



HAL
open science

Efficient Elimination of Redundancies in Polyhedra using Raytracing

Alexandre Maréchal, Michaël Périn

► **To cite this version:**

Alexandre Maréchal, Michaël Périn. Efficient Elimination of Redundancies in Polyhedra using Raytracing. Verification, Model Checking, and Abstract Interpretation (VMCAI), David Monniaux; Ahmed Bouajjani, Jan 2017, Paris, France. hal-01385653v2

HAL Id: hal-01385653

<https://hal.science/hal-01385653v2>

Submitted on 4 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Efficient Elimination of Redundancies in Polyhedra using Raytracing¹²

Alexandre Maréchal[‡], Michaël Périn[§]

Verimag Research Report n^o 6

October 2016

*This paper is an extended version of an article published in VMCAI'17 [15].

†This work was partially supported by the **European Research Council** under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement nr. 306595 "STATOR".

‡alex.marechal@imag.fr

§michael.perin@imag.fr

Reports are downloadable at the following address

<http://www-verimag.imag.fr>

Unité Mixte de Recherche 5104 CNRS - Grenoble INP - UGA

Bâtiment IMAG
Université Grenoble Alpes
700, avenue centrale
38401 Saint Martin d'Hères
France
tel : +33 4 57 42 22 42
fax : +33 4 57 42 22 22
<http://www-verimag.imag.fr/>

Efficient Elimination of Redundancies in Polyhedra using Raytracing⁵⁶

Alexandre Maréchal*, Michaël Périn†

Université Grenoble-Alpes, VERIMAG, F-38000 Grenoble, France
CNRS, VERIMAG, F-38000 Grenoble, France

October 2016

Abstract

Polyhedra are used in verification and automatic parallelization to capture linear relations between variables. A polyhedron can be represented as constraints, generators or both in the double description framework. Whatever the representation, most polyhedral operators spend a significant amount of time to maintain minimal representations. To minimize a polyhedron in constraints-only representation, the redundancy of each constraint must be checked with respect to others. Each of these redundancy tests generally implies solving a linear programming (LP) problem using the simplex algorithm. We present an algorithm that replaces most LP problem resolutions by distance computations. The geometric intuition is simple: consider ray traces starting from a point within the polyhedron and orthogonal to its faces. A face first encountered by one of these rays is an actual face of the polyhedron. It is therefore an irredundant constraint. Since this procedure is incomplete, LP problem resolutions are required for the remaining undetermined constraints. Experiments show that our algorithm drastically reduces the number of calls to the simplex, resulting in a considerable speed improvement. In addition, our algorithm generates by construction a certificate for each constraint: redundancy is established by exhibiting a nonnegative linear combination of constraints that yields the removed constraint, whereas irredundancy is shown by exhibiting a point outside of the polyhedron which is excluded only by the kept constraint. To follow the geometric interpretation, the algorithm is explained in terms of constraints but it can also be used to minimize generators.

Keywords: Polyhedra, Minimization, Linear Programming

Reviewers: Nicolas Halbwachs

[¶]This paper is an extended version of an article published in VMCAI'17 [15].

^{||}This work was partially supported by the [European Research Council](#) under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement nr. 306595 "STATOR".

*alex.marechal@imag.fr

†michael.perin@imag.fr

Contents

1	Redundancy in Polyhedra	2
2	Notations	4
3	Certifying a Minimization of Polyhedra	4
4	An Efficient Minimization Algorithm	5
4.1	The Frontier Detection Criterion	6
4.2	Irredundancy Certificates	9
4.3	Minimizing Generators	9
4.4	Using Floating Points in Raytracing	10
5	Experiments	10
6	Conclusion & Future Work	14
A	Proofs	16

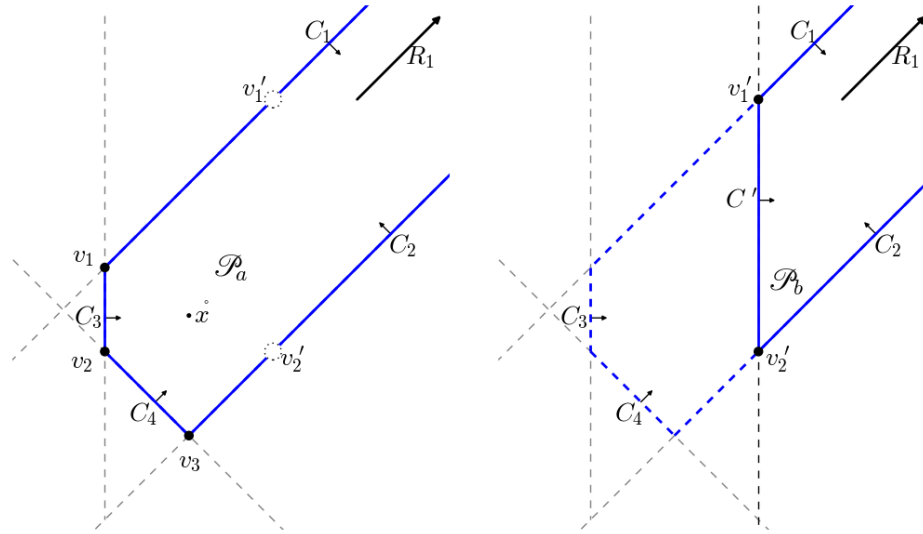


Figure 1: Emergence of redundant constraints and generators in polyhedra.

1 Redundancy in Polyhedra

Convex polyhedra are used in static analysis [2] and automatic parallelization [6] to capture linear inequalities of the form $\sum_{i=1}^n a_i x_i \leq b$ relating the program variables x_1, \dots, x_n .³ A polyhedron \mathcal{P} can be defined as the set of points $\mathbf{x} = (x_1, \dots, x_n)$ that satisfy a system of inequalities $\mathbf{A}\mathbf{x} \leq \mathbf{b}$. The ℓ^{th} row of the augmented matrix $[\mathbf{A} \mid -\mathbf{b}]$ is a vector $\mathbf{C}_\ell = (a_{\ell 1}, \dots, a_{\ell n}, -b_\ell)$ which encodes the constraint $\sum_{i=1}^n a_{\ell i} x_i \leq b_\ell$. A constraint \mathbf{C}_ℓ defines the bounding hyperplane normal to $(a_{\ell 1}, \dots, a_{\ell n})$ and shifted by b_ℓ (see Fig. 1). Alternatively the same set of points can be defined as the convex combination of generators (vertices and rays), i.e. $\left\{ \mathbf{x} \mid \mathbf{x} = \sum_{i=1}^v \beta_i \mathbf{v}_i + \sum_{i=1}^r \lambda_i \mathbf{R}_i, \beta_i, \lambda_i \geq 0, \sum \beta_i = 1 \right\}$ where \mathbf{R}_i 's and \mathbf{v}_i 's denote respectively rays and vertices. Fig. 1 shows two polyhedra \mathcal{P}_a and \mathcal{P}_b defined in double description as the set of constraints $\{\mathbf{C}_1 : x_2 - x_1 \leq 1, \mathbf{C}_2 : x_2 - x_1 \geq -2, \mathbf{C}_3 : x_1 \geq 1, \mathbf{C}_4 : x_1 + x_2 \geq 2\}$ and the set of generators $\{\mathbf{v}_1 : (1, 2), \mathbf{v}_2 : (1, 1), \mathbf{v}_3 : (2, 0), \mathbf{R}_1 : (1, 1)\}$ for \mathcal{P}_a ; respectively as $\{\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}' : x_1 \geq 3\}$ and $\{\mathbf{v}'_1 : (3, 4), \mathbf{v}'_2 : (3, 1), \mathbf{R}_1\}$ for \mathcal{P}_b .

The addition of new constraints or generators introduces redundancies which must be removed to reduce memory consumption and avoid useless computations in subsequent operations. In constraints-only representation, redundant constraints tend to grow exponentially during the computation of a projection by Fourier-Motzkin elimination [18]. For a description by generators the same pitfall occurs when a polyhedron is sliced with a constraint [8]. The emergence of redundancies is illustrated by Fig. 1: when constraint \mathbf{C}' is added into \mathcal{P}_a to form \mathcal{P}_b , constraints \mathbf{C}_3 and \mathbf{C}_4 become redundant. Conversely, the addition of points $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ into \mathcal{P}_b generates \mathcal{P}_a and makes \mathbf{v}'_1 and \mathbf{v}'_2 redundant.

Characterization of Redundancy. A ray \mathbf{R}_k is redundant if it is a nonnegative combination of the other rays and a point \mathbf{v}_k is redundant if it is a convex combination of the other generators, i.e. $\mathbf{v}_k = \sum_i \beta_i \mathbf{v}_i + \sum_i \lambda_i \mathbf{R}_i$ for some $\beta_i, \lambda_i \geq 0$ with $\sum \beta_i = 1$ and $\beta_k = 0$. Back to our running example, the equations $\mathbf{v}'_1 = 1 \times \mathbf{v}_1 + 2 \times \mathbf{R}_1$ and $\mathbf{v}'_2 = 1 \times \mathbf{v}_3 + 1 \times \mathbf{R}_1$ prove the redundancy of \mathbf{v}'_1 and \mathbf{v}'_2 in \mathcal{P}_a . Therefore, these equations account for *certificates of redundancy*.

Intuitively, a constraint is redundant if it is useless, in the sense that adding it does not change the geometrical space delimited by the polyhedron. Formally, a constraint \mathbf{C}_k is redundant if it is a nonnegative combination of other constraints. As we did for generators, we can find equations that prove the redundancy of \mathbf{C}_3 and \mathbf{C}_4 in \mathcal{P}_b , but we need to consider the tautological constraint $\mathbf{C}_0 : 1 \geq 0$ as being part of the system in order to exactly fit the constant b of the redundant constraint.

³We only deal with convex polyhedra. For readability, we will omit the adjective *convex* in the following.

Example. The equations $C_3 = 1 \times C' \oplus 2 \times C_0$ and $C_4 = 2 \times C' \oplus 1 \times C_2 \oplus 2 \times C_0$ are called the *Farkas decomposition* of C_3 and C_4 . They act as certificates of redundancy. Indeed, $C_3 : x_1 \geq 1 \equiv (x_1 \geq 3) \oplus 2 \times (1 \geq 0)$ and $C_4 : x_1 + x_2 \geq 2 \equiv 2 \times (x_1 \geq 3) \oplus (x_2 - x_1 \geq -2) \oplus 2 \times (1 \geq 0)$ where $(l \geq r) \oplus (l' \geq r') \stackrel{\text{def}}{=} l + l' \geq r + r'$.

If only one representation is available – as generators or as constraints – discovering redundancy requires solving linear programming (LP) problems of the form “does there exist nonnegative scalars satisfying some linear equations?”:

$$\exists \lambda_0, \dots, \lambda_p \geq 0, \mathbf{C}_k = \sum_{i=0, i \neq k}^p \lambda_i \mathbf{C}_i \quad (1) \text{ for constraints}$$

$$\exists \lambda_1, \dots, \lambda_r \geq 0, \mathbf{R}_k = \sum_{i=1, i \neq k}^r \lambda_i \mathbf{R}_i \quad (2) \text{ for rays}$$

$$\begin{aligned} \exists \beta_1, \dots, \beta_v, \lambda_1, \dots, \lambda_r \geq 0, \mathbf{v}_k &= \sum_{i=1, i \neq k}^v \beta_i \mathbf{v}_i + \sum_{i=1}^r \lambda_i \mathbf{R}_i \quad (3) \text{ for vertices} \\ &\wedge \sum_{i=1, i \neq k}^v \beta_i = 1 \end{aligned}$$

Polyhedral Cones. The way to reconcile the two definitions of redundancy is to switch to *polyhedral cones* to get a homogeneous system of constraints and only rays as generators. The trick for changing a polyhedron \mathcal{P} into a cone is to associate an extra variable η to the constant term \mathbf{b} as follows [19]: $\mathbf{A}\mathbf{x} \leq \mathbf{b} \equiv \eta(\mathbf{A}\mathbf{x}) \leq \eta\mathbf{b} \equiv \mathbf{A}(\eta\mathbf{x}) - \eta\mathbf{b} \leq 0 \equiv [\mathbf{A}] - \mathbf{b} \begin{pmatrix} \eta\mathbf{x} \\ \eta \end{pmatrix} \leq 0$ for any $\eta > 0$. It can be proved [9] that $\mathbf{x} \in \mathbb{Q}^n$ belongs to \mathcal{P} if and only if $\begin{pmatrix} \mathbf{x} \\ 1 \end{pmatrix} \in \mathbb{Q}^{n+1}$ belongs to the cone $\{\mathbf{x}' \in \mathbb{Q}^{n+1} \mid \mathbf{A}'\mathbf{x}' \leq 0\}$ where $\mathbf{A}' = [\mathbf{A}] - \mathbf{b}$. Using this transformation, operators on polyhedra can be implemented as computations on their associated cones producing a cone that, once intersected with the hyperplane $\eta = 1$, is the expected polyhedron. We switch back to polyhedra in illustrations as they are easier to draw. Considering cones simplifies the presentation: the constant term of constraints and the vertices disappear from definitions. Then, we end up with the same definition of redundancy for constraints (1) and for generators (2): a vector is redundant if it is a nonnegative combination of the others vectors. A constraint \mathbf{C}_k (resp. a ray \mathbf{R}_k) is redundant if

$$\begin{aligned} \exists \lambda_1, \dots, \lambda_p \geq 0, \mathbf{C}_k &= \sum_{i=1}^p \lambda_i \mathbf{C}_i \wedge \lambda_k = 0 \quad (1) \text{ for constraints} \\ \exists \lambda_1, \dots, \lambda_r \geq 0, \mathbf{R}_k &= \sum_{i=1}^r \lambda_i \mathbf{R}_i \wedge \lambda_k = 0 \quad (2) \text{ for rays} \end{aligned}$$

There is a technicality about the tautological constraint $\mathbf{C}_0 : 1 \geq 0$ in cones.

Example. Consider the constraint $\mathbf{C}_1 : x_1 \geq 1$ and $\mathbf{C}_2 : x_1 \geq 2$. Clearly \mathbf{C}_2 is redundant as $\mathbf{C}_2 = \mathbf{C}_1 \oplus \mathbf{C}_0$. The cone associated to the polyhedron $\mathcal{P} = \{\mathbf{C}_1, \mathbf{C}_2\}$ is $\text{Cone}(\mathcal{P}) = \{\eta x_1 - 1\eta \geq 0, \eta x_1 - 2\eta \geq 0, \eta > 0\} \equiv \{y_1 - y_2 \geq 0, y_1 - 2y_2 \geq 0, y_2 > 0\}$ with $\eta = y_2$ and $x_1 = \frac{y_1}{\eta}$.

Deciding Redundancy. The redundant/irredundant status of a constraint or a ray depends on the satisfiability of an existential problem (1,2) involving linear equations but also inequalities ($\bigwedge_i \lambda_i \geq 0$). Thus, such a problem does not fall within the realm of linear algebra but in that of LP for which the simplex algorithm is a standard solver [5]. In practice, the simplex performs much better than its theoretical exponential complexity – but still remains a costly algorithm. So, much research has been devoted to identifying many cases where the simplex can be avoided. Wilde [19] and Lassez *et al.* [13] suggest several *fast redundancy-detection criteria* before switching to the general LP problem:

- **The quasi-syntactic redundancy test** considers pairs of constraints and looks for single constraint redundancies of the form $\mathbf{C}' = \lambda\mathbf{C}$ with $\lambda > 0$, e.g. $\mathbf{C}' : 4x_1 - 6x_2 \geq 2$ is redundant with respect to $\mathbf{C} : x_1 - 3x_2 \geq 1$ since $\mathbf{C}' = 2 \times \mathbf{C}$.
- **The bound shifting test** exploits the implication $\sum_{i=1}^n a_i x_i \leq b \implies \sum_{i=1}^n a_i x_i \leq b'$ if $b \leq b'$. Hence, when the coefficients of two constraints \mathbf{C} and \mathbf{C}' only differ on b and b' with $b \leq b'$ then \mathbf{C}' is redundant and the certificate is $\mathbf{C}' = \mathbf{C} \oplus (b' - b) \times \mathbf{C}_0$ where \mathbf{C}_0 is the tautology $0 \leq 1$.

- **The combination of single variable inequalities** such as $x_1 \leq b_1$ and $x_2 \leq b_2$ entails for instance the redundancy of $C : 2x_1 + 3x_2 \leq b$ with $2b_1 + 3b_2 \leq b$. The corresponding certificate is $C = 2 \times (x_1 \leq b_1) \oplus 3 \times (x_2 \leq b_2) \oplus (2b_1 + 3b_2 - b) \times (0 \leq 1)$.

While these criteria can detect certain redundancies at low-cost, in this paper we investigate the other side of redundancy: we provide a *fast criterion to detect irredundant constraints*. The combination of the two approaches limits the usage of the simplex to constraints that are neither decided by our criteria nor by those of Wilde and Lassez *et al.*

Contributions. We present an algorithm that replaces most LP problem resolutions by distance computations. It is detailed in §4, after introducing useful notations in §2. The geometric intuition of our irredundancy criterion is simple: consider ray traces starting from a point within the polyhedron and orthogonal to its bounding hyperplanes. The hyperplane first encountered by one of these rays is an actual face of the polyhedron. It is therefore an irredundant constraint. Since this procedure is incomplete, LP problem resolutions are required for the remaining undetermined constraints. Experiments of §5 show that our algorithm drastically reduces the number of calls to the simplex, resulting in a considerable speed improvement. In addition, our algorithm generates certificates of correctness, precision and minimality which make it usable in a certified static analyzer. Certificates are presented in §3. To follow the geometric interpretation, the algorithm is explained below in terms of constraints but it can similarly be used to minimize generators. We conclude in §6 by a discussion of the potential benefit of integrating our algorithm in the double description framework.

2 Notations

Vectors and matrices are written in boldface to be distinguished from scalars, e.g. $\mathbf{0}$ is a vector of 0. For clarity and without loss of generality the rest of the paper will focus on *polyhedral cones over rationals*. A polyhedral cone \mathcal{P} of p constraints on n variables (x_1, \dots, x_n) is a conjunction (written as a set) of homogeneous linear constraints $\{C_1, \dots, C_p\}$ of the form $C_\ell : \sum_{i=1}^n a_{\ell i} x_i \leq 0$ where $C_\ell = (a_{\ell 1}, \dots, a_{\ell n})$. The inner product of vectors offers a convenient notation $\langle C_\ell, \mathbf{x} \rangle \leq 0$ for that inequality. Then, the cone \mathcal{P} corresponds to a matrix inequality $\mathbf{A}\mathbf{x} \leq \mathbf{0}$ where the rows of \mathbf{A} are the vectors C_1, \dots, C_p . Finally, $\{C_1, \dots, C_p\}$, $\bigwedge_{\ell=1}^p \langle C_\ell, \mathbf{x} \rangle \leq 0$ or $\mathbf{A}\mathbf{x} \leq \mathbf{0}$ are three equivalent ways of denoting a polyhedral cone \mathcal{P} . We use $\llbracket \mathcal{P} \rrbracket$ to specifically refer to the set of points defined by \mathcal{P} . Given a cone $\mathcal{P} : \mathbf{A}\mathbf{x} \leq \mathbf{0}$, the same system with a strict inequality defines $\overset{\circ}{\mathcal{P}}$, the interior of \mathcal{P} , and $\overset{\circ}{\mathbf{x}}$ denotes a point of $\llbracket \overset{\circ}{\mathcal{P}} \rrbracket \stackrel{\text{def}}{=} \{\mathbf{x} \mid \mathbf{A}\mathbf{x} < \mathbf{0}\}$.

3 Certifying a Minimization of Polyhedra

Our minimization algorithm is part of the Verimag Polyhedra Library (VPL) which operates on rational polyhedra in constraints-only representation. It was originally designed by Fouilhé *et al.* [7] as an abstract domain for the VERASCO certified static analyzer whose soundness is proved in COQ [12]. VERASCO can collaborate with an external library in OCAML such as the VPL, provided that it produces certificates of correctness, allowing a COQ-checker to verify the results computed in OCAML. In this section we recall the algorithm used in the original VPL for minimizing a *polyhedral cone represented as a set of constraints*. It is the standard algorithm but extended to produce on-the-fly *certificates of correctness, precision and minimality*. We recall the *fundamental theorem of linear inequalities* due to Farkas (1894) which ensures the existence of such certificates. Revisiting this theorem with a geometrical interpretation reveals an efficient way to determine irredundant constraints, which will be the key of our algorithm (§4).

Minimizing a cone \mathcal{P} consists in removing all redundant constraints such that the result, \mathcal{P}_M , represents the same geometrical space, *i.e.* $\llbracket \mathcal{P} \rrbracket = \llbracket \mathcal{P}_M \rrbracket$. Two certificates are needed to prove that equality: (1) one for the inclusion $\llbracket \mathcal{P} \rrbracket \subseteq \llbracket \mathcal{P}_M \rrbracket$ which guarantees *the correctness of the minimization* and (2) another one for $\llbracket \mathcal{P}_M \rrbracket \subseteq \llbracket \mathcal{P} \rrbracket$ which justifies its *precision*. A third certificate (3) ensures *the minimality of the result* showing that all constraints of \mathcal{P}_M are irredundant.

Certificate (1) must prove that each point of $\llbracket \mathcal{P} \rrbracket$ belongs to $\llbracket \mathcal{P}_M \rrbracket$. In the particular case of minimization, inclusion (1) is trivial because \mathcal{P}_M is obtained by only removing constraints from \mathcal{P} , which necessarily

leads to a larger set of points. By contrast, the existence of certificates (2) and (3) is not straightforward but the consequence of the following theorem, which we rephrased in our constraint terminology to ease its interpretation.

Theorem 1 (Fundamental theorem of linear inequalities [17, 7.1 p.85]). *Let C_1, \dots, C_p and C' be vectors in a n -dimensional space. Then,*

- (I) *either C' is **redundant** and there exists a Farkas decomposition of C' that is a nonnegative linear combination of linearly independent vectors from C_1, \dots, C_p , i.e. $C' = \lambda_1 C_1 + \dots + \lambda_p C_p$ for some scalars $\lambda_1, \dots, \lambda_p \geq 0$.*
- (II) *or C' is **irredundant** and there exists a n -dimensional vector w such that $\langle C', w \rangle > 0$ and $\langle C_1, w \rangle, \dots, \langle C_p, w \rangle \leq 0$.*

The standard algorithm (Algorithm 1) exploits the redundancy criterion (I) of the theorem which was already illustrated in §1 Example 1. The existence of a *Farkas decomposition* of C' is decided by solving a LP problem. If the simplex algorithm returns a solution λ then the pair (C', λ) is recorded as a *certificate of precision* (2) which proves that the removed constraint was indeed redundant. To get rid of all the redundancies, Algorithm 1 needs one execution of the simplex algorithm for each constraint.

Given an existential LP problem, the simplex can return either a solution or an explanation of the lack of solution. The proof of Theorem 1 and the simplex algorithm have strong connections which result in an interesting feature of the VPL simplex: calling $\text{simplex}(\exists \lambda_i \geq 0, C' = \sum_i \lambda_i C_i)$ returns either $\text{SUCCESS}(\lambda)$ or $\text{FAILURE}(w)$ such that $\langle C', w \rangle > 0 \wedge_i \langle C_i, w \rangle \leq 0$.⁴ This feature is a consequence of Theorem 1 and requires no additional computation.

When the simplex returns $\text{FAILURE}(w)$, the irredundancy criterion (II) of the theorem tells that C' is irredundant and must be kept in the set of constraints. Algorithm 1 builds the *certificate of minimality* (3) by associating a *witness point* to each constraint of the minimized polyhedron \mathcal{P}_M .

While the standard algorithm focuses on criterion (I), we revisit the theorem paying attention to the geometrical interpretation of criterion (II): when a constraint C' is irredundant, its associated bounding hyperplane is a *frontier* of the polyhedron separating the inside from the outside. Part (II) of the theorem ensures that we can exhibit a witness point w , outside of $\llbracket \mathcal{P} \rrbracket$, satisfying all constraints of \mathcal{P} except C' . The rest of the paper is dedicated to an algorithm that efficiently discovers such witness points.

Algorithm 1: The standard minimization algorithm (used in VPL 0.1)

Input : A set of constraints $\{C_1, \dots, C_p\}$.

Output: \mathcal{P}_M = the irredundant constraints of $\{C_1, \dots, C_p\}$
 (R, I) = the redundancy and irredundancy certificates

$\mathcal{P}_M \leftarrow \{C_1, \dots, C_p\}$

for C' *in* $\{C_1, \dots, C_p\}$ **do**

switch $\text{simplex} \left(\exists \lambda_i \geq 0, C' = \sum_{C_i \in \mathcal{P}_M \setminus C'} \lambda_i C_i \right)$ **do**

case $\text{SUCCESS}(\lambda)$: **do** $R \leftarrow R \cup \{(C', \lambda)\}$; $\mathcal{P}_M \leftarrow \mathcal{P}_M \setminus C'$
case $\text{FAILURE}(w)$: **do** $I \leftarrow I \cup \{(C', w)\}$

return (\mathcal{P}_M, R, I)

4 An Efficient Minimization Algorithm

Building up on the geometric interpretation of Theorem 1, we present a new minimization algorithm for polyhedral cones that brings two major improvements: it reduces the number of calls to the simplex algo-

⁴Conversely, $\text{simplex}(\exists w, \langle C', w \rangle > 0 \wedge_i \langle C_i, w \rangle \leq 0)$ returns either $\text{SUCCESS}(w)$ or $\text{FAILURE}(\lambda)$ such that $C' = \sum_i \lambda_i C_i$.

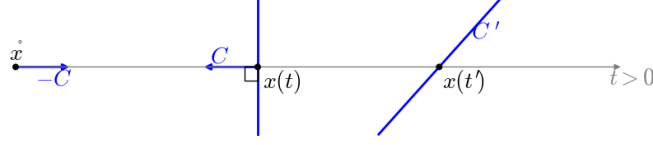


Figure 2: The ray starting at the interior point \hat{x} and orthogonal to a constraint C meets C and possibly others constraints.

rithm and limits the constraints they involve. The key idea of the algorithm is to trace rays starting from a point in the interior of the cone. The first hyperplane encountered by a ray is a frontier of the polyhedron, thus an irredundant constraint. Unfortunately, with a limited number of rays, some frontiers can be missed depending on the cone and the position of the interior point. This *raytracing procedure* is thus incomplete and LP problem resolutions are still required for the remaining undetermined constraints.

While the simplex algorithm is used in the standard minimization to discover Farkas decompositions, we rather use it to get closer to a witness point, and only when all previous rays failed to prove the irredundancy of a constraint. Of course, if the constraint is redundant, the simplex algorithm returns no point at all but an explanation of its failure which is nothing else than a Farkas decomposition proving the redundancy.

4.1 The Frontier Detection Criterion

We now detail the process of finding witness points by raytracing. We consider a cone \mathcal{P} with a nonempty interior. Then, there exists a point \hat{x} in $\overset{\circ}{\mathcal{P}}$. The basic operation of our algorithm consists in sorting the constraints of \mathcal{P} with respect to the order in which they are hit by a *ray*, *i.e.* a half-line starting at the interior point \hat{x} and extending along a given direction d .

Consider the constraint $\langle C, x \rangle \leq 0$. The hyperplane of the constraint is $\{x \mid \langle C, x \rangle = 0\}$, *i.e.* the set of points orthogonal to vector C . The ray starting at \hat{x} and extending in direction d is the set of points $\{x(t) \mid x(t) = \hat{x} + t \times d, t \geq 0\}$. Let us assume that the ray hits the C -hyperplane at point x_c . Then, there exists $t_c \geq 0$ such that $x_c = \hat{x} + t_c \times d$ and so, $x_c - \hat{x} = t_c \times d$. Therefore, the distance $\|\hat{x} - x_c\|$ is just a scaling by $|t_c|$ of the norm $\|d\|$ which does not depend on C . Hence, *by computing $|t_c|$ for each constraint we will be able to know in which order the constraints are hit by the ray*. Prior to computing t_c we check if the ray can hit the constraint, *i.e.* $\langle C, d \rangle \neq 0$. Then, we use the fact that $x_c \in \{x \mid \langle C, x \rangle = 0\}$ to get $t_c = -\frac{\langle C, \hat{x} \rangle}{\langle C, d \rangle}$. Indeed,

$$0 = \langle C, x_c \rangle = \langle C, \hat{x} + t_c \times d \rangle = \langle C, \hat{x} \rangle + t_c \times \langle C, d \rangle.$$

Hence, the basic operation of our raytracing algorithm consists in two evaluations of each constraint C of \mathcal{P} at \hat{x} and d in order to compute the scalar t_c . Let us explain how we exploit this information to discover actual frontiers of \mathcal{P} .

Note that any direction could be used to sort the constraints with respect to the order of intersection by a ray. We choose successively for d the opposite direction of the normal vector of each bounding hyperplane of \mathcal{P} . This heuristic ensures that each hyperplane will be hit by at least one ray. As illustrated by Fig. 2, a direction $d \stackrel{\text{def}}{=} -C$ necessarily intersects the C -hyperplane and may potentially cross many other constraints for some values of t . Considering a direction $d_i = -C_i$, we sort the intersected hyperplanes with respect to the increasing order of the scalar t , which is proportional to the distance between the interior point \hat{x} and the intersection point $x(t)$ of an hyperplane and the ray $\text{RAY}(\hat{x}, d_i)$. We obtain a sorted *intersection list* of pairs (t, S_t) where S_t is the set of the (possibly many) constraints vanishing at $x(t)$. If a constraint C is not hit by the ray (because $\langle C, d_i \rangle = 0$), then C is not added to the intersection list. The head pair provides the constraints which are encountered first by the ray. At the heart of our algorithm is the following proposition: “If the head of an intersection list is a pair $(t, \{C\})$ with a *single constraint*, then C is a frontier of \mathcal{P} ; otherwise we cannot conclude from this list.” This will be proved in §4.2 (Proposition 1) when we will come to the generation of witness points.

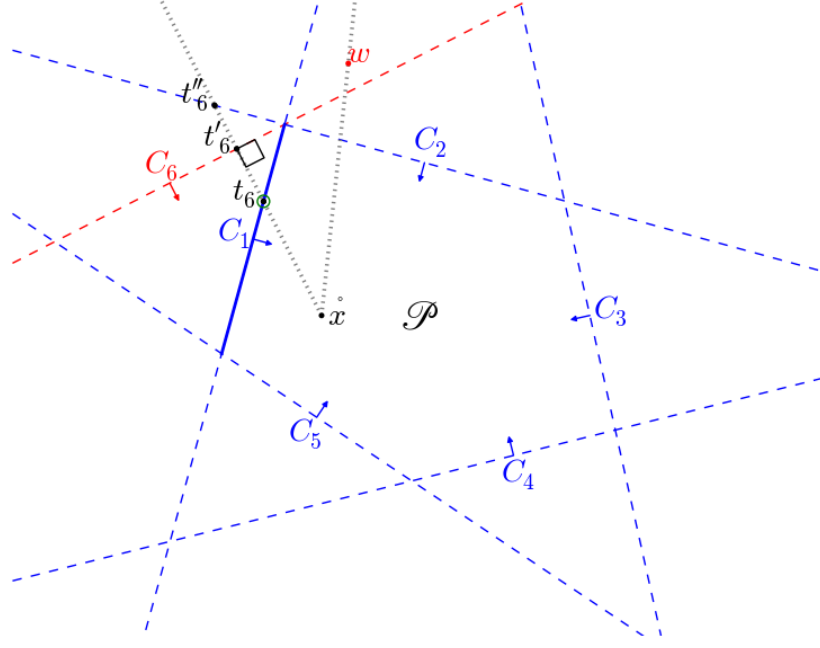


Figure 3: Detection of some frontiers of a polyhedron by looking at their intersections with rays starting from an interior point \hat{x} and orthogonal to a constraint. The thick lines are the discovered frontiers, confirmed by the doubly-circled intersection points.

Example. Here are the sorted intersection lists obtained for the 6-constraints polyhedron of Fig. 3. The list I_i records the constraints met along $\text{RAY}(\hat{x}, -C_i)$ from \hat{x} orthogonally to the hyperplane of C_i . It satisfies $t_i < t'_i < t''_i < t'''_i$.

$$\begin{aligned}
 I_1 &= [(t_1, \{C_1\}); (t'_1, \{C_5, C_6\}); (t''_1, \{C_2\})] & I_2 &= [(t_2, \{C_2\}); (t'_2, \{C_6\}); (t''_2, \{C_3\}); (t'''_2, \{C_1\})] \\
 I_3 &= [(t_3, \{C_3\}); (t'_3, \{C_2\}); (t''_3, \{C_4\})] & I_4 &= [(t_4, \{C_5\}); (t'_4, \{C_4\}); (t''_4, \{C_3\})] \\
 I_5 &= [(t_5, \{C_5\}); (t'_5, \{C_1, C_4\})] & I_6 &= [(t_6, \{C_1\}); (t'_6, \{C_6\}); (t''_6, \{C_2\})]
 \end{aligned}$$

These lists reveal that C_1 , C_2 , C_3 and C_5 are frontiers of \mathcal{P} ; C_1 and C_5 are even confirmed twice. Our criterion fails to decide the status of C_4 and C_6 because, in any of the considered directions, they are never encountered first. This situation is legitimate for the redundant constraint C_6 but also happens for C_4 even if it is a frontier of \mathcal{P} .

At this point (line 10 of Algorithm 2), we run the simplex to determine the irredundancy of the remaining constraints. In order to keep LP problems as small as possible, we build them incrementally as follows. Consider an undetermined constraint C_i and let I_i be the intersection list resulting from the direction $d_i = -C_i$. We pose a LP problem to find a point x'_i satisfying $\langle C_i, x'_i \rangle > 0 \wedge \langle C', x'_i \rangle \leq 0$, where C' is the single constraint that appears at the head of I_i . As said earlier, C' is a frontier because it is the first hyperplane encountered by the ray. We illustrate the algorithm on the case of a single head constraint as it is the most frequent one. If the head set contains several constraints we cannot know which one is a frontier, thus we add all of them in the LP problem (lines 13-14 of Algorithm 2). We distinguish two cases depending on the satisfiability of the existential LP problem: If the problem of line 15 is unsatisfiable, the simplex returns $\text{FAILURE}(\lambda)$, C_i is redundant with respect to C' and the Farkas decomposition of C_i is $\lambda \times C'$. Otherwise, the simplex exhibits a point x'_i which satisfies $\langle C_i, x'_i \rangle > 0 \wedge \langle C', x'_i \rangle \leq 0$. Here, we cannot conclude on C_i 's redundancy since x'_i is a witness showing that C_i is irredundant with respect to C' alone, but C_i could still be redundant with respect to the other constraints.

To check the irredundancy of C_i , we launch a new ray $\text{RAY}(\hat{x}, x'_i - \hat{x})$ from \hat{x} to x'_i in direction $d = x'_i - \hat{x}$. As before, we compute the intersection list of this ray with all the constraints but this time we know for sure that C_i will precede C' in the list. As this property is a pure technicality, it is given in Appendix A. Then, we analyze the head of the list: if C_i is the *single* first element, then it is a frontier.

Otherwise the first element, say C'' , is added to the LP problem, which is now asked for a point x''_i such that $\langle C_i, x''_i \rangle > 0 \wedge \langle C', x''_i \rangle \leq 0 \wedge \langle C'', x''_i \rangle \leq 0$ resulting in a new RAY($\hat{x}, x''_i - \hat{x}$). The way we choose rays guarantees that the previous constraints C', C'', \dots will always be hit after C_i by the next ray. Therefore, ultimately the constraint C_i will be hit first by a ray, or it will be proved redundant. Termination is guaranteed because the first constraint struck by the new ray is either C_i and we are done, or a not already considered constraint and there is a finite number of constraints in \mathcal{P} . Observe that this algorithm builds incremental LP problems which contain only frontiers that were between \hat{x} and the hyperplane of C_i at some step.

Example 2 (continued). In the above example, we found out that C_1, C_2, C_3 and C_5 were frontiers. To determine the status of C_4 , we solve the LP problem $\exists x'_4, \langle C_4, x'_4 \rangle > 0 \wedge \langle C_5, x'_4 \rangle \leq 0$ because C_5 is the head of I_4 . The simplex finds such a point x'_4 and the next step is to compute the intersection list corresponding to RAY($\hat{x}, x'_4 - \hat{x}$). This list will reveal C_4 as an actual frontier.

Similarly, the intersection list I_6 of the example suggests to solve the LP problem $\exists x'_6, \langle C_6, x'_6 \rangle > 0 \wedge \langle C_1, x'_6 \rangle \leq 0$ to launch a new ray toward C_6 . This problem is satisfiable and the simplex returns SUCCESS(x'_6). Then, we compute the intersection list corresponding to RAY($\hat{x}, x'_6 - \hat{x}$) and this time the head of the list is C_2 . We thus add C_2 to the previous LP problem and call the simplex on $\exists x''_6, \langle C_6, x''_6 \rangle > 0 \wedge \langle C_1, x''_6 \rangle \leq 0 \wedge \langle C_2, x''_6 \rangle \leq 0$. This problem has no solution: the simplex returns FAILURE($\lambda = (1, 1)$) showing that C_6 is redundant and its Farkas decomposition is $C_6 = 1 \times C_1 + 1 \times C_2$.

Algorithm 2: Raytracing algorithm

Input : A set of constraints $\mathcal{P} = \{C_1, \dots, C_p\}$; a point $\hat{x} \in \hat{\mathcal{P}}$
Output : \mathcal{P}_M : minimized version of \mathcal{P}
Data : $LP[i]$: LP problem associated to C_i ; $I[i]$: intersection list of C_i
Function: `intersectionList($d, \{C_1, \dots, C_q\}$)` returns the intersection list obtained by intersecting $\{C_1, \dots, C_q\}$ with ray d

```

1 Function updateFrontiers ( $I[i], \mathcal{P}_M, \mathcal{P}$ )
2   if head ( $I[i]$ ) = ( $t_F, \{F\}$ ) then
3      $\mathcal{P}_M \leftarrow \mathcal{P}_M \cup \{F\}$ 
4      $\mathcal{P} \leftarrow \mathcal{P} \setminus F$ 
5   return ( $\mathcal{P}_M, \mathcal{P}$ )

6  $\mathcal{P}_M \leftarrow \emptyset$ ;  $LP \leftarrow \text{arrayOfSize}(p)$ ;  $I \leftarrow \text{arrayOfSize}(p)$ 

7 for  $C_i$  in  $\mathcal{P}$  do /* First step of raytracing with orthogonal rays */
8    $I[i] \leftarrow \text{intersectionList}(\text{RAY}(\hat{x}, -C_i), \mathcal{P})$ 
9    $(\mathcal{P}_M, \mathcal{P}) \leftarrow \text{updateFrontiers}(I[i], \mathcal{P}_M, \mathcal{P})$ 

10 while  $\mathcal{P} \neq \emptyset$  do
11   for  $C_i$  in  $\mathcal{P}$  do
12     ( $t, S$ )  $\leftarrow$  head( $I[i]$ )
13     for  $C$  in  $S$  do
14        $LP[i] \leftarrow LP[i] \wedge \langle C, x'_i \rangle \leq 0$ 
15     switch simplex ( $\exists x'_i, \langle C_i, x'_i \rangle > 0 \wedge LP[i]$ ) do
16       case SUCCESS ( $x'_i$ ): do
17          $I[i] \leftarrow \text{intersectionList}(\text{RAY}(\hat{x}, x'_i - \hat{x}), \mathcal{P} \cup \mathcal{P}_M)$ 
18          $(\mathcal{P}_M, \mathcal{P}) \leftarrow \text{updateFrontiers}(I[i], \mathcal{P}_M, \mathcal{P})$ 
19       case FAILURE ( $\lambda$ ): do  $\mathcal{P} \leftarrow \mathcal{P} \setminus C_i$  /*  $C_i$  is redundant */
20
21 return  $\mathcal{P}_M$ 

```

4.2 Irredundancy Certificates

Let us explain how we compute witness points from the intersection lists defined in the previous section. From now on, we will denote a constraint by \mathbf{F} if it is a frontier and by \mathbf{C} when we do not know if it is a redundant constraint or an actual frontier. Let us come back to the list of the intersections of constraints of \mathcal{P} with a ray $\{\mathbf{x}(t) \mid \mathbf{x}(t) = \hat{\mathbf{x}} + t \times \mathbf{d}, t \geq 0\}$ for a direction \mathbf{d} .

Proposition 1. *If the head of an intersection list contains a single constraint \mathbf{F} , then we can build a witness point satisfying the irredundancy criterion of Theorem 1 which proves that \mathbf{F} is a frontier:*

(a) *For a list $[(t_{\mathbf{F}}, \{\mathbf{F}\})]$, we take the witness $\mathbf{w}_a = \hat{\mathbf{x}} + (t_{\mathbf{F}} + 1) \times \mathbf{d}$*

(b) *For a list $[(t_{\mathbf{F}}, \{\mathbf{F}\}); (t', S'); \dots]$ with at least two pairs, we define the witness $\mathbf{w}_b = \hat{\mathbf{x}} + \frac{t_{\mathbf{F}} + t'}{2} \times \mathbf{d}$.*

Proof. Let us prove that these witness points attest that \mathbf{F} is an irredundant constraint. According to Theorem 1, it amounts to proving that $\bigwedge_{\mathbf{C} \in \mathcal{P} \setminus \mathbf{F}} \langle \mathbf{C}, \mathbf{w} \rangle \leq 0 \wedge \langle \mathbf{F}, \mathbf{w} \rangle > 0$ for \mathbf{w}_a (resp. \mathbf{w}_b).

Let us first study the sign of $\langle \mathbf{F}, \mathbf{x}(t) \rangle$ at point $\mathbf{x}(t) = \hat{\mathbf{x}} + t \times \mathbf{d}$. Note that $\langle \mathbf{F}, \mathbf{x}(t) \rangle = \langle \mathbf{F}, \hat{\mathbf{x}} + t \times \mathbf{d} \rangle \stackrel{(\dagger)}{=} \langle \mathbf{F}, \hat{\mathbf{x}} \rangle + t \times \langle \mathbf{F}, \mathbf{d} \rangle$. By construction, $\langle \mathbf{F}, \mathbf{x}(t_{\mathbf{F}}) \rangle = 0$ then, by equation (\dagger) , $-\langle \mathbf{F}, \hat{\mathbf{x}} \rangle = t_{\mathbf{F}} \times \langle \mathbf{F}, \mathbf{d} \rangle$. Recall that $t_{\mathbf{F}} \geq 0$, $\langle \mathbf{F}, \mathbf{d} \rangle \neq 0$ since the ray hits \mathbf{F} and $\langle \mathbf{F}, \hat{\mathbf{x}} \rangle < 0$ because $\hat{\mathbf{x}} \in \mathcal{P}$. Thus, $\langle \mathbf{F}, \mathbf{d} \rangle$ and $t_{\mathbf{F}}$ are necessarily positive. Consequently, for a frontier \mathbf{F} found in a direction \mathbf{d} , $\langle \mathbf{F}, \mathbf{x}(t) \rangle = \langle \mathbf{F}, \hat{\mathbf{x}} \rangle + t \times \langle \mathbf{F}, \mathbf{d} \rangle$ is positive for any $t > t_{\mathbf{F}}$. Hence, in case (a) $\langle \mathbf{F}, \mathbf{w}_a \rangle \stackrel{\text{def}}{=} \langle \mathbf{F}, \mathbf{x}(t_{\mathbf{F}} + 1) \rangle > 0$ and in case (b) $\langle \mathbf{F}, \mathbf{w}_b \rangle \stackrel{\text{def}}{=} \langle \mathbf{F}, \mathbf{x}(\frac{t_{\mathbf{F}} + t'}{2}) \rangle > 0$ since $t_{\mathbf{F}} < \frac{t_{\mathbf{F}} + t'}{2} < t'$.

Let us now study the sign of $\langle \mathbf{C}, \mathbf{x}(t) \rangle$ for constraints other than \mathbf{F} :

(a) Consider the list $[(t_{\mathbf{F}}, \{\mathbf{F}\})]$. By construction, it means that no other constraint \mathbf{C} of \mathcal{P} is struck by the RAY $(\hat{\mathbf{x}}, \mathbf{d})$, i.e. whatever the value $t \geq 0$, the sign of $\langle \mathbf{C}, \mathbf{x}(t) \rangle = \langle \mathbf{C}, \hat{\mathbf{x}} \rangle + t \times \langle \mathbf{C}, \mathbf{d} \rangle$ does not change. As $\langle \mathbf{C}, \mathbf{x}(t=0) \rangle = \langle \mathbf{C}, \hat{\mathbf{x}} \rangle < 0$ because $\hat{\mathbf{x}} \in \mathcal{P}$, we can conclude that $\forall t \geq 0, \langle \mathbf{C}, \mathbf{x}(t) \rangle < 0$. Thus, in particular, $\langle \mathbf{C}, \mathbf{w}_a \rangle \stackrel{\text{def}}{=} \langle \mathbf{C}, \mathbf{x}(t_{\mathbf{F}} + 1) \rangle < 0$ for any $\mathbf{C} \in \mathcal{P} \setminus \mathbf{F}$.

(b) Consider now the list $[(t_{\mathbf{F}}, \{\mathbf{F}\}); (t', S'); \dots]$. A constraint \mathbf{C} that appears in the set S' vanishes at point $\mathbf{x}(t')$ with $t' > t_{\mathbf{F}} \geq 0$. The previous reasoning (\dagger) (on \mathbf{F}) based on equation $\langle \mathbf{C}, \mathbf{x}(t) \rangle = \langle \mathbf{C}, \hat{\mathbf{x}} \rangle + t \times \langle \mathbf{C}, \mathbf{d} \rangle$ is valid for \mathbf{C} , hence proving $\langle \mathbf{C}, \mathbf{d} \rangle > 0$. Thus, $\langle \mathbf{C}, \mathbf{x}(t) \rangle$ is negative for $t < t'$ (zero for $t = t'$ and positive for $t' < t$). Finally, $\langle \mathbf{C}, \mathbf{w}_b \rangle \stackrel{\text{def}}{=} \langle \mathbf{C}, \mathbf{x}(\frac{t_{\mathbf{F}} + t'}{2}) \rangle < 0$ since $\frac{t_{\mathbf{F}} + t'}{2} < t'$. The same reasoning applies to any other pair (t, S_i) in the tail of the list.

□

Fig. 4(\mathcal{P}_b) shows the irredundancy witness points w_1, w_2, w'_1, w'_2 of constraints C_1, C_2 and C' . The irredundancy of C' is confirmed three times by different rays respectively orthogonal to C' , C_3 and C_4 , leading to witnesses w'_1 (twice) and w'_2 . [Our algorithm could take advantage of this situation: if we pick up the ray orthogonal to C_3 (or C_4) before that of C' , constraint C' will be accidentally confirmed and thus the ray orthogonal to C' does not need to be considered.]←

4.3 Minimizing Generators

So far, to ease the understanding, we presented the raytracing for constraints-only polyhedra, but it works as well for generators. Indeed, we manipulated constraints as vectors and all our explanations and proofs are based on inner product. Moreover, Theorem 1 is not limited to constraints, it holds for any vector space and can be rephrased for generators. This time the irredundancy certificate for a generator \mathbf{g}' is a vector \mathbf{n} such that $\langle \mathbf{g}_1, \mathbf{n} \rangle, \dots, \langle \mathbf{g}_p, \mathbf{n} \rangle \leq 0$ and $\langle \mathbf{g}', \mathbf{n} \rangle > 0$. Such a vector defines a hyperplane orthogonal to \mathbf{n} , i.e. $\{\mathbf{x} \mid \langle \mathbf{n}, \mathbf{x} \rangle = 0\}$. It is called a *separating hyperplane* because it isolates generator \mathbf{g}' from the other ones. Fig. 4(\mathcal{P}_a) shows the separating hyperplanes defined by $\mathbf{n}_1, \mathbf{n}_2, \mathbf{n}_3$ and \mathbf{n}_4 . They respectively justify the irredundancy of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ and \mathbf{R}_1 in \mathcal{P}_a . The hyperplane defined by \mathbf{n}_4 vouches for the irredundancy of \mathbf{R}_1 because it satisfies $\langle \mathbf{n}_4, \mathbf{g} \rangle \geq 0$ for all $\mathbf{g} \in \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}'_1, \mathbf{v}'_2\}$ and $\langle \mathbf{n}_4, \mathbf{R}_1 \rangle < 0$.

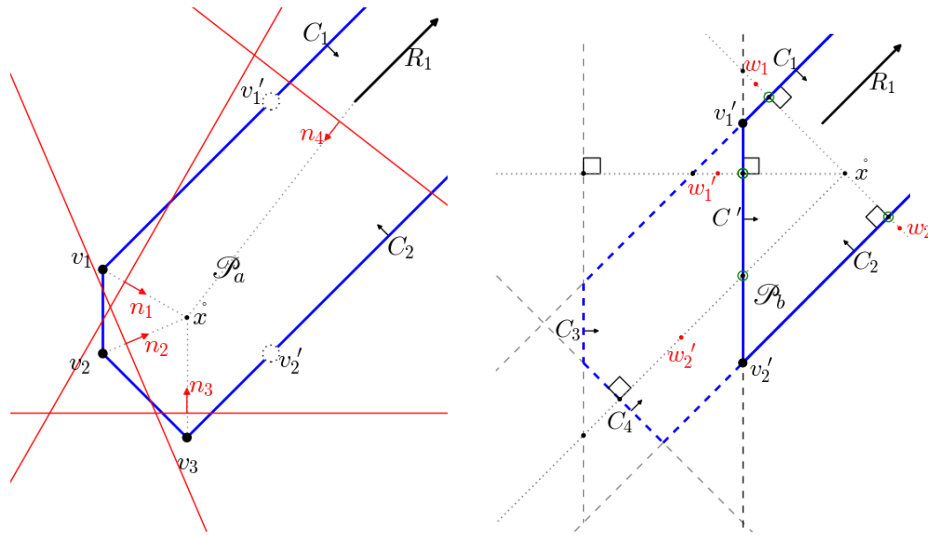


Figure 4: Irredundancy witnesses for generators of \mathcal{P}_a and constraints of \mathcal{P}_b .

4.4 Using Floating Points in Raytracing

It is possible to make raytracing even more efficient by using floating points instead of rationals. Thereby, we experimented floating points in both LP problem resolutions and distance computations. The rational coefficients of constraints are translated into floating points. It introduces a loss in precision which does not jeopardize the result because the certificate checking controls the minimization process. However, we must pay attention to the generation of *exact* (*i.e.* rational) certificates from floating point computations. The solution we propose differs depending on the kind of certificate.

Witness Points. Checking a certificate of irredundancy consists in evaluating the sign of $\langle C_i, w \rangle$ for all constraints C_i of \mathcal{P} with the provided witness point w . A witness point w must then be given with rational coefficients to avoid sign errors if $\langle C_i, w \rangle$ is too close to 0. Thus, the witness point $w^{\mathbb{F}}$ obtained with floating point computations is translated into a rational one $w^{\mathbb{Q}}$, without loss of precision (each floating point $0.d_1\dots d_m 10^e$ is changed into a rational $\frac{d_1\dots d_m \cdot 10^e}{10^m}$). Then we check the irredundancy certificate with $w^{\mathbb{Q}}$ and the rational version of the constraints. If the verification passes, then $w^{\mathbb{Q}}$ is indeed a witness point. In the rare case of failure, using the exact simplex of the VPL on the LP problem will fix the approximation error by providing a rational witness point.

Farkas Decompositions. To prove a redundancy we need to exhibit the Farkas decomposition of the redundant constraint. To obtain an exact decomposition from the floating LP solution, we record which constraint is actually part of the decomposition.⁵ Then, we run the exact simplex on a LP problem involving only those constraints to retrieve the exact Farkas decomposition.

5 Experiments

This section is devoted to the comparison of three minimization algorithms:

- **The Standard Minimization Algorithm (SMA).** The standard Algorithm 1 of §3 is available in the VPL since version 0.1. It works on rationals and can generate certificates of precision, minimality and correctness. The VPL implementation carries an optimization: the LP problem of Algorithm 1 is built only once before testing the redundancy of each constraint. Initially it is composed of all

⁵What is needed from the floating point solution is the set of basic variables and an ordering of the nonnull λ_i coefficients to speed up the search in exact simplex.

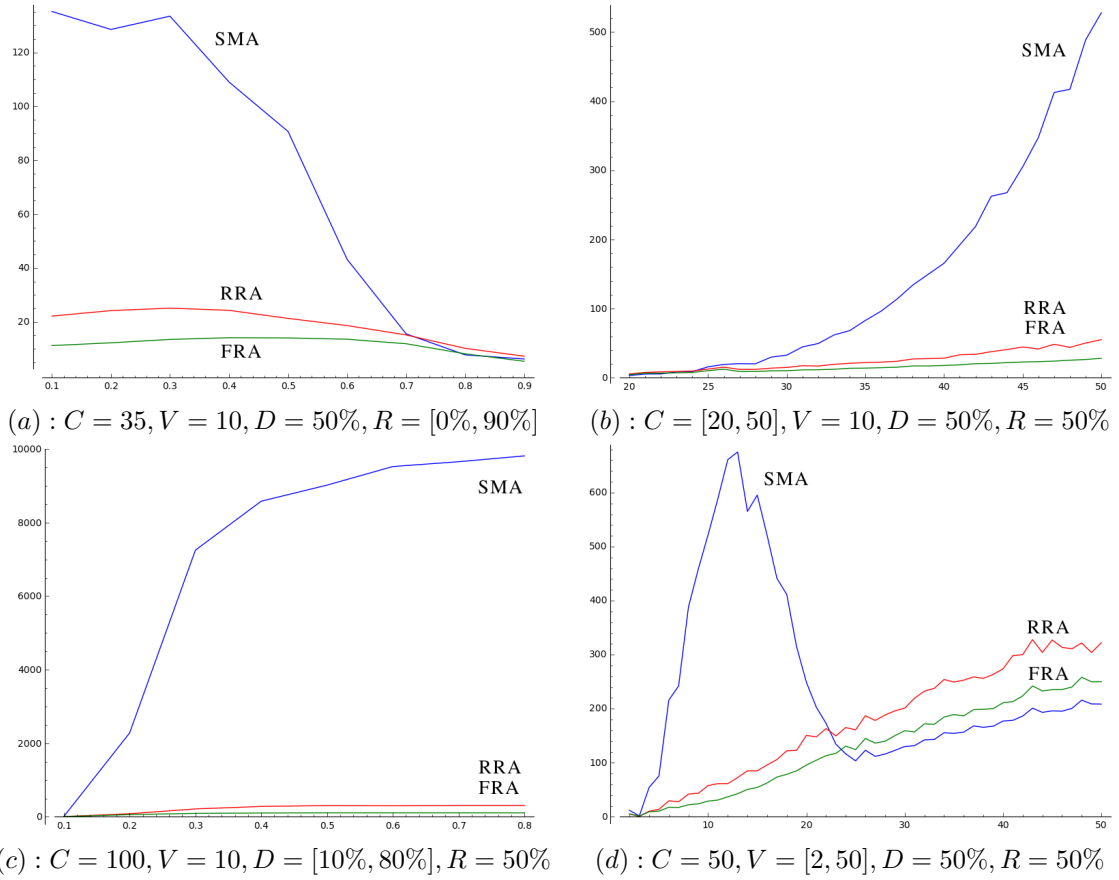


Figure 5: Execution time in milliseconds of SMA (blue), RRA (red) and FRA (green) depending on respectively (a) redundancy, (b) number of constraints, (c) density and (d) number of variables.

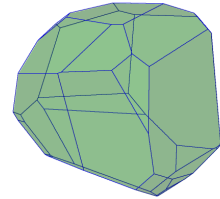
the constraints $\{C_1, \dots, C_n\}$. Then, for each $C_i \in \{C_1, \dots, C_n\}$, it checks C_i 's redundancy with respect to other constraints present in the LP problem. If C_i is shown redundant, it is removed from the LP problem and the algorithm continues with C_{i+1} . It means that the redundancy of the following constraints C_{i+i}, \dots, C_n will be tested without taking C_i into account.

- **The Rational Raytracing Algorithm (RRA).** RRA and SMA use the same LP solver, thus comparing their running time is relevant to estimate the efficiency of raytracing with respect to the standard algorithm.
- **The Floating point Raytracing Algorithm (FRA).** FRA implements raytracing with floating points as explained in §4.4. LP problems are solved by the GNU LP Kit which provides a simplex algorithm on floating points.

These three algorithms are all implemented in the current version (0.2) of the VPL.⁶ For computing the exact Farkas decomposition that proves a constraint's irredundancy, the three algorithms ultimately rely on the VPL simplex in rational. They use the same datastructures (e.g. for constraints), allowing more reliable timing comparisons between them. Moreover, they share the same pre-processing step of finding a point within the polyhedron interior. This point is obtained by solving a LP problem checking the emptiness of the polyhedron with strict inequalities. The time measurements given below include this step but not the reconstruction of exact certificates from floating point ones.

⁶<https://github.com/VERIMAG-Polyhedra>

Benchmarks. Throughout the paper, we focused on cones to simplify both notations and explanations. However, our algorithm works for general convex polyhedra and we build our experiments as follows. To compare the three algorithms, we asked them to minimize polyhedra that were generated randomly from four parameters that will be detailed further: the number of variables ($V \in [2, 50]$), the number of constraints ($C \in [2, 50]$), the redundancy rate ($R \in [0\%, 90\%]$) and the density rate ($D \in [10\%, 80\%]$). Each constraint is created by giving a random integer between -100 and 100 to the coefficient of each variable, within the density rate. All constraints are attached the same constant bound ≤ 20 . Such polyhedra have a convex potatoid shape, shown on the right hand side. We do not directly control the number of generators but we count them using the APRON interface [11] to polyhedral libraries in double description. Among all our measurements, the number of generators ranged from 10 to 6400 and this number grows polynomially in the number of constraints. This covers a wide variety of polyhedra and our experiments show that raytracing is always more efficient.



The representativeness of the polyhedra encountered in practice is a recurrent issue in experiments. It is difficult to answer as the polyhedra generated depend on the verification tool and the program under analysis. Their dimension can range from four variables, e.g. in a tuned analysis of a C program, to a thousand in analyzes of LUSTRE programs. Moreover, it is biased by the fact that verification tools pay attention to limit their use of polyhedra to a small number of variables and constraints, as polyhedra are known to be costly. They switch to interval domains if these limits are exceeded.

Redundancy Rate. The effect of redundancy on execution time is displayed on Fig. 5(a). These measures come from the minimization of polyhedra with 10 variables and 35 constraints, and a redundancy rate ranging from 0% to 90% of the number of constraints. To generate a redundant constraint, we randomly pick two constraints and produce a nonnegative combination of them. We took care of avoiding redundancies that can be discarded by the fast detection criteria of §1. The graph clearly shows that raytracing has a big advantage on polyhedra with few redundancies. This phenomenon was expected: raytracing is good at detecting irredundancy at low-cost. SMA becomes similar to raytracing when the redundancy rate is high. This is explained by the implementation details given in previous paragraphs: when a redundant constraint is found, it is removed from the LP problem. Thus, if the redundancy rate reaches a very high level, the LP problem becomes smaller and smaller at each iteration, lowering the impact of using floating points. Moreover, the heuristic used by our algorithm never hits if almost all constraints are redundant, which makes the raytracing computations useless. To be fair between raytracing and the standard algorithm, we set the redundancy rate at 50% in the other experiments.

Number of Constraints. Fig. 5(b) measures the minimization time depending on the number of constraints for polyhedra with 10 variables. FRA and RRA scale better with respect to the number of constraints than SMA: experiments show that when C ranges from 20 to 50 constraints, SMA has a quadratic evolution compared to raytracing algorithms.

Density Rate. The density of a polyhedron is the (average) rate of nonnull coefficients within a constraint. For instance, a density of 60% with 10 variables means that on average, constraints have 6 nonnull coefficients. Fig. 5(c) shows the execution time for 10-dimensional polyhedra with 100 constraints, where the density rate D goes from 10% to 80%. The raytracing algorithms are almost insensitive to density, whereas the execution time of the standard algorithm blows up with density. Actually, having a lot of nonnull coefficients in constraints tends to create huge numerators and denominators because a pivot in the simplex performs many combinations of constraints. The blow up does not happen in RRA because LP problems are much smaller in the raytracing algorithms.

Number of Variables. The effect of the dimension on execution time is shown on Fig. 5(d). Whereas raytracing seems linearly impacted by the dimension, SMA has a behaviour that may look a bit strange. After a dramatic increase of execution time, the curve falls down when the dimension reaches the number of irredundant constraints, which is half the given number of constraints as the redundancy rate equals 50%. SMA finally joins and sticks to FRA curve. This phenomenon may be explained if we take a look at the number of generators. Indeed, when $V \geq C$, the number of generators becomes close to the number of constraints, as shown on Fig. 6. Recall that the simplex algorithm travels from one vertex to another until

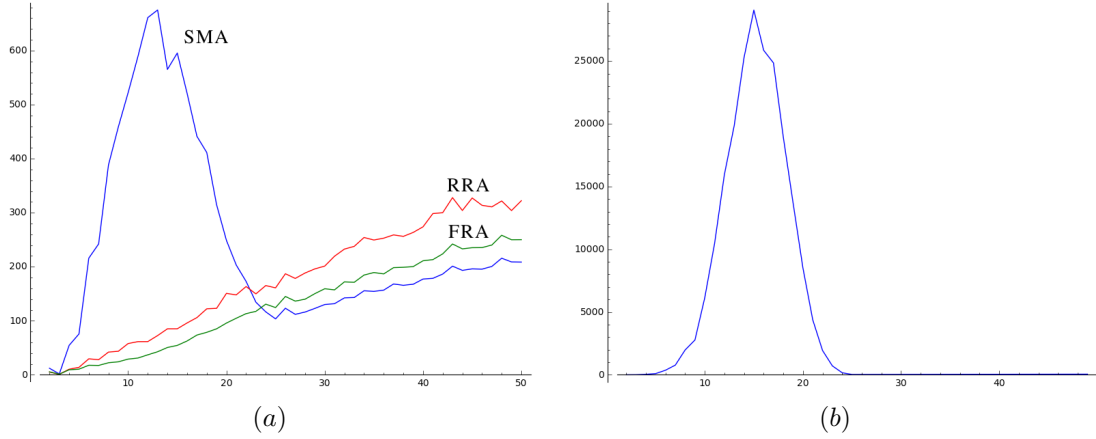


Figure 6: (a) : $C = 50, V = [2, 50], D = 50\% R = 50\%$; (b): number of generators associated to polyhedra tested in (a).

Table 1: Time measures of the three minimization algorithms SMA, RRA and FRA for different values of variables and constraints.

Var		5 constraints			10 constraints			25 constraints			50 constraints			100 constraints		
		SMA	RRA	FRA	SMA	RRA	FRA	SMA	RRA	FRA	SMA	RRA	FRA	SMA	RRA	FRA
2	# lp	5	2	2	10	5	5	25	18	18	50	42	42	100	90	90
	lp size	4	3	3	7	3	3	16	4	3	29	4	3	53	5	4
	time(ms)	0.12	0.19	0.46	0.54	0.51	0.75	3.1	1.7	2.1	12.4	5.4	6.1	65.4	23.3	25.0
	speed up	-	0.65	0.26	-	1.1	0.73	-	1.9	1.5	-	2.3	2.0	-	2.8	2.6
5	# lp	5	2	2	10	6	6	25	15	15	50	35	35	100	78	78
	lp size	4	3	3	8	3	3	20	5	4	37	6	5	75	7	6
	time(ms)	0.13	0.31	0.36	0.60	0.83	0.85	8.5	3.4	3.0	74.3	12.7	10.3	571	57.6	41.6
	speed up	-	0.42	0.36	-	0.73	0.71	-	2.5	2.8	-	5.8	7.2	-	9.9	13.7
10	# lp	5	2	2	10	5	5	25	12	12	50	27	27	100	58	58
	lp size	4	4	3	8	4	4	19	6	6	38	9	7	74	13	10
	time(ms)	0.24	0.30	0.31	0.64	1.1	1.0	11.2	8.7	6.9	639	59.7	31.4	9s	331	117
	speed up	-	0.82	0.77	-	0.59	0.64	-	1.3	1.6	-	10.7	20.4	-	28.3	80.1

the optimal value is reached. If the number of generators is low, few pivots are then needed to solve the LP problem. This makes SMA competitive even with more LP problems to solve.

Table 1 shows results for several values of dimension and number of constraints. Again, each cell of this table gives the average values resulting from the minimization of 50 convex potatoids, with a density and a redundancy both fixed at 50%. For each pair (number of variables \times number of constraints), Table 1 gives the number of LP problems that were solved and their size (*i.e.* the number of constraints they involve) on average. It contains also the computation time of the minimization in milliseconds and the speed up of raytracing compared to SMA. Results of Table 1 show that for small polyhedra, either in dimension or in number of constraints, raytracing does not help. Indeed, for such small LP problems, the overhead of our algorithm is unnecessary and leads to time losses. Raytracing becomes interesting for larger polyhedra, where the speed improvement is significant. For instance, FRA is 44.5 times faster with 10 variables and 100 constraints than SMA. The gain can be explained by the number of LP problems solved and their average size, noticeably smaller in raytracing than in SMA. As expected, raytracing is faster with floating points.

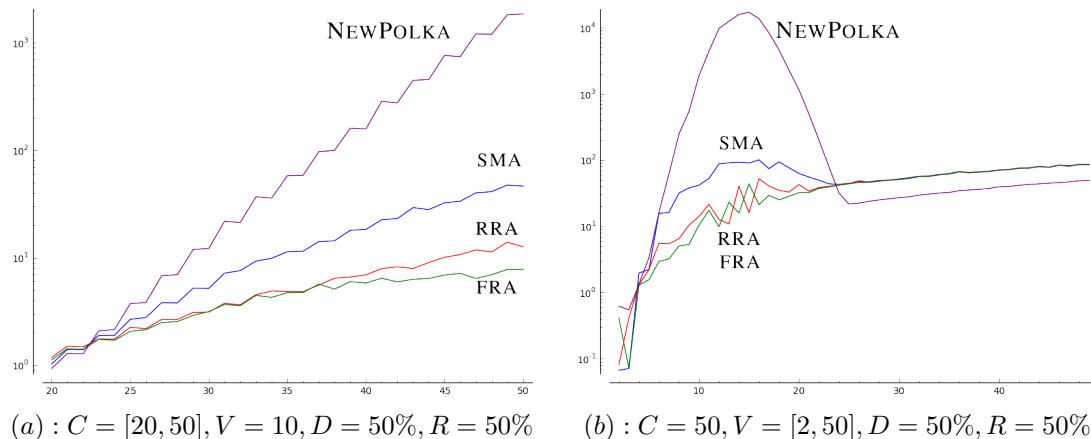


Figure 7: Execution time in milliseconds (log scale) of SMA (blue), RRA (red), FRA (green) and NEWPOLKA (purple) depending on respectively (a) the number of constraints and (b) the number of variables.

We also compare our algorithms with the NEWPOLKA, a double description library available in APRON.⁷ As an illustration, Fig. 7 adds the minimization time computed with APRON for the same tests as Fig. 5(b) and (d), with a log scale. Not surprisingly the minimization time of APRON is significantly larger than that of SMA, FRA and RRA, as it must first compute all generators using Chernikova’s algorithm [4]. The linear behaviour of APRON in log scale (Fig. 7(a)) shows the exponential complexity of this conversion from constraints to generators. The behaviour of APRON in Fig. 7(b) is the same one as SMA and is explained by the number of generators, as mentioned above.

6 Conclusion & Future Work

In this paper, we present a new algorithm to minimize the representation of a polyhedron, available in the VPL. It is based on raytracing and provides an efficient irredundancy check in which LP executions are replaced by distance computations. The *raytracing procedure* is incomplete and LP problem resolutions are still required for deciding the redundancy of the remaining constraints. However, our algorithm reduces not only the number of LP problems solved along the minimization, but also their size by an incremental approach. Moreover, it is usable for polyhedra in single representation, as constraints or as generators. It can be used either with rational or floating coefficients. In both cases, it can produce certificates of correctness, precision and minimality.

Parallelizing. Our raytracing algorithm is well-suited to parallelization: computing the intersection lists could be done by as many threads as rays. These computations boil down to matrix multiplications for which there exist efficient libraries, e.g. the LAPACK library [1]. Actually, to fully benefit from parallelism, the algorithm should be implemented in C because OCAML does not support native concurrency yet. Exploiting multi-cores, the number of ray traces could be greatly increased, and applying the raytracing principle from several interior points would allow us to discover frontiers even more easily.

Redundancy in the Double Description Framework (DDF). Our algorithm has been designed to minimize polyhedra in single representation, but the principle of raytracing can be reused in the double description framework, where it could *quickly detect irredundant constraints*. Redundancy is easier to detect when the two representations of a polyhedron are available. Let the pair $(\mathcal{C}, \mathcal{G})$ denote the set of constraints and the set of generators of a polyhedron in \mathbb{Q}^n and $(\mathcal{C}_M, \mathcal{G}_M)$ be its minimal version. A constraint $C \in \mathcal{C}$ is irredundant if it is saturated by at least n irredundant generators, i.e. $\exists g_1, \dots, g_n \in \mathcal{G}_M, \langle C, g_i \rangle = 0$. Similarly, a generator $g \in \mathcal{G}$ is irredundant if it is the intersection of at least n irredundant constraints i.e. $\exists C_1, \dots, C_n \in \mathcal{C}_M, \langle C_i, g \rangle = 0$. Think for instance of a line in 2D being defined by two points and a

⁷<http://apron.cri.enscm.fr/library/>

point being the intersection of at least two lines. The principle of the minimization algorithm is the following [10]: build the *boolean saturation matrix* \mathcal{S} of size $|\mathcal{C}| \times |\mathcal{G}|$ defined by $\mathcal{S}[\mathcal{C}][\mathcal{g}] := (\langle \mathcal{C}, \mathcal{g} \rangle = 0)$, then iteratively remove constraints (and the corresponding rows of \mathcal{S}) which are insufficiently saturated and do the same for generators (and columns of \mathcal{S}) until reaching a stable matrix. The remaining constraints and generators form the minimal version $(\mathcal{C}_M, \mathcal{G}_M)$ which mutually justify the irredundancy of each other. This algorithm is appealing compared to its counterpart in single representation but the number of evaluation of $\langle \mathcal{C}, \mathcal{g} \rangle$ is huge when each variable x_i ranges in an interval $[l_i, u_i]$. Such a product of intervals can be represented by $2n$ constraints (two inequalities $l_i \leq x_i \wedge x_i \leq u_i$ per variable) which corresponds to 2^n vertices [3].⁸ Therefore, the size of \mathcal{S} is $n2^{n+1}$. To limit the computations, the saturation matrix is not fully constructed. Let us summarize the improved algorithm [19]: (1) Some constraints are removed by the *fast redundancy detection* recalled in §1. (2) The irredundant generators of \mathcal{G}_M are constructed from the remaining constraints using Chernikova’s algorithm with some optimized adjacency criteria [14, 8, 20]. The adjacency criterion ensures that the construction cannot produce redundant generators [16]. (3) Finally, the saturation matrix is built to remove the constraint redundancies but a row is only completed if the constraint never finds enough saturating generators, otherwise the computation of the row is interrupted.

We believe that our orthogonal raytracing phase can be used at Step (3) to *quickly discover irredundant constraints*, which therefore do not have to be confirmed by the saturation matrix. The cost of this initial raytracing is reasonable: \mathcal{C} rays and $2 \times |\mathcal{C}|$ evaluations per ray resulting in $2 \times |\mathcal{C}|^2$ computations of inner products. It could therefore benefit to minimization in the DDF especially when $|\mathcal{C}| \ll |\mathcal{G}|$ as in hypercubes.

References

- [1] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users’ Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999. 6
- [2] R. Bagnara, P. M. Hill, and E. Zaffanella. Applications of polyhedral computations to the analysis and verification of hardware and software systems. *Theoretical Computer Science*, 410(46):4672–4691, 2009. 1
- [3] F. Benoy, A. King, and F. Mesnard. Computing convex hulls with a linear solver. *TPLP: Theory and Practice of Logic Programming*, 5(1-2):259–271, 2005. 6, 8
- [4] N. V. Chernikova. Algorithm for discovering the set of all the solutions of a linear programming problem. *USSR Computational Mathematics and Mathematical Physics*, 1968. 5
- [5] V. Chvatal. *Linear Programming*. Series of books in the Mathematical Sciences. W. H. Freeman, 1983. 1
- [6] P. Feautrier and C. Lengauer. Polyhedron model. In *Encyclopedia of Parallel Computing*, volume 1, pages 1581–1592. Springer, 2011. 1
- [7] A. Fouilhé, D. Monniaux, and M. Périn. Efficient certificate generation for the abstract domain of polyhedra. In *Static Analysis Symposium*, 2013. 3
- [8] K. Fukuda and A. Prodon. Double description method revisited. In *Combinatorics and Computer Science*, volume 1120 of *LNCS*, page 91–111. Springer, 1996. 1, 2
- [9] A. J. Goldman and A. W. Tucker. Polyhedral convex cones. In *Linear inequalities and related systems*, volume 38 of *Annals of Mathematics Studies*, pages 19–40. Princeton University Press, 1956. 1
- [10] N. Halbwachs. *Détermination automatique de relations linéaires vérifiées par les variables d’un programme*. PhD thesis, Université scientifique et médicale de Grenoble, 1979. (in french). 6

⁸The opposite phenomena ($2n$ vertices corresponding to 2^n constraints) also exists but hardly ever occurs in practice [3].

- [11] B. Jeannet and A. Miné. APRON: A library of numerical abstract domains for static analysis. In *Computer Aided Verification (CAV)*, volume 5643 of *LNCS*, pages 661–667, 2009. 5
- [12] J.-H. Jourdan, V. Laporte, S. Blazy, X. Leroy, and D. Pichardie. A formally-verified C static analyzer. In *ACM Principles of Programming Languages (POPL)*, pages 247–259. ACM Press, January 2015. 3
- [13] J.-L. Lassez, T. Huynh, and K. McAloon. Simplification and elimination of redundant linear arithmetic constraints. In *Constraint Logic Programming*, pages 73–87. MIT Press, 1993. 1
- [14] H. Le Verge. A note on Chernikova’s algorithm. Research Report RR-1662, INRIA, 1992. 2
- [15] A. Maréchal and M. Périn. Efficient Elimination of Redundancies in Polyhedra by Raytracing. In *Verification, Model Checking, and Abstract Interpretation (VMCAI)*, volume 10145 of *LNCS*, pages 367–385. Springer, 2017. *, ¶
- [16] T. S. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall. The double description method. In *Contributions to the theory of games*, volume 2 of *Annals of Mathematics Studies*, pages 51–73. Princeton University Press, 1953. 2
- [17] A. Schrijver. *Theory of linear and integer programming*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1999. 1
- [18] A. Simon and A. King. Exploiting Sparsity in Polyhedral Analysis. In *Static Analysis Symposium (SAS)*, volume 3672 of *LNCS*, pages 336–351. Springer-Verlag, 2005. 1
- [19] D. K. Wilde. A library for doing polyhedral operations. Master’s thesis, Oregon State University, Corvallis, Oregon, December 1993. Also published as IRISA Technical Report PI 785, Rennes, France (1993). 1, 1, 6
- [20] N. Y. Zolotykh. New modification of the double description method for constructing the skeleton of a polyhedral cone. *Computational Mathematics and Mathematical Physics*, 52(1):146–156, 2012. 2

A Proofs

Proposition 2. *Let C_i, C' be two constraints of a polyhedron \mathcal{P} , and let $\hat{x} \in \mathring{\mathcal{P}}$. Let x'_i be a point such that $\langle C_i, x'_i \rangle > 0$ and $\langle C', x'_i \rangle \leq 0$. Then $\text{RAY}(\hat{x}, x'_i - \hat{x})$ intersects C_i at some point $x(t_{C_i})$. Moreover, assume it crosses C' at $x(t_{C'})$, then $t_{C_i} < t_{C'}$.*

Proof. Because $\langle C_i, \hat{x} \rangle < 0$ and $\langle C_i, x'_i \rangle > 0$ then $\text{RAY}(\hat{x}, x'_i - \hat{x})$ necessarily crosses C_i , say at $x(t_{C_i})$. Then, $\langle C_i, x(t_{C_i}) \rangle = 0 = \underbrace{\langle C_i, \hat{x} \rangle}_{<0} + \underbrace{t_{C_i}}_{>0} \times \langle C_i, x'_i - \hat{x} \rangle$. Thus, $\langle C_i, x'_i - \hat{x} \rangle > 0$ (‡).

We supposed that $\text{RAY}(\hat{x}, x'_i - \hat{x})$ intersects C' at $x(t_{C'})$ then, similarly,
 $\langle C', x(t_{C'}) \rangle = 0 = \underbrace{\langle C', \hat{x} \rangle}_{<0} + \underbrace{t_{C'}}_{>0} \times \langle C', x'_i - \hat{x} \rangle$. Thus, $\langle C', x'_i - \hat{x} \rangle > 0$ (‡').

The computation of intersection with a $\text{RAY}(\hat{x}, x'_i - \hat{x})$ is explained in §4.1. It gives

$$t_{C'} \stackrel{\text{def}}{=} -\frac{\langle C', \hat{x} \rangle}{\langle C', x'_i - \hat{x} \rangle} \text{ and } t_{C_i} \stackrel{\text{def}}{=} -\frac{\langle C_i, \hat{x} \rangle}{\langle C_i, x'_i - \hat{x} \rangle}.$$

Let us now show that $0 < t_{C'} - t_{C_i}$.

$$\begin{aligned} t_{C'} - t_{C_i} &= -\frac{\langle C', \hat{x} \rangle}{\langle C', x'_i - \hat{x} \rangle} + \frac{\langle C_i, \hat{x} \rangle}{\langle C_i, x'_i - \hat{x} \rangle} = \frac{-\langle C', \hat{x} \rangle \langle C_i, x'_i - \hat{x} \rangle + \langle C_i, \hat{x} \rangle \langle C', x'_i - \hat{x} \rangle}{\langle C', x'_i - \hat{x} \rangle \langle C_i, x'_i - \hat{x} \rangle} \\ &= \frac{-\langle C', \hat{x} \rangle \langle C_i, x'_i \rangle + \langle C_i, \hat{x} \rangle \langle C', x'_i \rangle}{\langle C', x'_i - \hat{x} \rangle \langle C_i, x'_i - \hat{x} \rangle} \\ &= \frac{-\langle C', \hat{x} \rangle \langle C_i, x'_i \rangle + \langle C_i, \hat{x} \rangle \langle C', x'_i \rangle}{\langle C', x'_i - \hat{x} \rangle \langle C_i, x'_i - \hat{x} \rangle} \end{aligned}$$

Let us now study the sign of each term:

1. The denominator $\langle \mathbf{C}', \mathbf{x}'_i - \hat{\mathbf{x}} \rangle \langle \mathbf{C}_i, \mathbf{x}'_i - \hat{\mathbf{x}} \rangle$ is positive, as a consequence of (\dagger') and (\ddagger) .
2. By definition of the interior point $\hat{\mathbf{x}}$, $\langle \mathbf{C}_i, \hat{\mathbf{x}} \rangle < 0$ and $\langle \mathbf{C}', \hat{\mathbf{x}} \rangle < 0$. By construction of \mathbf{x}'_i , $\langle \mathbf{C}_i, \mathbf{x}'_i \rangle > 0$ and $\langle \mathbf{C}', \mathbf{x}'_i \rangle \leq 0$. Thereby, the numerator is positive as $-\underbrace{\langle \mathbf{C}', \hat{\mathbf{x}} \rangle}_{<0} \underbrace{\langle \mathbf{C}_i, \mathbf{x}'_i \rangle}_{>0} > 0$
and $\underbrace{\langle \mathbf{C}_i, \hat{\mathbf{x}} \rangle}_{<0} \underbrace{\langle \mathbf{C}', \mathbf{x}'_i \rangle}_{\leq 0} \geq 0$.

Finally, we have proved that the fraction is positive, hence $t_{\mathbf{C}_i} < t_{\mathbf{C}'}$.

□