbf{u}' \rangle \leftarrow$ empty set of constraints
 - 24: **while** $|R| < n$ **do**
 - 25: extract $\langle \mathbf{w}, c, d \rangle$ from Q with maximal priority
 - 26: **if** \mathbf{w} is linearly independent from R **then**
 - 27: add $\langle \mathbf{w}, c, d \rangle$ to $\langle R, \mathbf{l}', \mathbf{u}' \rangle$
 - 28: **end if**
 - 29: **end while**
 - 30: **return** $\langle R, \mathbf{l}', \mathbf{u}' \rangle$
-




Figure 4: Union of parallelotopes \mathcal{P}_1 and \mathcal{P}_2 .

Definition 8. Given two parallelotopes $\mathcal{P}_A = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}_B = \langle B, \mathbf{j}, \mathbf{k} \rangle$, let $\langle A, \mathbf{j}', \mathbf{k}' \rangle = \text{rot}_A(\mathcal{P}_B)$. We define $\mathcal{P}_A \nabla \mathcal{P}_B = \langle A, \mathbf{l}', \mathbf{u}' \rangle$ where for each $i \in \{1, \dots, n\}$, we have that:

$$l'_i = \begin{cases} -\infty & \text{if } j'_i < l_i \\ l_i & \text{otherwise} \end{cases} \quad u'_i = \begin{cases} \infty & \text{if } k'_i > u_i \\ u_i & \text{otherwise.} \end{cases}$$

The problem with this widening is that it cannot change the constraint matrix. For this reason, it is better to use it with delayed widening. Even a small delay allows the constraint matrix in loops to stabilize to a good value. Then, when widening kicks in, constraints cannot change anymore.

We introduce an alternative widening operator which tries to change the constraint matrix while avoiding infinite ascending chains and that will be used in the rest of the paper.

Definition 9. Given two parallelotopes $\mathcal{P}_A = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}_B = \langle B, \mathbf{j}, \mathbf{k} \rangle$, we define

$$\mathcal{P}_A \bar{\nabla} \mathcal{P}_B = \begin{cases} \text{rot}_B(\mathcal{P}_A) \nabla \mathcal{P}_B & \text{if } \text{rot}_B(\mathcal{P}_A) < \mathcal{P}_B \\ \mathcal{P}_A \nabla \mathcal{P}_B & \text{otherwise} \end{cases}$$

The above definition works because applying the rot operator does not decrease the number of infinite bounds in the parallelotope.

Theorem 18. *The operator $\bar{\nabla}$ is a widening.*

4.7. Narrowing

First of all, we give a definition of narrowing which, as we will discuss later, slightly differs from the standard one.

Definition 10 (Narrowing). Given a lattice (C, \subseteq) (the concrete domain) and a partial order (A, \leq) (the abstract domain) with monotone concretization map $\gamma : A \rightarrow C$, a *narrowing* is a binary operator $\Delta : A \times A \rightarrow A$ such that:

- $a_1 \Delta a_2 \leq a_1$;
- $\gamma(a_1) \cap \gamma(a_2) \subseteq \gamma(a_1 \Delta a_2)$;
- for each sequence a_0, \dots, a_i, \dots of elements of A , the sequence defined as $b_0 = a_0$ and $b_i = b_{i-1} \Delta a_i$ for $i > 0$ is definitively stationary.

This generalizes the standard narrowing in [15] to the case when the second argument is not bigger than the first one. The generalization is needed since abstract operators on parallelotopes are not monotone, hence the arguments to narrowing may not be in the correct relation as expected in [15]. Another generalization of narrowing, designed to work with non-monotonic abstract operators, recently appeared in [16]. However, we find the latter not completely satisfactory, since correctness of the descending iterations does not depend on the definition of narrowing alone, but on the interaction between narrowing and the abstract operators.

The first condition in Definition 10 ensures that an iteration sequence built with narrowing is decreasing, the second condition means that Δ is a correct approximation of the concrete meet, while the last condition enforces termination. When $a_2 \leq a_1$, then $\gamma(a_2) \cap \gamma(a_1) = \gamma(a_2)$, and the condition $\gamma(a_1) \cap \gamma(a_2) \subseteq \gamma(a_1 \Delta a_2)$ is enforced by the stronger condition $a_2 \leq a_1 \Delta a_2$ of the standard narrowing.

A narrowing satisfying the definition above may be used to improve the result of the ascending phase of the analysis:

Theorem 19. *In the hypothesis of Definition 10, let \bar{F} be a (possibly non-monotone) abstract operator $\bar{F} : A \rightarrow A$ which is a correct abstraction of the monotone operator $F : C \rightarrow C$. Given $a \in A$, let $c \in C$ be a fixpoint of F such that $c \subseteq \gamma(a)$. We define the iteration sequence $y_0 = a$, $y_{i+1} = y_i \Delta \bar{F}(y_i)$. Then*

- *the iteration sequence $\{y_i\}_{i \in \mathbb{N}}$ is decreasing and definitively stationary;*
- *for each $i \in \mathbb{N}$, y_i is a correct approximation of c .*

Our definition of narrowing ensures that the descending phase terminates on a correct abstraction, which might not be a post-fixpoint of the set of semantic equations.

For parallelotopes \mathcal{P}_A and \mathcal{P}_B , the narrowing is defined by first computing $\mathcal{P}' = \text{rot}_A(\mathcal{P}_B)$ and then refining with \mathcal{P}' the unbounded constraints in \mathcal{P}_A .

Definition 11. Given two parallelotopes $\mathcal{P}_A = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}_B = \langle B, \mathbf{j}, \mathbf{k} \rangle$, let $\langle A, \mathbf{j}', \mathbf{k}' \rangle = \text{rot}_A(\mathcal{P}_B)$. We define $\mathcal{P}_A \Delta \mathcal{P}_B = \langle A, \mathbf{l}', \mathbf{u}' \rangle$ where for each $i \in \{1, \dots, n\}$, we have that:

$$l'_i = \begin{cases} j'_i & \text{if } l_i = -\infty \\ l_i & \text{otherwise} \end{cases} \quad u'_i = \begin{cases} k'_i & \text{if } u_i = +\infty \\ u_i & \text{otherwise} \end{cases}$$

Theorem 20. *The operator Δ for parallelotopes is a narrowing according to Definition 10.*

5. Parametric parallelotopes

In Sections 4.5 we have designed a parametric algorithm for the abstract union depending on the flag *favorAxes*. Changing the flag we modify the priorities assigned to interval constraints, thus obtaining different parallelotopes during the analysis. We now show that the resulting three abstract union operations produce quite different results.

We have implemented a prototype of the abstract domain of parallelotopes in the static analyzer Jandom [17] and did some preliminary test on the ALICE² benchmarks [18] (plus some additional test programs which can be found in [17]). The test-suite comprises a total of 105 models with 326 program points. The models have at most 11 different program points and 10 variables.

The domain of parallelotopes Par has been implemented with the standard union, namely *favorAxes* = 0. The domain $\text{Par}_{+\text{axes}}$ is tuned with *favorAxes* = 1, thus assigning to the interval constraints the highest priority, so that they are always preferred. On the contrary, in the domain $\text{Par}_{-\text{axes}}$ we use *favorAxes* = -1, so producing parallelotopes which hardly use interval constraints.

Figure 5 shows the results of the experimental comparison between the standard parallelotope Par , $\text{Par}_{+\text{axes}}$ and $\text{Par}_{-\text{axes}}$. The experimental results show that in about one third of the experiments we get incomparable results. This is due to the fact that we deal with a non-template numerical domain, and thus the shape of the result may considerably vary when changing the abstract operations. The experiments show that the domain with the best performance in precision is clearly $\text{Par}_{+\text{axes}}$, which improves over the results of both Par and $\text{Par}_{-\text{axes}}$. This result was largely expected, since variable initialization is an important source of information which should be carried on during the analysis. Thus, when the abstract domain of parallelotopes is used alone, the preferred version should certainly be $\text{Par}_{+\text{axes}}$.

On the other hand, we show in the next section that things change when considering a reduced product of $\text{Par}_{-\text{axes}}$ with intervals which dramatically improves the analysis precision.

6. Reduced product

Although the abstract domain of parallelotopes is very expressive and can potentially represent any linear invariant, it is limited by the reduced number of constraints allowed in the constraint matrix. In order to overcome this limit, we strongly believe that parallelotopes should be used in combination with another simple, even non-relational, numerical abstract domain. In particular, the combination of a non-template, fully relational domain with a rich expressiveness,

²<http://alice.cri.enscm.fr/>




Figure 5: Comparing parallelotopes with different flags *favorAxes*. On the left we compare Par with Par_{-axes}, on the right Par with Par_{+axes}.

such as parallelotopes, and a template, weakly relational domain, able to carry on the basic information at least on interval constraints, allows parallelotopes to exploit its full potential.

In this section we define a simple reduced product between parallelotopes and intervals. We denote by **Box** the domain of intervals with the standard operations [2]. Since intervals are parallelotopes (whose constraint matrix is the identity matrix), we abuse notation and use the same denotations for the concretization map and the abstract operations.

6.1. Representation of reduced product

The (reduced) product of the domains **Par** and **Box** is defined as

$$\text{Par} \sqcap \text{Box} = \{\langle P, B \rangle \mid P \in \text{Par}, B \in \text{Box}\}$$

and the concretization map is simply given by

$$\gamma(\langle P, B \rangle) = \gamma(P) \cap \gamma(B) .$$

6.2. Abstract operators on the reduced product

The abstract operations on **Par** \sqcap **Box** are obtained by combining the operations of the original domains, with the help of a *reduction* operator responsible for transferring information between the two domains.

We first need to introduce a new abstract operator of *weak intersection* on parallelotopes which is defined asymmetrically by first performing a rotation of the second argument on the constraint matrix of the first one.




Figure 6: Weak intersection of parallelotopes P_1 and P_2 .

Definition 12 (Weak intersection). Given two parallelotopes $\mathcal{P}_1 = \langle A, \mathbf{l}, \mathbf{u} \rangle$, $\mathcal{P}_2 \in \text{Par}_n$, let $\langle A, \mathbf{l}', \mathbf{u}' \rangle = \text{rot}_A(\mathcal{P}_2)$. We define the weak intersection of \mathcal{P}_1 and \mathcal{P}_2 as:

$$\mathcal{P}_1 \cap^\alpha \mathcal{P}_2 = \langle A, \mathbf{l}'', \mathbf{u}'' \rangle$$

where $l''_i = \max(l_i, l'_i)$ and $u''_i = \min(u_i, u'_i)$.

The idea of the weak intersection is to preserve the constraint matrix of the first parallelotope. Algorithm 6 shows the weak intersection operator.

Example 7. Given the parallelotopes

$$\mathcal{P}_1 = \{\mathbf{x} \in \mathbb{R}^n \mid 0 \leq -x_1 + 3x_2 \leq 2, 3 \leq x_1 + 2x_2 \leq 8\}$$

$$\mathcal{P}_2 = \{\mathbf{x} \in \mathbb{R}^n \mid 2 \leq x_1 + x_2 \leq 4, -2 \leq -x_1 + x_2 \leq 2\}$$

depicted in Figure 6, $\mathcal{P}_1 \cap^\alpha \mathcal{P}_2$ is the parallelotope filled with both vertical and horizontal lines.

Proposition 21. *The operator \cap^α is a correct approximation of the concrete intersection. It is γ -complete when the two arguments are defined over the same constraint matrix.*

We can now define the reduction operator in a standard way, by using the weak intersection operator.

Definition 13 (Reduction). Given a parallelotope $\mathcal{P} \in \text{Par}_n$ and a box $B \in \text{Box}$, we define the reduction operator $\text{red} : \text{Par} \sqcap \text{Box} \rightarrow \text{Par} \sqcap \text{Box}$ as:

$$\text{red}(\langle P, B \rangle) = \langle P \cap^\alpha B, B \cap^\alpha P \rangle$$

Proposition 22. *The reduction operator is correct, i.e., given any $P \in \text{Par}$ and $B \in \text{Box}$, we have that $\gamma(\text{red}(\langle P, B \rangle)) = \gamma(\langle P, B \rangle)$.*

Algorithm 6 The weak intersection operator

Require: $P_A = \langle A, \mathbf{l}, \mathbf{u} \rangle, P_B = \langle B, \mathbf{j}, \mathbf{k} \rangle \in \text{Par}_n$

```

1:  $\langle A, \mathbf{l}', \mathbf{u}' \rangle \leftarrow \text{rot}_A(P_B)$ 
2: for all  $i \in \{1, \dots, n\}$  do
3:    $l'_i = \max(l'_i, l_i)$ 
4:    $h'_i = \min(h'_i, h_i)$ 
5:   if  $(l'_i > h'_i)$  then
6:     return  $\epsilon$ 
7:   end if
8: end for
9: return  $\langle A, \mathbf{l}', \mathbf{u}' \rangle$ 

```

Since weak intersection is not relatively optimal, it happens that multiple iterations of the reduction operator may improve the precision of the result [19], at the cost of an increase in the computational cost of the operation.

All the operators $\odot^{P,B}$ on $\text{Par} \sqcap \text{Box}$ except widening and narrowing are simply defined componentwise, by first performing the operation on the original domain and then applying the reduction operator:

$$\langle P_1, B_1 \rangle \odot^{P,B} \langle P_2, B_2 \rangle = \text{red}(\langle P_1 \odot^P P_2, B_1 \odot^B B_2 \rangle) \quad (2)$$

Corollary 23. *All the operations on $\text{Par} \sqcap \text{Box}$ defined according to Equation 2 are correct, namely:*

$$\gamma(\langle P_1, B_1 \rangle \odot^{P,B} \langle P_2, B_2 \rangle) \supseteq \gamma(\langle P_1, B_1 \rangle) \odot \gamma(\langle P_2, B_2 \rangle)$$

The widening and narrowing operators cannot be defined using the reduction, since it may prevent termination. Thus we use a standard definition for reduced product omitting the reduction step.

Definition 14 (Widening and narrowing). Given $P_1, P_2 \in \text{Par}$ and $B_1, B_2 \in \text{Box}$, the abstract widening and narrowing on $\text{Par} \sqcap \text{Box}$ are defined as

$$\langle P_1, B_1 \rangle \nabla^{P,B} \langle P_2, B_2 \rangle = \langle (P_1 \bar{\nabla}^P P_2), (B_1 \nabla^B B_2) \rangle$$

$$\langle P_1, B_1 \rangle \Delta^{P,B} \langle P_2, B_2 \rangle = \langle (P_1 \Delta^P P_2), (B_1 \Delta^B B_2) \rangle$$

7. Evaluating the reduced product

In the previous section, we have defined the reduced product $\text{Par} \sqcap \text{Box}$. According to Section 5, we can easily design a parametric reduced product along the lines of $\text{Par} \sqcap \text{Box}$, by simply tuning the flag *favorAxes*, so obtaining the reduced products $\text{Par}_{-\text{axes}} \sqcap \text{Box}$ and $\text{Par}_{+\text{axes}} \sqcap \text{Box}$. We have implemented a prototype of the reduced products in the static analyzer Jandom and performed the test on the ALICe benchmarks. We have performed two different kinds of




Figure 7: Comparing Par with different versions of reduced product

evaluations. Firstly, we have compared the parametric versions of the reduced product with the domain Par, in order to validate our intuition that the domain $\text{Par}_{-axes} \cap \text{Box}$ is generally more precise than the other versions. Secondly, we have tested the domain $\text{Par}_{-axes} \cap \text{Box}$ versus the domains Par, Par_{+axes} and Par_{-axes} , in order to show the increment in precision that we gain when using the reduced product.

7.1. Parametric reduced products

We have compared the domain of parallelotope Par to the reduce products $\text{Par} \cap \text{Box}$, $\text{Par}_{-axes} \cap \text{Box}$ and $\text{Par}_{+axes} \cap \text{Box}$. In addition, for completeness, we have also implemented a naive analysis, denoted by $\text{Par} \cap \text{Box}$, obtained by separately performing the analysis with parallelotopes and intervals, and then taking the intersection of the results only once at the end of the two analyses.

The results of the experiments are given in Figure 7. From the experiments it emerges that the domain $\text{Par}_{-axes} \cap \text{Box}$ is more precise and more expressive, since it is able to compute more and different invariants with respect to the other domains, according to the two last columns of the bar chart which show the program points with a better behavior (147) and those where we can produce (at least a) new invariants (55). We strongly believe that the ability to found new invariants is specially important.




Figure 8: Comparing the reduced product $\text{Par}_{-axes} \sqcap \text{Box}$ with Par_{+axes} , Par and Par_{-axes} .

7.2. Reduced product vs parallelotopes

We have compared the best domain resulting from the previous section, $\text{Par}_{-axes} \sqcap \text{Box}$, with the different versions of parallelotopes Par , Par_{+axes} and Par_{-axes} , in order to show the increment in precision obtained by considering the combination with intervals. In all the tests, we have used a delayed widening and narrowing, fixing the delay to 5, i.e. we avoid applying the widening and narrowing operators during the first 5 iterations of any loop.

The comparison in Figure 8 shows that the reduced product highly improves over the results of the analysis when using the parallelotopes alone.

8. Experimental comparison with other domains

We have compared the domain $\text{Par}_{-axes} \sqcap \text{Box}$ with the polyhedra [1], octagons [3] and intervals [2] abstract domains. The tests have been conducted in the Jandom static analyzer. The comparison is mostly performed in precision, since parallelotopes and intervals are implemented natively in Jandom, while for polyhedra and octagons we use the PPL (Parma Polyhedra Library) [20] and comparing the execution times would be unfair.

We have conducted two different evaluations. First, we have compared the domain $\text{Par}_{-axes} \sqcap \text{Box}$ with the polyhedra, octagons and intervals domains. For each model in the benchmarks and for each program point, we have evaluated




Figure 9: Comparing $\text{Par}_{\text{-axes}} \sqcap \text{Box}$ with polyhedra, octagons and intervals. The first polyhedra domain is equipped with the standard polyhedra widening [1], the second one with the widening in [21].

both the shapes and the bounds of the invariants found by the analyses. In the experiments we use two different polyhedra domains, the first equipped with the standard polyhedra widening [1], the second one with the widening in [21].

The results in Figure 9 show that in about one half of the program points we obtain the same results as polyhedra and octagons. In the comparison with octagons, we obtain a better result in 38 program points (out of 326), we get worse results in 67 program points, and incomparable results in 62 program points. Thus, in about one third of the cases (100 program points), we get a better or different result than octagons. A similar conclusion holds for the comparison with polyhedra.

Even if octagons seem to be better than our domain in many program points, this is mainly due to the fact that octagons can represent a greater number of constraints. In fact, it is worth noting that this comparison measures not only the ability of the domain to discover better bounds or new invariants, but also the total number of constraints found by each domain. Of course, this is unfair with respect to the domain $\text{Par}_{\text{-axes}} \sqcap \text{Box}$, which can represent at most $2n$ constraints, while octagons can represent n^2 constraints, and the polyhedra




Figure 10: Comparing $\text{Par}_{\text{axes}} \sqcap \text{Box}$ with polyhedra, octagons and intervals on the interval constraints.

potentially any number. Things change if we are interested in (a limited number of) specific constraints, for instance interval constraints, which can be expressed in all the domains. A precise interval analysis is very helpful in many contexts, therefore it is interesting to compare the relative precision of different domains with respect to interval properties. Thus, we have compared the analysis results projecting them on the domain of intervals. The results are in Figure 10.

In this case, the domain $\text{Par}_{\text{axes}} \sqcap \text{Box}$ dramatically outperforms the other ones. This evaluation shows that, if we are interested in interval bounds only, then the domain $\text{Par}_{\text{axes}} \sqcap \text{Box}$ is strictly more precise than octagons in 69 program points (21.1%), octagons is better in 5 cases (1.5%), the results are incomparable in only 1 program point (0.3%), and we get the same result on the other program points.

Surprisingly, we get a similar result also when comparing $\text{Par}_{\text{axes}} \sqcap \text{Box}$ with polyhedra. In fact, our domain is strictly more precise than polyhedra [1] in 77 program points (23.6%), polyhedra is better in 13 cases (3.9%) and the results are incomparable in 6 program points (1.8%). Similar results hold when we use the widening in [21]. We believe that this improvement is mainly due to the widening on polyhedra, which is responsible for losing precision with respect to

Delay	3	4	5	6	7
equals	245	234	230	231	225
$\text{Par}_{\text{-axes}} \sqcap \text{Box}$ is better	58	73	77	79	85
polyhedra [1] is better	16	12	13	12	13
incomparable	7	7	6	4	3

Table 1: Comparing $\text{Par}_{\text{-axes}} \sqcap \text{Box}$ with polyhedra [1] using various widening and narrowing delays.

Domain	Time
$\text{Par}_{\text{-axes}} \sqcap \text{Box}$	3.7
polyhedra [21]	2.6
polyhedra [1]	2.4
octagons	1.8
intervals	0.2

Table 2: Comparing analysis time. Results are in seconds.

more abstract domains.

Our experiments also show that the widening and narrowing delays may greatly change the results of the comparison. In Table 1 we compare $\text{Par}_{\text{-axes}} \sqcap \text{Box}$ with polyhedra [1] on the interval constraints using delays from 3 to 7. For instance, when choosing the delay 5, we obtain the same result for 230 program points, $\text{Par}_{\text{-axes}} \sqcap \text{Box}$ is better in 77 program points, polyhedra is better in 13 program points, and in 6 we get incomparable results.

Finally, we have briefly compared the efficiency of our domain with respect to polyhedra, octagons and intervals. We believe that this is not completely fair because we use the PPL [20] for polyhedra and octagons, which is carefully engineered and optimized, so comparing their efficiency against our implementation is difficult. Moreover, the PPL is written in C++, while our analyzer Jandom [17], where we have implemented parallelotopes, box and their reduced product, is written in Scala, so that we also pay the overhead of the Java virtual machine. Table 2 shows some preliminary results. Analysis time is computed on the full set of benchmarks and it is in seconds. Because of the above consideration, and since there is still space for improvements on the abstract operators, we believe that this result is actually very promising.

9. Conclusions

We have proposed a new, non-template abstract domain based on parallelotopes. The domain of parallelotopes can represent any linear constraint, thus retaining the expressivity of the polyhedra abstract domain, with reasonable computational costs. We have provided the full set of abstract operators and have defined a reduced product with intervals. Operations are mostly performed exploiting the inversion of the constraint matrix and, unlike template polyhedra, do not resort to linear programming.

We have experimentally validated the precision of parallelotopes versus the domains of polyhedra and octagons. In many cases we compute more invariants or invariants with tighter bounds. In particular, if we are interested in inferring interval constraints, the reduced product of parallelotopes and intervals greatly improves over polyhedra and octagons, still having a reasonable computational cost similar to octagons.

We believe that the ability to potentially represent any constraint is the key of the parallelotope domain, and makes it a valid alternative to template polyhedra.

We think that the domain of parallelotopes can be further improved by exploring alternative heuristics used in some abstract operators, which need an extensive experimental evaluation. Also the implementation of all the abstract operators could be enhanced, avoiding to recompute the inverse matrix (or solving a full linear system) when only few rows of the constraint matrix vary.

Appendix (Proofs)

Theorem 1. *For every parallelotope P there is a representation \mathcal{P} such that $\gamma(\mathcal{P}) = P$.*

Proof. If P is a parallelotope, there is $A \in \text{GL}(n)$ and $\mathbf{l}, \mathbf{u} \in \bar{\mathbb{R}}^n$ such that (1) holds. If P is empty, its representation is ϵ and $\gamma(\epsilon) = \emptyset$ by definition. If P is non-empty, then $\mathbf{l} \leq \mathbf{u}$ and for all $i \in \{1, \dots, n\}$, $l_i \neq +\infty$ and $u_i \neq -\infty$. Hence $\langle A, \mathbf{l}, \mathbf{u} \rangle$ is a representation of P . \square

Theorem 2 (Parallelotopes and Representations). *If $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ is a representation of a parallelotope in \mathbb{R}^n , the following holds:*

1. $\dim(\mathcal{P})$ is equal to $n - |E(\mathcal{P})|$; ³
2. $\mathbf{v} \neq 0$ is a ray in \mathcal{P} iff for each $i \in \{1, \dots, n\}$, $l_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \geq 0$ and $u_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \leq 0$;
3. \mathcal{P} is bounded iff all the bounds are finite;
4. $\mathbf{v} \neq 0$ is a line in \mathcal{P} iff it is orthogonal to all the constrained rows.

Proof. We start proving the first point. Assume without loss of generality that $E(\mathcal{P}) = \{1, \dots, m\}$. Since the rows of A are linearly independent, we know from linear algebra that the solution set of the first m constraints in \mathcal{P} is a flat of dimension $n - m$. Therefore $\dim(\mathcal{P}) \leq n - m$.

Now we prove the opposite inequality. Let us define vectors $\bar{\mathbf{l}}, \bar{\mathbf{u}} \in \mathbb{R}^n$ such that

$$\bar{l}_i = \begin{cases} l_i & \text{if } l_i \in \mathbb{R} \\ u_i - 1 & \text{if } l_i \notin \mathbb{R} \text{ and } u_i \in \mathbb{R} \\ -1 & \text{otherwise} \end{cases} \quad \bar{u}_i = \begin{cases} u_i & \text{if } u_i \in \mathbb{R} \\ l_i + 1 & \text{if } u_i \notin \mathbb{R} \text{ and } l_i \in \mathbb{R} \\ +1 & \text{otherwise} \end{cases}$$

³As an example of the problems which arise in considering representations with $\mathbf{l} \not\leq \mathbf{u}$, this property is false in that case, since \mathcal{P} is empty and $\dim(\mathcal{P}) = -1$ regardless of the number of equality rows.

Note that $\mathbf{l} \leq \bar{\mathbf{l}} \leq \bar{\mathbf{u}} \leq \mathbf{u}$ and $\bar{l}_i < \bar{u}_i$ for each $i > m$. For each $i > m$ we define $\mathbf{u}^i = \bar{\mathbf{l}}[i \mapsto \bar{u}_i]$, and note that $\bar{\mathbf{l}} \leq \bar{\mathbf{u}}^i \leq \mathbf{u}$. To ease notation, we also define $\mathbf{u}^{n+1} = \bar{\mathbf{l}}$. We prove that the set $\{\mathbf{u}^i\}_{i=m+1}^{n+1}$ is affinely independent. Assume otherwise, i.e. $\sum_{i=m+1}^{n+1} \lambda_i \mathbf{u}^i = \mathbf{0}$ with $\sum_{i=m+1}^{n+1} \lambda_i = 0$ and $\lambda_j \neq 0$. We have $\sum_{i=m+1}^{n+1} \lambda_i \mathbf{u}_j^i = (\sum_{k \neq j} \lambda_k) \bar{l}_j + \lambda_j \bar{u}_j = -\lambda_j \bar{l}_j + \lambda_j \bar{u}_j = 0$. However, since $\bar{u}_j \neq \bar{l}_j$, this is only possible if $\lambda_j = 0$, which contradicts our hypothesis. Then $\{\mathbf{u}^i\}_{i=m+1}^{n+1}$ is affinely independent.

Since A is invertible, it is possible to find points $\{\mathbf{v}^i\}_{i=m+1}^{n+1}$ such that $A\mathbf{v}^i = \mathbf{u}^i$. By definition of \mathbf{u}^i we have that \mathbf{v}^i is a point in \mathcal{P} . Since $\{\mathbf{u}^i\}_{i=m+1}^{n+1}$ is a set of affinely independent vectors, the same is true for $\{\mathbf{v}^i\}_{i=m+1}^{n+1}$. Therefore $\dim(\mathcal{P}) \geq n - m$ and this proves the result.

In order to prove the second point, assume $\mathbf{v} \neq \mathbf{0}$ is a ray in \mathcal{P} . If $l_i \in \mathbb{R}$ and $\mathbf{a}_{i*} \cdot \mathbf{v} = z < 0$, consider a generic $\mathbf{x} \in P$. Then it is possible to find $k \geq 0$ such that $\mathbf{a}_{i*} \cdot (\mathbf{x} + k\mathbf{v}) = \mathbf{a}_{i*} \cdot \mathbf{x} + kz < l_i$, which implies $\mathbf{x} + k\mathbf{v} \notin P$. This is absurd, hence $\mathbf{a}_{i*} \cdot \mathbf{v} \geq 0$. The case when $u_i \in \mathbb{R}$ is similar. For the converse implication, assume $\mathbf{v} \neq \mathbf{0}$ such that for each $i \in \{1, \dots, n\}$, $l_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \geq 0$ and $u_i \in \mathbb{R}$ implies $\mathbf{a}_{i*} \cdot \mathbf{v} \leq 0$. We need to prove that \mathbf{v} is a ray. Given $\mathbf{x} \in P$ and $\alpha \geq 0$, consider the point $\mathbf{w} = \mathbf{x} + \alpha\mathbf{v}$. Then, $A\mathbf{w} = A\mathbf{x} + \alpha A\mathbf{v}$, where $\mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}$ by definition. Given $i \in \{1, \dots, n\}$, if $l_i \in \mathbb{R}$ then $\mathbf{a}_{i*} \cdot \mathbf{v} = z \geq 0$, hence $\mathbf{a}_{i*} \cdot \mathbf{w} = \mathbf{a}_{i*} \cdot \mathbf{x} + \alpha z \geq \mathbf{a}_{i*} \cdot \mathbf{x} \geq l_i$. The same holds for the case $u_i \in \mathbb{R}$, hence $\mathbf{w} \in P$.

We come to the third point. It is clear that \mathcal{P} is bounded iff it has no rays. Assume there is a row \mathbf{a}_{i*} with $u_i = +\infty$, and we show that \mathcal{P} has a ray. We write \mathbf{a}_{i*} as $\mathbf{v} + \mathbf{w}$ where \mathbf{v} is in the space V of the rows of A different from i , and \mathbf{w} in V^\perp . Note that $\mathbf{w} \neq \mathbf{0}$, otherwise A is not invertible, and $\mathbf{a}_{i*} \cdot \mathbf{w} = \mathbf{w} \cdot \mathbf{w} > 0$. Given a point $\mathbf{x} \in P$, we have that $\mathbf{a}_{j*} \cdot (\mathbf{x} + \lambda\mathbf{w}) = \mathbf{a}_{j*} \cdot \mathbf{x}$ for $j \neq i$, while $\mathbf{a}_{i*} \cdot (\mathbf{x} + \lambda\mathbf{w}) \geq \mathbf{a}_{i*} \cdot \mathbf{x}$. Hence \mathbf{w} is a ray. The same happens if $l_i = -\infty$. On the converse, assume \mathcal{P} has a ray \mathbf{v} . If all bounds are finite, by the point 2 of this Lemma, we have that $A\mathbf{v} = \mathbf{0}$, i.e., $\mathbf{v} = \mathbf{0}$, which is not possible since a ray, by definition, cannot be zero.

Finally, if $\mathbf{v} \neq \mathbf{0}$ is orthogonal to all the constrained rows, then, according to the second point, both \mathbf{v} and $-\mathbf{v}$ are rays, hence \mathbf{v} is a line. On the contrary, assume $\mathbf{v} \neq \mathbf{0}$ is a line and \mathbf{a}_{i*} has a real bound, let us say $l_i \in \mathbb{R}$. Since \mathbf{v} is a ray, we have $\mathbf{a}_{i*} \cdot \mathbf{v} \leq 0$, and since $-\mathbf{v}$ is a ray, we have $\mathbf{a}_{i*} \cdot \mathbf{v} \geq 0$, hence \mathbf{v} is orthogonal to \mathbf{a}_{i*} . \square

Corollary 3. *The linearity space of a parallelotope \mathcal{P} is V^\perp where V is the linear space generated by all the constrained rows of \mathcal{P} . It is generated by the orthogonal projections of the unconstrained rows onto V^\perp .*

Proof. The fact that the linearity space of a parallelotope is V^\perp is an immediate consequence of the last point of the previous theorem.

For each unconstrained row \mathbf{a}_{i*} in \mathcal{P} , consider its projection \mathbf{w}_i over V^\perp . Since it is a vector in V^\perp , it is a line. Moreover, all the \mathbf{w}_i 's are linearly independent. Actually, we have that $\mathbf{a}_{i*} = \mathbf{w}_i + \mathbf{v}_i$ with $\mathbf{v}_i \in V$ and $\mathbf{w}_i \in V^\perp$. Since $\mathbf{v}_i \in V$, we have $\mathbf{v}_i = \sum_j \beta_{ij} \mathbf{a}_{j*}$ where j ranges over the constrained

rows. If $\sum_i \alpha_i \mathbf{w}_i = 0$, we have that $\sum_i \alpha_i \mathbf{a}_{i*} - \alpha_i \sum_j \beta_{ij} \mathbf{a}_{j*} = 0$, which would mean that A is singular. \square

Theorem 4. *Given the representation $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and a vector $\mathbf{c} \in \mathbb{R}^n$, we have that*

$$\inf_{\mathbf{x} \in \mathcal{P}} \mathbf{c} \cdot \mathbf{x} = \inf_{\mathbf{l} \leq \mathbf{y} \leq \mathbf{u}} \mathbf{c}^T A^{-1} \mathbf{y}.$$

The computational complexity of the minimization operation is $O(n^3)$.

Proof. By replacing $A\mathbf{x}$ with \mathbf{y} we may transform the linear programming problem $\inf_{\mathbf{x} \in \langle A, \mathbf{l}, \mathbf{u} \rangle} \mathbf{c} \cdot \mathbf{x} = \inf_{\mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}} \mathbf{c}^T \mathbf{x}$ into $\inf_{\mathbf{l} \leq \mathbf{y} \leq \mathbf{u}} \mathbf{c}^T A^{-1} \mathbf{y}$. The computational complexity is bounded by the cost for computing $\mathbf{c}^T A^{-1}$, which may be found out by solving for \mathbf{z} in the system of linear equations $\mathbf{z}A = \mathbf{c}^T$. \square

Theorem 5. *Given a parallelotope \mathcal{P} and $A' \in \text{GL}(n)$, $\text{rot}_{A'}(\mathcal{P})$ is the least parallelotope definable over A' which contains \mathcal{P} , i.e., $\text{rot}_{A'}(\mathcal{P}) = \alpha_{A'}(\gamma(\mathcal{P}))$. The computational complexity of $\text{rot}_{A'}$ is $O(n^3)$.*

Proof. If $\mathcal{P} = \epsilon$ the result is obvious. Otherwise, let $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$. If we want to compute $\alpha_{A'}(\gamma(\mathcal{P}))$, it is enough to find out, for each row \mathbf{a}'_{i*} in A' , the minimum and maximum value of $\mathbf{a}'_{i*} \mathbf{x}$ for $\mathbf{x} \in \gamma(\mathcal{P})$. These will be the values for l'_i and u'_i respectively. Note that $\inf_{\mathbf{x} \in \mathcal{P}} \mathbf{a}'_{i*} \mathbf{x} = \inf_{\mathbf{l} \leq \mathbf{y} \leq \mathbf{u}} \mathbf{a}'_{i*} A^{-1} \mathbf{y}$ by Theorem 4. Moreover $\mathbf{a}'_{i*} A^{-1} = \mathbf{b}_{i*}$, which proves that $\alpha_{A'}(\gamma(\mathcal{P})) = \text{rot}_{A'}(\mathcal{P})$. The computational cost is bounded by the cost for computing B , which may be found out as the solution on X of the equation $XA = A'$. \square

Theorem 6. *Given two representations \mathcal{P} and \mathcal{P}' , we have that $\mathcal{P} \leq \mathcal{P}'$ iff $\gamma(\mathcal{P}) \subseteq \gamma(\mathcal{P}')$. Moreover, \leq is a pre-order. The computational complexity of deciding \leq is $O(n^3)$.*

Proof. Given the Galois connection between $\alpha_{A'}$ and γ , we have $\gamma(\mathcal{P}) \subseteq \gamma(\mathcal{P}')$ iff $\alpha_{A'}(\gamma(\mathcal{P})) = \text{rot}_{A'}(\mathcal{P}) \leq \mathcal{P}'$ which is equivalent to $\mathcal{P} \leq \mathcal{P}'$. The fact that \leq is a pre-order immediately follows from its characterization in terms of γ and \subseteq . The computational complexity of deciding \leq is clearly bounded by the complexity of rot which is $O(n^3)$. \square

Theorem 7. *The operation $\text{assign}^\alpha(i, \mathbf{c}, b)$ is correct and γ -complete. The computational complexity is $O(mn)$ where m is the number of non-zero components in \mathbf{c} .*

Proof. It trivially holds that $\text{assign}(i, \mathbf{c}, b)(\gamma(\epsilon)) = \text{assign}(i, \mathbf{c}, b)(\emptyset) = \emptyset = \gamma(\epsilon)$. Given $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle \in \text{Par}_n$, we have

$$\begin{aligned} & \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P})) \\ &= \{ \mathbf{x} [i \mapsto \mathbf{c} \cdot \mathbf{x} + b] \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u} \} \\ &= \{ K\mathbf{x} + be^i \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u} \} \end{aligned}$$

where $K = I_n + \mathbf{e}^i(\mathbf{c} - \mathbf{e}^i)^T$. Note that K is invertible and $K^{-1} = I_n - \frac{1}{c_i}\mathbf{e}^i(\mathbf{c} - \mathbf{e}^i)^T$. Therefore, by the change of variable $\mathbf{y} = K\mathbf{x} + b\mathbf{e}^i$, we have

$$\begin{aligned} \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P})) &= \{\mathbf{y} \mid \mathbf{l} \leq AK^{-1}(\mathbf{y} - b\mathbf{e}^i) \leq \mathbf{u}\} \\ &= \{\mathbf{y} \mid \mathbf{l} + bAK^{-1}\mathbf{e}^i \leq AK^{-1}\mathbf{y} \leq \mathbf{u} + bAK^{-1}\mathbf{e}^i\}. \end{aligned}$$

Since $AK^{-1} = A - \frac{1}{c_i}\mathbf{a}_{*i}(\mathbf{c} - \mathbf{e}^i)^T$ we have $AK^{-1}\mathbf{e}^i = \mathbf{a}_{*i} - \frac{1}{c_i}\mathbf{a}_{*i}(c_i - 1) = \frac{1}{c_i}\mathbf{a}_{*i}$. Hence

$$\begin{aligned} \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P})) &= \left\{ \mathbf{y} \mid \mathbf{l} + \frac{b}{c_i}\mathbf{a}_{*i} \leq \left(A - \frac{1}{c_i}\mathbf{a}_{*i}(\mathbf{c} - \mathbf{e}^i)^T \right) \mathbf{y} \leq \mathbf{u} + \frac{b}{c_i}\mathbf{a}_{*i} \right\} \\ &= \gamma(\text{assign}^\alpha(i, \mathbf{c}, b))(\mathcal{P}). \end{aligned}$$

This proves that $\text{assign}^\alpha(i, \mathbf{c}, b)$ is γ -complete. The computational complexity of this operation is dominated by the cost of computing all non-zero products in $\mathbf{a}_{*i}(\mathbf{c} - \mathbf{e}^i)^T$, which is $O(mn)$. This is true even if we use a dense representation for vectors and matrices, since $\frac{1}{c_i}\mathbf{a}_{*i}(\mathbf{c} - \mathbf{e}^i)^T$ may be subtracted from A without materializing it. \square

Proposition 8. *Given two parallelotopes $\langle A, \mathbf{l}, \mathbf{u} \rangle \leq \langle A', \mathbf{l}', \mathbf{u}' \rangle$, we have that*

$$\text{forget}(i)(\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)) \subseteq \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle) \text{ iff } \forall j (a'_{ji} \neq 0 \rightarrow (l'_j = -\infty \wedge u'_j = +\infty)).$$

Proof. If $a'_{ji} \neq 0 \rightarrow (l'_j = -\infty \wedge u'_j = +\infty)$ for all j , it follows immediately that for each $v \in \mathbb{R}$ and $\mathbf{x} \in \langle A', \mathbf{l}', \mathbf{u}' \rangle$, $\mathbf{x}[i \mapsto v] \in \langle A', \mathbf{l}', \mathbf{u}' \rangle$, since the variable x_i is not used in $\mathbf{l}' \leq A'\mathbf{x} \leq \mathbf{u}'$. Since $\langle A, \mathbf{l}, \mathbf{u} \rangle \leq \langle A', \mathbf{l}', \mathbf{u}' \rangle$, it follows that $\text{forget}(i)(\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)) \subseteq \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle)$.

Now, let $\text{forget}(i)(\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)) \subseteq \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle)$. By contradiction, assume there exists an index j such that $a'_{ji} \neq 0$ and $l'_j \in \mathbb{R}$. Let $\mathbf{x} \in \langle A, \mathbf{l}, \mathbf{u} \rangle$ and let $v \in \mathbb{R}$ such that $va'_{ji} < l'_j - \mathbf{a}'_{j*}\mathbf{x}$. It follows that $\mathbf{a}'_{j*}(\mathbf{x} + v\mathbf{e}^i) = \mathbf{a}'_{j*}\mathbf{x} + va'_{ji} < l'_j$, and thus $\mathbf{x} + v\mathbf{e}^i \notin \langle A', \mathbf{l}', \mathbf{u}' \rangle$. However, $\mathbf{x} + v\mathbf{e}^i \in \text{forget}(i)(\gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle))$, which is a contradiction. The case that there exists an index j such that $a'_{ji} \neq 0$ and $u'_j \in \mathbb{R}$ is analogous. \square

Theorem 9. *The operator $\text{forget}^\alpha(i)$ described in Algorithm 1 is correct and relatively optimal. The computational complexity is $O(n^2)$.*

Proof. The algorithm begins at line 1 by computing a set J of indexes. Each element $j \in J$ is the index of a constrained row \mathbf{a}_{j*} of the constraint matrix A , where the i -th component a_{ji} is non-zero. If J is empty, by Proposition 8, $\gamma(\mathcal{P}) = \text{forget}(i)(\gamma(\mathcal{P}))$ and therefore we return \mathcal{P} . The result is γ -complete, and thus relatively optimal.

Otherwise, lines from 6 to 14 of the algorithm choose an index $r \in R$, according to a given heuristic. The loop at lines 15–21 finds a new constraint matrix

A' and builds at the same time the parallelotope $\langle A', \mathbf{l}', \mathbf{u}' \rangle = \text{rot}_{A'}(P)$. In particular, the r -th constraint of the parallelotope ($l_r \leq \mathbf{a}_{*r} \leq u_r$) is combined with the j -th constraint ($l_j \leq \mathbf{a}_{j*} \leq u_j$) to obtain a new constraint $l'_j \leq \mathbf{a}'_{j*} \leq u'_j$ which is implied by the first two. The set of rows $\{\mathbf{a}'_{j*} \mid j \in J\}$ generates the same vector space as $\{\mathbf{a}_{j*} \mid j \in J\}$, therefore A' is invertible. Moreover, after the loop $a'_{ji} = 0$ for each $j \in J \setminus \{r\}$.

From Proposition 8, we know that the smallest parallelotope whose constraint matrix is A' and which contains $\text{forget}(i)(\gamma(\mathcal{P}))$ is given by turning the constrained rows which contain a non-zero entry in the i -th position into unconstrained rows. There is only one such a row in A' , which is row r . The bounds for this row are changed in the final return statement.

The complexity is bounded by the cost of determining all the linear combinations in the loop at lines 15–21. \square

Theorem 10. *If \mathbf{e}^i is a line in \mathcal{P} or $J_0 \neq \emptyset$, then it holds that $\gamma(\text{forget}^\alpha(i)(\mathcal{P})) = \text{forget}(i)(\gamma(\mathcal{P}))$.*

Proof. If \mathbf{e}^i is a line in $\gamma(\mathcal{P})$, then $\text{forget}(i)(\gamma(\mathcal{P})) = \gamma(\mathcal{P})$. Moreover, since \mathbf{e}^i is orthogonal to the constrained rows in \mathcal{P} (Theorem 2), it means that J in line 1 of the algorithm is empty, hence $\text{forget}^\alpha(i)(\mathcal{P}) = \mathcal{P}$.

Otherwise, assume the set J_0 in line 6 of the algorithm is not empty and the row \mathbf{a}_{r*} chosen in line 9 is an equality constraint, with bounds $l_r = u_r = \alpha$.

We now prove that the steps 15–21 in the algorithm do not change the parallelotope. Consider the constraint $c_1 \equiv l'_j \leq \mathbf{a}'_{j*} \cdot \mathbf{x} \leq u'_j$ computed during the for loop. We prove that this constraint together with the r -th constraint of \mathcal{P} , i.e. $c \equiv \mathbf{a}_{r*} \cdot \mathbf{x} = \alpha$, is equivalent to the combination of the r -th and j -th constraint of \mathcal{P} , i.e. $c_2 \equiv l_j \leq \mathbf{a}_{j*} \cdot \mathbf{x} \leq u_j$. Given what we have said in the proof of correctness, we only need to check that the j -th constraint of \mathcal{P} is a logical consequence of the new constraint and the r -th constraint in \mathcal{P} . Assume without loss of generality that $a_{ri} > 0$ and $a_{ji} > 0$. Then

$$c_1 \equiv a_{ji}\alpha - a_{ri}u_j \leq (a_{ji}\mathbf{a}_{r*} - a_{ri}\mathbf{a}_{j*}) \cdot \mathbf{x} \leq a_{ji}\alpha - a_{ri}l_j .$$

It is easy to check that $(a_{ji}c - c_1)/a_{ri} = c_2$. Since this holds for all $j \in J \setminus \{r\}$, we have that, at the end of the loop, $\langle A', \mathbf{l}', \mathbf{u}' \rangle$ and \mathcal{P} are representations for the same parallelotope.

The last step is removing the bounds in the row r , which is the only constrained row with a non-null i -th coefficient.

Now, given a generic point $\mathbf{x} \in \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$, we need to prove that $\mathbf{x} \in \text{forget}(i)(\gamma(\mathcal{P}))$. It is obvious that $\mathbf{y} = \mathbf{x} + \lambda \mathbf{e}^i \in \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$ for each $\lambda \in \mathbb{R}$. Note that, for each constrained row \mathbf{a}'_{j*} with $j \neq r$, we have that $\mathbf{a}'_{j*} \cdot \mathbf{y} = \mathbf{a}'_{j*} \cdot \mathbf{x}$, while $\mathbf{a}_{r*} \cdot \mathbf{y} = \mathbf{a}_{r*} \cdot \mathbf{x} + \lambda a_{ri}$. By choosing $\lambda = \alpha - (\mathbf{a}_{r*} \cdot \mathbf{x})/a_{ri}$, we have $\mathbf{y} \in \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle) = \gamma(\mathcal{P})$. Since $\mathbf{x} = \mathbf{y} - \lambda \mathbf{e}^i$ and $\mathbf{y} \in \gamma(\mathcal{P})$, then $\mathbf{x} \in \text{forget}(i)(\gamma(\mathcal{P}))$. \square

Lemma 24. *Consider the polyhedron $P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}\}$ such that $A \in \mathbb{R}(m, n)$ has maximal row rank, $\mathbf{a}_{*r} = \mathbf{1}$ (the column vector with 1 in all*

its components) and $\mathbf{l}, \mathbf{u} \in \bar{\mathbb{R}}^n$ with $\mathbf{l} < \mathbf{u}$. The set of constraints generated by Fourier-Motzkin elimination is $\{(\mathbf{a}_{i*} - \mathbf{a}_{j*}) \cdot \mathbf{x} \leq u_i - l_j \mid i, j \in \{1, \dots, m\}\}$. None of these constraints is redundant.

Proof. First of all, considering $\mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}$ as a short form of $A\mathbf{x} \leq \mathbf{u}$ and $-A\mathbf{x} \leq -\mathbf{l}$ and applying Fourier-Motzkin elimination, we obtain the set

$$\{(\mathbf{a}_{i*} - \mathbf{a}_{j*}) \cdot \mathbf{x} \leq u_i - l_j \mid i, j \in \{1, \dots, n\}\} .$$

Without loss of generality, consider the inequality $(\mathbf{a}_{1*} - \mathbf{a}_{2*}) \cdot \mathbf{x} \leq u_1 - l_2$ and assume $u_1 - l_2$ is finite. We want to prove that this constraint is not redundant.

By Farkas Lemma, this inequality is redundant iff there are $\lambda_{ij} \geq 0$ such that $\lambda_{12} = 0$ and

$$\mathbf{a}_{1*} - \mathbf{a}_{2*} = \sum_{i,j} \lambda_{ij} (\mathbf{a}_{i*} - \mathbf{a}_{j*}) \quad (3)$$

and

$$\sum_{i,j} \lambda_{ij} (u_i - l_j) \leq u_1 - l_2 .$$

We may rewrite $\sum_{i,j} \lambda_{ij} (\mathbf{a}_{i*} - \mathbf{a}_{j*})$ as $\sum_{i,j} \lambda_{ij} \mathbf{a}_{i*} - \sum_{i,j} \lambda_{ij} \mathbf{a}_{j*} = \sum_{i,j} \lambda_{ij} \mathbf{a}_{i*} - \sum_{i,j} \lambda_{ji} \mathbf{a}_{i*}$. Since the \mathbf{a}_{i*} are linearly independent, we have that (3) holds iff

$$\sum_j \lambda_{ij} - \sum_j \lambda_{ji} = \begin{cases} 1 & \text{if } i = 1 \\ -1 & \text{if } i = 2 \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Let $\alpha_i = \sum_j \lambda_{ji}$ and $\beta_i = \sum_j \lambda_{ij}$. Equation 4 means that $\beta_1 = 1 + \alpha_1$, $\beta_2 = -1 + \alpha_2$ and $\alpha_i = \beta_i$ for each $i > 2$. Using these relationships, we may rewrite $\sum \lambda_{ij} (u_i - l_j) \leq u_1 - l_2$ as follows:

$$\begin{aligned} \sum_{i,j} \lambda_{ij} (u_i - l_j) &= \sum_i \left(\sum_j \lambda_{ij} \right) u_i - \sum_i \left(\sum_j \lambda_{ji} \right) l_i \\ &= \sum_i (\beta_i u_i - \alpha_i l_i) \\ &= (\alpha_1 + 1)u_1 - \alpha_1 l_1 + \sum_{i>2} (\alpha_i u_i - \alpha_i l_i) + \beta_2 u_2 - (1 + \beta_2)l_2 \\ &= u_1 - l_2 + \sum_{i \neq 2} \alpha_i (u_i - l_i) + \beta_2 (u_2 - l_2) . \end{aligned}$$

Since $u_i > l_i$, $\lambda_{ij} \geq 0$ and $\sum_{i,j} \lambda_{ij} (u_i - l_j) \leq u_1 - l_2$, then $\alpha_i = 0$ for each $i \neq 2$ and $\beta_2 = 0$. This entails that the only possibly positive λ 's are the λ_{j2} with $j > 2$. Let $h > 2$ such that $\lambda_{h2} > 0$, then $\beta_h = \alpha_h > 0$, which is a contradiction. Therefore all the λ_{ij} are zero, i.e. $\mathbf{a}_{1*} = \mathbf{a}_{2*}$ which is not possible since A is of maximal row rank. Therefore $\mathbf{a}_{1*} - \mathbf{a}_{2*} \leq u_1 - l_2$ is not redundant. \square

Corollary 25. Consider the polyhedron $P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}\}$ such that $A \in \mathbb{R}(m, n)$ has maximal row rank, \mathbf{a}_{*r} has no zero components and $\mathbf{l}, \mathbf{u} \in \mathbb{R}^n$ with $\mathbf{l} < \mathbf{u}$. Then, none of the constraints obtained by Fourier-Motzkin elimination of P is redundant.

Proof. Consider the matrix A' such that $a'_{ij} = a_{ij}/a_{ir}$, and vectors \mathbf{l}' and \mathbf{u}' such that

$$l'_i = \begin{cases} l_i/a_{ir} & \text{if } a_{ir} > 0; \\ u_i/a_{ir} & \text{otherwise;} \end{cases} \quad u'_i = \begin{cases} u_i/a_{ir} & \text{if } a_{ir} > 0; \\ l_i/a_{ir} & \text{otherwise.} \end{cases}$$

Then $P = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l}' \leq A'\mathbf{x} \leq \mathbf{u}'\}$. We may apply Lemma 24 and note that each constraint $(\mathbf{a}'_{i*} - \mathbf{a}'_{j*}) \cdot \mathbf{x} \leq u'_i - l'_j$ is equivalent to one of the constraints obtained by Fourier-Motzkin elimination of P . For example, if $a_{ir} > 0$ and $a_{jr} < 0$, constraints $\mathbf{a}_{i*} \cdot \mathbf{x} \leq u_i$ and $\mathbf{a}_{j*} \cdot \mathbf{x} \leq u_j$ are combined into $(-a_{jr}\mathbf{a}_{i*} + a_{ir}\mathbf{a}_{j*}) \cdot \mathbf{x} \leq -a_{jr}u_i + a_{ir}u_j$, which is equivalent to $(\mathbf{a}_{i*}/a_{ij} - \mathbf{a}_{j*}/a_{jr}) \cdot \mathbf{x} \leq u_i/a_{ir} - u_j/a_{jr}$, i.e., $(\mathbf{a}'_{i*} - \mathbf{a}'_{j*}) \cdot \mathbf{x} \leq u'_i - l'_j$. The other cases are similar. \square

Proposition 11. Given a parallelotope $\langle A, \mathbf{l}, \mathbf{u} \rangle$ and a column index i , assume that $l_j < u_j$ for each constrained row j such that $a_{ji} \neq 0$. Consider the polyhedron obtained by Fourier-Motzkin elimination of the i -th variable. None of its constraints with finite bounds is redundant.

Proof. Let J be defined as in Algorithm 1, i.e. $J = \{j \mid a_{ji} \neq 0, l_j \neq -\infty \vee u_j \neq +\infty\}$. We have that the result of Fourier-Motzkin elimination is the polyhedron

$$\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l}_{-J} \leq A_{-J*}\mathbf{x} \leq \mathbf{u}_{-J}\} \cap \{\mathbf{x} \in \mathbb{R}^n \mid D\mathbf{x} \leq \mathbf{d}\}$$

where $D\mathbf{x} \leq \mathbf{d}$ is the set of constraints determined by Fourier-Motzkin elimination from the polyhedron $\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l}_J \leq A_{J*}\mathbf{x} \leq \mathbf{u}_J\}$. We need to prove that all the constraints with finite bounds are not redundant. All finitely bounded constraints of the kind $l_j \leq \mathbf{a}_{j*} \cdot \mathbf{x}$ or $\mathbf{a}_{j*} \cdot \mathbf{x} \leq u_j$ for $j \notin J$ are not redundant simply because \mathbf{a}_{j*} is linearly independent from all the rows in D and in $A_{-(J \cup \{j\})*}$. Therefore, it is not possible for such a constraint to be a positive linear combination of other constraints.

Moreover, since the rows of D are linearly independent from the rows of A_{-J*} , if a constraint generated by D is redundant, then it should be a positive linear combination of other constraints generated by D . This is not possible by Corollary 25. \square

Lemma 26. Let $P \subseteq P' \subseteq P''$ be convex sets, and H an hyper-plane. If $P \cap H$ and $P'' \cap H$ are faces of dimension k , then $P' \cap H$ is a face of dimension k .

Proof. By definition, if C is a convex set and H is a hyper-plane, $C \cap H$ is a face when it is non-empty and C is entirely on one side of H . Since $P \cap H$ is non-empty and $P' \supseteq P$, then $P \cap H$ is not-empty. Moreover, since P'' is entirely on one side of H and $P' \subseteq H$, then P' is entirely on one side of H . Then $P' \cap H$ is a face. Moreover, since $P \cap H \subseteq P' \cap H \subseteq P'' \cap H$ and both $P \cap H$ and $P'' \cap H$ have dimension k , also $P' \cap H$ has dimension k . \square

Theorem 12. *The operator $\text{forget}^\alpha(i)$ is minimal.*

Proof. Let $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ be a parallelotope representation and $P = \gamma(\mathcal{P}) \subseteq \mathbb{R}^n$. If we are in the case where $\text{forget}^\alpha(i)$ is γ -complete, then it is also minimal. Otherwise, assume \mathcal{P} has a equality rows and u unconstrained rows. Therefore, $\dim(P) = n - a$ and $\dim(\text{lin}(P)) = u$. If j is an equality row, then $\mathbf{a}_{ji} = 0$, otherwise we fall in the case where $\text{forget}^\alpha(i)$ is γ -complete.

Consider now the polyhedron $\text{forget}(i)(P)$. By the properties of P discussed above, we have that $\dim(\text{forget}(i)(P)) = \dim(P)$, $\text{aff. hull}(\text{forget}(i)(P)) = \text{aff. hull}(P)$ and $\dim(\text{lin}(\text{forget}(i)(P))) = u + 1$. Assume $c \equiv \mathbf{v} \cdot \mathbf{x} \leq u$ is a constraint generated by Fourier-Motzkin elimination and $H = \{\mathbf{x} \mid \mathbf{v} \cdot \mathbf{x} \leq u\}$. Since c is not redundant, by a property of polyhedron we have that $\text{forget}(i)(P) \cap H$ is a facet of $\text{forget}(i)(P)$, i.e., a face of dimension $\dim(P) - 1$.

If $\mathcal{P}' = \text{forget}^\alpha(i)(\mathcal{P}) = \langle A', \mathbf{l}', \mathbf{u}' \rangle$ and $P' = \gamma(\mathcal{P}')$ we have that \mathcal{P}' has a equality rows and $u + 1$ unconstrained rows. The constraints corresponding to the equality rows are unchanged from those in \mathcal{P} , hence $\dim(P') = \dim(P)$, $\text{aff. hull}(P') = \text{aff. hull}(P)$ and $\dim(\text{lin}(P')) = u + 1$. The other constraints are a selection of constraints obtained by Fourier-Motzkin elimination. Since these constraints were non-redundant in $\text{forget}(i)(P)$ by Proposition 11, they are also non-redundant in \mathcal{P}' . Hence if $\mathbf{a}'_{j*} \cdot \mathbf{x} \leq u_j$ is a bounded constraint in \mathcal{P}' and $H = \{\mathbf{x} \mid \mathbf{a}'_{j*} \cdot \mathbf{x} \leq u_j\}$, then $P' \cap H$ is a facet of P' . The same holds for constraints of the kind $l_j \leq \mathbf{a}'_{j*} \cdot \mathbf{x}$.

Now assume \mathcal{P}'' is a correct approximation of $\text{forget}(i)(P)$ such that $\mathcal{P}'' \leq \mathcal{P}'$ and $P'' = \gamma(\mathcal{P}'')$. Therefore, $\dim(P'') = \dim(P') = \dim(\text{forget}(i)(P))$, $\text{aff. hull}(P'') = \text{aff. hull}(P') = \text{aff. hull}(\text{forget}(i)(P))$ and $\dim(\text{lin}(P'')) = u + 1$. We want to prove that in \mathcal{P}'' should explicitly appear all the constraints which appear in \mathcal{P}' . Without loss of generality, consider the constraint $\mathbf{a}'_{j*} \cdot \mathbf{x} \leq u_j$ in \mathcal{P}' and $H = \{\mathbf{x} \mid \mathbf{a}'_{j*} \cdot \mathbf{x} \leq u_j\}$. We know that both $P' \cap H$ and $\text{forget}(i)(P) \cap H$ are faces of dimension $\dim(P') - 1$. By Lemma 26, $P'' \cap H$ is a face of P'' of the same dimension, and since $\dim(P'') = \dim(P')$ then $P'' \cap H$ is a facet.

For a polyhedron, and therefore also for a parallelotope, facets exactly corresponds to the inequality constraints in its definition. Therefore, $\mathbf{a}'_{j*} \cdot \mathbf{x} \leq u_j$ (or an equivalent multiple) should appear in \mathcal{P}'' . This holds for all constrained inequality rows of A' and for constraints determined by either lower or upper bounds. Since a rows in \mathcal{P}'' are needed for the equality constraints which force $\text{aff. hull}(P'') = \text{aff. hull}(P')$ and $u + 1$ row should remain unconstrained, so that $\dim(\text{lin}(P'')) = \dim(\text{lin}(P'))$, we have that all remaining constraints in \mathcal{P}'' are multiples of constraints in \mathcal{P}' , and therefore $P'' = P'$. \square

Theorem 13. *The operator $\text{assign}^\alpha(i, c, b)$ described in Algorithm 2 is correct. The computational complexity is $O(n^2)$.*

Proof. We begin by proving that lines 2–5 of the algorithm change the representation $\mathcal{P}' = \langle A', \mathbf{l}', \mathbf{u}' \rangle$ into another representation for the same parallelotope. First of all, we need to check that the new vectors $\mathbf{v}_s = \mathbf{a}'_{s*} - (a'_{si}/a'_{ji})\mathbf{a}'_{j*}$ correspond to unbounded directions in $\gamma(\mathcal{P}')$, i.e., that the linear form $\mathbf{x} \mapsto \mathbf{v}_s \cdot \mathbf{x}$ is unbounded for $\mathbf{x} \in \gamma(\mathcal{P}')$. Note that this does not mean that \mathbf{v}_s is a line.

Since A' is invertible, \mathbf{v}_s cannot be obtained as a linear combination of the constrained rows of A' , hence $\mathbf{v}_s = \lambda_1 \mathbf{w}_1 + \lambda_2 \mathbf{w}_2$ where \mathbf{w}_2 is a line and \mathbf{w}_1 is orthogonal to the linearity space of \mathcal{P}' . Hence, $\inf\{\mathbf{v}_s \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}'\} = -\infty$ and $\sup\{\mathbf{v}_s \cdot \mathbf{x} \mid \mathbf{x} \in \mathcal{P}'\} = +\infty$. Moreover, the operations at line 4 are *basic row operations*, which are known to preserve the rank of a matrix.

At line 6, \mathbf{a}'_{j^*} is the only row in A' with a non-zero i -th component. If we replace it with another row which has a non-empty i -th component, such as $\mathbf{e}^i - \mathbf{c}$, the matrix remains invertible. Moreover, $\mathbf{e}^i - \mathbf{c}$ is an unbounded direction in $\text{forget}^\alpha(i)(\langle A, \mathbf{l}, \mathbf{u} \rangle)$. Therefore, at the end of line 6, we still have $\gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle) = \gamma(\text{forget}^\alpha(i)(\langle A, \mathbf{l}, \mathbf{u} \rangle))$.

Lines 7 and 8 in the algorithm amount to building the parallelotope representation $\alpha_{A'}(\text{assign}(i, \mathbf{c}, b)(\gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle))) = \alpha_{A'}(\text{assign}(i, \mathbf{c}, b)(\gamma(\text{forget}^\alpha(i)(\langle A, \mathbf{l}, \mathbf{u} \rangle)))$. Since $\text{forget}^\alpha(_)$ is correct, this proves the correctness of the non-invertible assignment.

The computational complexity of the algorithm is dominated by the cost of the loop at lines 3–5, which is $O(n^2)$. \square

Theorem 14. *Consider $\mathbf{c} \in \mathbb{R}^n$, $b \in \mathbb{R}$ and $c_i = 0$. If \mathcal{P} is a parallelotope in \mathbb{R}^n and \mathbf{e}^i is a line in \mathcal{P} or $J_0 \neq \emptyset$ in Algorithm 1 then $\gamma(\text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P})) = \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P}))$.*

Proof. Assume $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\langle A'', \mathbf{l}'', \mathbf{u}'' \rangle = \text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P})$. Consider a generic $\mathbf{x} \in \gamma(\text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P}))$. If j is the row determined at line 2 of Algorithm 2, then $\mathbf{a}''_{j^*} \mathbf{x} = b$, which means that \mathbf{x} satisfies the equation $x_i = \mathbf{c} \cdot \mathbf{x} + b$. At the end of line 5 of Algorithm 2, we have values of A' , \mathbf{l}' and \mathbf{u}' which only differ from A'' , \mathbf{l}'' and \mathbf{u}'' for the j -th constraint, which directly comes from the result of $\text{forget}^\alpha(i)$. From the definition of $\text{forget}^\alpha(i)$, it turns out that $l'_j = -\infty$ and $u'_j = +\infty$, hence $\mathbf{x} \in \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle)$. Since lines 2–5 in the algorithm do not change the parallelotope, we have $\mathbf{x} \in \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$. By Theorem 10, $\mathbf{x} \in \text{forget}(i)(\gamma(\mathcal{P}))$, and since \mathbf{x} satisfies the equation $x_i = \mathbf{c} \cdot \mathbf{x} + b$, then also $\mathbf{x} \in \text{assign}(i, \mathbf{c}, b)(\text{forget}(i)(\gamma(\mathcal{P}))) = \text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P}))$. \square

Theorem 15. *The operator $\text{assign}^\alpha(i, \mathbf{c}, b)$ with $c_i = 0$ is minimal.*

Proof. Let $\mathcal{P}'' = \langle A'', \mathbf{l}'', \mathbf{u}'' \rangle$ be a correct approximation of $\text{assign}(i, \mathbf{c}, b)(\mathcal{P})$, $\mathcal{P}' = \langle A', \mathbf{l}', \mathbf{u}' \rangle = \text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P})$ and assume $\mathcal{P}'' \leq \mathcal{P}'$.

Note that $\gamma(\mathcal{P}')$ lies on the hyper-plane $(\mathbf{e}^i - \mathbf{c}) \cdot \mathbf{x} = b$. Since $\mathcal{P}'' \leq \mathcal{P}'$, the same holds for $\gamma(\mathcal{P}'')$. It is easy to check that this implies

$$\gamma(\mathcal{P}'') = \text{assign}(i, \mathbf{c}, b)(\text{forget}(i)(\gamma(\mathcal{P}''))) . \quad (5)$$

Since \mathcal{P}'' lies in the hyper-plane $(\mathbf{e}^i - \mathbf{c}) \cdot \mathbf{x} = b$, one of its equality rows must have a non-zero i -th component. The same also holds for \mathcal{P}' . By Theorem 10 we get:

$$\gamma(\text{forget}^\alpha(i)(\mathcal{P}')) = \text{forget}(i)(\gamma(\mathcal{P}')) , \quad (6)$$

$$\gamma(\text{forget}^\alpha(i)(\mathcal{P}'')) = \text{forget}(i)(\gamma(\mathcal{P}'')) . \quad (7)$$

Moreover, it is clear that $\text{forget}(i) \circ \text{assign}(i, \mathbf{c}, b) = \text{forget}(i)$ and $\gamma \circ \text{forget}^\alpha(i) \circ \text{assign}^\alpha(i, \mathbf{c}, b) = \gamma \circ \text{forget}^\alpha(i)$. Therefore, by (7),

$$\begin{aligned} \gamma(\text{forget}^\alpha(i)(\mathcal{P}'')) &= \text{forget}(i)(\gamma(\mathcal{P}'')) \supseteq \\ &\supseteq \text{forget}(i)(\text{assign}(i, \mathbf{c}, b)(\gamma(\mathcal{P}))) = \text{forget}(i)(\gamma(\mathcal{P})) \end{aligned}$$

which means that $\text{forget}^\alpha(i)(\mathcal{P}'')$ is a correct approximation of $\text{forget}(i)(\mathcal{P})$. Moreover, by (6) and (7),

$$\begin{aligned} \gamma(\text{forget}^\alpha(i)(\mathcal{P}'')) &= \text{forget}(i)(\gamma(\mathcal{P}'')) \subseteq \text{forget}(i)(\gamma(\mathcal{P}')) = \\ &= \gamma(\text{forget}^\alpha(i)(\mathcal{P}')) = \gamma(\text{forget}^\alpha(i)(\text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P}))) = \gamma(\text{forget}^\alpha(i)(\mathcal{P})) \end{aligned} ,$$

hence $\gamma(\text{forget}^\alpha(i)(\mathcal{P}'')) \subseteq \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$. However, by Theorem 12, we have that $\text{forget}^\alpha(i)(\mathcal{P})$ is a minimal correct approximation of $\text{forget}(i)(\mathcal{P})$. Hence, we have $\gamma(\text{forget}^\alpha(i)(\mathcal{P}'')) = \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$, and together with (7) this gives $\text{forget}(i)(\gamma(\mathcal{P}'')) = \gamma(\text{forget}^\alpha(i)(\mathcal{P}))$. Therefore, by (5),

$$\gamma(\mathcal{P}'') = \text{assign}(i, \mathbf{c}, b)(\text{forget}(i)(\gamma(\mathcal{P}''))) = \text{assign}(i, \mathbf{c}, b)(\gamma(\text{forget}^\alpha(i)(\mathcal{P})))$$

and since e^i is a line in $\text{forget}^\alpha(i)(\mathcal{P})$, by Theorem 14 we have that $\gamma(\mathcal{P}'') = \gamma(\text{assign}^\alpha(i, \mathbf{c}, b)(\text{forget}^\alpha(i)(\mathcal{P}))) = \gamma(\text{assign}^\alpha(i, \mathbf{c}, b)(\mathcal{P}))$. \square

Theorem 16. *The operator $\text{refine}^\alpha(\mathbf{c}, b)$ described in Algorithm 3 is correct and relatively optimal. If $\mathbf{c} \notin \text{lin}(\gamma(\mathcal{P}))^\perp$, then $\text{refine}^\alpha(\mathbf{c}, b)(\mathcal{P})$ is γ -complete. The computational complexity is $O(n^3)$.*

Proof. Assume $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$, and let $\mathbf{y} \in \mathbb{R}^n$ be the solution of the equation $A^T \mathbf{y} = \mathbf{c}$. First of all, we prove γ -optimality when $\mathbf{c} \notin \text{lin}(\gamma(\mathcal{P}))^\perp$. In this case there is an index j that satisfies the condition in line 2 of the algorithm: since $\mathbf{c} \notin \text{lin}(\gamma(\mathcal{P}))^\perp$, then \mathbf{c} is not a linear combination of the constrained rows in \mathcal{P} . Steps 3 and 4 introduce the new constraint and, due to the choice of j , the modified matrix A is invertible. Let $A', \mathbf{l}', \mathbf{u}'$ be the values of variables A, \mathbf{l} and \mathbf{u} at line 5, i.e., $\langle A', \mathbf{l}', \mathbf{u}' \rangle = \text{refine}^\alpha(\mathbf{c}, b)(\mathcal{P})$. Then $\mathbf{x} \in \gamma(\langle A', \mathbf{l}', \mathbf{u}' \rangle)$ iff $\mathbf{x} \in \gamma(\langle A, \mathbf{l}, \mathbf{u} \rangle)$ and $\mathbf{c} \cdot \mathbf{x} \leq b$ iff $\mathbf{x} \in \text{refine}(\mathbf{c}, b)(\gamma(\mathcal{P}))$. This proves that, when $\mathbf{c} \notin \text{lin}(\gamma(\mathcal{P}))^\perp$, $\text{refine}^\alpha(\mathbf{c}, b)$ is γ -complete.

We now prove correctness when $\mathbf{c} \in \text{lin}(\gamma(\mathcal{P}))^\perp$. Note that, in this case,

$$\begin{aligned} \text{refine}(\mathbf{c}, b)(\mathcal{P}) &= \{\mathbf{x} \mid \mathbf{l} \leq A\mathbf{x} \leq \mathbf{u} \wedge \mathbf{c} \cdot \mathbf{x} \leq b\} = \\ &= \{A^{-1}\mathbf{z} \mid \mathbf{l} \leq \mathbf{z} \leq \mathbf{u} \wedge \mathbf{c} \cdot (A^{-1}\mathbf{z}) \leq b\} = \\ &= \{A^{-1}\mathbf{z} \mid \mathbf{l} \leq \mathbf{z} \leq \mathbf{u} \wedge ((A^T)^{-1}\mathbf{c}) \cdot \mathbf{z} \leq b\} = \\ &= \{A^{-1}\mathbf{z} \mid \mathbf{l} \leq \mathbf{z} \leq \mathbf{u} \wedge \mathbf{y} \cdot \mathbf{z} \leq b\} = \\ &= \{A^{-1}\mathbf{z} \mid \mathbf{z} \in \text{refine}(\mathbf{y}, b)(\gamma(\langle I_n, \mathbf{l}, \mathbf{u} \rangle))\} = \\ &= \{\mathbf{x} \mid A\mathbf{x} \in \text{refine}(\mathbf{y}, b)(\gamma(\langle I_n, \mathbf{l}, \mathbf{u} \rangle))\} \subseteq \\ &= \{\mathbf{x} \mid A\mathbf{x} \in \gamma(\text{refine}^\alpha(\mathbf{y}, b)(\langle I_n, \mathbf{l}, \mathbf{u} \rangle))\} \end{aligned}$$

If $\text{refine}^\alpha(\mathbf{y}, b)(\langle I_n, \mathbf{l}, \mathbf{u} \rangle) = \langle I_n, \mathbf{l}', \mathbf{u}' \rangle$, the last set in the previous equation is equal to $\gamma(\langle A, \mathbf{l}', \mathbf{u}' \rangle)$. Since $\langle A, \mathbf{l}', \mathbf{u}' \rangle$ is the result computed in lines 7–8 of the algorithm, we have that Algorithm 3 is correct. Moreover, since $\text{refine}^\alpha(\mathbf{c}, b)$ on boxes is optimal, the algorithm is also relatively optimal.

Finally, the computational complexity of the algorithm is bounded by the time needed to solve the system of linear equations in line 1, which is $O(n^3)$. \square

Theorem 17. *The abstract union operator described in Algorithm 5 is correct and relatively optimal. The computational complexity is $O(n^4)$.*

Proof. We first prove that the result of Algorithm 5 is correct.

Given two parallelotopes $\mathcal{P}_A, \mathcal{P}_B \in \text{Par}_n$ and a vector $\mathbf{v} \in \mathbb{R}^n$, Algorithm 4 computes a triple $\langle l, u, p \rangle \in \mathbb{R}^3$. By line 22 in Algorithm 4 we know that $l = \min(\inf\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_A\}, \inf\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_B\})$, and thus $l = \inf\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_A \cup \mathcal{P}_B\}$. Symmetrically we have that $u = \sup\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_A \cup \mathcal{P}_B\}$.

In Algorithm 5, we first fill a priority queue Q with the candidate constraints. The lines 2-6 ensure that all the rows of the matrices A and B are added to Q . By line 4, we insert in Q the constraints with their lower and upper bounds, as computed by Algorithm 4. With the lines 7-22 we insert in Q the candidate constraints which form an inversion, with their lower and upper bounds computed by Algorithm 4. The while loop in lines 24-29 extracts n linearly independent elements from Q which form the constraint matrix. First note that the queue Q contains at least n such linear forms. In fact, all the n constraints of A are in the queue, and they are linearly independent. Thus, lines 25-28 which extract n linearly independent elements from Q always succeed.

Moreover, for each $\langle \mathbf{v}, c, d \rangle$, in the queue, we know by Algorithm 4 that $c = \inf\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_A \cup \mathcal{P}_B\}$, and $d = \sup\{\mathbf{v} \cdot \mathbf{x} \mid x \in \mathcal{P}_A \cup \mathcal{P}_B\}$. Thus, $\gamma(\{\mathbf{x} \in \mathbb{R}^n \mid c \leq \mathbf{v} \cdot \mathbf{x} \leq d\}) \supseteq \gamma(\mathcal{P}_A) \cup \gamma(\mathcal{P}_B)$. By line 27 we know that the resulting matrix R is formed by n linear forms all enjoying this property. Then for the algorithm output $\langle R, \mathbf{l}', \mathbf{u}' \rangle$ it holds that $\gamma(\{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{l}' \leq R\mathbf{x} \leq \mathbf{u}'\}) = \bigcap_{i \in \{1, \dots, n\}} \gamma(\{\mathbf{x} \in \mathbb{R}^n \mid l'_i \leq \mathbf{r}_{i*} \cdot \mathbf{x} \leq u'_i\})$ which, by the previous property, $\supseteq \gamma(\mathcal{P}_A) \cup \gamma(\mathcal{P}_B)$. Thus the union operator is correct.

Relative optimality immediately follows from the fact that Algorithm 4 computes the tightest bounds in lines 1-4.

The computational complexity of Algorithm 4 is $O(n^3)$ by Theorem 4, and all the other operators on vectors have a lower complexity. But the complexity is $O(n^3)$ only the first time, since we need to compute the inverse of the two matrices. Successive calls have a complexity of $O(n^2)$.

Since Algorithm 4 is called n times in lines 2-5, n times in line 6, at most n^2 times in the loop in 7-19, and the same for lines 20, 21 and 22, we have that the complexity of all the vector operations is at most $O(n^4)$. The other part comes from the complexity for insert and extract operations in the priority queue. We perform at most $O(n^2)$ insert operations in lines 7-22, and at most $O(n^2)$ extract operations in line 25 (since the number of elements in the queue is at most $O(n^2)$), and thus the complexity is bounded by $O(n^4)$. \square

Lemma 27. *Given the parallelotope $\mathcal{P} = \langle A, \mathbf{l}, \mathbf{u} \rangle$ in \mathbb{R}^n and $A' \in \text{GL}(n)$, let $\mathcal{P}' = \text{rot}_{A'}(\mathcal{P}) = \langle A', \mathbf{l}', \mathbf{u}' \rangle$. Then, the sum of the number of infinite components in \mathbf{l}' and \mathbf{u}' is greater or equal then the sum of the number of infinite components in \mathbf{l} and \mathbf{u} .*

Proof. Each row of A' may be written as a linear combination of the rows in A . Assume we have

$$\mathbf{a}'_{j*} = \sum_{i=1}^n \lambda_i \mathbf{a}_{i*}$$

with $\lambda_k > 0$ and assume $u_k = +\infty$. Then either $\sup_{\mathbf{x} \in \mathcal{P}} \mathbf{a}'_{j*} \cdot \mathbf{x} = +\infty$ or $\inf_{\mathbf{x} \in \mathcal{P}} \mathbf{a}'_{j*} \cdot \mathbf{x} = -\infty$. Consider the set R of all the rows in A with an infinite upper bound. Then, there are at least $|R|$ rows in A' whose linear combinations involve rows in R . Actually, if this is not the case, it means that there are strictly more than $n - |R|$ linearly independent rows in A' which may be written as linear combination of $n - |R|$ rows in A , and this is impossible. Therefore, to each infinite upper bound in \mathcal{P} corresponds an infinite lower or upper bound in \mathcal{P}' . The same happens for the lower bounds. Therefore, the set of infinite bounds in \mathcal{P}' is no less than the set of infinite bounds in \mathcal{P} . \square

Theorem 18. *The operator $\bar{\nabla}$ is a widening.*

Proof. The fact that $\bar{\nabla}$ is an upper bound immediately derives from the fact that ∇ is an upper bound. Now consider a sequence $\mathcal{P}_0, \dots, \mathcal{P}_i, \dots$ of parallelotopes. We do not require the sequence to be increasing. We need to prove that the sequence $\{\mathcal{P}_i^w\}$ defined as $\mathcal{P}_0^w = \mathcal{P}_0$ and $\mathcal{P}_i^w = \mathcal{P}_{i-1}^w \bar{\nabla} \mathcal{P}_i$ for $i > 0$ is definitively stationary. Let $\mathcal{P}_i = \langle A_i, \mathbf{l}_i, \mathbf{u}_i \rangle$ and $\mathcal{P}_i^w = \langle A_i^w, \mathbf{l}_i^w, \mathbf{u}_i^w \rangle$.

We show that for each i , either $\mathcal{P}_i^w = \mathcal{P}_{i+1}^w$ or the number of infinite components in \mathbf{l}_i^w and \mathbf{u}_i^w is strictly smaller than the number of infinite components in \mathbf{l}_{i+1}^w and \mathbf{u}_{i+1}^w . This is obviously true when $\mathcal{P}_{i+1}^w = \mathcal{P}_i^w \nabla \mathcal{P}_i$. Now we consider the case $\text{rot}_{A_{i+1}}(\mathcal{P}_i^w) < \mathcal{P}_{i+1}$, which implies $\mathcal{P}_{i+1}^w = \text{rot}_{A_{i+1}}(\mathcal{P}_i^w) \nabla \mathcal{P}_{i+1}$. By Lemma 27 we now that the number of infinite bounds in $\text{rot}_{A_{i+1}}(\mathcal{P}_i^w)$ is greater or equal than the number of infinite bounds in \mathcal{P}_i^w . Since \mathcal{P}_{i+1} is strictly greater than $\text{rot}_{A_{i+1}}(\mathcal{P}_i^w)$ and we apply the ∇ widening, at least one finite bound is turned into an infinite bound. \square

Theorem 19. *In the hypothesis of Definition 10, let \bar{F} be a (possibly non-monotone) abstract operator $\bar{F} : A \rightarrow A$ which is a correct abstraction of the monotone operator $F : C \rightarrow C$. Given $a \in A$, let $c \in C$ be a fixpoint of F such that $c \subseteq \gamma(a)$. We define the iteration sequence $y_0 = a$, $y_{i+1} = y_i \Delta \bar{F}(y_i)$. Then*

- *the iteration sequence $\{y_i\}_{i \in \mathbb{N}}$ is decreasing and definitively stationary;*
- *for each $i \in \mathbb{N}$, y_i is a correct approximation of c .*

Proof. The fact that $\{y_i\}_{i \in \mathbb{N}}$ is decreasing and definitively stationary immediately follows from the first and third properties in the definition of narrowing. Now we prove that if $c \subseteq \gamma(y_i)$, i.e., y_i is a correct approximation of c , then $c \subseteq \gamma(y_{i+1})$. By the second property of narrowing, $\gamma(y_{i+1}) \supseteq \gamma(y_i) \cap \gamma(\bar{F}(y_i))$.

By correctness of \bar{F} , we have $\gamma(\bar{F}(y_i)) \supseteq F(\gamma(y_i)) \supseteq F(c) = c$, and $\gamma(y_{i+1}) \supseteq c$. \square

Theorem 20. *The operator Δ for parallelotopes is a narrowing according to Definition 10.*

Proof. Since $\mathcal{P}_A \Delta \mathcal{P}_B$ is built on the same constraint matrix of \mathcal{P}_A by possibly restricting the bounds, then $\mathcal{P}_A \Delta \mathcal{P}_B \leq \mathcal{P}_A$. Now, let \mathcal{P}_A , \mathcal{P}_B and $\text{rot}_A(\mathcal{P}_B)$ as in Definition 11. Assume $\mathbf{x} \in \gamma(\mathcal{P}_A) \cap \gamma(\mathcal{P}_B)$. Then, for each row i in A , we have $l_i \leq \mathbf{a}_{i*}\mathbf{x} \leq u_i$ and $j'_i \leq \mathbf{a}_{i*}\mathbf{x} \leq k'_i$, hence $\max(l_i, j'_i) \leq \mathbf{a}_{i*}\mathbf{x} \leq \min(u_i, k'_i)$. Since $l'_i \leq \max(l_i, j'_i)$ and $u'_i \geq \min(u_i, k'_i)$, then $l'_i \leq \mathbf{a}_{i*}\mathbf{x} \leq u'_i$, hence $\mathbf{x} \in \gamma(\mathcal{P}_A \Delta \mathcal{P}_B)$. Termination is guaranteed because bounds may be enlarged at most $2n$ times. \square

Proposition 21. *The operator \cap^α is a correct approximation of the concrete intersection. It is γ -complete when the two arguments are defined over the same constraint matrix.*

Proof. We first prove that when $\mathcal{P}_1 = \langle A, \mathbf{l}, \mathbf{u} \rangle$ and $\mathcal{P}_2 = \langle A, \mathbf{l}', \mathbf{u}' \rangle$ are defined over the same constraint matrix, then \cap_α is γ -complete. In this case the rot operation is a no-op, hence $\text{rot}_A(\mathcal{P}_2) = \mathcal{P}_2$. Let $\mathbf{x} \in \gamma(\mathcal{P}_1) \cap \gamma(\mathcal{P}_2)$. This is equivalent to $\mathbf{l} \leq A\mathbf{x} \leq \mathbf{u}$ and $\mathbf{l}' \leq A\mathbf{x} \leq \mathbf{u}'$. If we use \mathbf{l}'' and \mathbf{u}'' as in the definition of weak intersection, we have the equivalent property $\mathbf{l}'' \leq A\mathbf{x} \leq \mathbf{u}''$. If $\mathbf{l}'' \leq \mathbf{u}''$, this means $\mathbf{x} \in \gamma(\langle A, \mathbf{l}'', \mathbf{u}'' \rangle)$, otherwise the intersection is empty. In both cases $\mathbf{x} \in \gamma(\mathcal{P}_1) \cap^\alpha \gamma(\mathcal{P}_2)$ iff $\mathbf{x} \in \gamma(\mathcal{P}_1 \cap^\alpha \mathcal{P}_2)$.

Now, let us consider the general case. It is immediate to check that if $\mathcal{P}_1 = \langle A, \mathbf{l}, \mathbf{u} \rangle$, then $\mathcal{P}_1 \cap^\alpha \mathcal{P}_2 = \mathcal{P}_1 \cap^\alpha \text{rot}_A(\mathcal{P}_2)$. By the γ -completeness proved above, $\gamma(\mathcal{P}_1 \cap^\alpha \mathcal{P}_2) = \gamma(\mathcal{P}_1) \cap \gamma(\text{rot}_A(\mathcal{P}_2)) \supseteq \gamma(\mathcal{P}_1) \cap \gamma(\mathcal{P}_2)$, i.e., \cap^α is a correct abstraction of \cap . \square

Proposition 22. *The reduction operator is correct, i.e., given any $P \in \text{Par}$ and $B \in \text{Box}$, we have that $\gamma(\text{red}(\langle P, B \rangle)) = \gamma(\langle P, B \rangle)$.*

Proof. We need to prove that, given any $P \in \text{Par}$ and $B \in \text{Box}$, we have that $\gamma(\text{red}(\langle P, B \rangle)) \supseteq \gamma(P) \cap \gamma(B)$.

First note that, given any $P \in \text{Par}$ and matrix $A \in \text{GL}(n)$, we have that $P \leq \text{rot}_A(P)$. Thus, the weak intersection on parallelotopes correctly approximates the concrete intersection, namely

$$\gamma(P_1 \cap^\alpha P_2) \supseteq \gamma(P_1) \cap \gamma(P_2)$$

and therefore

$$\gamma(\text{red}(\langle P, B \rangle)) \supseteq \gamma(P) \cap \gamma(B)$$

which proves the result. \square

Corollary 23. *All the operations on $\text{Par} \sqcap \text{Box}$ defined according to Equation 2 are correct, namely:*

$$\gamma(\langle P_1, B_1 \rangle \odot^{P, B} \langle P_2, B_2 \rangle) \supseteq \gamma(\langle P_1, B_1 \rangle) \odot \gamma(\langle P_2, B_2 \rangle)$$

Proof. It immediately follows from Proposition 22. \square

References

- [1] P. Cousot, N. Halbwachs, Automatic discovery of linear restraints among variables of a program, in: POPL '78: Proceedings of the 5th ACM SIGACT-SIGPLAN symposium on Principles of programming languages, ACM Press, New York, NY, USA, 84–97, doi:10.1145/512760.512770, 1978.
- [2] P. Cousot, R. Cousot, Static determination of dynamic properties of programs, in: Proceedings of the Second International Symposium on Programming, Dunod, Paris, France, 106–130, 1976.
- [3] A. Miné, The Octagon Abstract Domain, Higher-Order and Symbolic Computation 19 (1) (2006) 31–100, ISSN 1388-3690, doi:10.1007/s10990-006-8609-1.
- [4] R. Clarisó, J. Cortadella, The octahedron abstract domain, Science of Computer Programming 64 (2007) 115–139, ISSN 0167-6423, doi:10.1016/j.scico.2006.03.009.
- [5] J. M. Howe, A. Simon, Logahedra: A New Weakly Relational Domain, in: Z. Liu, A. P. Ravn (Eds.), Automated Technology for Verification and Analysis, 7th International Symposium, ATVA 2009, Macao, China, October 14–16, 2009. Proceedings., vol. 5799 of *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, ISBN 978-3-642-04760-2, 306–320, doi:10.1007/978-3-642-04761-9_23, 2009.
- [6] A. Simon, A. King, J. M. Howe, Two Variables per Linear Inequality as an Abstract Domain, in: M. Leuschel (Ed.), Logic Based Program Synthesis and Transformation 12th International Workshop, LOPSTR 2002, Madrid, Spain, September 17–20, 2002. Revised Selected Papers, vol. 2664 of *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, ISBN 978-3-540-40438-5, 71–89, doi:doi://10.1007/3-540-45013-0_7, 2003.
- [7] S. Sankaranarayanan, H. B. Sipma, Z. Manna, Scalable Analysis of Linear Systems using Mathematical Programming, in: R. Cousot (Ed.), Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Paris, France, January 17–19, 2005. Proceedings, vol. 3385 of *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, ISBN 978-3-540-24297-0, 25–41, doi:10.1007/b105073, 2005.
- [8] M. Colón, S. Sankaranarayanan, Generalizing the template polyhedral domain, in: G. Barthe (Ed.), Programming Languages and Systems: 20th European Symposium on Programming, ESOP 2011, held as part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011. Proceedings, vol. 6602 of *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, ISBN 978-3-642-19717-8, 176–195, 2011.
- [9] G. Amato, M. Parton, F. Scozzari, Deriving Numerical Abstract Domains via Principal Component Analysis, in: R. Cousot, M. Martel (Eds.), 17th

- International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010, Proceedings, vol. 6337 of *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, 134–150, doi:10.1007/978-3-642-15769-1, 2010.
- [10] G. Amato, M. Parton, F. Scozzari, Discovering invariants via Simple Component Analysis, *Journal of Symbolic Computation* 47 (12) (2012) 1533–1560, doi:10.1016/j.jsc.2011.12.052.
- [11] G. Amato, F. Scozzari, The abstract domain of parallelotopes, in: J. Midtgaard, M. Might (Eds.), Proceedings of the Fourth International Workshop on Numerical and Symbolic Abstract Domains, NSAD 2012, vol. 287 of *Electronic Notes in Theoretical Computer Science*, Elsevier, 17–28, doi:10.1016/j.entcs.2012.09.003, 2012.
- [12] A. Schrijver, *Theory of Linear and Integer Programming*, John Wiley & Sons, 1986.
- [13] P. Cousot, R. Cousot, Abstract Interpretation Frameworks, *Journal of Logic and Computation* 2 (4) (1992) 511–549, doi:10.1093/logcom/2.4.511.
- [14] S. Sankaranarayanan, M. Colón, H. B. Sipma, Z. Manna, Efficient Strongly Relational Polyhedral Analysis, in: E. A. Emerson, K. S. Namjoshi (Eds.), VMCAI, vol. 3855 of *Lecture Notes in Computer Science*, Springer, ISBN 3-540-31139-4, 111–125, 2006.
- [15] P. Cousot, R. Cousot, Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints, in: POPL '77: Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages, ACM Press, New York, NY, USA, 238–252, doi:10.1145/512950.512973, 1977.
- [16] P. Cousot, Abstract Induction by Extrapolation and Interpolation, in: D. D’Souza, A. Lal, K. G. Larsen (Eds.), Verification, Model Checking, and Abstract Interpretation. 16th International Conference, VMCAI 2015, Mumbai, India, January 12-14, 2015, Proceedings, vol. 8931 of *Lecture Notes in Computer Science*, Springer, 19–42, doi:10.1007/978-3-662-46081-8_2, 2015.
- [17] G. Amato, F. Scozzari, Jandom, <https://github.com/jjandom-devel/Jandom>, 2016.
- [18] V. Maisonneuve, O. Hermant, F. Irigoin, ALICe: A framework to improve affine loop invariant computation, in: 5th Workshop on INvariant Generation, 2014.
- [19] P. Granger, Improving the Results of Static Analyses Programs by Local Decreasing Iteration, in: R. Shyamasundar (Ed.), Foundations of Software Technology and Theoretical Computer Science. 12th Conference New Delhi, India, December 18–20, 1992 Proceedings, *Lecture Notes in Computer Science*, Springer, 68–79, doi:10.1007/3-540-56287-7_95, 1992.

- [20] R. Bagnara, P. M. Hill, E. Zaffanella, The Parma Polyhedra Library: Toward a Complete Set of Numerical Abstractions for the Analysis and Verification of Hardware and Software Systems, *Science of Computer Programming* 72 (1–2) (2008) 3–21, ISSN 0167-6423, doi: 10.1016/j.scico.2007.08.001.
- [21] R. Bagnara, P. M. Hill, E. Ricci, E. Zaffanella, Precise widening operators for convex polyhedra, *Science of Computer Programming* 58 (1) (2005) 28–56, ISSN 0167-6423, doi:10.1016/j.scico.2005.02.003.