

Essential Convexity and Complexity of Semi-Algebraic Constraints

Manuel Bodirsky¹, Peter Jonsson², and Timo von Oertzen³

¹ CNRS/LIX, École Polytechnique, 91128 Palaiseau, France.

E-mail: bodirsky@lix.polytechnique.fr

² Department of Computer and System Science, Linköpings Universitet

SE-581 83 Linköping, Sweden. Email: petej@ida.liu.se

³ Max-Planck-Institute for Human Development, Königin-Luise-Strasse 5, 14195 Berlin, Germany,
and University of Virginia, Department of Psychology, Charlottesville, USA.

Email: vonoertzen@mpib-berlin.mpg.de

Abstract. Let Γ be a structure with a finite relational signature and a first-order definition in $(\mathbb{R}; *, +)$ with parameters from \mathbb{R} , that is, a relational structure over the real numbers where all relations are semi-algebraic sets. In this article, we study the computational complexity of constraint satisfaction problem (CSP) for Γ : the problem to decide whether a given primitive positive sentence is true in Γ . We focus on those structures Γ that contain the relations \leq , $\{(x, y, z) \mid x + y = z\}$ and $\{1\}$. Hence, all CSPs studied in this article are at least as expressive as the feasibility problem for linear programs. The central concept in our investigation is *essential convexity*: a relation S is essentially convex if for all $a, b \in S$, there are only finitely many points on the line segment between a and b that are not in S . If Γ contains a relation S that is not essentially convex and this is witnessed by *rational* points a, b , then we show that the CSP for Γ is NP-hard. Furthermore, we characterize essentially convex relations in logical terms. This different view may open up new ways for identifying tractable classes of semi-algebraic CSPs. For instance, we show that if Γ is a first-order expansion of $(\mathbb{R}; +, 1, \leq)$, then the CSP for Γ can be solved in polynomial time if and only if all relations in Γ are essentially convex (unless $P=NP$).

1 Introduction

Linear Programming is a computational problem of outstanding theoretical and practical importance. It is known to be computationally equivalent to the problem to decide whether a given set of linear (non-strict) inequalities is *feasible*, i.e., defines a non-empty set:

Linear Program Feasibility

INPUT: A finite set of variables V ; a finite set of linear inequalities of the form $a_1x_1 + \dots + a_kx_k \leq a_0$ where $x_1, \dots, x_k \in V$ and a_0, \dots, a_k are rational numbers where the numerators and denominators are represented in binary.

QUESTION: Does there exist an $x \in \mathbb{R}^{|V|}$ that satisfies all inequalities?

This problem can be viewed as a *constraint satisfaction problem*, where the allowed constraints are linear inequalities with rational coefficients, and the question is whether there is an assignment of real values to the variables such that all the constraints are satisfied. For formal definitions of concepts related to constraint satisfaction, we refer the reader to Section 2.1. It is obvious that

this problem cannot be formulated with a finite constraint language; however, we will later on (Proposition 12) see that the feasibility problem for linear programs is polynomial-time equivalent to the constraint satisfaction problem for the structure

$$\Gamma_{\text{lin}} := (\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \leq, \{1\}) .$$

It is well-known that linear programming can be solved in polynomial time; moreover, several algorithms are known that are efficient also in practice. In this article, we study how far Γ_{lin} can be expanded such that the corresponding constraint satisfaction problem remains polynomial-time solvable. An important class of relations that generalizes the class of relations defined by linear inequalities is the class of all *semi-algebraic* relations, i.e., relations that have a first-order definition over $(\mathbb{R}; *, +)$ using parameters from \mathbb{R} . By the fundamental theorem of Tarski and Seidenberg, it is known that a relation $S \subseteq \mathbb{R}^n$ is semi-algebraic if and only if it has a quantifier-free first-order definition in $(\mathbb{R}; *, +, \leq)$ using parameters from \mathbb{R} . Geometrically, we can view semi-algebraic sets as finite unions of finite intersections of the solution sets of strict and non-strict polynomial inequalities.

We propose a framework for systematically studying the computational complexity of expansions of Γ_{lin} by semi-algebraic relations. In this framework, a constraint satisfaction problem is given by a (fixed and finite) constraint language Γ . All the constraints in the input of such a feasibility problem must be chosen from this constraint language Γ (a formal definition can be found in Section 2.1). This way of parameterizing constraint satisfaction problems by their constraint language has proved to be very fruitful for finite domain constraint satisfaction problems [1, 7–9, 14]. Since the constraint language is finite, the computational complexity of such a problem does not depend on how the constraints are represented in the input. We believe that the very same approach is very promising for studying the complexity of problems in real algebraic geometry. In Section 6 we will discuss a connection between some of the CSPs with semi-algebraic constraint languages and open problems in convex geometry and semidefinite programming.

One of the key reasons why linear program feasibility can be decided in polynomial time is that the feasible regions of linear inequalities are *convex*. Convexity is not a necessary condition for tractability of semi-algebraic constraint satisfaction problems, though. It is, for instance, well-known that linear program feasibility can also be decided in polynomial time when some of the input constraints are *disequalities*, i.e., constraints of the form $a_1x_1 + \dots + a_kx_k \neq a_0$ for rational values a_0, \dots, a_k . However, we show that if $\Gamma_{\text{lin}} \subseteq \Gamma$ and Γ contains a relation S with rational $a, b \in S$ such that on the line segment L between a and b there are infinitely many points that are not in S , then the CSP(Γ) is NP-hard. This motivates the notion of *essential convexity*: a set $S \subseteq \mathbb{R}^k$ is *essentially convex* if for all $p, q \in S$ there are only finitely many points on the line between p and q that are not in S . One of our central results is a logical characterization of essentially convex semi-algebraic relations in Section 4. This characterization can be used to show several results that are briefly described next.

A relation is called *semi-linear* if it has a first-order definition with rational parameters⁴ in the structure $(\mathbb{R}; +, \leq)$. From the perspective of constraint satisfaction, the set of semi-linear relations is a rich set. For example, every relation $S \subseteq \mathbb{Q}^k$ with finitely many elements is semi-linear; thus, every finitary relation on a finite set can be viewed as a semi-linear relation. In Section 5.1, we show that when we add a finite number of semi-linear relations to Γ_{lin} , then the resulting language

⁴ We deviate from model-theoretic terminology as it is used e.g. in [24] in that we only allow rational and not arbitrary real parameters in first-order definitions. Our definition coincides with the definition of semi-linear sets given in e.g. [12, 13].

either has a polynomial-time or an NP-hard constraint satisfaction problem. This result is useful for studying optimization problems: note that linear programming can be viewed as optimizing a linear function over the feasible points of a set of linear inequalities. This view suggests an immediate generalization: optimize a linear function over the feasible points of an instance of a constraint satisfaction problem for semi-linear constraint languages. We completely classify the complexity of this problem in Section 5.2.

Another application concerns temporal reasoning. A *temporal constraint language* Γ is a structure $(\mathbb{R}; R_1, \dots, R_l)$ with a first-order definition in $(\mathbb{R}; <)$. Many computational problems in artificial intelligence and scheduling can be modeled as constraint satisfaction problems for temporal constraint languages. The complexity of the CSP for temporal constraint languages Γ has been completely classified recently [6]; there are 9 tractable classes of temporal constraint satisfaction problems. Often, temporal languages are extended with some mechanism for expressing *metric* time, i.e., the ability to assign numerical values to variables and performing some kind of arithmetic calculations [11]. It has been observed that many metric languages Γ are semi-linear and satisfy $\Gamma_{\text{lin}} \subseteq \Gamma$, and if such a language is polynomial-time solvable, then it is a subclass of the so-called Horn-DLR class [21]. Our result shows that this is not a coincidence: whenever Γ is not a subclass of Horn-DLR, then the $\text{CSP}(\Gamma)$ is NP-hard.

2 Preliminaries

2.1 Constraint Satisfaction Problems

A first-order formula⁵ is called *primitive positive* (pp) if it is of the form

$$\exists x_1, \dots, x_n. (\psi_1 \wedge \dots \wedge \psi_m)$$

where ψ_i are atomic formulas, i.e., formulas of the form $x = y$ or $S(x_{i_1}, \dots, x_{i_k})$ where S is the relation symbol for a k -ary relation in Γ . We call such a formula a *pp-formula*, and as usual a pp-formula without free variables is called a pp-sentence.

Let $\Gamma = (D; S_1, \dots, S_l)$ be a structure with domain D and a finite relational signature. The *constraint satisfaction problem for Γ* ($\text{CSP}(\Gamma)$ in short) is the computational problem to decide whether a given primitive positive sentence Φ involving relation symbols for the relations in Γ is true in Γ . The conjuncts in a pp-sentence Φ are also called the *constraints* of Φ , and to emphasize the connection between the structure Γ and the constraint satisfaction problem, we typically refer to Γ as a *constraint language*. By choosing an appropriate constraint language Γ , many computational problems that have been studied in the literature can be formulated as $\text{CSP}(\Gamma)$ (see e.g. [5, 9]).

When studying the complexity of different CSPs, it is often useful to be able to derive new relations from old. If $\Gamma = (D; S_1, \dots, S_l)$ is a relational structure and $S \subseteq D^k$ is a relation, then (Γ, S) denotes the expansion $(D; S, S_1, \dots, S_l)$ of the structure Γ by the relation S . We say that an n -ary relation S is *pp-definable* in Γ if there exists a pp-formula ϕ with free variables x_1, \dots, x_n such that $(x_1, \dots, x_n) \in S$ iff $\phi(x_1, \dots, x_n)$ holds in Γ . The following simple but important result explains the importance of pp-definability for constraint satisfaction problems.

Lemma 1 (Jeavons et al. [20]). *Let Γ be a relational structure, and let S be pp-definable over Γ . Then $\text{CSP}(\Gamma, S)$ is polynomial-time equivalent to $\text{CSP}(\Gamma)$.*

⁵ Our terminology is standard; all notions that are not explicitly introduced can be found in standard textbooks, e.g., in [19].

2.2 Semi-algebraic and semi-linear relations

We say that a relation $S \subseteq D^n$ is *first-order definable* in a structure Γ with domain D if there exists a formula $\phi(x_1, \dots, x_n)$ using universal and existential quantification, disjunction, conjunction, negation, and atomic formulas over Γ (where x_1, \dots, x_n denote the free variables in ϕ) such that $\phi(a_1, \dots, a_n)$ is true over Γ if and only if $(a_1, \dots, a_n) \in S$. We always admit equality when building atomic formulas, i.e., we have atomic formulas of the form $t_1 = t_2$ for terms t_1, t_2 formed from function symbols for Γ and variables. We say that S is *first-order definable in Γ with parameters from A* , for $A \subseteq D$, if additionally we can use constant symbols for the elements of A in the first-order definition of S .

A set $S \subseteq \mathbb{R}^n$ is called *semi-algebraic* if it has a first-order definition in $(\mathbb{R}; *, +)$ using parameters from \mathbb{R} . Note that the order \leq of the real numbers is first-order definable in $(\mathbb{R}; *, +)$, since

$$a \leq b \Leftrightarrow \exists c. b = a + c * c.$$

We need some basic algebraic and topological concepts and facts.

Definition 2 (Section 3.1 in [2]). A set $S \subseteq \mathbb{R}^n$ is open if it is the union of open balls, i.e., if every point of S is contained in an open ball contained in S . A set $S \subseteq \mathbb{R}^n$ is closed if its complement is open. The closure of a set S , denoted \bar{S} , is the intersection of all closed sets containing S . Equivalently, $\bar{S} = \{x \in \mathbb{R}^n \mid \forall r > 0 \exists y \in S. (y - x)^2 < r^2\}$. A point p in S is a boundary point if for every $\epsilon > 0$, the n -dimensional open ball with radius ϵ around p contains at least one point in S and one point not in S . The set of boundary points is denoted ∂S . The interior of S , denoted by S° , is $S \setminus \partial S$.

Note that the interior of S consists of all $p \in S$ such that there exists an $\epsilon > 0$ with the following property: the n -dimensional open ball with radius ϵ around p is contained in S . Also note that a finite union of closed sets is closed.

Proposition 3 (Proposition 2.2.2. in [4]). The closure of a semi-algebraic relation is semi-algebraic.

We use the notion of *dimension* $\dim(S) \in \mathbb{N}$ of a semi-algebraic set S as defined in [4].

Definition 4 (Section 2.8 in [4]). Let $S \subseteq \mathbb{R}^k$ be a semi-algebraic set, and let $\mathcal{P}(S)$ be the ring of polynomial functions on S , i.e., the ring of functions $S \rightarrow \mathbb{R}$ which are the restriction of a polynomial. Then the dimension of S , denoted by $\dim(S)$, is the maximal length of chains of prime ideals of $\mathcal{P}(S)$, i.e., the maximal d such that there exist distinct prime ideals I_0, I_1, \dots, I_d of $\mathcal{P}(S)$ with $I_0 \subset I_1 \subset \dots \subset I_d$.

To work with this definition of dimension, we need some more concepts.

Definition 5 (see [4]). Let $S \subseteq \mathbb{R}^k$ and $T \subseteq \mathbb{R}^l$ be semi-algebraic sets. A function $f : S \rightarrow T$ is semi-algebraic if the set $\{(x_1, \dots, x_k, y_1, \dots, y_l) \mid f(x_1, \dots, x_k) = (y_1, \dots, y_l)\}$ is a semi-algebraic subset of \mathbb{R}^{k+l} .

As usual, bijective functions $f : S \rightarrow T$ such that $S' \subseteq S$ is open if and only if $f(S') \subseteq T$ is open are called *homeomorphisms*.

Lemma 6 (Propositions 2.8.5, 2.8.9, and 2.8.13 in [4]). Let $S \subseteq \mathbb{R}^n$ be semi-algebraic.

- If $S = S_1 \cup S_2$ then $\dim(S) = \max(\dim(S_1), \dim(S_2))$.
- If there is a semi-algebraic homeomorphism from S to $(0, 1)^d$, then $\dim(S) = d$.
- $\dim(\bar{S} \setminus S) < \dim(S)$.

In particular, if $S \subseteq T$, then $\dim(S) \leq \dim(T)$.

A set $V \subseteq \mathbb{R}^n$ is called an *(algebraic) variety* if it can be defined as a conjunction of the form $p_1 = 0 \wedge \dots \wedge p_m = 0$ where p_1, \dots, p_m are polynomials in the variables x_1, \dots, x_n with coefficients from \mathbb{R} . We allow terms in polynomials to have degree zero.

Lemma 7. *Let $V \subseteq \mathbb{R}^n$ be a variety and let $L \subseteq \mathbb{R}^n$ be a line. If infinitely many points of L are in V , then $L \subseteq V$.*

Proof. Let V be defined by $p_1(x_1, \dots, x_n) = 0 \wedge \dots \wedge p_m(x_1, \dots, x_n) = 0$, and let l_1, \dots, l_n be univariate linear polynomials such that $L = \{(l_1(x), \dots, l_n(x)) \mid x \in \mathbb{R}\}$. For each p_i , the univariate polynomial $p_i(l_1(x), \dots, l_n(x))$ equals 0 infinitely often. So it is always 0, and it follows that every point on L satisfies $p_1 = 0 \wedge \dots \wedge p_m = 0$. \square

Theorem 8 (Tarski and Seidenberg; Proposition 5.2.2 in [4]). *Every first-order formula over $(\mathbb{R}; *, +, \leq)$ with parameters from \mathbb{R} is equivalent to a quantifier-free formula with parameters from \mathbb{R} .*

By an *interval* we mean either an open, half-open, or closed interval with more than one element. An ordered structure $(D; \leq, \dots)$ is *o-minimal* (see [23], Definition 3.1.18) if for any first-order definable $S \subseteq D$ with parameters from D there are finitely many intervals I_1, \dots, I_m with endpoints in $D \cup \{\pm\infty\}$ and a finite set $D_0 \subseteq D$ such that $S = D_0 \cup I_1 \cup \dots \cup I_m$. The following is an easy and well-known consequence of Theorem 8.

Theorem 9 (see e.g. [23]). *Let R_1, \dots, R_n be semi-algebraic relations. Then $(\mathbb{R}; \leq, R_1, \dots, R_n)$ is o-minimal.*

A set $S \subseteq \mathbb{R}^n$ is called *semi-linear* if it has a first-order definition in $(\mathbb{R}; +, \leq)$ with parameters from \mathbb{Q} ; we also call first-order formulas over $(\mathbb{R}; +, \leq)$ with parameters from \mathbb{Q} *semi-linear*. It has been shown in [12, 13] that it is decidable whether a given first-order formula over $(\mathbb{R}; *, +, \leq)$ with parameters from \mathbb{Q} defines a semi-linear relation or not. A set $V \subseteq \mathbb{R}^n$ is called a *linear set* if it can be defined as a conjunction of the form $p_1 \geq 0 \wedge \dots \wedge p_m \geq 0$ where p_1, \dots, p_m are linear polynomials in the variables x_1, \dots, x_n with coefficients from \mathbb{Q} . It is not hard to see that every semi-linear relation S can be viewed as a finite union of linear sets. We also have quantifier elimination for semi-linear relations.

Theorem 10 (Ferrante and Rackoff [15]). *Every semi-linear relation has a quantifier-free definition over $(\mathbb{R}; +, -, \leq)$ with parameters from \mathbb{Q} .*

2.3 Definability of Rational Expressions

The following elementary lemma will be needed for the observation that the feasibility problem for linear programs is polynomial-time equivalent to $\text{CSP}(T_{\text{lin}})$; it is also essential for the hardness proofs in Section 3 and for proving the dichotomy result for metric temporal constraint reasoning.

Lemma 11. *Let $n_0, \dots, n_l \in \mathbb{Q}$ be arbitrary rational numbers. Then the relation $\{(x_1, \dots, x_l) \mid n_1x_1 + \dots + n_lx_l = n_0\}$ is pp-definable in $(\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\})$. Furthermore, the pp-formula that defines the relation can be computed in polynomial time.*

Proof. We first note that we can assume that n_0, \dots, n_l are integers. To see this, suppose that the rational coefficients n_0, \dots, n_l are represented as pairs of integers $(a_0, b_0), \dots, (a_l, b_l)$ where a_i denotes the nominator and b_i the denominator. Let $c = \prod_{i=0}^l b_i$ and create a new sequence of coefficients $n'_0, \dots, n'_l = (a_0 \cdot c / b_0, 1), \dots, (a_l \cdot c / b_l, 1)$. The resulting equation is obviously equivalent. It is also clear that it only takes polynomial time to compute these coefficients.

Before the actual proof, we note that $x = 0$ is pp-definable by $x + x = x$, and we therefore freely use the terms 0 and 1 in pp-definitions. Similarly, $x = -1$ is pp-definable by $x + 1 = 0$. The proof is by induction on l . We first show how to express equations of the form $n_1x_1 + n_2x_2 = x_3$. By setting x_2 to -1 and x_3 to 0, this will solve the case $l = 1$. For positive n_1, n_2 , the formula $n_1x_1 + n_2x_2 = x_3$ is equivalent to

$$\begin{aligned} \exists u_1, \dots, u_{n_1}, v_1, \dots, v_{n_2}. \quad & u_1 = x_1 \wedge \bigwedge_{i=1}^{n_1-1} x_1 + u_i = u_{i+1} \\ & \wedge v_1 = x_2 \wedge \bigwedge_{i=1}^{n_2-1} x_2 + v_i = v_{i+1} \\ & \wedge u_{n_1} + v_{n_2} = x_3. \end{aligned}$$

However, this formula is exponential in the representation size of n_1 and n_2 , and cannot be used in polynomial-time reductions.

Let $\text{bit}(n, i)$ denote the i -th lowest bit in the binary representation of an integer n and $1 \leq i \leq \lfloor \log n \rfloor + 1$. The following formula is equivalent to the previous one (we are still in the case that both n_1 and n_2 are positive) and has polynomial length in the representation size of n_1 and n_2 .

$$\begin{aligned} \exists \bar{a}, \bar{b}, \bar{c}, \bar{d}. \quad & a_1 = x_1 \wedge \bigwedge_{i=1}^{\lfloor \log n_1 \rfloor} a_i + a_i = a_{i+1} \wedge \\ & b_1 = x_2 \wedge \bigwedge_{i=1}^{\lfloor \log n_2 \rfloor} b_i + b_i = b_{i+1} \wedge \\ & c_1 = \text{bit}(n_1, 1) \wedge \bigwedge_{i=1}^{\lfloor \log n_1 \rfloor} \text{bit}(n_1, i+1)a_i + c_i = c_{i+1} \wedge \\ & d_1 = \text{bit}(n_2, 1) \wedge \bigwedge_{i=1}^{\lfloor \log n_2 \rfloor} \text{bit}(n_2, i+1)b_i + d_i = d_{i+1} \wedge \\ & c_{\lfloor \log n_1 \rfloor + 1} + d_{\lfloor \log n_2 \rfloor + 1} = x_3 \end{aligned}$$

If $l = 2$, and $n_1 = 0$ or $n_2 = 0$, then the proof is similar. If n_1 and n_2 have different signs, we replace the conjunct $c_{\lfloor \log n_1 \rfloor + 1} + d_{\lfloor \log n_2 \rfloor + 1} = x_3$ in the formula above appropriately by $c_{\lfloor \log n_1 \rfloor + 1} + x_3 = d_{\lfloor \log n_2 \rfloor + 1}$ or $d_{\lfloor \log n_2 \rfloor + 1} + x_3 = c_{\lfloor \log n_1 \rfloor + 1}$. If both n_1 and n_2 are negative, then we use the pp-definition $\exists x'_3. (-n_1x_1 - n_2x_2 = x'_3 \wedge x'_3 + x_3 = 0)$.

Equalities of the form $n_1x_1 + n_2x_2 = n_0$ can be defined by $\exists x_3.(n_1x_1 + n_2x_2 = x_3 \wedge x_3 = n_0)$. Now suppose that $l > 2$. By the inductive assumption, there is a pp-definition $\phi_1(x_1, x_2, u)$ for $n_1x_1 + n_2x_2 + u = n_0$ and a pp-definition $\phi_2(x_3, \dots, x_l, u)$ for $n_3x_3 + \dots + n_lx_l = u$. Then $\exists u.(\phi_1 \wedge \phi_2)$ is a pp-definition for $n_1x_1 + \dots + n_lx_l = n_0$.

We next prove that the pp-definition can be computed in time $O(p(m))$ time where p is a fixed polynomial and m denotes the number of bits needed to represent the input. Let $T(l, m)$ denote the maximum time needed to compute the definition for the equations containing l variables and where m bits are needed for representing the input. We assume without loss of generality that $l \leq m$. If $1 \leq l \leq 3$, then $T(l, m) \leq p'(m)$ for some polynomial p' since we can compute $\text{bit}(\cdot, \cdot)$ in polynomial time. If $l > 3$, then $T(l, m) \leq T(3, m) + T(l-2, m) + p''(m)$ for some polynomial p'' by the inductive construction of the definition. Assume for simplicity that l is even; the case when l is odd can be solved analogously. By unfolding the recursive definition of T , we see that

$$\begin{aligned} T(l, m) &\leq \frac{l}{2}T(3, m) + p'(m) + \sum_{j=1}^{l/2} p''(m) \leq l \cdot p'(m) + p'(m) + l \cdot p''(m) \\ &\leq m \cdot p'(m) + p'(m) + m \cdot p''(m) \end{aligned}$$

which is a polynomial in m . □

By extending the previous result to inequalities, we prove that $\text{CSP}(\Gamma_{\text{lin}})$ and linear program feasibility are polynomial-time equivalent problems. The dichotomy for metric temporal reasoning follows immediately by combining this result and Theorem 29.

Proposition 12. *The linear program feasibility problem is polynomial-time equivalent to $\text{CSP}(\Gamma_{\text{lin}})$.*

Proof. It is clear that an instance of $\text{CSP}(\Gamma_{\text{lin}})$ can be seen as a linear program feasibility problem, since the three different relations in the constraint language, $x + y = z$, $x = 1$, $x \leq y$, are linear.

For the opposite direction, let Φ be an arbitrary instance of the linear program feasibility problem. Given a linear equality $L(x_1, \dots, x_k) \equiv c_1x_1 + \dots + c_kx_k = c_0$, let $\phi_{L(x_1, \dots, x_k)}$ denote the pp-definition of $L(x_1, \dots, x_k)$ in $(\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\})$ obtained in Lemma 11. Construct an instance Ψ of $\text{CSP}(\Gamma_{\text{lin}})$ by replacing each occurrence of a linear inequality constraint $c_1x_1 + \dots + c_lx_l \leq c_0$ in Φ by a $\phi_{c_1x_1 + \dots + c_lx_l - y = 0} \wedge y \leq c_0$; use fresh variables for y and for the existentially quantified variables introduced by ϕ_L . The resulting formula Ψ can be rewritten as a primitive positive sentence over Γ without increasing its length and, by Lemma 11, the length of Ψ is polynomial in the length of Φ . Since Φ is satisfiable if and only if Ψ is satisfiable, this shows that the problems are polynomial-time equivalent. □

3 Hardness

We consider relations that give rise to NP-hard CSPs in this section. We first need some definitions: a relation $S \subseteq \mathbb{R}^k$ is *convex* if for all $p, q \in S$, S contains all points on the line segment between p and q . We say that a relation $S \subseteq \mathbb{R}^k$ *excludes an interval* if there are $p, q \in S$ and real numbers $0 < \delta_1 < \delta_2 < 1$ such that $p + (q - p)y \notin S$ whenever $\delta_1 \leq y \leq \delta_2$. Note that we can assume that δ_1, δ_2 are rational numbers, since we can choose any two distinct rational numbers $\gamma_1 < \gamma_2$ between δ_1 and δ_2 instead of δ_1 and δ_2 .

Definition 13. We say that $S \subseteq \mathbb{R}^n$ is essentially convex if for all $p, q \in S$ there are only finitely many points on the line segment between p and q that are not in S .

If S is not essentially convex, and if p and q are such that there are infinitely many points on the line segment between p and q that are not in S , then p and q witness that S is not essentially convex. The following is a direct consequence of Theorem 9, and we will use it in the following without further reference.

Corollary 14. If S is a semi-algebraic relation that is not essentially convex, then S excludes an interval. If S is an essentially convex semi-algebraic relation, and a, b are two distinct points from S , then the line segment between a and b contains an interval I with $I \subseteq S$.

The next proposition will be used several times in the sequel; it clarifies the relation between finite unions of varieties and essentially convex relations.

Proposition 15. Let W be a finite union of varieties $V_1, \dots, V_k \subseteq \mathbb{R}^n$, and let $C \subseteq W$ be essentially convex. Then, there is an $i \leq k$ such that $C \subseteq V_i$.

Proof. Let $J \subseteq \{1, \dots, k\}$ be minimal such that $C \subseteq \bigcup_{i \in J} V_i$. If $|J| = 1$, then there is nothing to show. So suppose for contradiction that there are distinct $i, j \in J$. Then there must be points $a, b \in C$ such that $a \in V_i$ and $a \notin V_l$ for all $l \in J \setminus \{i\}$, and $b \in V_j$ and $b \notin V_l$ for all $l \in J \setminus \{j\}$. By essential convexity of C and Corollary 14, the line segment L between a and b must contain an interval I that lies in C . Since J is finite, there must be $l \in J$ such that infinitely many points on I are from V_l . By Lemma 7, all points on the line through a and b are from V_l ; this contradicts the choice of a and b . \square

The rest of the section is divided into two parts. We first prove that if $S \subseteq \mathbb{R}^k$ is a semi-algebraic relation that is not essentially convex and this is witnessed by two rational points p and q , then $\text{CSP}((\Gamma_{\text{lin}}, S))$ is NP-hard. In the second part, we prove that if $S \subseteq \mathbb{R}^k$ is a semi-linear relation that is not essentially convex, then this is witnessed by rational points and, consequently, $\text{CSP}((\Gamma_{\text{lin}}, S))$ is NP-hard.

3.1 Semialgebraic relations and rational witnesses

We begin with the special case when S is a unary relation. The hardness proof is by a reduction from $\text{CSP}(\{0, 1\}; R_{1/3})$ where

$$R_{1/3} = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}.$$

This NP-complete problem is also called POSITIVE ONE-IN-THREE 3SAT [16, LO4], which is the variant of ONE-IN-THREE 3SAT where we have the extra requirement that in all input instances of the problem, no clause contains a negated literal.

Lemma 16. Let $S \subseteq \mathbb{R}$ be a unary relation. If S excludes an interval and this is witnessed by rational points p and q , then $\text{CSP}((\Gamma_{\text{lin}}, S))$ is NP-hard.

Proof. We know that there are rational numbers $0 < \delta_1 < \delta_2 < 1$ such that $p + (q - p)y \notin S$ whenever $\delta_1 \leq y \leq \delta_2$. Let

$$a = \sup\{\delta_2 - \delta_1 \mid 0 < \delta_1 < \delta_2 < 1 \text{ and } [p + (q - p)\delta_1, p + (q - p)\delta_2] \cap S = \emptyset\},$$

i.e., the least upper bound on the length (scaled to the interval $[0, 1]$) of excluded intervals between p and q . Choose rational numbers δ_1, δ_2 such that

- there exists $y \in [\delta_1 - d, \delta_1]$ such that $p + (q - p)y \in S$; and
- there exists $y \in [\delta_2, \delta_2 + d]$ such that $p + (q - p)y \in S$.
- $S \cap [p + (q - p)\delta_1, p + (q - p)\delta_2] = \emptyset$.

where $d = (\delta_2 - \delta_1)/5$. It is easy to see that such δ_1, δ_2 exist; we simply need to find δ_1, δ_2 such that $S \cap [p + (q - p)\delta_1, p + (q - p)\delta_2] = \emptyset$ and $\delta_2 - \delta_1$ is sufficiently close to a . Clearly, for any $\epsilon > 0$, there exist suitable δ_1, δ_2 such that $a - (\delta_2 - \delta_1) < \epsilon$.

Now, define $p' = p + (q - p)(\delta_1 - d)$, $q' = p + (q - p)(\delta_2 + d)$, and

$$U(y) \equiv \exists z. (z = p' + (q' - p')y \wedge S(z) \wedge 0 \leq y \leq 1).$$

Observe that U is pp-definable in $\Gamma_{\text{lin}} \cup \{S\}$ by Lemma 11 combined by the fact that p' and q' are rational numbers. We claim that U contains at least one point in the interval $[0, d']$, at least one point in the interval $[1 - d', 1]$, and no points in the interval $[d', 1 - d']$ where $d' = 1/7$. Let us consider the interval $[0, d']$. The point (expressed in p and q) corresponding to $y = 0$ is p' (which equals $p + (q - p)(\delta_1 - d)$) while the point corresponding to $y = 1/7$ is

$$\begin{aligned} p' + \frac{(q' - p')}{7} &= p + (q - p)(\delta_1 - d) + \frac{p + (q - p)(\delta_2 + d) - p - (q - p)(\delta_1 - d)}{7} \\ &= p + (q - p)(\delta_1 - d) + \frac{(q - p)((\delta_2 + d) - (\delta_1 - d))}{7} \\ &= p + (q - p)(\delta_1 - d) + \frac{(q - p)(\delta_2 - \delta_1 + 2d)}{7} \\ &= p + (q - p)(\delta_1 - d) + \frac{(q - p)(5d + 2d)}{7} \\ &= p + (q - p)(\delta_1 - d) + (q - p)d \\ &= p + (q - p)\delta_1 \end{aligned}$$

We know that the choice of δ_1 and δ_2 implies that there is at least one point in S on the line segment between $p + (q - p)(\delta_1 - d)$ and $p + (q - p)\delta_1$. The other two cases can be proved similarly.

We show NP-hardness by a polynomial-time reduction from $\text{CSP}(\{\{0, 1\}; R_{1/3}\})$. Let ϕ be an arbitrary instance of this problem and let V denote the set of variables appearing in ϕ . Construct a formula

$$\psi \equiv \bigwedge_{v \in V} U(v) \wedge \bigwedge_{R_{1/3}(v_i, v_j, v_k) \in \phi} v_i + v_j + v_k \geq \frac{6}{7} \wedge \bigwedge_{R_{1/3}(v_i, v_j, v_k) \in \phi} v_i + v_j + v_k \leq \frac{11}{7}.$$

Lemma 11 implies that ψ is pp-definable in $(\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\}, \leq, U)$ (and, consequently, pp-definable in $(\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\}, \leq, S)$) and the formula can be constructed in polynomial time. We now verify that the formula ψ has a solution if and only if ϕ has a solution.

Assume that there exists a satisfying truth assignment $f : V \rightarrow \{0, 1\}$ to the formula ϕ . We construct a solution g for ψ as follows: arbitrarily choose a point $t_0 \in [0, d']$ such that $t_0 \in U$ and a point $t_1 \in [1 - d', 1]$ such that $t_1 \in U$. Let $g(v) = t_0$ if $f(v) = 0$ and $g(v) = t_1$, otherwise. Clearly, this assignment satisfies every literal of the type $U(v)$. Each literal $v_i + v_j + v_k \geq 6/7$ is satisfied, too, since $g(v_i) + g(v_j) + g(v_k) = 2 \cdot t_0 + t_1 \geq 2 \cdot 0 + (1 - d') = 1 - d' = 6/7$. Similarly, each literal $v_i + v_j + v_k \leq 11/7$ is also satisfied: $g(v_i) + g(v_j) + g(v_k) = 2 \cdot t_0 + t_1 \leq 2 \cdot d' + 1 = 9/7$.

Assume now instead that there exists a satisfying assignment $g : V \rightarrow \mathbb{R}$ for the formula ψ . Each variable obtains a value that is in either the interval $[0, d']$ or in the interval $[1 - d', 1]$. If a variable is assigned a value in $[0, d']$, then we consider this variable ‘false’, i.e., having the truth value 0; analogously, variables assigned values in $[1 - d', 1]$ are considered ‘true’.

We continue by looking at an arbitrary literal $R_{1/3}(v_i, v_j, v_k)$ in ϕ and its corresponding inequalities (1) $v_i + v_j + v_k \geq 6/7$ and (2) $v_i + v_j + v_k \leq 11/7$. If all three variables are assigned values within $[0, d']$, then their sum is at most $3d' = 3/7$ which violates inequality (1). If two of the variables appear within $[1 - d', 1]$, then their sum is at least $0 + 2(1 - d') = 12/7$ which violates inequality (2); naturally, this inequality is violated if all three variables appear within $[1 - d', 1]$, too. If exactly one variable appears within $[1 - d', 1]$, then the sum of the variables is at least $1 - d' = 6/7$ and at most $1 + 2d' = 9/7$. We see that both inequality (1) and (2) are satisfied. Hence, we can define a satisfying assignment $f : V \rightarrow \{0, 1\}$ for ϕ :

$$f(v) = \begin{cases} 0 & \text{if } 0 \leq g(v) \leq d' \\ 1 & \text{otherwise} \end{cases}$$

This concludes the proof. \square

It is now straightforward to lift Lemma 16 to relations with arbitrary arities.

Lemma 17. *Let $S \subseteq \mathbb{R}^k$ be a semi-algebraic relation that is not essentially convex, and this is witnessed by two rational points $p = (p_1, \dots, p_k)$ and $q = (q_1, \dots, q_k)$. Let Γ be the structure (Γ_{lin}, S) . Then, $\text{CSP}(\Gamma)$ is NP-hard.*

Proof. Define

$$U(y) \equiv \exists \bar{z}. \bigwedge_{i=1}^k z_i = p_i + (q_i - p_i)y \wedge S(\bar{z}) \wedge 0 \leq y \leq 1$$

where $\bar{z} = (z_1, \dots, z_k)$. By Corollary 14, U excludes an interval and $\text{CSP}(\Gamma)$ is NP-hard by Lemma 16 since U is pp-definable in Γ . \square

Remark 18. If S is not essentially convex and this is witnessed by non-rational points only, then the problem $\text{CSP}(\Gamma)$ for $\Gamma = (\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\}, \leq, S)$ might still be solvable in polynomial time. Consider for instance the binary relation

$$S = \{(x, y) \mid (|x + y| \leq 1) \wedge (y = \sqrt{2}x \rightarrow |x + y| = 1)\}.$$

Clearly, S is not essentially convex; however, the only witnesses are $(\sqrt{2} - 1, 2 - \sqrt{2})$ and $(-\sqrt{2} + 1, -2 + \sqrt{2})$ (see Figure 1).

We show that $\text{CSP}(\Gamma)$ can be solved in polynomial time. To see this, define

$$S' = \{(x, y) \mid (|x + y| \leq 1) \wedge (x \neq 0 \vee y \neq 0)\}$$

and $\Delta = (\mathbb{R}; \{(x, y, z) \mid x + y = z\}, \{1\}, \leq, S')$. We first prove that a primitive positive sentence is true in Γ if and only if it is true in Δ . Clearly, if a primitive positive sentence is true in Γ , then it is also true in Δ , since the relations in Δ are supersets of the corresponding relations in Γ . Conversely, suppose that Φ is primitive positive and true in Δ . Let α be an assignment of the variables of Φ that satisfies all conjuncts in Φ . Since Δ is semi-linear, we can assume that α is




Fig. 1. Illustration of relation S

rational (see Lemma 19). The only relations that are different in Γ and in Δ are the relations S and S' . Since $S' \setminus S$ contains irrational points only, the assignment α shows that Φ is also true in Δ . Finally, $\text{CSP}(\Delta)$ can be solved in polynomial time by the results in Section 5.1 (note that the constraint $|x + y| \leq 1$ is equivalent to a conjunction of four linear inequalities). \square

3.2 Semilinear relations

In the previous section, we showed that there exists a relation S that is not essentially convex, this is witnessed by non-rational points only, and $\text{CSP}(\{(x, y, z) \mid x + y = z\}, \{1\}, \leq, S)$ is polynomial-time solvable. If we restrict ourselves to semi-linear relations, then this phenomenon cannot occur. By semi-linearity, all equalities and inequalities are of degree one and use rational parameters only. In particular, the boundary sets of excluded areas are rationally defined, and every excluded interval is witnessed by rational points. This section gives a formal confirmation of this fact.

Lemma 19. *Every non-empty semi-linear relation S contains at least one rational point.*

Proof. Assume first that S is a non-empty unary relation such that $S \cap \mathbb{Q} = \emptyset$. If S contains infinitely many points, then it also contains an interval due to ω -minimality of S ; this contradicts that $S \cap \mathbb{Q} = \emptyset$. So we assume that S contains a finite number of points. Consider the unary relation $S' = \{\min(S)\}$ and note that it can be pp-defined in (Γ_{lin}, S) by $S'(x) \equiv S(x) \wedge x \leq p$ where p denotes a suitably chosen rational number. By Theorem 10, S' has a quantifier-free definition ϕ over $(\mathbb{R}; +, -, \leq)$ with parameters from \mathbb{Q} , and we can without loss of generality assume that ϕ is in disjunctive normal form, and contains a single disjunct since $|S'| = 1$. Assume without loss of generality that every conjunct of this disjunct of ϕ is of one of the following forms: $x \geq c$, $x \leq c$, or $x \neq c$ (where c denotes some rational number). Let $a = \max\{c \mid (x \geq c) \in \phi\}$ and $b = \min\{c \mid (x \leq c) \in \phi\}$. If $a = b$ then $S' = \{a\}$ and we have a contradiction since a is a rational number. If $a < b$, then S' contains an infinite number of points (regardless of the number of disequality constraints in ϕ) and we have a contradiction once again.

Assume now that $\text{ar}(S) = d > 1$. Arbitrarily choose a point $s = (s_1, \dots, s_d) \in S$ with a maximum number of rational components. Assume without loss of generality that $s_1, \dots, s_{k'}, k' < k$

are rational points and consider the unary relation

$$U(x_k) \equiv \exists x_1, \dots, x_{k-1}. (S(x_1, \dots, x_k) \wedge x_1 = s_1 \wedge \dots \wedge x_{k'} = s_{k'}).$$

We get a contradiction since $U \cap \mathbb{Q} = \emptyset$, U is non-empty, and U is semi-linear. \square

Corollary 20. *Let $S \subseteq \mathbb{R}^k$ be a semi-linear relation and let $s \in S$ be arbitrary. Then, every open k -dimensional ball B around s of radius $\epsilon > 0$ contains a rational point in S .*

Proof. If there is an ϵ such that B does not contain any rational point in S , then there is a linear set P within B such that $S \cap P$ only contains irrational points. This contradicts Lemma 19. \square

A *hyperplane* is a set $V = \{x \in \mathbb{R}^k \mid p(x) = 0\}$ where p is a linear term such that $\emptyset \subset V \subset \mathbb{R}^k$ (this makes sure that the degree of p is one). We do not require that the coefficients in p are rational; it is important to note that this differs from the definition of a linear set. If all coefficients appearing in p are rational, then we say that the hyperplane is rational.

Lemma 21. *If T is a semi-linear relation that is not essentially convex, then this is witnessed by rational points, and $\text{CSP}((\Gamma_{lin}, T))$ is NP-hard.*

Proof. If there are rational witnesses of the fact that T is not essentially convex, then NP-hardness follows from Lemma 17 and we are done.

Assume now that there exists a relation T that is not essentially convex but T lacks rational witnesses. Arbitrarily choose such a T with minimal arity k . We first consider the case when $k = 1$. Arbitrarily choose witnesses $p, q \in T$. By o -minimality, there are finitely many intervals I_1, \dots, I_m with endpoints in $\mathbb{R} \cup \{\pm\infty\}$ and a finite set $D_0 \subseteq \mathbb{R}$ such that $T = D_0 \cup \bigcup_{i=1}^m I_i$. Now, apply the following process to D_0 and I_1, \dots, I_m .

- if there is a point $d \in D_0$ and an interval I_j , $1 \leq j \leq m$, such that d is in $\overline{I_j}$, then remove d from D_0 and replace I_j with $I_j \cup \{d\}$;
- repeat until D_0 is not changed.

After these modifications, the sets I_1, \dots, I_m are still (open, half-open, or closed) intervals, and for every point $d \in D_0$, there exists an $\epsilon_d > 0$ such that $[d - \epsilon_d, d + \epsilon_d] \cap T = \{d\}$.

Assume without loss of generality that $p \notin \mathbb{Q}$. If $p \in D_0$, then choose rational numbers p^-, p^+ such that $p - \epsilon_p < p^- < p < p^+ < p + \epsilon_p$; this is always possible since the rationals are a dense subset of the reals. Consider the semi-linear relation

$$T'(x) \equiv T(x) \wedge p^- \leq x \leq p^+$$

and note $T' = \{p\}$. However, p is not a rational number which contradicts Lemma 19. We may thus assume that $p \notin D_0$ and that p is a member of an interval $I \in \{I_1, \dots, I_m\}$. Arbitrarily choose one rational point $p' \in I$; once again, this is possible since the rationals are a dense subset of the reals. Note that p', q witness that T is not essentially convex. If $q \in \mathbb{Q}$, then we are done so we assume that $q \notin \mathbb{Q}$. We see that $q \notin D_0$ by reasoning as above. Consequently, q is a member of an interval $J \in \{I_1, \dots, I_m\}$ and $I \neq J$. Finally choose a rational point $q' \in J$ and note that p', q' are rational points witnessing that T is not essentially convex.

Assume instead that $k > 1$. Let \mathcal{S}_k denote the set of relations S that satisfy 1, 2, and 3:

1. S is a semi-linear relation of arity k ,
2. S is not essentially convex, and
3. for every pair of witnesses that S is not essentially convex, at least one is irrational.

We now conclude the proof by considering two different cases.

Case 1: There exists an $S \in \mathcal{S}_k$ and a finite set of rational hyperplanes H_1, \dots, H_h such that $S \subseteq \bigcup_{j=1}^h H_j$. Choose the hyperplanes such that h is minimal. Let $v, w \in S$ be arbitrarily chosen witnesses for the fact that S excludes an interval, and let I denote this interval.

Suppose first that $h = 1$, i.e., that there is a single hyperplane H such that $S \subseteq H$. Obviously, $x = (x_1, \dots, x_k) \in H \Leftrightarrow c_1 x_1 + \dots + c_k x_k = c_0$ for some rational constants c_0, \dots, c_k . We assume without loss of generality that at least one c_i , say c_k , is non-zero. Define the relation S' by

$$S'(x_1, \dots, x_{k-1}) \equiv \exists y. (S(x_1, \dots, x_{k-1}, y) \wedge y = \frac{c_0 - c_1 x_1 - \dots - c_{k-1} x_{k-1}}{c_k}).$$

Let $v' = (v_1, \dots, v_{k-1})$ and $w' = (w_1, \dots, w_{k-1})$, and note that v', w' are witnesses of an excluded interval in S' . If S' lacks rational witnesses of essential non-convexity, then the fact that S' has arity $k-1$ contradicts the choice of T . Hence, S' has two rational witnesses $t = (t_1, \dots, t_{k-1})$ and $u = (u_1, \dots, u_{k-1})$. This implies that

$$t' = \left(t_1, \dots, t_{k-1}, \frac{c_0 - c_1 t_1 - \dots - c_{k-1} t_{k-1}}{c_k} \right)$$

and

$$u' = \left(u_1, \dots, u_{k-1}, \frac{c_0 - c_1 u_1 - \dots - c_{k-1} u_{k-1}}{c_k} \right)$$

are rational witnesses for S , which leads to a contradiction.

Next, suppose that $h \geq 2$. Let $H'_1 = S \cap (H_1 \setminus \bigcup_{j=2}^h H_j)$ and $H'_2 = S \cap (H_2 \setminus \bigcup_{j \in \{1, 3, \dots, h\}} H_j)$. By the minimal choice of h , H'_1 and H'_2 are non-empty. Furthermore, they are semi-linear so we can choose rational points $p_i \in H'_i$, $1 \leq i \leq 2$, by Lemma 19. We now claim that at most a finite number of points on the line segment between p_1 and p_2 lie in S . Suppose to the contrary that infinitely many points lie on the line segment. Then, there must be one H_i , $i \geq 1$, such that infinitely many points from H_i lie on the line segment. Hence, H_i (since it is a variety) must contain the entire line by Lemma 7. This leads to a contradiction since p_1 and p_2 are chosen so that $|\{p_1, p_2\} \cap H_j| \leq 1$, $1 \leq j \leq h$. Thus, we have found rational witnesses for essential non-convexity of S and obtained a contradiction since $S \in \mathcal{S}_k$.

Case 2: There is no $S \in \mathcal{S}_k$ such that there exists a finite set of rational hyperplanes H_1, \dots, H_h and $S \subseteq \bigcup_{j=1}^h H_j$. Arbitrarily choose $S \in \mathcal{S}_k$, let $v, w \in S$ be arbitrarily chosen witnesses for the fact that S excludes an interval, and let I denote such an interval.

If there exists a rational hyperplane H such that $\{v, w\} \subseteq S \cap H$, then the semi-linear relation

$$S'(x_1, \dots, x_k) \equiv S(x_1, \dots, x_k) \wedge H(x_1, \dots, x_k)$$

excludes an interval and this is witnessed by v and w . Obviously, $S' \in \mathcal{S}_k$ and $S' \subseteq H$. This contradicts the assumptions for this case so we assume that $\{v, w\}$ (and consequently I) do not lie on any rational hyperplane.

Next, we prove a couple of facts.

Fact 1: $I \subseteq \bar{S} \setminus S$. We show that there is no point $e \in I$ and an $\epsilon > 0$ such that the open k -dimensional ball B around e with radius ϵ satisfies $B \cap S = \emptyset$. Assume to the contrary that there is a point $e \in I$ satisfying this condition. By Corollary 20, there exist rational points in S arbitrary close to v and w . Thus, one can find rational points $v', w' \in S$ such that the line segment L between v' and w' passes through B and $L' = L \cap B$ has non-zero length. In other words, v' and w' are rational witnesses of an excluded interval and we have obtained a contradiction.

Fact 2: There exists a finite set $\{H_1, \dots, H_h\}$ of rational hyperplanes such that $\bar{S} \setminus S \subseteq \bigcup_{i=1}^h H_i$. Let ϕ be a first-order definition of S and let $\psi = D_1 \vee \dots \vee D_n$ be a quantifier-free definition of S in disjunctive normal form; such a ψ exists due to Theorem 10. Note that every parameter appearing in ψ is rational: initially, every parameter in ϕ is rational, the quantifier elimination does not add any irrational parameters, and the conversion to disjunctive normal form does not introduce any new parameters. Let l_1, \dots, l_m denote the literals appearing in ϕ . For each $l_i \equiv p(x_1, \dots, x_k) r 0$ (where $r \in \{\leq, <, =, \neq, >, \geq\}$), create a hyperplane $H_i = \{(x_1, \dots, x_k) \subseteq \mathbb{R}^k \mid p(x_1, \dots, x_k) = 0\}$. In other words, we let the boundary of the subspace defined by l_i define a hyperplane H_i . It is now easy to see that $\bar{S} \setminus S \subseteq \partial S \subseteq \bigcup_{i=1}^m H_i$. Furthermore, every hyperplane H_1, \dots, H_m is rational.

We are now ready to prove the second case of the proof. By Fact 1, $I \subseteq \bar{S} \setminus S$. The set $\bar{S} \setminus S$ is a subset of $\bigcup_{i=1}^h H_i$ where H_1, \dots, H_h are rational hyperplanes by Fact 2. Hence, I is a subset of some H_i by Proposition 15, a contradiction. \square

4 Essentially Convex Relations

Before we present a logical characterization of essentially convex semi-algebraic relations, we give examples that show that two more naive syntactic restrictions of first-order formulas are not powerful enough for defining all essentially convex semi-algebraic relations. Both of those restrictions are motivated by classes of essential convex semi-linear relations that have appeared in the literature, cf. [21]. When S is a subset of \mathbb{R}^n , we write $\neg S$ for the complement of S , i.e., for $\mathbb{R}^n \setminus S$.

We start with an example that shows that not every essentially convex semi-algebraic relation can be defined by conjunctions of first-order formulas of the form

$$p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$$

where p_1, \dots, p_k are polynomials with coefficients from \mathbb{R} , and where ϕ defines a convex set. It is easy to see that every relation that can be defined by such a conjunction is essentially convex.

See Figure 2, left side. The figure shows a 1-dimensional variety $C \subseteq \mathbb{R}^2$, given as $\{(p(t), q(t)) \mid t \in \mathbb{R}\}$ for polynomials p and q . The figure also shows two marked segments S_1, S_2 on the curve C . The marked segments are chosen such that one end point is contained in interior of the convex hull of the other three end points of the segments.

Let S be the set $\neg C \cup S_1 \cup S_2$. Clearly, S is essentially convex. Now, suppose for contradiction that S has a definition ψ as described above. Let H be the convex hull of $S_1 \cup S_2$. The crucial observation is that the set $G := (H \cap C) \setminus (S_1 \cup S_2)$ is infinite. Since no point from G is in S , there must be a conjunct $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$ in ψ that excludes infinitely many points from G . In particular, the variety V defined by $p_1 = \dots = p_k = 0$ contains infinitely many points from C . As in the proof of Lemma 7, one can see that V must contain C . Hence, all points in $S_1 \cup S_2$ must satisfy ϕ ; but in this case, all points in G satisfy $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$, a contradiction.




Fig. 2. Examples of essentially convex relations.

Our first example might motivate the following notion of definability: we consider conjunctions of formulas of the form $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$ such that for every conjunction of linear equalities ψ that implies $p_i = 0$ for all $i \leq k$, the set defined by $\phi \wedge \psi$ is convex. The set described above can indeed be defined by such a formula. Similarly as before, it is also easy to see that all relations that can be defined in such a way are essentially convex. However, we again have an example of a semi-algebraic essentially convex relation that cannot be defined by such a conjunction.

See Figure 2, right side. The figure shows the boundary B of a doubly infinite cone with apex o . On the boundary, there is a straight line segment L through o , and a circle C that cuts L . Let S be the set $\neg B \cup C \cup (L \setminus \{o\})$. It can be verified that S is essentially convex. However, we claim that there is no conjunction as described above that defines S . The reason is that when $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$ is such that $p_1 = 0 \wedge \dots \wedge p_k = 0$ describes B , and if ϕ contains $C \cup (L \setminus \{o\})$, then it must also contain o in order that $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$ meets the required condition. However, the set $\neg B \cup C \cup L$ is not essentially convex since o and points from the cycle exclude an interval.

The correct definition of formulas that correspond to essentially convex sets has to take these examples into account. We call these formulas *convex Horn formulas*. Basically, a convex Horn formula is a conjunction of implications $p_1 = \dots = p_k = 0 \rightarrow \phi$ such that the premise defines a variety V , and the formula ϕ is again convex Horn when restricted to any convex subset of V . Formally, we have the following definition.

Definition 22. *The set of convex Horn formulas is the smallest set of first-order formulas such that*

- all formulas defining convex closed semi-algebraic relations over $(\mathbb{R}; +, *, \leq)$ with parameters from \mathbb{R} are convex Horn;
- Suppose that p_1, \dots, p_k are polynomials, ϕ is a first-order formula that defines a set $U \subseteq \mathbb{R}^n$, and for every semi-algebraic convex set C contained in the set defined by $p_1 = \dots = p_k = 0$, the

set $C \cap U$ can be defined by a convex Horn formula, and has strictly smaller dimension than the set defined by $\psi \equiv (p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi)$. Then ψ is also convex Horn.

– Finite conjunctions of convex Horn formulas are convex Horn.

For example, the formula $(x^2 - y \neq 0) \wedge (y \geq 1)$, describing the half-plane above $y = 1$ with exception of the standard parabola, is convex Horn. Every convex set C contained in the set defined by $x^2 - y = 0$ consists of at most one point, and hence is 0-dimensional and can be defined by a convex Horn formula.

We can prove properties about the set of all convex Horn formulas by induction over the *level* of a convex Horn formula, which is defined as follows. The level of a formula that defines a convex closed semi-algebraic relation is 0. Now, suppose we have already defined convex Horn formulas of level smaller than i , and let ψ be a convex Horn formula that does not have level smaller than i . Then ψ has level i if it is the finite conjunction of formulas $\psi' \equiv (p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi)$ such that for every semi-algebraic convex set C contained in the set defined by $p_1 = \dots = p_k = 0$, the intersection of C with the set defined by ϕ is convex Horn, has level at most $i - 1$, and strictly smaller dimension than the set defined by ψ' . Since the intersection of sets of dimension n has at most dimension n , it follows directly from the definition of convex Horn formulas that the level of a convex Horn formula ϕ is bounded by the dimension of the set defined by ϕ .

Looking back at the formula $(x^2 - y \neq 0) \vee (y \geq 1)$, we claim that it is convex Horn of level one: every convex set contained in the parabola consists of at most one point, and can hence be described by a convex Horn formula of level zero. For another example, consider $(x^2 - y \neq 0) \vee (z \geq 0)$, i.e., the same parabola in three dimensions on one side side of the x - y -plane. Each convex subset of $x^2 - y = 0$ is a point, a straight line, or a line segment in the z direction and can again be described by a level zero convex Horn formula.

We are now ready to logically define essentially convex semi-algebraic sets via convex Horn formulas. This is done in two steps; we first prove (in Proposition 23) that every set defined by a semi-algebraic convex Horn formula is essentially convex. The rest of the section is devoted to proving the other direction—the final result can be found in Theorem 27.

Proposition 23. *Any set S defined by a convex Horn formula ψ over $(\mathbb{R}; *, +, \leq)$ is essentially convex.*

Proof. Our proof is by induction over the level of ψ . Let m denote the number of free variables in ψ . If the level of ψ is 0, then $S = \{x \in \mathbb{R}^m \mid \psi(x)\}$ is a closed convex set and, in particular, essentially convex.

Assume that all relations defined by convex Horn formulas of level $< i$ are essentially convex. Arbitrarily choose a convex Horn formula $\psi \equiv p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi$ with level i . Define $S = \{x \in \mathbb{R}^m \mid \psi(x)\}$, $V = \{x \in \mathbb{R}^m \mid p_1(x) = \dots = p_k(x) = 0\}$, and $U = \{x \in \mathbb{R}^m \mid \phi(x)\}$. Since ψ is level- i convex Horn, we know that for every semi-algebraic convex set C such that $C \subseteq V$, the set $C \cap U$ can be defined by a convex Horn formula of level smaller than i and $\dim(C \cap U) < \dim(S)$.

Suppose for contradiction that there are $a, b \in S$ and an infinite set I of points on the line segment L between a and b that is not contained in S . In particular, $I \subseteq V$ since $S = \neg V \cup U$. By Lemma 7, all points on the line through a and b are in V . However, a and b are in S and therefore in U so $L \cap U$ is not essentially convex. We now note that L is a semi-algebraic convex set that is a subset of V so $L \cap U$ can be defined by a convex Horn formula of level smaller than i . Consequently, $L \cap U$ is essentially convex by the inductive assumption which leads to a contradiction.

Finally, suppose that ψ is a finite conjunction of convex Horn formulas of level at most i . Since the intersection of finitely many essentially convex relations is essentially convex, we are done also in this case. \square

Next, we need some preparations for the proof of the converse implication (Theorem 27): we show that semi-algebraic relations can be defined by a special type of formula (Lemma 25), and that the closure \bar{S} of an essentially convex relation S is convex (Lemma 26).

Definition 24. *Let S be a semi-algebraic relation. We say that a first-order formula ϕ is a standard definition of S if*

- $\phi(x_1, \dots, x_k)$ defines $S \subseteq \mathbb{R}^k$ over $(\mathbb{R}; *, +, -, \leq)$ with parameters from \mathbb{R} ;
- ϕ is in quantifier-free conjunctive normal form;
- if we remove any literal from ϕ , then the resulting formula is not equivalent to ϕ ; and
- all literals are of the form $t \leq 0$ or $t \neq 0$.

Lemma 25. *Every semi-algebraic relation S has a standard definition. If S is even semi-linear, then it has a standard definition ϕ that does not involve the function symbol for multiplication and irrational parameters.*

Proof. By Theorem 8, we know that S has a quantifier-free definition over $(\mathbb{R}; *, +, \leq)$ with parameters from \mathbb{R} , and it is clear that such a definition can be rewritten in conjunctive normal form ϕ . Replace a clause α in ϕ with a literal of the form $a < b$ by two clauses α_1 and α_2 obtained from α by replacing $a < b$ by $a \leq b$ and by $a \neq b$, respectively. In the same way we can eliminate occurrences of $a = b$ from ϕ using \leq . Literals of the form $a \leq b$ ($a \neq b$) can then be replaced by $a - b \leq 0$ (and $a - b \neq 0$, respectively). Finally, we remove literals from ϕ as long as the resulting formula is equivalent to the original formula.

If S is semi-linear, then by Theorem 10 we know that S has a quantifier-free definition over $(\mathbb{R}; +, -, \leq)$ with parameters from \mathbb{Q} , and it is then clear that the formula constructed from ϕ as above will be a standard definition of S without the function symbol for multiplication and irrational parameters. \square

Lemma 26. *The closure \bar{S} of an essentially convex relation S is convex.*

Proof. Let $a, b \in \bar{S}$. We will show that all points c on the line segment between a and b are in \bar{S} . We have to show that for every $\epsilon > 0$ there is a point c' in S such that the distance between c and c' is smaller than ϵ . Since $a \in \bar{S}$ and $b \in \bar{S}$, there are points $a' \in S$ and $b' \in S$ that are closer than $\epsilon/2$ to a and b , respectively. Let L be the line from a' to b' . It is clear that there are infinitely many points on L that are at distance less than ϵ from c . Hence, since a' and b' are in S and S is essentially convex, there must be one such point in S , and we are done. \square

Theorem 27. *A semi-algebraic relation $S \subseteq \mathbb{R}^n$ is essentially convex if and only if it has a convex Horn definition. Moreover, when S is even semi-linear then S has a semi-linear convex Horn definition.*

Proof. We have already seen in Proposition 23 that every relation defined by a convex Horn formula is essentially convex. We now show the more difficult implication of the statement. Let ϕ be a standard definition of S . The proof is by induction on the dimension d of S . For $d = 0$, the set S consists of at most one point, and the statement is trival. Otherwise, if $|S| \geq 2$, then by essential convexity, Corollary 14, and Lemma 6 we have that $\dim(S) \geq 1$.

For $d > 0$, we will construct two formulas ϕ_1, ϕ_2 such that ϕ is equivalent to $\phi_1 \wedge \phi_2$. Thereafter, we will show that ϕ_1 is equivalent to a conjunction of convex Horn formulas, and that ϕ_2 defines a closed convex relation (and consequently is convex Horn). Since finite conjunctions of convex Horn formulas are also convex Horn, ϕ is then equivalent to a convex Horn formula.

We begin by writing all clauses of ϕ as $\alpha \rightarrow \beta$ where α is either ‘true’ or a conjunction of polynomial equalities and β is either ‘false’ or a disjunction of inequalities. This is always possible since a clause

$$(p_1 \leq 0 \vee \cdots \vee p_k \leq 0 \vee q_1 \neq 0 \vee \cdots \vee q_m \neq 0)$$

is logically equivalent to

$$(q_1 = 0 \wedge \cdots \wedge q_m = 0) \rightarrow (p_1 \leq 0 \vee \cdots \vee p_k \leq 0).$$

Next, we rewrite all clauses $\alpha \rightarrow \beta$ where α is not equal to ‘true’, as $(\alpha \rightarrow (\beta \wedge \phi))$. Let ϕ_1 be the conjunction of all the implications of the type $(\alpha \rightarrow (\beta \wedge \phi))$ and ϕ_2 the conjunction of the remaining implications, i.e., those of the type $(\text{true} \rightarrow \beta)$. The formula $\phi_1 \wedge \phi_2$ is clearly equivalent to the formula ϕ .

We begin by studying the formula ϕ_1 . Let $\alpha \rightarrow (\beta \wedge \phi)$ be a clause from ϕ_1 , let V be the variety defined by α , and let U be the set defined by $\beta \wedge \phi$. Observe that $U \subseteq S$. We now show that the intersection of the set U with a semi-algebraic convex set $C \subseteq V$ can be defined by a convex Horn formula. We make two claims about the set $U \cap C$:

Claim 1. $U \cap C$ is essentially convex. For arbitrary points $a, b \in U \cap C$, let L_{ab} denote the line segment from a to b , and $X_{ab} = \{x \in L_{ab} \mid x \notin U \cap C\}$. Suppose for contradiction that there exist $a, b \in U \cap C$ such that X_{ab} is an infinite set. Since $a, b \in C$ and C is convex, $L_{ab} \subseteq C$ which implies that $X_{ab} = \{x \in L_{ab} \mid x \notin U\}$. Moreover, $C \subseteq V$ so $X_{ab} \subseteq V$. Now recall that $a, b \in S$ since $U \cap C \subseteq U \subseteq S$: thus, there are infinitely many points (those that are in X_{ab}) between $a, b \in S$ that are in V but not in U . This shows that no point in X_{ab} satisfies $\alpha \rightarrow (\beta \wedge \phi)$, and $X_{ab} \cap S = \emptyset$. This fact contradicts the essential convexity of S .

Claim 2. $U \cap C$ has smaller dimension than S . Let T be the set $S \setminus (U \cap C)$. It suffices to show that $U \cap C$ is a subset of $\bar{T} \setminus T$, because Lemma 6 asserts that $\dim(\bar{T} \setminus T) < \dim(T) \leq \dim(S)$.

The set S must contain a point p that is not in V , because if $S \subseteq V$ then we could replace the clause of ϕ that was re-written to $\alpha \rightarrow (\beta \wedge \phi)$ by β and obtain a formula that is equivalent to ϕ ; this contradicts the assumption that ϕ is a standard definition of S .

To show that $(U \cap C) \subseteq \bar{T} \setminus T$, let x be an arbitrary point in $U \cap C$. Only finitely many points on the line segment between p and x can be from $(U \cap C) \subseteq V$, because otherwise Proposition 15 implies that V must contain the entire line between x and p , including p , a contradiction. Also the set S contains all but finitely many points on the line segment between p and x : this is by essential convexity of S , since $x \in U \cap C \subseteq S$ and $p \in S$. Hence, we can choose a sequence of points from $T = S \setminus (U \cap C)$ on this line segment that approaches x , which shows that $x \in \bar{T}$.

Since $U \cap C$ is semi-algebraic, essentially convex, and has smaller dimension than S , it follows by the inductive assumption that it can be defined by a convex Horn formula. Thus, ϕ_1 is equivalent to a finite conjunction of convex Horn formulas.

We claim that ϕ_2 defines a closed convex set D . This follows from Lemma 26, since D is in fact the closure of S . To see this, observe that D is clearly a closed set, D contains S , and hence $\bar{S} \subseteq \bar{D} = D$. To prove that $D \subseteq \bar{S}$, let y be from $D \setminus S$. Consider the clauses $\alpha_1 \rightarrow (\beta_1 \wedge \phi), \dots, \alpha_l \rightarrow (\beta_l \wedge \phi)$ of ϕ_1 , and let V_i , for $1 \leq i \leq l$, be the variety $\{x \in \mathbb{R}^k \mid x \text{ satisfies } \alpha_i\}$. There must be a point q in

S that is not contained in the set $W = \bigcup_{i \leq l} V_i$; otherwise, Proposition 15 implies that there exists an $i \leq l$ such that $S \subseteq V_i$. In other words, all points in S satisfy α_i . This is in contradiction to the assumption that ϕ is a standard definition of S , since in this case the formula ϕ is equivalent to the formula where the clause of ϕ that has been rewritten to $\alpha_i \rightarrow (\beta_i \wedge \phi)$ is replaced by β_i . Only finitely many points on the line segment L between q and y can be from W , because otherwise Lemma 7 implies that W contain the entire line between y and q , including q , a contradiction. Hence, $y \in \bar{S}$.

Finally, consider the case that S is semi-linear. By Lemma 25, we can choose ϕ to be a standard definition which is semi-linear (and only uses parameters in \mathbb{Q}). Then the proof above leads to a semi-linear convex Horn definition of S . \square

5 Applications

5.1 Semi-linear constraint languages

We will now show that a finite semi-linear expansion Γ of Γ_{lin} has a polynomial-time tractable constraint satisfaction problem if and only if all relations of Γ are essentially convex.

Recall that a relation is semi-linear if it has a first-order definition in $(\mathbb{R}; +, 1, \leq)$. A quantifier-free first-order formula in CNF is called *Horn-DLR* [21] (where ‘DLR’ stands for *disjunctive linear relations*) if its clauses are of the form

$$p_1 \neq 0 \vee \dots \vee p_k \neq 0$$

or of the form

$$p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee p_0 \leq 0$$

where p_0, p_1, \dots, p_k are linear terms with rational coefficients. A semi-linear relation is called *Horn-DLR* if it can be defined by a Horn-DLR formula. The constraint satisfaction problem for Horn-DLRs can be solved in polynomial time.

Theorem 28 (see [10, 21, 22]). *Let Γ be a structure with domain \mathbb{R} whose relations are Horn-DLR. Then $\text{CSP}(\Gamma)$ is in P .*

In this section, we show the following.

Theorem 29. *Let $\Gamma = (\Gamma_{\text{lin}}, S_1, \dots, S_l)$ be a constraint language such that S_1, \dots, S_l are semi-linear relations. Then, either each relation S_1, \dots, S_l is Horn-DLR and $\text{CSP}(\Gamma)$ is in P , or $\text{CSP}(\Gamma)$ is NP-complete.*

In order to prove Theorem 29, we need to characterize convex and essentially convex semi-linear relations. This is done in Lemma 30 and Theorem 31, respectively.

Let P_1, \dots, P_n be (possibly unbounded) polyhedra defined such that $P_i = \{x \in \mathbb{R}^k \mid A_i x \leq b_i\}$. Bemporad et al. [3] define the *envelope* of P_1, \dots, P_n ($\text{env}(P_1, \dots, P_n)$) as the polyhedron

$$\{x \in \mathbb{R}^k \mid A'_1 x \leq b'_1, \dots, A'_n x \leq b'_n\}$$

where $A'_i x \leq b'_i$ is the subsystem of $A_i x \leq b_i$ obtained by removing all the inequalities not valid for the other polyhedrons $P_1, \dots, P_{i-1}, P_{i+1}, P_n$. We note that if P_1, \dots, P_n are defined with coefficients from \mathbb{Q} , then $\text{env}(P_1, \dots, P_n)$ can be described by rational coefficients, too. By combining Theorem 3 with Remark 1 in Bemporad et al., it follows that $\bigcup_{i=1}^n P_i$ is convex if and only if $\bigcup_{i=1}^n P_i = \text{env}(P_1, \dots, P_n)$.

Lemma 30. *A closed semi-linear relation S is convex if and only if it has a primitive positive definition in $(\mathbb{R}; +, 1, \leq)$.*

Proof. It is straightforward to verify that relations with a primitive positive definition in $(\mathbb{R}; +, 1, \leq)$ are convex; each relation defines a convex set and the intersection of convex sets is convex itself.

For the converse, let $\phi = \psi_1 \vee \dots \vee \psi_m$ be a quantifier-free definition of S over $(\mathbb{R}; +, -, \leq)$ with parameters from \mathbb{Q} , written in disjunctive normal form. If there is a disjunct ψ_i that contains a literal $p \neq q$, for linear terms p and q , then split the disjunct into two; one containing $p - q < 0$ and one containing $p - q > 0$. By repeating this process, every literal $p \neq q$ can be removed. Similarly, every literal $p = q$ can be replaced by $p - q \leq 0 \wedge q - p \leq 0$. Thus, we may assume that every literal appearing in the ψ_i is of the type $p \leq 0$ or $p < 0$, for a linear term p . Let D_1, \dots, D_m be the sets defined by ψ_1, \dots, ψ_m , respectively.

Now recall that the topological closure operator preserves finite unions, i.e., $\bar{D}_1 \cup \dots \cup \bar{D}_m = \overline{D_1 \cup \dots \cup D_m}$. Hence,

$$S = D_1 \cup \dots \cup D_m \subseteq \bar{D}_1 \cup \dots \cup \bar{D}_m = \overline{D_1 \cup \dots \cup D_m} = \bar{S} = S$$

and $S = \bar{D}_1 \cup \dots \cup \bar{D}_m$. We now note that if $P = \{x \in \mathbb{R}^k \mid Ax \leq b, Cx < d\}$ and $P \neq \emptyset$, then $\bar{P} = \{x \in \mathbb{R}^k \mid Ax \leq b, Cx \leq d\}$, cf. Case (i) of Proposition 1.1 in [17]. Thus, each $\bar{D}_i = \{x \in \mathbb{R}^k \mid A_i x \leq b_i\}$ for some rational A_i, b_i . Furthermore, $S = \bigcup_{i=1}^m \bar{D}_i$ is convex so $S = \text{env}(\bar{D}_1, \dots, \bar{D}_m)$ by Bemporad et al.'s result. This implies that $S = \{x \in \mathbb{R}^k \mid C_1 x \leq d_1, \dots, C_m x \leq d_m\}$ for some rational matrices C_1, \dots, C_m and rational vectors d_1, \dots, d_m . It is easy to see that each $C_i x \leq d_i$ is pp-definable in $(\mathbb{R}; +, 1, \leq)$ by the same technique as in the proof of Lemma 11, and this concludes the proof. \square

Theorem 31. *A semi-linear relation S is essentially convex if and only if S is Horn-DLR.*

Proof. We first prove that every Horn-DLR relation S is essentially convex. Let ϕ be a Horn-DLR definition of S . Suppose for contradiction that there are $a, b \in S$ and an infinite set I of points on the line segment L between a and b is not contained in S . Since ϕ has finitely many conjuncts, there is a conjunct ψ in ϕ that is false for an infinite subset I' of I . If ψ is of the form $p_1 \neq 0 \vee \dots \vee p_k \neq 0$, then all points in I' satisfy $p_1 = \dots = p_k = 0$. By Lemma 7, the entire line L satisfies $p_1 = \dots = p_k = 0$. This contradicts the assumption that $a \in L$ and $b \in L$ satisfy ϕ . If ψ is of the form $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee p_0 \leq 0$, then all points in I' satisfy $p_1 = \dots = p_k = 0$ and $p_0 > 0$. Again by Lemma 7 we find that both a and b must satisfy $p_1 = \dots = p_k = 0$. Since a, b also satisfy ψ , we conclude that both points satisfy $p_0 \leq 0$. But then also all points in L must satisfy $p_0 \leq 0$, which contradicts the fact that the points in I' satisfy $p_0 > 0$.

The other direction of the statement can be derived from Theorem 27 as follows. Let S be an essentially convex semi-linear relation. By Theorem 27, S has a semi-linear convex Horn definition ψ . We prove by induction on the level of ψ that ψ is Horn-DLR. If the level is 0, then S is closed and convex and the claim follows from Lemma 30. Now suppose that ψ has level $i > 0$ and is of the form $p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \psi'$, where $p_1 = \dots = p_k = 0$ defines a set V , and ψ' defines a set U such that for every semi-algebraic convex set $C \subseteq V$ the set $C \cap U$ has a convex Horn definition of level strictly smaller than i . Since ψ is semi-linear, the terms p_1, \dots, p_k are linear. Hence, V is convex, and by taking $C := V$ in the statement above we see that $V \cap U$ has a convex Horn definition of level strictly smaller than i . By the inductive assumption, $V \cap U$ has a definition by a Horn-DLR formula ϕ with clauses ϕ_1, \dots, ϕ_m . Then $\phi' = \bigwedge_{1 \leq i \leq m} (p_1 \neq 0 \vee \dots \vee p_k \neq 0 \vee \phi_i)$ is clearly Horn-DLR. We claim that ϕ' defines S . First suppose that $a \in \neg V$. In this case, a clearly

satisfies ϕ' and this is justified by the fact $\neg V \subseteq S$. Suppose instead that $a \in V$. Then a satisfies ϕ' if and only if it satisfies ϕ , and since ϕ defines $V \cap U$ this is the case if and only if a satisfies ψ' .

Finally, the statement holds if ψ is the conjunction of finitely many convex Horn formulas (which are Horn-DLR by inductive assumption). \square

Proof (of Theorem 29). If all relations of Γ are Horn-DLR, then $\text{CSP}(\Gamma)$ can be solved in polynomial time (Theorem 28). Otherwise, if there is a relation S from Γ that is not Horn-DLR, then Theorem 31 shows that S is not essentially convex, and NP-hardness of $\text{CSP}(\Gamma)$ follows by Lemma 21.

So we only have to show that $\text{CSP}(\Gamma)$ is in NP. Let Φ be an arbitrary instance of $\text{CSP}(\Gamma)$. By Theorem 10 every relation of Γ has a quantifier-free definition in conjunctive normal form over $(\mathbb{R}; +, -, \leq)$ with rational parameters. One can now non-deterministically guess *one* literal from each clause of in the defining formula for each constraint and verify – in polynomial-time by Theorem 28 – that all the selected literals are simultaneously satisfiable. \square

5.2 Generalized linear programming

In this section, we study generalizations of the following problem.

Linear Programming (LP)

INPUT: A finite set of variables V , a vector $c \in \mathbb{Q}^{|V|}$, a number $M \in \mathbb{Q}$, and a finite set of linear inequalities of the form $a_1x_1 + \dots + a_nx_n \leq a_0$ where $x_1, \dots, x_n \in V$ and $a_0, \dots, a_n \in \mathbb{Q}$. All rationals are given by numerators and denominators represented in binary.

QUESTION: Is there a vector $x \in \mathbb{R}^{|V|}$ that satisfies the inequalities and $c^T x \geq M$?

We generalize LP as follows. Let Γ be a structure $(\Gamma_{\text{lin}}, R_1, \dots, R_m)$ such that R_1, \dots, R_m are semi-linear relations.

Generalized Linear Programming for Γ (GLP(Γ))

INPUT: A finite set of variables V , a vector $c \in \mathbb{Q}^{|V|}$, a number $M \in \mathbb{Q}$, and a finite set Φ of expressions of the form $R(x_1, \dots, x_k)$ where R is a relation from Γ and $x_1, \dots, x_k \in V$.

QUESTION: Is there a vector $x \in \mathbb{R}^{|V|}$ that satisfies the constraints and $c^T x \geq M$?

This can indeed be viewed as a generalization of LP because of Proposition 12: the problem LP is polynomial-time equivalent to the problem $\text{GLP}(\Gamma_{\text{lin}})$.

Theorem 32. *Let $\Gamma = (\mathbb{R}; \Gamma_{\text{lin}}, R_1, \dots, R_l)$ be a structure with semi-linear relations R_1, \dots, R_l . Then, either each R_i is Horn-DLR and $\text{GLP}(\Gamma)$ is in P, or $\text{GLP}(\Gamma)$ is NP-hard.*

Proof. If there is an R_i that is not Horn-DLR, then the relation R_i is not essentially convex by Theorem 31, and $\text{CSP}((\Gamma_{\text{lin}}, R_i))$ is NP-hard by Theorem 29. Clearly, $\text{GLP}(\Gamma)$ is NP-hard, too.

Assume instead that each R_i is Horn-DLR. We present an algorithm that actually solves a more general problem that includes $\text{GLP}(\Gamma)$. Let Φ be an arbitrary satisfiable Horn-DLR formula⁶ over variable vector $\bar{x} = (x_1, \dots, x_n)$ and let c be a rational n -vector.

⁶ Note that if we are given an instance of $\text{CSP}((\Gamma_{\text{lin}}, R_1, \dots, R_l))$, then it can be transformed into an equivalent Horn-DLR formula in polynomial time since there is only a finite number of relations in the given structure. Hence, there is no loss of generality in considering Horn-DLR formulas instead of CSP instances. Also note that the resulting formula is (up to a multiplicative constant depending on the structure) of the same size as the CSP instance.

We assume additionally that $\Phi \wedge D$ is satisfiable for every disequality literal (i.e., literal $p(\bar{x}) \neq a$) D appearing in Φ . If $\Phi \wedge D$ is not satisfiable, then every occurrence of D in Φ can be removed without changing the set defined by the formula. Furthermore, this check can be carried out in polynomial time by Theorem 28. Hence, we may assume that Φ has this additional property (which we will refer to as $(*)$) without loss of generality. Let $\Phi = \Phi' \wedge \Phi''$ where Φ' consists of the clauses not containing any disequality literal $p(\bar{x}) \neq a$, and Φ'' consists of the remaining clauses.

Given Φ , our algorithm returns one of the following three answers:

- ‘unbounded’: for every $K \in \mathbb{Q}$, there exists a solution y such that $c^T y \geq K$;
- ‘optimum: K ’: there exists a $K \in \mathbb{Q}$ and a solution y such that $c^T y = K$, but there is no solution y' such that $c^T y' > K$;
- ‘optimum is arbitrarily close to K ’: there exists a $K \in \mathbb{Q}$ such that there is no solution y satisfying $c^T y \geq K$, but for every $\epsilon > 0$ there is a solution y' with $c^T y' \geq K - \epsilon$.

We claim that the following algorithm solves the task described above in polynomial time.

Step 1. Maximize $c^T \bar{x}$ over Φ' (by using some polynomial-time algorithm for linear programming). Let K denote the optimum. If $K = \infty$, then return ‘unbounded’ and stop.

Step 2. Check whether $\Phi \wedge c^T \bar{x} = K$ is satisfiable. Note that $c^T \bar{x} = K$ has a primitive positive definition in Γ_{lin} , which furthermore can be computed in polynomial time by Lemma 11. Therefore this check can be reduced to deciding satisfiability of Horn-DLRs. If $\Phi \wedge c^T \bar{x} = K$ is satisfiable, then return ‘optimum: K ’. If this is not the case, then return ‘optimum is arbitrarily close to K ’.

We first show that the algorithm runs in polynomial time. **Step 1** takes polynomial time since maximizing $c^T \bar{x}$ over Φ' is equivalent to solving a linear program with size polynomially bounded in the size of Φ . Finally, **Step 2** takes polynomial time due to Lemma 11 and Theorem 28.

Next, we prove the correctness of the algorithm. Correctness is obvious if the algorithm answers ‘optimum: K ’ in **Step 2**. For the remaining cases, we need to make a couple of observations. Define $S = \{x \in \mathbb{R}^n \mid x \text{ satisfies } \Phi\}$ and $S' = \{x \in \mathbb{R}^n \mid x \text{ satisfies } \Phi'\}$. Let D_1, \dots, D_m denote the disequality literals appearing in Φ . Let $H_i = \{x \in \mathbb{R}^n \mid x \text{ does not satisfy } D_i\}$, $1 \leq i \leq m$, and note that each H_i is a hyperplane.

Observation 1. The formula $\Phi^- \equiv \Phi \wedge D_1 \wedge \dots \wedge D_m$ is satisfiable.

Otherwise, $S \subseteq H_1 \cup \dots \cup H_m$. The set S is essentially convex and each H_m is a variety, so there exists an $1 \leq i \leq m$ such that $S \subseteq H_i$ by Proposition 15. Consequently, $\Phi \wedge D_i$ is not satisfiable which contradicts the fact that Φ has property $(*)$.

Observation 2. For every $\epsilon > 0$ there is a $y \in S$ satisfying $|c^T w - c^T y| < \epsilon$.

Let $d(\cdot, \cdot)$ denote the Euclidean distance in \mathbb{R}^n , i.e., $d(a, b) = \sqrt{\sum_{i=1}^n (a_i - b_i)^2}$, and $\|\cdot\|$ the corresponding norm, i.e., $\|a\| = \sqrt{a^T a}$.

Arbitrarily choose a point z that satisfies Φ^- ; this is always possible by Observation 1. Consider the line segment L between z and w . Note the following: if H is a hyperplane in \mathbb{R}^n , then either H intersects L in at most one point or $L \subseteq H$. Also note that $w, z \in S'$ and S' is convex so $L \subseteq S'$. Arbitrarily choose a clause $C \in \Phi''$ and assume $C = (p_1(\bar{x}) \neq 0 \vee \dots \vee p_k(\bar{x}) \neq 0 \vee p_0(\bar{x}) \leq 0)$. Assume that there exist two distinct points $a, b \in L$ such that $p_1(a) = p_1(b) = 0$. If so, then every point $c \in L$ satisfies $p_1(c) = 0$. This is not possible since $z \in L$ satisfies $D_1 \wedge \dots \wedge D_m$, and in particular $p_1(z) \neq 0$. Hence, at most one point $c \in L$ satisfies $p(c) = 0$, and c is the only point in L

that potentially does not satisfy the clause C . This implies that only finitely many points in L do not satisfy Φ'' , and it follows that for every $\delta > 0$ there is a point $y \in S$ such that $d(w, y) < \delta$.

We proceed by showing that if $w, y \in \mathbb{R}^n$ and $d(w, y) = d$, then $|c^T w - c^T y| \leq \|c\| \cdot d$. This follows from the Cauchy-Schwarz inequality (that is, $|a^T b| \leq \|a\| \cdot \|b\|$ for vectors a, b in \mathbb{R}^n):

$$|c^T w - c^T y| = |c^T(w - y)| \leq \|c\| \cdot \|w - y\| = \|c\| \cdot d(w, y) = \|c\| \cdot d.$$

To find a vector y that satisfies $|c^T w - c^T y| < \epsilon$, we simply choose $y \in S$ such that $d(w, y) < \frac{\epsilon}{\|c\|}$; we know that such a y exists by the argument above.

If the algorithm outputs ‘unbounded’ in **Step 2**, then arbitrarily choose a sufficiently large number k and note that there exists a vector $w \in S'$ such that $c^T w \geq k$. By Observation 2, there exists a vector $y \in S$ such that $|c^T w - c^T y| < 1$. Hence, S has unbounded solutions, too.

Assume finally that the algorithm answers ‘optimum is arbitrarily close to K ’ in **Step 3**; Observation 2 immediately proves correctness in this case, too. \square

6 Open Problems

The most prominent open question is whether there are essentially convex relations S with a first-order definition in $(\mathbb{R}; *, +)$ such that $\text{CSP}((\Gamma_{\text{lin}}, S))$ is NP-hard. Resolving this question is probably difficult, since the following closely related problem is of unknown computational complexity:

Feasibility of Convex Polynomial Inequalities

INPUT: A set of variables V , a set of polynomial inequalities each of which defining a convex set; the coefficients of the polynomials are rational numbers where the numerators and denominators are represented in binary.

QUESTION: Is there a point in $\mathbb{R}^{|V|}$ that satisfies all inequalities?

One may note that the problems we have considered could be easier since they are defined over finite constraint languages. On the other hand, convexity is much more restrictive than essential convexity; moreover, we are only given polynomial inequalities (there are convex semi-algebraic relations that cannot be defined as the intersection of convex polynomial inequalities). Still, the computational complexity of the Feasibility Problem of Convex Polynomial Inequalities is open.

Convex semi-algebraic relations are of particular interest in the quest for efficiently solvable semi-algebraic constraint languages because of a conjectured link to semidefinite programming. Every semidefinite representable set is convex and semi-algebraic. Recently, Helton, Vinnikov and Nie showed that the converse statement is true in surprisingly many cases and conjectured that it remains true in general [18].

Acknowledgements

We would like to thank Johan Thapper for pointing out inaccuracies in an earlier version of the article, Frank-Olaf Schreyer for helpful discussions in the early stages of this work, and the referees for their detailed remarks. Manuel Bodirsky has received funding from the *European Research Council* under the European Community’s Seventh Framework Programme (FP7/2007-2013 Grant

Agreement no. 257039). Peter Jonsson is partially supported by the *Center for Industrial Information Technology* (CENIIT) under grant 04.01 and by the *Swedish Research Council* (VR) under grants 2006-4532 and 621-2009-4431.

References

1. L. Barto and M. Kozik. Constraint satisfaction problems of bounded width. In *Proceedings of FOCS*, pages 595–603, 2009.
2. S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry, 2nd edition*. Springer-Verlag, 2009.
3. A. Bemporad, K. Fukuda, and F. Torrisi. Convexity recognition of the union of polyhedra. *Computational Geometry*, 18(3):141–154, 2001.
4. J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*. Springer-Verlag, 1998.
5. M. Bodirsky. Constraint satisfaction problems with infinite templates. In H. Vollmer, editor, *Complexity of Constraints (a collection of survey articles)*, pages 196–228. Springer, LNCS 5250, 2008.
6. M. Bodirsky and J. Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2), 2009. An extended abstract appeared in the proceedings of STOC’08.
7. A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of LICS’03*, pages 321–330, Ottawa, Canada, 2003.
8. A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.
9. A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.
10. D. Cohen, P. Jeavons, P. Jonsson, and M. Koubarakis. Building tractable disjunctive constraints. *Journal of the ACM*, 47(5):826–853, 2000.
11. T. Drakengren and P. Jonsson. Computational complexity of temporal constraint problems. In *Handbook of Temporal Reasoning in Artificial Intelligence*, pages 197–218. Elsevier, 2005.
12. F. Dumortier, M. Gyssens, L. Vandeurenzen, and D. V. Gucht. On the decidability of semilinearity for semialgebraic sets and its implications for spatial databases. *Journal on Computer and System Sciences*, 58(3):535–571, 1999.
13. F. Dumortier, M. Gyssens, L. Vandeurenzen, and D. V. Gucht. On the decidability of semilinearity for semialgebraic sets and its implications for spatial databases - corrigendum. *Journal on Computer and System Sciences*, 59(3):557–562, 1999.
14. T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1999.
15. J. Ferrante and C. Rackoff. A decision procedure for the first order theory of real addition with order. *SIAM Journal on Computing*, 4(1):69–76, 1975.
16. M. Garey and D. Johnson. *A guide to NP-completeness*. CSLI Press, Stanford, 1978.
17. M. A. Goberna, V. Jornet, and M. M. L. Rodríguez. On linear systems containing strict inequalities. *Linear Algebra and its Applications*, 360:151 – 171, 2003.
18. J. W. Helton and J. Nie. Sufficient and necessary conditions for semidefinite representability of convex hulls and sets. *SIAM Journal on Optimization*, 20(2):759–791, 2009.
19. W. Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.
20. P. Jeavons, D. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.
21. P. Jonsson and C. Bäckström. A unifying approach to temporal constraint reasoning. *Artificial Intelligence*, 102(1):143–155, 1998.
22. M. Koubarakis. Tractable disjunctions of linear constraints: Basic results and applications to temporal reasoning. *Theoretical Computer Science*, 266:311–339, 2001.
23. D. Marker. *Model Theory: An Introduction*. Springer, New York, 2002.

24. D. Marker, Y. Peterzil, and A. Pillay. Additive reducts of real closed fields. *Journal of Symbolic Logic*, 57(1):109–117, 1992.