

NP-membership for the boundary-boundary art-gallery problem

Jack Stade

November 2025

Abstract

The boundary-boundary art-gallery problem asks, given a polygon P representing an art-gallery, for a minimal set of guards that can see the entire boundary of P (the wall of the art gallery), where the guards must be placed on the boundary. We show that this art-gallery variant is in NP. In order to prove this, we develop a constraint-propagation procedure for continuous constraint satisfaction problems where each constraint involves at most 2 variables.

The X-Y variant of the art-gallery problem is the one where the guards must lie in X and need to see all of Y. Each of X and Y can be either the vertices of the polygon, the boundary of the polygon, or the entire polygon, giving 9 different variants. Previously, it was known that X-vertex and vertex-Y variants are all NP-complete and that the point-point, point-boundary, and boundary-point variants are $\exists\mathbb{R}$ -complete [Abrahamsen, Adamaszek, and Miltzow, JACM 2021][Stade, SoCG 2025]. However, the boundary-boundary variant was only known to lie somewhere between NP and $\exists\mathbb{R}$.

The X-vertex and vertex-Y variants can be straightforwardly reduced to discrete set-cover instances. In contrast, we give example to show that a solution to an instance of the boundary-boundary art-gallery problem sometimes requires placing guards at irrational coordinates, so it unlikely that the problem can be easily discretized.

Contents

1	Introduction	1
1.1	Background and related work	1
1.2	Our results	3
2	The boundary-boundary art-gallery problem and $\mathcal{M} - 2\text{SAT}$	3
3	The continuous inference system	6
3.1	Inference with function composition	8
4	Bounded-depth refutations	11
4.1	Infinite compositions	11
4.2	Compressing paths	13
5	Algorithms for $\mathcal{M} - 2\text{SAT}$	14
6	Irrational Coordinates	17
7	Acknowledgments	19
8	Conclusion	23

1 Introduction

Given a plane polygon P , two points x and y in P are said to be visible to each other if the line segment between x and y is contained in P . The art-gallery problem [22] is a decision problem that asks, given a polygon P (called the art-gallery), whether there are n points in P (called the guards) such that every point in P is visible to at least one of the guards.

For $X, Y \in \{\text{Vertex, Boundary, Point}\}$, the X-Y art-gallery problem [5] is a variant of the art-gallery problem where the guards are restricted to lie in X , but only need to see Y . The original variant (as described above) is called the point-point variant under this system. The variant where the guards only need to see the boundary of the polygon is called the point-boundary variant.

All of these variants are at least NP-hard [17, 16], but it isn't clear that they are necessarily contained in NP. The coordinates of the guards can be arbitrary real numbers, and some instances require guards to be placed at irrational coordinates in order to be guarded optimally, even when the vertices of P are given as integers (see [1]).

The problem ETR (short for Existential Theory of the Reals) asks whether a statement of form:

$$\exists x_1, \dots, x_n \in \mathbb{R} : \Phi(x_1, \dots, x_n)$$

is true, where Φ is a boolean formula involving the signs of integer-coefficient polynomials in the variables $x_1, \dots, x_n$. The complexity class $\exists\mathbb{R}$ consists of problems that can be polynomial-time reduced to ETR. It is straightforward to see that $\text{NP} \subseteq \exists\mathbb{R}$. In 1988, Canny [10] established that $\exists\mathbb{R} \subseteq \text{PSPACE}$. It is unknown whether either inclusion is strict.

In 2018, Abrahamsen, Adamaszek and Miltzow [2] showed that the point-point and boundary-point art-gallery variants are $\exists\mathbb{R}$ -complete, meaning in particular that they are not contained in NP unless $\text{NP} = \exists\mathbb{R}$. Stade [28] later proved the point-boundary variant is also $\exists\mathbb{R}$ -complete. In contrast, all the vertex-Y and X-vertex variants are easily seen to be in NP. We show that the boundary-boundary variant is in NP, completing the classification of the complexities of X-Y art-gallery variants (at least up to $\text{NP} \stackrel{?}{=} \exists\mathbb{R}$). This is summarized in Table 1.

Variant	Complexity
Vertex-Y	NP[17]
X-Vertex	NP[17]
Point-Point	$\exists\mathbb{R}$ [2]
Boundary-Point	$\exists\mathbb{R}$ [2]
Point-Boundary	$\exists\mathbb{R}$ [28]
Boundary-Boundary	NP[this paper]

Table 1: Complexity of X-Y art-gallery variants

1.1 Background and related work

Biniaz, Maheshwari, Merrild, Mitchell, Odak, Polishchuk, Robson, Rysgaard, Schou, Shermer, Spalding-Jamieson, Svenning, and Zheng [8] recently gave a polynomial time algorithm for the *contiguous art gallery problem*, a variant of the point-boundary art-gallery where each guard can be responsible only for a connected section of the polygon boundary.

In addition to art-galleries, many other important problems are known to be $\exists\mathbb{R}$ -complete, notable examples including packing [4], drawing graphs with fixed edge lengths [23], realizability of line arrangements [21, 27], and training neural networks [3, 7]. A recent survey by Cardinal, Miltzow and Schaefer [24] lists roughly 150 $\exists\mathbb{R}$ -complete problems.

A constraint satisfaction problem (CSP) is a decision problem that asks whether there is an assignment of some variables that satisfies a set of constraints. A 2CSP is a CSP where each constraint involves at most 2 variables.

A *continuous constraint satisfaction problem* (CCSP) is a constraint satisfaction problem where the variables take values over a continuous domain, typically $\mathbb{R}$ or a subset thereof. The constraints are typically represented by equalities and inequalities involving polynomials in the variables. A 2CCSP is a CCSP where each constraint involves at most 2 variables.

Discrete CSPs are typically in NP and are frequently NP-complete. Similarly, CCSPs are typically in $\exists\mathbb{R}$ and are frequently $\exists\mathbb{R}$ -complete. Miltzow and Schmierrmann [19] showed that CCSPs typically become $\exists\mathbb{R}$ -hard if constraints of form $x + y = z$ can occur along with some nonlinear constraints. They conjecture that it is necessary to have a constraints that depends on 3 or more variables in order for a CCSP to be $\exists\mathbb{R}$ -hard, as current techniques seem to be unable to show $\exists\mathbb{R}$ -hardness for any 2CCSPs.

It has been known since at least 1967 [15] that boolean 2CSPs can be decided in polynomial time. On the other hand, ternary 2CSPs are already NP-hard [6]. In 2000, Beckert, Hahnle and Manya [6] gave a polynomial-time algorithm for solving 2CSPs where the variables lie in a lattice and the constraints are disjunctions of terms of form $x \geq c$ or $x \leq c$. Charatonik and Wrona [11] later gave a faster algorithm for these instances.

One very important example of a CCSP is linear programming. In 1983, Megiddo [18] gave a strongly polynomial time algorithm for determining feasibility of a linear programming instance where each constraint involves at most 2 variables. In 2023, Dadush, Koh, Natura, Olver, and Végé [12] gave a strongly polynomial time algorithm for determining optimality for these instances. The best known algorithms for general linear programming run in only weakly polynomial time.

A common technique for solving CSPs is *constraint propagation*: given an instance Γ , we can derive new constraints from the existing ones. A method for producing new constraints giving existing ones is called an *inference rule*. If we can derive the unsatisfiable clause $\emptyset$ from the original instance using some set of inference rules, then we conclude that the instance is unsatisfiable. A derivation of $\emptyset$ is called a *refutation* of the instance. A set of inference rules for a CSP is *refutation complete* if every unsatisfiable instance has a refutation.

In 1960, Davis and Putman [13] showed that the inference rule:

$$\frac{x \vee C \quad \neg x \vee D}{C \vee D}$$

is refutation-complete for boolean satisfiability. The notation here describes an inference rule (called binary resolution) that produces a constraint $C \vee D$ given constraints $x \vee C$ and $\neg x \vee D$, where x is a variable and C and D are (possibly empty) constraints. In general, there is no guarantee that a refutation has polynomial size or that one can be found efficiently if it exists. However, resolution leads to a polynomial-time algorithm for boolean 2SAT (see Krom 1967 [15]).

In 2000, Beckert, Hahnle and Manya [6] extended binary resolution to the case where the variables take values in a finite lattice N , and the constraints are disjunctions of terms of form $x \leq c$ or $x \geq c$, where x is a variable and $c \in N$. Specifically, they show that the inference rule:

$$\frac{(x \leq c) \vee C \quad (x \geq d) \vee D}{C \vee D} \quad \text{where } d > c$$

is refutation complete for these CSPs. In the case where each constraint involves at most 2 variables, they show that there is a polynomial-time algorithm for finding refutations if they exist.

A family of techniques for solving CCSPs is quantifier elimination: given some subset $S \subset \{x_1, \dots, x_n\}$, we can try to find conditions that determine whether a given assignment of the variables in S extends to a satisfying assignment of the instance. The variables in $\{x_1, \dots, x_n\} \setminus S$ are said to be eliminated. If we can eliminate all the variables, then we can decide the truth value of the instance.

In 1948, Tarski [29] gave the first quantifier-elimination procedure for ETR. Given a formula:

$$\exists x_n : \Phi(x_1, \dots, x_n)$$

he shows how to compute an equivalent quantifier-free formula $\Psi(x_1, \dots, x_{n-1})$. The formula Ψ represents a sort of algorithm for determining if an assignment of the variables $x_1, \dots, x_{n-1}$ can be extended to an assignment of $x_1, \dots, x_n$ that satisfies Φ . The procedure can be repeated to eliminate any subset of the variables, though each elimination causes an exponential blow-up in the size of the formula.

Even if computing it is infeasible, the existence of a quantifier elimination can sometimes be used to bound the complexity of solutions if they exist. For example, if we eliminate all but one variable x using Tarski's quantifier elimination procedure for ETR, then we obtain a formula only involving polynomials in x , showing that if a satisfying assignment exists, then one exists where x is an algebraic number. More usefully, Schaefer and Stefankovič [25] use the existence of certain a certain quantifier-elimination to show that ETR is $\exists\mathbb{R}$ -complete even when the variables are restricted to the range $[-1, 1]$.

1.2 Our results

We show that the boundary-boundary art-gallery problem is in NP. The proof is by a nondeterministic reduction to a particular 2CCSP problem that we call $\mathcal{M} - 2\text{SAT}$, which is a 2CCSP where the constraints are given by (piecewise) fractional-linear functions. Most of the paper is devoted to showing that $\mathcal{M} - 2\text{SAT}$ is in NP. The strategy is to show that, if a solution exists, then there is a solution where the variables have form $p + q\sqrt{r}$, where p , q and r are rational numbers with a polynomially-bounded number of bits. Such solutions can be verified efficiently.

In order to prove this, we define a broad class of monotone CCSPs that contains $\mathcal{M} - 2\text{SAT}$. In Section 3, we develop an inference system for this class based on function composition, and show that it is refutation-complete. This is proved by using a compactness argument to show that it is sufficient to look for refutations over some finite subset of $\mathbb{R}$. The results of [6] imply that binary resolution is refutation complete over this discrete subset, and we show that a refutation using binary resolution over a finite subset of the domain can be lifted to one using function composition.

The nature of a compactness argument means that we don't get any useful bound on the length of refutations needed by our system. So in Section 4 we extend our inference system to allow infinite compositions and n -way minimum operations. In the new inference system, we prove that it is sufficient to consider refutations that in some sense have logarithmic depth.

In Section 5, we use these results to give two upper bounds for $\mathcal{M} - 2\text{SAT}$. First, we show that $\mathcal{M} - 2\text{SAT} \in \text{NP}$. Second, we show that there is a quasi-polynomial-time algorithm for $\mathcal{M} - 2\text{SAT}$.

The problem $\mathcal{M} - 2\text{SAT}$ will contain a range constraint $x \in [a, b]$ for each variable. The construction in Section 4 is used to construct a sort of circuit that determines the satisfiability of the instance given the values of a and b . If we set $a = b = c$, then this circuit determines whether $x = c$ extends to a satisfying assignment of the instance, giving a method of quantifier elimination. We can bound the complexity of arithmetic operations performed by the circuit, giving us a bound on the complexity of solutions needed. We obtain $\mathcal{M} - 2\text{SAT} \in \text{NP}$.

The quasi-polynomial-time algorithm for $\mathcal{M} - 2\text{SAT}$ is essentially straightforward to obtain from the results of Section 4.

It is easy to construct instances of $\mathcal{M} - 2\text{SAT}$ where a solution requires irrational coordinates, but it is not immediately clear if such an instance can be realized as an instance of the boundary-boundary art-gallery problem. In Section 6, we give an example to show that a solution to the boundary-boundary art-gallery problem can require irrational coordinates.

2 The boundary-boundary art-gallery problem and $\mathcal{M} - 2\text{SAT}$

Let $\mathcal{F}$ be a class of continuous increasing bijections $\mathbb{R} \rightarrow \mathbb{R}$. Then an instance of $\mathcal{F} - 2\text{SAT}$ consists of a set of variables $v_1, \dots, v_n$, a closed interval $\text{range}(v_i)$ for each variable, and a set of constraints of form $x \leq f(y)$ where x, y are *literals*, that is elements of $\{\pm v_1, \dots, \pm v_n\}$. The problem asks whether there is an assignment of real values $v_i \in \text{range}(v_i)$ that satisfy all the constraints. We write $\text{range}(-v_i) = -\text{range}(v_i)$.




Figure 1: Some examples of constraints of the form $x \leq f(y)$ for different f .

Figure 1 shows some examples of the types of sets that can occur. Even though the variable v_i will only every take values in $\text{range}(v_i)$, it is helpful to have the functions defined on all of $\mathbb{R}$ so that inverses and function compositions are well-defined.

In many cases, the constraint $x \leq f(y)$ already restricts x or y to a smaller interval. We could try to restrict the variables to this range as a pre-processing step. However, restricting the range of one variable might then required restricting the range of another variable, and this can lead to an infinite loop. This process does converge in some sense, but it isn't clear that we can compute the limit.

If the functions in $\mathcal{F}$ are very complicated, for example polynomial circuits, then there is a little hope of being able to decide $\mathcal{F} - 2\text{SAT}$ efficiently. The specific class of functions that we are most interested in are *increasing continuous piecewise fractional linear*, that is (increasing, continuous) functions that are piecewise of form:

$$f(x) = \frac{ax + b}{cx + d}$$

where a, b, c and d are integers. We let $\mathcal{M}$ be this class of functions, so $\mathcal{M} - 2\text{SAT}$ is the 2CCSP using functions from $\mathcal{M}$ as constraints.

Lemma 1. *The boundary-boundary art-gallery art-gallery problem reduces in NP to $\mathcal{M} - 2\text{SAT}$*

Proof. If P is an art-gallery with n vertices, then the boundary of P can be easily guarded with n (or even $\lceil \frac{n}{2} \rceil$) boundary guards. In a guarding configuration of P , the portion of the boundary that can be seen by any given guard is composed of at most n connected regions. This is because any two different connected




Figure 2: Left: the portion of the boundary seen by a single guard is a union of at most n connected pieces, since any two different pieces are separated by a vertex of the polygon. There are at most n guards, so in a guarding configuration, each edge is covered by a union of at most n^2 intervals that are each visible to a single guard. Right: we guess that a guard at position $\mathbf{ax} + \mathbf{b}$ sees the interval from $\mathbf{cy} + \mathbf{d}$ to $\mathbf{cz} + \mathbf{d}$. In order for this to be true, the vertex $\mathbf{v}$ of P must be on the left side of the $\mathbf{ax} + \mathbf{b}$ to $\mathbf{cy} + \mathbf{d}$, creating a fractional-linear constraint.

components of the boundary region must be separated by a vertex of P , as illustrated in Figure 2 (left). So in a guarding configuration, each edge of P is covered by a set of at most n^2 intervals that are each guarded by a single guard.

First, we nondeterministically guess the structure of a guarding configuration. That is, for each guard, we guess which edge of P it is on, and for each edge of P we guess a partition of it into at most n^2 intervals, and for each of these intervals we guess which guard should guard it. The exact locations of each guard and the endpoints of the guarding intervals are not guessed, but will become variables in an instance of $\mathcal{M} - 2\text{SAT}$.

For each edge of P , choose vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^2$ such that points on that edge have form $\mathbf{ax} + \mathbf{b}$ for $x \in [0, 1]$. We now write down constraints on the positions of the guards and the endpoints of guarding intervals that determine whether an assignment of the positions represents a valid guarding configuration. Figure 2 (right) illustrates an example.

In order for the guard at position $\mathbf{ax} + \mathbf{b}$ to see the line segment from $\mathbf{cy} + \mathbf{d}$ to $\mathbf{cz} + \mathbf{d}$, we need to know that the segments from $\mathbf{ax} + \mathbf{b}$ to $\mathbf{cy} + \mathbf{d}$ or $\mathbf{cz} + \mathbf{d}$ are contained in the polygon P . For example, the vertex marked $\mathbf{v}$ in the figure should be on the left side of the line from $\mathbf{ax} + \mathbf{b}$ to $\mathbf{cy} + \mathbf{d}$. This condition can be written:

$$(\mathbf{ax} + \mathbf{b} - \mathbf{v}) \times (\mathbf{cy} + \mathbf{d} - \mathbf{v}) \geq 0 \quad (1)$$

Here $\times$ represents cross product. The cross product of two vectors $\mathbf{p}$ and $\mathbf{q}$ in $\mathbb{R}^2$ is a scalar value given by the determinant of the matrix with columns $\mathbf{p}$ and $\mathbf{q}$. The condition (1) is equivalent to:

$$\mathbf{a} \times (\mathbf{cy} + \mathbf{d} - \mathbf{v})x \geq -(\mathbf{b} - \mathbf{v}) \times (\mathbf{cy} + \mathbf{d} - \mathbf{v})$$

Let $a = -(\mathbf{b} - \mathbf{v}) \times \mathbf{c}$, $b = -(\mathbf{b} - \mathbf{v}) \times (\mathbf{d} - \mathbf{v})$, $c = \mathbf{a} \times \mathbf{c}$ and $d = \mathbf{a} \times (\mathbf{d} - \mathbf{v})$. Depending on the sign of $\mathbf{a} \times (\mathbf{cy} + \mathbf{d} - \mathbf{v})$ (that is, $cy + d$), we want to create the constraint either $x \geq \frac{ay+b}{cy+d}$ or $x \leq \frac{ay+b}{cy+d}$. In the example in Figure 2 (right), $cy + d$ is always positive, but in general it could have different signs for different values of y .

We define a function f that is equal to $\frac{ay+b}{cy+d}$ when $cy + d \geq 0$ and $0 \leq ay + b \leq cy + d$ (so $\frac{ay+b}{cy+d} \in [0, 1]$ when y satisfies these conditions). The conditions $cy + d \geq 0$ and $0 \leq ay + b \leq cy + d$ are convex, so the set of values of y satisfying them is an interval if it is non-empty. If this interval is nonempty, then extend f to

a monotone function $\mathbb{R} \rightarrow \mathbb{R}$ by attaching lines of slope ± 1 outside of the interval, and add the constraint $x \geq f(y)$.

Similarly, let g be a monotone function $\mathbb{R} \rightarrow \mathbb{R}$ that is equal to $\frac{ay+b}{cy+d}$ when $cy+d \leq 0$ and $cy+d \leq ay+b \leq 0$ (assuming that such values of y exist), and add the constraint $x \leq g(y)$. The constraints $x \geq f(y)$ and $x \leq g(y)$ are equivalent to (1).

We want the constraints to have form $x \leq f(y)$ for an increasing function f . If $ad - bc$ is positive, then f and g are increasing, and we can write $x \geq f(y)$ and $-x \leq \tilde{f}(-y)$, where $\tilde{f}(z) = -f(-z)$. If $ad - bc$ is negative, then f and g are decreasing, and we can write $x \geq f(y)$ as $-x \leq -f(y)$ and $x \leq g(y)$ and $x \leq \hat{g}(-y)$ where $\hat{g}(z) = g(-z)$. If $ad - bc = 0$, then $\frac{ay+b}{cy+d}$ doesn't depend on y , and we adjust the range of x appropriately.

All of the necessary constraints can be written in a similar way, and they will all be piecewise-fractional-linear. We obtain an instance of $\mathcal{M} - 2\text{SAT}$. $\square$

So if we can show that $\mathcal{M} - 2\text{SAT}$ is in NP, then we will show that the boundary-boundary art-gallery problem is in NP (the certificate consists of values of the guesses needed in the proof of Lemma 1 and a certificate for the $\mathcal{M} - 2\text{SAT}$ instance generated). Most of the rest of this paper will be devoted to proving that $\mathcal{M} - 2\text{SAT}$ is in NP. This is not immediately clear, since irrational values are sometimes required. For example, the instance:

$$x \leq \frac{x+p}{x+1}, x \leq \frac{x-p}{-x+1}$$

is satisfiable for p positive, but the only solutions are $x = \pm\sqrt{p}$, which can be irrational. What we will show is that if an instance is satisfiable, then there is a solution where each variable is of form $p + q\sqrt{r}$, where p , q and r are rational numbers with at most polynomially many bits.

3 The continuous inference system

In this section, we construct an inference system for $\mathcal{F} - 2\text{SAT}$ instances and prove that it is refutation complete.

Throughout this section, we work with n real variables $v_1, \dots, v_n$, where v_i is restricted to a compact interval $\text{range}(v_i) \subset \mathbb{R}$. In this section, a *formula* is a conjunction *clauses* that are each a conjunction of at most 2 terms of form $v_i \leq c$ or $v_i \geq c$ for some $c \in \mathbb{R}$. The empty clause $\emptyset$ is unsatisfiable. A *continuous formula* is a conjunction of finitely many *constraints*, where a constraint is of form $x \leq f(y)$ for some literals x and y , where f is a continuous strictly increasing bijection $\mathbb{R} \rightarrow \mathbb{R}$.

Theorem 2. (Beckert, Hahnle and Many, [6]) *Let N be a lattice (that is, a partially ordered set where every pair of elements has a unique least upper bound and greatest lower bound), let $v_1, \dots, v_n$ be variables taking values in N , and let Γ be a conjunction of (finitely many) clauses each containing literals of form $x_i \leq c$ or $x_i \geq c$ for $c \in N$. If Γ is unsatisfiable, then the empty clause can be derived from Γ using the binary resolution rule:*

$$\frac{(x_1 \leq a) \vee T_1 \quad (x_1 \geq b) \vee T_2}{T_1 \vee T_2} \text{ where } a < b$$

The authors of [6] restricted to the case where N is finite, but the proof works for N infinite as long as the number of clauses is finite (we can just restrict to the finite sublattice generated by the elements appearing in the Γ). In the case where $N \subset \mathbb{R}$ is a compact interval, we can use a compactness argument to extend to the case where the number of clauses is infinite.

Lemma 3. *For each real variable v_i , let $\text{range}(v_i) \subset \mathbb{R}$ be a compact interval. Suppose Ξ is a formula consisting of a (possibly infinite) conjunction of disjunctions of terms $v_i \leq c$ or $v_i \geq c$ with $c \in \text{range}(v_i)$. Then $\emptyset$ can be derived from Ξ with (finitely many) applications of binary resolution.*

Proof. Let $D \subset \mathbb{R}^n$ be the set of values $(v_1, \dots, v_n)$ with each v_i in $\text{range}(v_i)$. Since each $\text{range}(v_i)$ is compact, D is a product of compact spaces, so is compact.

For each clause C in Ξ , the set of points $S(C)$ in D satisfying C is closed. Since Ξ is unsatisfiable, the intersection:

$$\bigcap \{S(C) : C \text{ a clause in } \Xi\}$$

is empty. So:

$$\bigcup \{D \setminus S(C) : C \text{ a clause in } \Xi\} = D \setminus \bigcap \{S(C) : C \text{ a clause in } \Xi\} = D$$

Since each $D \setminus S(C)$ is open and D is compact, there is a finite set of the $D \setminus S(C)$ that cover D , and so the intersection of these $S(C)$ is empty. That is, there is a finite subformula Ξ' of Ξ that is not satisfiable.

By Theorem 2, binary resolution can be used to derive $\emptyset$ from Ξ' . Since Ξ' is a subformula of Ξ , this gives a derivation of $\emptyset$ from Ξ . $\square$

We now write $x \leq f(y)$ in terms of infinitely many clauses involving terms of form $x \leq c$ or $y \geq d$.

Lemma 4. *Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous increasing bijection. Values x and y satisfy $x \leq f(y)$ if and only if they satisfy $(x \leq f(c)) \vee (y \geq c)$ for each $c \in \mathbb{R}$.*

Proof. First, suppose $x > f(y)$. By continuity of f , there is some $c > y$ such that $x > f(c)$. So x and y do not satisfy $(x \leq f(c)) \vee (y \geq c)$.

Now suppose $x \leq f(y)$ and let $c \in \mathbb{R}$. If $y < c$, then $f(y) < f(c)$, so $x \leq f(c)$. So at least one of $(x \leq f(c))$ or $(y \geq c)$ is satisfied. $\square$

We next want to define an infinite set of clauses $C(x, y, f)$ associated to the constraint $x \leq f(y)$. In order to eventually use Theorem 2, we don't want to have a clause of form $x \leq c$ or $x \geq c$ unless $c \in \text{range}(x)$. So we take all the clauses $(x \leq f(c)) \vee (y \geq c)$ whenever $f(c)$ is in $\text{range}(x)$ and c is in $\text{range}(y)$. So $C(x, y, f)$ contains the clauses:

$$\{(x \leq f(c)) \vee (y \geq c) : c \in \text{range}(y) \cap f^{-1}(\text{range}(x))\}$$

We also need some additional clauses corresponding to the endpoints of $\text{range}(x)$ and $\text{range}(y)$. If $f(\max(\text{range}(y))) \in \text{range}(x)$, then we include $x \leq f(\max(\text{range}(y)))$. If $f^{-1}(\min(\text{range}(x))) \in \text{range}(y)$, then we include $y \geq f^{-1}(\min(\text{range}(x)))$.

If $f(\max(\text{range}(y))) < \min(\text{range}(x))$ or $f^{-1}(\min(\text{range}(x))) > \max(\text{range}(y))$, then we add the empty clause $\emptyset$ to $C(x, y, f)$, since the constraint $x \leq f(y)$ is unsatisfiable for $x \in \text{range}(x)$ and $y \in \text{range}(y)$.

Lemma 5. *Values x and y with $x \in \text{range}(x)$ and $y \in \text{range}(y)$ satisfy $x \leq f(y)$ if and only if they satisfy all the constraints in $C(x, y, f)$.*

Proof. First suppose x and y satisfy $x \leq f(y)$. By Lemma 4, all the $(x \leq f(c)) \vee (y \geq c)$ clauses are satisfied. Since $y \in \text{range}(y)$, we have $x \leq f(\max(\text{range}(y)))$. Since $x \in \text{range}(x)$, we have $f(y) \geq \min(\text{range}(x))$, so $y \geq f^{-1}(\min(\text{range}(x)))$.

Now suppose that x and y are such that $x > f(y)$. By Lemma 4, $(x \leq f(c)) \vee (y \geq c)$ fails for some $c \in \mathbb{R}$. If this value of c is in $\text{range}(y) \cap f^{-1}(\text{range}(x))$, then we are done.

Suppose $(x \leq f(c)) \vee (y \geq c)$ fails for some $c \notin \text{range}(y)$. Since $y \geq c$, we have $c > \max(\text{range}(y))$. Since $x > f(c)$, $x > f(\max(\text{range}(y)))$, which fails to satisfy $x \leq f(\max(\text{range}(y)))$.

Similarly, if $(x \leq f(c)) \vee (y \geq c)$ fails for some $f(c) \notin \text{range}(x)$, then $f(c) < \min(\text{range}(x))$ so $y < f^{-1}(\min(\text{range}(x)))$. $\square$

If Γ is a continuous formula, then we define the associated formula Ξ to be formula consisting of the disjunction of all the clauses in $C(x, y, f)$ for each x, y, f such that $x \leq f(y)$ is a constraint in Γ .




Figure 3: The case where the third inference rule is needed. There appear to be points that satisfy both constraints, but none where $x = -(-x)$. The other inference rules don't know about the relationship between x and $-x$ as literals, so aren't powerful enough to refute this case.

3.1 Inference with function composition

The *continuous inference rules* are:

$$\frac{x \leq f(y) \quad y \leq g(z)}{x \leq f(g(z))}$$

$$\frac{x \leq f(y)}{\emptyset} \text{ where } f(\max(\text{range}(y))) < \min(\text{range}(x))$$

$$\frac{x \leq f(-x) \quad -x \leq g(x)}{\emptyset} \text{ where } \exists c : f(-c) < c \text{ and } g(c) < -c$$

Lemma 6. *Let Γ be a continuous formula and let Γ' be a continuous formula that can be obtained from Γ by one of the continuous inference rules. Then Γ is satisfiable if and only if Γ' is.*

Proof. Γ' contains all the clauses in Γ , so Γ' is unsatisfiable if Γ is.

If Γ is satisfiable, then there is some assignment of the variables satisfying all the clauses. We want to show that the same variable assignment satisfies Γ' . If x, y and z satisfy $x \leq f(y)$ and $y \leq g(z)$, then by monotonicity of f , $f(y) \leq f(g(z))$, so $x \leq f(g(z))$. So the first inference rule is valid.

If $x \leq f(y)$ and $y \in \text{range}(y)$, then $y \leq \max(\text{range}(y))$ so $x \leq f(\max(\text{range}(y)))$. So if $f(\max(\text{range}(y))) < \min(\text{range}(x))$ then x is not in $\text{range}(x)$. So if Γ is satisfiable then the second inference rule can never be applied.

If $x \leq f(-x)$ and $-x \leq g(x)$ and there is some c such that $f(-c) < c$ and $g(c) < -c$, then $0 \leq f(-x) - x$ and $0 \leq g(x) + x$. $f(-x) - x$ is increasing in $-x$ and $g(x) + x$ is increasing in x , so $-x > -c$ and $x > c$, which is a contradiction. This case is shown in Figure 3. So if Γ is satisfiable then the third inference rule can never be applied. □

If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly increasing continuous bijection, then f has a well-defined continuous inverse. The constraint $x \leq f(y)$ could equivalently be written $f^{-1}(x) \leq y$. Define $\tilde{f} : \mathbb{R} \rightarrow \mathbb{R}$ by $\tilde{f}(x) = -f^{-1}(-x)$. So $\tilde{f}$ is a strictly increasing continuous bijection and $x \leq f(y)$ is equivalent to $-y \leq \tilde{f}(-x)$. Note that $C(x, y, f)$ contains exactly the same clauses as $C(-y, -x, \tilde{f})$. We say that a continuous formula Γ is *symmetric* if, for every constraint of the form $x \leq f(y)$, the constraint $-y \leq \tilde{f}(-x)$ also appears.

Lemma 7. *If Γ is symmetric and Γ' can be derived from Γ , then a symmetric formula containing Γ' can be obtained from Γ' .*

Proof.

$$\widetilde{f \circ g} = \widetilde{g} \circ \widetilde{f}$$

□

We can now state the main result of this section:

Lemma 8. *Let Γ be an unsatisfiable symmetric continuous formula. Then the empty constraint can be derived from Γ by the continuous inference rules.*

To prove Lemma 8, we will use Lemma 3 to show that if Γ is unsatisfiable, then the empty clause can be derived from the constraints in the $C(x, y, f)$ by binary resolution. If a clause C can be derived from clauses in $C(x, y, f)$ and $C(y, z, g)$, then we would like to argue that C is contained in $C(x, z, f \circ g)$. However, this is not quite true.

For a clause C , we say that another clause C' is *stronger than* C if the set of values satisfying C' is contained in the set of values satisfying C . Roughly speaking, the strategy to prove Lemma 8 is to show that, if a clause C can be derived from $C(x, y, f)$ and $C(y, z, g)$, then a stronger clause C' is contained in $C(x, z, f \circ g)$.

However, this creates a problem. The inference system in [6] implicitly allows a clause of form $(x \leq c) \vee (x \leq c)$ to be contracted to $x \leq c$. However, the stronger clause $(x \leq c) \vee (x \leq d)$ for $d < c$ can't be contracted in this way. This is a problem for us because it means that derivations behave badly with respect to strengthening clauses. A derivation starting with a formula including $(x \leq c) \vee (x \leq c)$ might need fewer steps than one where $(x \leq c) \vee (x \leq c)$ is replaced by $(x \leq c) \vee (x \leq d)$ for $d < c$. We instead use an inference system with *strong contractions*, where $(x \leq c) \vee (x \leq d)$ to be contracted to $x \leq c$ whenever $c \geq d$.

Lemma 9. *Let Γ be a formula that can be refuted in k steps of binary resolution, allowing strong contractions. If Γ' is obtained from Γ by replacing a clause C with a stronger clause C' , then there is a derivation of the empty clause from Γ' requiring the same number of steps, again allowing strong contractions.*

Proof. Straightforward. □

Proof of Lemma 8. Let Ξ be formula consisting of the disjunction of all the $C(x, y, f)$ for each $x \leq f(y)$ that occurs in Γ . By Lemma 5, Ξ is unsatisfiable, so by Lemma 3, there is a derivation of the empty clause from Ξ . We will show that if a clause C that can be obtained from Ξ by binary resolution, then we can derive some $x \leq f(y)$ from Γ such that $C(x, y, f)$ contains a clause stronger than C . By induction and Lemma 9, this will show that the empty clause can be derived from Γ .

Each clause in Ξ has form either $(x \leq f(c)) \vee (y \geq c)$, $x \leq f(\max(\text{range}(y)))$, or $x \geq f^{-1}(\min(\text{range}(y)))$ for some constraint $x \leq f(y)$ in Γ . The clause $x \geq f^{-1}(\min(\text{range}(y)))$ can be written $-x \leq \widetilde{f}(\max(\text{range}(-y)))$. Since Γ is symmetric we will always use the later form of this clause.

Let C be a clause that can be derived directly from Ξ by binary resolution. There are three cases to consider, depending on whether C has 0, 1, or 2 terms.

If C has two terms, then it is produced by a step that can be written:

$$\frac{(y \leq a) \vee (w_1 \leq c_1) \quad (y \geq b) \vee (w_2 \leq c_2)}{(w_1 \leq c_1) \vee (w_2 \leq c_2)}$$

for $b > a$ and some literals x and z . Since Γ is symmetric, the clause of form $(w \leq f(c)) \vee (y \geq c)$ is also contained in $C(-y, -w, \widetilde{f})$ in the form $(-y \leq \widetilde{f}(-f(c))) \vee (-w \geq -f(c)) = (-y \leq -c) \vee (-w \geq -f(c))$. By choosing the constraints containing $(y \leq a)$ and $(y \geq b)$ appropriately, we can write the step as:

$$\frac{(y \leq g(d)) \vee (z \geq d) \quad (x \leq f(c)) \vee (y \geq c)}{(z \geq d) \vee (x \leq f(c))}$$

for some literals y and z . We can derive $x \leq (f \circ g)(z)$ from $x \leq f(y)$ and $y \leq g(z)$. The constraint $C(x, z, f \circ g)$ contains $(x \leq f(g(d))) \vee (z \geq d)$. Since $g(d) > c$, this is stronger than $(x \leq f(c)) \vee (z \geq d)$.

Next we consider the case when C has 1 term. So one of the clauses input into the step has 1 term and the other has 2 terms. There are two ways for Ξ to contain a clause with only one term. Clauses of form $x \leq f(\max(\text{range}(y)))$ appear directly in Ξ , but there are also contractions of a clauses $(y \leq f(c)) \vee (-y \geq c)$ contained in a constraint like $C(y, -y, f)$.

Consider a step starting with one clause of form $y \leq g(\max(\text{range}(y)))$ and one clause of form $(x \leq f(c)) \vee (y \geq c)$ or $(-y \leq f(c)) \vee (x \geq c)$. By replacing f with $\tilde{f}$, we may assume that the second clause has form $(x \leq f(c)) \vee (y \geq c)$, so this step can be written:

$$\frac{(x \leq f(c)) \vee (y \geq c) \quad (y \leq g(\max(\text{range}(z))))}{x \leq f(c)}$$

The constraint $x \leq (f \circ g)(z)$ can be derived from $x \leq f(y)$ and $y \leq g(z)$. Since $g(\max(\text{range}(z))) < c$, the clause $x \leq (f \circ g)(\max(\text{range}(z)))$ is stronger than $x \leq f(c)$.

Now we consider the case where a term like $y \leq a$ appears as a contraction of a clause like $(y \leq g(c)) \vee (-y \geq c) = (y \leq g(c)) \vee (y \leq -c)$ in $C(y, -y, h)$. Here a is the larger of $-c$ and $h(c)$. If $-c < h(c)$ then there is an equivalent clause $(y \leq \tilde{h}(d)) \vee (-y \geq d)$ in $C(-y, y, \tilde{h})$. Here $d = -g(c)$, so $-d > \tilde{h}(d)$. By choosing g to be h or $\tilde{h}$ appropriately, we have either $a = -c$ or $a = -d$ with $g(-a) \leq a$.

A step producing a 1-term clause that uses a contraction can therefore always be written:

$$\frac{(x \leq f(b)) \vee (y \geq b) \quad y \leq a}{x \leq f(b)}$$

From $x \leq f(y)$ and $y \leq g(-y)$ we derive $x \leq f(g(-y))$. From $-y \leq \tilde{f}(-x)$ we derive $x \leq f(g(\tilde{f}(-x)))$. $C(x, -x, f \circ g \circ \tilde{f})$ contains the clause $(x \leq f(g(\tilde{f}(-f(b)))) \vee (-x \geq -f(b)))$. Now $\tilde{f}(-f(b)) = -b < -a$, so $g(\tilde{f}(-f(b))) \leq a < b$. So $f(g(\tilde{f}(-f(b)))) \leq f(b)$, and so $(x \leq f(g(\tilde{f}(-f(b)))) \vee (-x \geq -f(b)))$ contracts to $(x \leq f(g(\tilde{f}(-f(b))))$, which is stronger than $x \leq f(b)$.

Finally, we consider the case where $C = \emptyset$. This step can always be written:

$$\frac{y \leq a \vee y \geq b}{\emptyset}$$

There are 3 cases to consider here, depending on whether 0, 1, or 2 of the inputs come from a contractions. First consider the case where neither inputs comes from a contraction, yielding:

$$\frac{x \leq f(\max(\text{range}(y))) \quad -x \leq g(\max(\text{range}(z)))}{\emptyset}$$

The constraint $-x \leq g(z)$ has dual $-z \leq \tilde{g}(x)$, and $-z \leq \tilde{g}(f(y))$ can be derived from $x \leq f(y)$ and $-z \leq \tilde{g}(x)$. Since $f(\max(\text{range}(y))) < -g(\max(\text{range}(z)))$, we have $\tilde{g}(f(\max(\text{range}(y)))) < -\max(\text{range}(z)) = \min(\text{range}(-z))$. So $\emptyset$ can be derived from $-z \leq \tilde{g}(f(y))$.

If one of the inputs comes from a contraction, then the step can be written:

$$\frac{x \leq f(\max(\text{range}(y))) \quad -x \leq c}{\emptyset}$$

where $-x \leq c$ comes from a clause $(-x \leq g(-c)) \vee (x \geq -c)$ in $C(-x, x, g)$ with $g(-c) \leq c$. $-x \leq g(f(y))$ can be derived from $-x \leq g(x)$ and $x \leq f(y)$. Since $f(\max(\text{range}(y))) < -c$, $g(f(\max(\text{range}(y)))) \leq c$. So we now have $x \leq f(y)$ and $-x \leq g(f(y))$ where $f(\max(\text{range}(y))) < -g(f(\max(\text{range}(y))))$, which reduces to the previous case.

Finally, we have the case where both clauses come from contractions. A step of this form can be written:

$$\frac{x \leq a \quad -x \leq -b}{\emptyset}$$

where $x \leq a$ is contained in $C(x, -x, f)$ and $-x \leq -b$ is contained in $C(-x, x, g)$. By symmetry, we can suppose that $f(-a) \leq a$ and $g(b) \leq -b$. We have $a < b$, so let $c \in (a, b)$. $f(-x) - x$ is strictly increasing in

$-x$ and $g(x) + x$ is strictly increasing in x . Since $-c < -a$, $c < b$ and $f(-a) - a \leq 0$ and $g(b) + b \leq 0$, we have $f(-c) - c < 0$ and $g(c) + c < 0$, so $f(-c) < c$ and $g(c) < -c$. So $\emptyset$ can be derived from $x \leq f(-x)$ and $-x \leq g(x)$. $\square$

4 Bounded-depth refutations

In general, there is no way to bound the length of the refutation produced by Lemma 8. However, if a refutation is too large, then we might be able to simplify it. For example, a refutation might generate many different constraints of form $x \leq f(y)$, in which case we can replace constraints $x \leq f_1(y), \dots, x \leq f_k(y)$ with a single constraint $x \leq \min(f_1, \dots, f_k)(y)$.

If a constraint of form $x \leq f(x)$ appears, then we can obtain $x \leq (f \circ \dots \circ f)(x)$. We may need many compositions of f in order to obtain a strong enough constraint, so this is another way that the length of a refutation can become very large. In these cases, we can compute (roughly) the limit of f^k as k goes to ∞ . As we show in Section 4.1, the constraint $x \leq f^\infty(x)$ and the tautology $x \leq x$ are together stronger than any constraint of form $x \leq f^k(x)$, and so we are saved from having to compute an arbitrarily large number of compositions of f .

In this section, we show that the addition of k -way minimum operations and infinite compositions allow us to produce a refutations where the depth of arithmetic operations required is at most logarithmic. For every pair of literals x and y , this refutation tries to calculate the function $\min_{x \rightarrow y}$ representing the infimum over all f such that the constraint $x \leq f(y)$ can be derived.

Given an $\mathcal{F}$ -2SAT instance Γ , let G be a graph with vertices representing literals $\pm v_i$ and directed edges representing constraints, so that an edge from x to y represents a constraint of form $x \leq f(y)$ in Γ . The constraints $x \leq f(y)$ that can be derived from Γ correspond to paths from x to y in G .

If $x \leq f(y)$ is a constraint that can be derived from Γ using the composition, minimum, and infinite composition operations, then the sequence of operations used to define f represents (in some sense) a regular expression matching paths in G , and f is the infimum of the constraints represented by the paths matched by the regular expression. Given vertices x and y in G , it is possible (see Brzozowski and McCluskey [9]) to construct a regular expression matching all directed paths from x to y . However, the regular expression produced can be as long as n^n , which is not useful for us.

So instead, we construct a regular expression that matches the paths from x to y of length at most $2n$, so in particular it matches any non-repeating paths. We let $f_{x \rightarrow y}$ be the function so that this regular expression constructs the constraint $x \leq f_{x \rightarrow y}(y)$. In Section 4.2, we show that, for a real number c , $\min_{x \rightarrow y}(c)$ is either $f_{x \rightarrow y}(c)$ or $f_{x \rightarrow z}(d)$ for some literal z , where d is an attracting fixed point of $f_{z \rightarrow z}$. In the case of $\mathcal{M}$ -2SAT, this will give a bound on the complexity of coordinates required for a satisfying assignment, proving NP-membership for $\mathcal{M}$ -2SAT.

For $\mathcal{M}$ -2SAT, the refutations considered in Section 4.2 can be explicitly constructed in time $n^{\mathcal{O}(\log(n))}$, allowing us to solve $\mathcal{M}$ -2SAT in quasi-polynomial time. In particular, $\mathcal{M}$ -2SAT isn't NP-hard unless the exponential-time hypothesis is false. It seems to be difficult to perform these calculations in polynomial time, since the regular expression representing by $f_{x \rightarrow y}$ is a monotone formula for solving directed st-connectivity. Such a formula must have size at least $n^{\Omega(\log(n))}$ (see Karchmer and Wigderson [14]).

4.1 Infinite compositions

If a constraint of form $x \leq f(x)$ appears, then we can obtain $x \leq (f \circ \dots \circ f)(x)$. We may need many compositions of f in order to obtain a strong enough constraint. Our strategy to handle this case is to compute (roughly) the limit of $f^k(x)$ as k goes to ∞ . As we show below, the constraint $x \leq f^\infty(x)$ and the tautology $x \leq x$ are together stronger than any constraint of form $x \leq f^k(x)$, so we are saved from having to compute an arbitrarily large number of compositions of f .

Given a function f , we now define precisely what we mean by the function f^∞ . Our precise definition is someone subtle, because we want f^∞ to be piecewise-constant. This might fail if f has infinitely many fixed points, for example if f is a piecewise linear function where one piece is $f(x) = x$.




Figure 4: Illustration of f^∞

If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly increasing continuous bijection, then we define $f^\infty : \mathbb{R} \rightarrow \mathbb{R} \cup \{\pm\infty\}$ by:

$$f^\infty(x) = \inf\{y \geq \sup\{z \leq x : z \leq f(z)\} : y > f(y)\}$$

here taking $\inf\{\emptyset\} = +\infty$ and $\sup\{\emptyset\} = -\infty$.

Equivalently, we could define:

$$f^\infty(x) = \lim_{\epsilon \rightarrow 0^+} \lim_{k \rightarrow \infty} (f + \epsilon)^k(x)$$

but the first definition will be easier to work with.

We now prove some useful properties of this construction.

Lemma 10. *If $f(x) < x$, then $f^\infty(x) \leq f(x)$.*

Proof. If z is such that $f(x) < z \leq x$ then $f(z) \leq f(x) < z$. So $\sup\{z \leq x : z \leq f(z)\} \leq f(x)$. Since $f(x) < x$, $f(f(x)) < f(x)$, so $f^\infty(x) \leq f(x)$. $\square$

Lemma 11. *$x \leq f(x)$ if and only if $x \leq f^\infty(x)$.*

Proof. If $x \leq f(x)$, then $\sup\{z \leq x : z \leq f(z)\} = x$, so $f^\infty(x) \geq x$.

If $x > f(x)$, then by Lemma 10, $f^\infty(x) \leq f(x)$, so $x > f^\infty(x)$. $\square$

Lemma 12. *For all $x \in \mathbb{R}$, $f^\infty(x) = \pm\infty$ or $f^\infty(x)$ is a locally maximal fixed point of f . That is, $f^\infty(x) + \epsilon$ is not a fixed point of f for all sufficiently small $\epsilon > 0$.*

Also $f^\infty \circ f^\infty = f^\infty$.

Proof. It is clear from the definition that $f(f^\infty(x) + \epsilon) > f(f^\infty(x))$ for arbitrarily small $\epsilon > 0$ (when it is finite). By continuity of f , $f^\infty(x)$ is a fixed point of f .

It is straightforward from the definition that $f^\infty(x) = x$ when x is a locally maximal fixed point of f . So $f^\infty \circ f^\infty = f^\infty$. $\square$

The values taken by f^∞ are always fixed points of f . Those fixed points of f that can occur as values of f^∞ are called *attracting points* of f .

4.2 Compressing paths

Given two literals x and y , any function f such that $x \leq f(y)$ can be derived from a continuous formula Γ can be written $f = f_1 \circ \dots \circ f_k$ for functions f_i from the original problem instance. We would like to compute the minimum (that is, infimum) of the possible functions f such that $x \leq f(y)$ can be derived.

Given a continuous formula Γ , we say a *path* from a literal x to a literal y is a list of constraints $x = x_1 \leq f_1(x_2), \dots, x_k \leq f_k(x_{k+1} = y)$. A path is non-repeating if $x_i \neq x_j$ for $i \neq j$ unless $\{i, j\} = \{1, k+1\}$ (so a path from x to itself can be non-repeating by this definition). The resolution of a path is the clause $x \leq (f_1 \circ \dots \circ f_k)(y)$. We will sometimes write $p(y)$ for $(f_1 \circ \dots \circ f_k)(y)$. The interior vertices of a path are the vertices $x_2, \dots, x_{k+1}$.

The number of paths from x to y is infinite in general, but we hope to use infinite compositions in order to reduce this to a finite size. But there can still be roughly $n!$ non-repeating paths between a given pair of literals.

We now introduce the minimum operation. If we have constraints $x \leq f_1(y), \dots, x \leq f_k(y)$, then we can replace them by $x \leq \min(f_1, \dots, f_k)(y)$. Note that an expression like $\min(f_1, g_1) \circ \dots \circ \min(f_n, g_n)$ can represent the minimum of an exponential number of paths in only polynomial size.

Given a continuous formula Γ , we produce a new continuous formula $H(\Gamma)$ by the following process:

- For each pair of constraints of form $x \leq f(y)$ and $y \leq g(z)$ with y not being the same literal as x or z , add the clause $x \leq f(g(z))$.
- For every pair of (not necessarily distinct) literals x and y , write $\{x \leq f_i(y) : i = 1, \dots, k\}$ for the set of constraints from x to y in Γ , and replace them with the single constraint $x \leq \min_i(f_i)(y)$.

Lemma 13. *$H(\Gamma)$ is satisfiable if and only if Γ is satisfiable.*

Proof. Whenever a constraint of form $x \leq f(y)$ appears in Γ , a constraint $x \leq g(y)$ appears in $H(\Gamma)$ where $g(c) \leq f(c)$ for all $c \in \mathbb{R}$. So if $H(\Gamma)$ is satisfiable then Γ is.

If $x \leq f_i(y)$ then $x \leq \min_i(f_i)(y)$. So if Γ is satisfiable, then $H(\Gamma)$ is. $\square$

Write $H^2(\Gamma) = H(H(\Gamma))$ and write $H^k(\Gamma)$ for the result after applying this process k times.

Lemma 14. *Let Γ be a continuous formula with n variables. Suppose there is a non-repeating path p from a literal x to a literal y in Γ . If $m \geq \log_2(n) + 1$ then $H^m(\Gamma)$ contains a constraint $x \leq f(y)$ with $f(c) \leq p(c)$ for all $c \in \mathbb{R}$.*

Proof. By induction, we show that for non-repeating path p from x to y in Γ of length at most 2^k , there is a constraint $C(x, y, f)$ in $H^k(\Gamma)$ that is stronger than the resolution of p . This is clear when $k = 0$.

Suppose this is true for some k and p is a path from x to y of length at most 2^{k+1} . Then p can be split into non-repeating paths p_1 from x to z and p_2 from z to y , each of length at most 2^k . By induction, $H^k(\Gamma)$ contains constraints $x \leq f(z)$ and $z \leq g(y)$ with $f \leq p_1$ and $g \leq p_2$. So $f(g(c)) \leq p(c)$ for $c \in \mathbb{R}$. Since z is not the same literal as x or y , $H^{k+1}(\Gamma)$ contains a constraint $x \leq h(y)$ where $h \leq f \circ g$.

A non-repeating path has length at most $2n$, so $H^m(\Gamma)$ contains constraints for all such paths if $m \geq \log_2(n) + 1$. $\square$

Now let $V(\Gamma)$ be the continuous formula obtained from $H^{\lceil \log_2(n)+1 \rceil}(\Gamma)$ by replacing every constraint of form $x \leq f(x)$ with $x \leq f^\infty(x)$. We say that a path in $V(\Gamma)$ is *tight* if it alternates between loops and edges that aren't loops.

Lemma 15. *Suppose there is a path p from x to y in Γ . Then for each $c \in \mathbb{R}$, there is a tight path q from x to y in $V(\Gamma)$ such that $q(c) \leq p(c)$.*

Proof. If p is non-repeating then we are done by Lemma 14.

Now induct on the length of p . A path of length 1 can never repeat. Suppose that p repeats a literal z . Split p at the first and second points where it hits z , so p splits into sub-paths $\{p_1, p_2, p_3\}$ where p_2 is a non-repeating cycle.

Let $c \in \mathbb{R}$. First suppose that $(p_2 \circ p_3)(c) \geq p_3(c)$. In this case, $(p_1 \circ p_2 \circ p_3)(c) \leq (p_1 \circ p_3)(c)$. The path $\{p_1, p_3\}$ is a path shorter than p , so $p_1(p_3(c))$ is at most $q(c)$ where q is a tight path.

Now suppose that $p_2(p_3(c)) < p_3(c)$. Since p_3 is shorter than p , $p_3(c)$ is at most $q_3(c)$ for some tight path q_3 . By Lemma 14, there is a loop ℓ in $H^{\lceil \log_2(m)+1 \rceil}(\Gamma)$ where $p_2(p_3(c)) \leq \ell(p_3(c))$. Since $p_2(p_3(c)) < p_3(c)$, $\ell(p_3(c)) < p_3(c)$, so $\ell^\infty(p_3(c)) \leq p_2(p_3(c))$ by Lemma 10. So $\ell^\infty(q_3(c)) \leq p_2(p_3(c))$. There is a tight path q_1 such that $q_1(\ell^\infty(q_3(c))) \leq p_1(\ell^\infty(q_3(c))) \leq p_1(p_2(p_3(c)))$.

The path (q_1, ℓ^∞, q_3) may not be tight since the paths q_1 and q_3 can have loops at the ends, so (q_1, ℓ^∞, q_3) might have multiple loops in a row. But $V(\Gamma)$ has only one loop per vertex, and $\ell^\infty \circ \ell^\infty = \ell^\infty$, so (q_1, ℓ^∞, q_3) is equivalent to a tight path.

So for each $c \in \mathbb{R}$, there is a tight path q in $V(\Gamma)$ with $q(c) \leq p(c)$. □

Write $x \leq f_{x \rightarrow y}(y)$ for the constraint from x to y in $H^{\lceil \log_2(n)+1 \rceil}(\Gamma)$ (if one exists) and define:

$$\min_{x \rightarrow y}(y) = \min\{p(y) : p \text{ is a tight path in } V(\Gamma)\}$$

We obtain the following:

Lemma 16. *For $c \in \mathbb{R}$, the value of $\min_{x \rightarrow y}(c)$ is one of:*

- $f_{x \rightarrow y}(c)$
- An attracting point of $f_{x \rightarrow x}$
- $f_{x \rightarrow z}(d)$ for some literal $z \neq x$, where d is an attracting point of $f_{z \rightarrow z}$

Proof. The values taken by a loop from z to itself in $V(\Gamma)$ are attracting points of $f_{z \rightarrow z}$, so the result is clear from the definition of a tight path. □

By Lemma 15, the functions $\min_{x \rightarrow y}$ determine the satisfiability of Γ :

Lemma 17. *Γ is satisfiable if and only if neither of the following happen:*

- There are literals x and y such that $\min_{x \rightarrow y}(\max(\text{range}(y))) < \min(\text{range}(x))$
- There is a literal x and a $c \in \text{range}(x)$ such that $\min_{x \rightarrow -x}(-c) < c$ and $\min_{-x \rightarrow x}(c) < -c$

Proof. If Γ is satisfiable, then $V(\Gamma)$ is satisfiable by Lemmas 11 and 13. Each of the cases implies that there is a derivation of $\emptyset$ from $V(\Gamma)$, which can't happen by Lemma 6.

Suppose Γ is unsatisfiable. By Lemma 8, there is a derivation of $\emptyset$ from Γ . For every constraint of form $x \leq f(y)$ that can be derived from Γ , the function f is the resolution of a path in Γ . By Lemma 15, these paths correspond to tight paths p or p and q in $V(\Gamma)$ satisfying one of the conditions. □

5 Algorithms for $\mathcal{M} - 2\text{SAT}$

Theorem 18. *$\mathcal{M} - 2\text{SAT}$ is in NP*

Proof. Let Γ be an instance of $\mathcal{M} - 2\text{SAT}$ with n variables where the each constraint in Γ is represented with at most β bits. Suppose that Γ is satisfiable. First, we compute the set of values of v_1 that can be extended to satisfying assignments.

The function $\min_{x \rightarrow y}$ only depend on the functions appearing in Γ and not on the ranges $\text{range}(v_i)$. By Since Γ is satisfiable, the second case in Lemma 17 never happens. So by Lemma 17 the problem has a satisfying assignment with $v_1 = c$ if and only if:

- $\min_{-v_1 \rightarrow x}(\max(\text{range}(x))) \geq -c$ and $\min_{v_1 \rightarrow x}(\max(\text{range}(x))) \geq c$ for each x a literal other than v_1 or $-v_1$

- $\min_{v_1 \rightarrow -v_1}(-c) \geq c$ and $\min_{-v_1 \rightarrow v_1}(c) \geq -c$
- $\min_{v_1 \rightarrow v_1}(c) \geq c$ (equivalently $\min_{-v_1 \rightarrow -v_1}(-c) \geq -c$)

Note that we don't need to check e.g. $\min_{x \rightarrow v_1}(c) \geq \min(\text{range}(x))$ because this is equivalent to $\min_{-v_1 \rightarrow -x}(\max(\text{range}(-x))) \geq -c$.

The set of values of c satisfying all these conditions is a union of closed intervals, where each endpoint of one of the intervals is one of:

- $-\min_{-v_1 \rightarrow x}(\max(\text{range}(x)))$ for x a literal other than x
- $\min_{v_1 \rightarrow x}(\max(\text{range}(x)))$ for x a literal other than x
- A point where $\min_{v_1 \rightarrow -v_1}(-c) = c$
- A point where $\min_{-v_1 \rightarrow v_1}(c) = -c$
- A locally minimal or locally maximal fixed point of $\min_{v_1 \rightarrow v_1}$

The constraints in Γ are piecewise fractional linear, that is they are piecewise of form:

$$\frac{ax + b}{cx + d}$$

where a, b, c and d are integers with at most β bits each.

Taking the minimum of some of these constraints increases only the number of pieces and not the complexity of the pieces themselves. Composing two functions of this form causes the bit-complexity of the coefficients to grow by at most a factor of 3, since:

$$\frac{a_1 \left(\frac{a_2 x + b_2}{c_2 x + d_2} \right) + b_1}{c_1 \left(\frac{a_2 x + b_2}{c_2 x + d_2} \right) + d_1} = \frac{(a_1 a_2 + b_1 c_2)x + (a_1 b_2 + b_1 d_2)}{(c_1 a_2 + d_1 c_2)x + (c_1 b_2 + d_1 d_2)}$$

In order to make Γ symmetric, we also need to compute $\tilde{f}$ for each constraint $x \leq f(y)$ in Γ . The inverse of:

$$f(x) = \frac{ax + b}{cx + d}$$

is:

$$f^{-1}(x) = \frac{dx - b}{-cx + d}$$

so computing inverses doesn't increase the bit complexity of a fractional linear function.

So $H(\Gamma)$ has constraints that are piecewise fractional-linear with coefficients having at most 3β bits each. So $H^{\lceil \log_2(n)+1 \rceil}(\Gamma)$ has constraints that are piecewise of the same form, where the coefficients are integers with at most $\mathcal{O}(n^{\log_2(3)}\beta)$ bits. Since the constraints should be continuous, the break points between two pieces are coincidence points of the functions on the those pieces.

By Lemma 16, the value of $\min_{v_1 \rightarrow x}(\max(\text{range}(x)))$ is either $f_{v_1 \rightarrow x}$, an attracting point of $f_{v_1 \rightarrow v_1}$, or $f_{v_1 \rightarrow z}(d)$ for some literal $z \neq v_1$ and d an attracting point of $f_{z \rightarrow z}$. The functions $f_{x \rightarrow y}$ are just the functions that appear in $H^{\lceil \log_2(n)+1 \rceil}(\Gamma)$, so are piecewise fractional-linear by the above discussion.

The equation:

$$\frac{ax + b}{cx + d} = x$$

yields a quadratic equation:

$$cx^2 + (d - a)x - b = 0$$

after clearing denominators. So a fractional-linear function is either the identity function $f(x) = x$ or has at most 2 fixed points. In particular, each attracting point of a piecewise fractional-linear function is an isolated fixed point of one of the pieces. So the attracting points of each $f_{x \rightarrow y}$ are of each form $p + q\sqrt{r}$, where p, q , and r are integers with a number of bits at most polynomial in n and β .

Applying a fractional linear map to a number of form $p + q\sqrt{r}$ yields a number of the same form, indeed:

$$\frac{a(p + q\sqrt{r}) + b}{c(p + q\sqrt{r}) + d} = \frac{((ap + b) + aq\sqrt{r})(cp + d - cq\sqrt{r})}{(cp + d)^2 + rq^2c^2} = \frac{(ap + b)(cp + d) - acq^2r}{(cp + d)^2 + rq^2c^2} + \frac{(ad - bc)}{(cp + d)^2 + rq^2c^2}q\sqrt{r}$$

Since the endpoints of the interval $\text{range}(x)$ are rational, we conclude that $\min_{v_1 \rightarrow x}(\max(\text{range}(x)))$ is a number of form $p + q\sqrt{r}$ where p, q , and r have at most polynomially many bits. The same is true of locally extremal fixed points of $\min_{v_1 \rightarrow v_2}$ and points where $\min_{v_1 \rightarrow -v_1}(-c) = c$ or $\min_{-v_1 \rightarrow v_1}(c) = -c$ (note that $\min_{v_1 \rightarrow -v_1}$ is increasing, so there is at most one point where $\min_{v_1 \rightarrow -v_1}(-c) = c$). In particular, we can choose a value of v_1 that has this form where v_1 extends to a satisfying assignment of Γ .

We now repeat this process to set values of $v_2, \dots, v_n$. The different now is that $\text{range}(v_1)$ may no longer be bounded by rational numbers, since we might have set v_1 to an irrational value. But this is fine, because applying $\min_{v_2 \rightarrow x}$ to a number of form $p + q\sqrt{r}$ still yields a number of the same form. Inductively, we conclude that if Γ has a satisfying assignment, then it has an assignment where all variables have form $p + q\sqrt{r}$, where p, q , and r are rational numbers with at most $\mathcal{O}(\beta n^{2 \log_2(3)})$ bits.

Given an assignment of variables of this form, checking a constraint in Γ requires at most a constant number of arithmetic operations with numbers that roots of polynomials of degree at most 2. These arithmetic operations can be performed in polynomial time (see e.g. Mishra and Pedersen [20]), so a solution can be checked in polynomial time, proving that $\mathcal{M} - 2\text{SAT}$ is in NP. $\square$

By Lemma 1, this also proves that the boundary-boundary art-gallery problem is in NP.

The boundary-boundary art-gallery problem is NP-hard, but the reduction in Lemma 1 is nondeterministic, so it doesn't imply anything about the complexity of $\mathcal{M} - 2\text{SAT}$. We show that there is a quasi-polynomial-time algorithm for $\mathcal{M} - 2\text{SAT}$, so it is not NP-hard unless the exponential-time hypothesis fails.

Theorem 19. *There is a quasi-polynomial-time algorithm for $\mathcal{M} - 2\text{SAT}$*

Proof. Let Γ be an instance of $\mathcal{M} - 2\text{SAT}$ with n variables where each constraint is a piecewise fractional-linear function where the coefficients have at most β bits each. For each pair of literals x and y , we can take the minimum over all constraints of form $x \leq f(y)$ in Γ , so assume that Γ has at most one constraint $x \leq f(y)$ for each such pair of literals. Let k be the maximum number of pieces in a constraint in Γ .

As shown in the proof of Theorem 19, the functions $f_{x \rightarrow y}$ are piecewise fractional linear where the coefficients of each piece are integers with at most $\mathcal{O}(n^{\log_2(3)}\beta)$ bits each. We also need a bound on the number of pieces. If f and g are monotone fractional-linear functions with at most k pieces each, then the composition $f \circ g$ has at most $2k$ pieces.

If $f_1, \dots, f_j$ are piecewise-fractional-linear functions with at most k pieces each, then $\min_i(f_i)$ has at most $(2j^3 + j)k$ pieces. To see this, split $\mathbb{R}$ at the endpoint of every piece in f_i . This splits $\mathbb{R}$ into at most $j(k-1) + 1$ chunks such that each f_i restricts to a single piece on each chunk. Two fractional linear functions intersect in at most two places, so the number of intersection points on each of these chunks is at most $2j^2$.

Each application of H involves computing minimums of sets of at most $2n$ functions, so each constraint in $H^{\lceil \log_2(n) + 1 \rceil}(\Gamma)$ has at most $n^{\mathcal{O}(\log(n))}k$ pieces, and computing all the functions $f_{x \rightarrow y}$ takes time $n^{\mathcal{O}(\log(n))}k\beta$.

In order to decide if Γ is decidable, we need to compute the functions $\min_{x \rightarrow y}$, which correspond to tight paths in $V(\Gamma)$. Each loop in $V(\Gamma)$ is the infinite composition of some $f_{x \rightarrow x}$. The number of values taken by $f_{x \rightarrow x}^\infty$ is at most twice the number of pieces in f .

Let G be a directed graph that has a vertex (x, c) whenever x is a literal in Γ and c is a value that can be taken by $f_{x \rightarrow x}^\infty$. G has an edge from (x, c) to (y, d) whenever $f_{x \rightarrow x}^\infty(f_{x \rightarrow y}(d)) = c$. For every pair of a literal x and a vertex (y, c) of G , define:

$$G(x \rightarrow (y, d)) = \min\{c : (x, c) \text{ is a vertex of } G \text{ and there is a path from } (x, c) \text{ to } (y, d) \text{ in } G\}$$

We now claim that:

$$\min_{x \rightarrow y}(c) = \min\{\text{id}, f_{x \rightarrow x}^\infty\} \circ \min(\{f_{x \rightarrow y}(c)\} \cup \{f_{x \rightarrow z}(G(z \rightarrow (w, (f_{w \rightarrow w}^\infty \circ f_{w \rightarrow y})(c)))) : z \text{ a literal}\})$$

To see this, let p be a tight path from x to y in $V(\Gamma)$. For each $c \in \mathbb{R}$, we want to show that $p(c)$ is one of $f_{x \rightarrow y}(c)$, $(f_{x \rightarrow x}^\infty \circ f_{x \rightarrow y})(c)$, $f_{x \rightarrow z}(d)$, or $(f_{x \rightarrow x}^\infty \circ f_{x \rightarrow z})(d)$, for some literal z and $d \in \mathbb{R}$ such that there is a literal w and path from (z, d) to $(w, (f_{w \rightarrow w}^\infty \circ f_{w \rightarrow y})(c))$ in G . This is by induction on the length of p .

If p has length 1, then it must be exactly $f_{x \rightarrow y}$. If p is longer, then there are two cases based on whether the first edge in p is a loop.

First, suppose we can write $p = f_{x \rightarrow x}^\infty \circ q$ for q a tight path from x to y . Since q is shorter than p and $f_{x \rightarrow x} \circ f_{x \rightarrow x} = f_{x \rightarrow x}$, we have that p is has one of the required forms.

Otherwise, we can write $p = f_{x \rightarrow z} \circ f_{z \rightarrow z} \circ q$ for $z \neq x$ and q a tight path from z to y . By induction, q has one of the types above. Since $f_{z \rightarrow z}^\infty \circ f_{z \rightarrow z}^\infty = f_{z \rightarrow z}^\infty$, we have that $p(c)$ is either $(f_{x \rightarrow z} \circ f_{z \rightarrow z}^\infty \circ f_{x \rightarrow y})(c)$ or $(f_{x \rightarrow z} \circ f_{z \rightarrow z}^\infty \circ f_{z \rightarrow z'})(d)$ where there is a path from (z', d) to $(w, (f_{w \rightarrow w}^\infty \circ f_{w \rightarrow y})(c))$ in G .

There is a path from $(z, (f_{z \rightarrow z}^\infty \circ f_{x \rightarrow y})(c))$ to itself in G , so the first case can be written $f_{x \rightarrow z}(d)$ with d, z satisfying the appropriate conditions. There is an edge from $(z, (f_{z \rightarrow z}^\infty \circ f_{z \rightarrow z'})(d))$ to (z', d) in G , so there is a path from $(z, (f_{z \rightarrow z}^\infty \circ f_{z \rightarrow z'})(d))$ to $(w, (f_{w \rightarrow w}^\infty \circ f_{w \rightarrow y})(c))$ in G . So the second case also has the appropriate form.

For each of $f_{x \rightarrow y}(c)$, $(f_{x \rightarrow x}^\infty \circ f_{x \rightarrow y})(c)$, $f_{x \rightarrow z}(d)$, and $(f_{x \rightarrow x} \circ f_{x \rightarrow z})(d)$ for z, d as above, there is a tight path in $V(\Gamma)$ that evaluates to that value at c . This proves the claim.

We can compute G in time $n^{\mathcal{O}(\log(n))}k\beta$, and so we can calculate each $G(x \rightarrow (y, d))$ in quasi-polynomial time. The above calculation shows that this lets us calculate the functions $\min_{x \rightarrow y}$. In order to determine if Γ is satisfiable, we just need to test if any of the conditions in Lemma 17 occur. Since each $\min_{x \rightarrow y}$ is a piecewise-fractional-linear (though not necessarily continuous) function, this is straightforward. The entire computation runs in time $n^{\mathcal{O}(\log(n))}k\beta$. $\square$

The minimum of k piecewise-fractional-linear functions with k pieces each is the lower envelope of an arrangement of nk arcs, where each pair of arcs has at most 2 intersection points. Results on Davenport-Schinzel sequences (see Sharir and Agarwal [26]) imply the number of pieces in the minimum is at most $\mathcal{O}(jk2^{\alpha(jk)})$, where α is the inverse Ackermann function. Unless k is much much larger than n , this improves the bound used in the proof of Theorem 19, but the result of the theorem does not change qualitatively.

6 Irrational Coordinates

In this section, we give a sketch of a construction of an instance of the boundary-boundary art-gallery problem where the unique solution with 3 guards requires placing the guards at irrational coordinates. The instance is illustrated in Figure 5.

The instance can be guarded with 3 guards. The positions of the guards represent variables x, y and z , and the 4 nooks create constraints $z \geq x$, $y \geq z$, $\frac{4x+2}{x+4} \geq y$ and $\frac{4x-2}{-x+4} \geq y$. The constraints $z \geq x$ and $y \geq z$ imply that $y \geq x$, so the only two solutions are $x = y = z = \pm\sqrt{2}$, as illustrated by Figure 6. Only the solution $x = y = z = \sqrt{2}$ satisfies the range constraints created by the polygon.

We now give a brief sketch of the techniques used to construct this polygon. The key ideas are adapted from the construction in [28]. There are 6 small triangular slits in the polygon boundary, and the only way to guard all of the slits with 3 guards total is to place one guard somewhere on each of the red segments shown in Figure 7. There are 4 nooks (highlighted in blue in Figure 7) that each create a continuous constraint between two of the guard positions. By the discussion in the proof of Lemma 1, the constraint created by each nook is fractional-linear.




Figure 5: Left: an art gallery whose boundary can be guarded with three guards on the boundary, but only if irrational coordinates are allowed. Right: guarding this example with 3 guards. There are 4 nooks that represent constraints between pairs of guards.




Figure 6: The three regions created by constraints $y \geq x$, $\frac{4x+2}{x+4} \geq y$ and $\frac{4x-2}{-x+4} \geq y$. The only two points in all three regions are $\pm(\sqrt{2}, \sqrt{2})$




Figure 7: The 6 slits on the boundary of the polygon can only be guarded with 3 guards if there is one guard on each of the red segments shown. The nooks that create continuous constraints are highlighted in blue.

Two fractional linear functions are identical if they agree on three different points (for f and g fractional linear, the identity $f(x) = g(x)$ becomes quadratic after clearing denominators). So in order to verify that a nook enforces the appropriate constraint, it is sufficient to check 3 points. We illustrate this by constructing a nook that enforces the constraint $\frac{4x+2}{x+4} \geq y$.

First, we choose two lines, so that points on those lines represent values of variables x and y , as illustrated in Figure 8. We choose label the points corresponding to $x = 0$, $x = 1$, and $x = 2$. These should be mapped to $y = \frac{1}{2}$, $y = \frac{6}{5}$, and $y = \frac{5}{3}$ by the nook. We draw a line through the points representing $x = 0$ and $y = \frac{1}{2}$, and choose a point P on this line, as in Figure 8.

Next, we choose a line ℓ through the point representing $x = 0$, as illustrated in Figure 9. We draw lines from the points representing $y = \frac{6}{5}$ and $y = \frac{5}{3}$ through P , and let points I and J respectively be the intersection points of these new lines with ℓ .

Next, we draw a line from I to $x = 1$ and from J to $x = 2$, and let Q be the intersection point of these lines, as in Figure 10. We create a nook from the points P , I , J and Q . Let g be the function from the line representing x to ℓ given by projecting through Q and let f be the function from ℓ to the line representing y given by projecting through P . If two guards are placed at points representing values of x and y with $x \in [1, 2]$ and $y \in [\frac{6}{5}, \frac{5}{3}]$, then they together can see all of the nook if and only if $f(g(x)) \leq y$, as illustrated in Figure 11.



The composition $f \circ g$ is fractional linear, so in order to verify that $f(g(x)) = \frac{4x+2}{x+4}$, we just need to check that $f(g(0)) = \frac{1}{2}$, $f(g(1)) = \frac{6}{5}$ and $f(g(2)) = \frac{5}{3}$. This can be easily seen from the figures.

Each of the 4 nooks in Figure 5 is constructed using a similar method. A detailed specification of the construction is shown in Figure 12.

7 Acknowledgments

I am very grateful to Lucas Meijer and Tillman Miltzow for discussions that helped to produce the example in Section 6.

This work is supported by Starting Grant 1054-00032B from the Independent Research Fund Denmark under the Sapere Aude research career programme, and by the Carlsberg Foundation, grant CF24-1929. The






Figure 10: The second step in the construction of a nook.




Figure 11: In order to see the parts of the nook that cannot be seen by a guard at position x , a guard on the vertical line should be placed at a position representing a y value of at most $f(g(x))$. The composition $f \circ g$ is a fractional linear function that maps 1 to $\frac{6}{5}$ and 2 to $\frac{5}{3}$ by construction. To see that it maps 0 to $\frac{1}{2}$, we notice that the line through $x = 0$ and Q intersects ℓ at the point $x = 0$, and the line through $x = 0$ and P intersects the vertical line at $y = \frac{1}{2}$ by construction.




Figure 12: Specification of the construction. The 6 slits are omitted, since the exact parameters of the slits are not important. Each of the purple points is a vertex of the polygon with explicitly-given coordinates. Each of the blue points is an explicitly-given point that is used to help construct the polygon (but not a vertex of the polygon). The red points are derived vertices, which are defined as intersections of lines through some of the other points. A guard represents a value of x if it is at point $(-x, 0)$, a value of y if it is at $(\frac{3}{2} + \frac{1}{2}y, 1 + y)$, or a value of z if it is at point $(\frac{3}{2} + \frac{1}{2}z, -1 - z)$.

8 Conclusion

We have shown that the boundary-boundary art-gallery problem is in NP, despite the fact that irrational coordinates are sometimes needed. Each of the X-Y art-gallery variants is now known to be either NP-complete or $\exists\mathbb{R}$ -complete.

We have described new techniques for approaching 2CCSPs, in particular giving a quasi-polynomial-time algorithm for $\mathcal{M} - 2\text{SAT}$ and showing that it is in NP. It is interesting to wonder if $\mathcal{M} - 2\text{SAT}$ can be solved in polynomial time.

We could also ask if more complicated 2CCSPs might also be in NP. The techniques from Section 3 can be applied to more complicated problems, but the numerical calculations required quickly become intractable in the Turing model.

Even if we allow a model of computation that can perform more complicated arithmetic operations, it is unclear whether general 2CCSPs become easy. Our algorithms for $\mathcal{M} - 2\text{SAT}$ exploit the fact that the class of fractional-linear functions is closed under composition. This seems to be in contrast to higher-degree polynomials, where the degree grows exponentially with the number of compositions. However, only a small fraction of all high-degree polynomials can be written as a composition of lower-degree polynomials, and it isn't obvious that the exponential degree increase leads to an exponential increase in combinatorial complexity. For example, it seems to be unknown whether the composition $f \circ g$ can have more than $\mathcal{O}(n)$ fixed points, when f and g are monotone polynomials of degree at most n .

References

- [1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. Irrational Guards are Sometimes Needed. In Boris Aronov and Matthew J. Katz, editors, *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:15, Dagstuhl, Germany, 2017. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SocG.2017.3>, doi:10.4230/LIPIcs.SocG.2017.3.
- [2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$ -complete. *Journal of the ACM*, 69(1):1–70, 2021. doi:10.1145/3486220.
- [3] Mikkel Abrahamsen, Linda Kleist, and Tillmann Miltzow. Training neural networks is $\exists\mathbb{R}$ -complete. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18293–18306. Curran Associates, Inc., 2021. URL: https://proceedings.neurips.cc/paper_files/paper/2021/file/9813b270ed0288e7c0388f0fd4ec68f5-Paper.pdf.
- [4] Mikkel Abrahamsen, Tillmann Miltzow, and Nadja Seiferth. Framework for $\exists\mathbb{R}$ -completeness of two-dimensional packing problems. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1014–1021, 2020. doi:10.1109/FOCS46700.2020.00098.
- [5] Akanksha Agrawal, Kristine V. K. Knudsen, Daniel Lokshantov, Saket Saurabh, and Meirav Zehavi. The Parameterized Complexity of Guarding Almost Convex Polygons. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2020/12161>, doi:10.4230/LIPIcs.SocG.2020.3.

- [6] B. Beckert, R. Hahnle, and F. Manya. The 2-sat problem of regular signed cnf formulas. In *Proceedings 30th IEEE International Symposium on Multiple-Valued Logic (ISMVL 2000)*, pages 331–336, 2000. doi:10.1109/ISMVL.2000.848640.
- [7] Daniel Bertschinger, Christoph Hertrich, Paul Jungeblut, Tillmann Miltzow, and Simon Weber. Training fully connected neural networks is er-complete. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA, 2023. Curran Associates Inc.
- [8] Ahmad Biniiaz, Anil Maheshwari, Magnus Christian Ring Merrild, Joseph S. B. Mitchell, Saeed Odak, Valentin Polishchuk, Eliot W. Robson, Casper Moldrup Rysgaard, Jens Kristian Refsgaard Schou, Thomas Shermer, Jack Spalding-Jamieson, Rolf Svenning, and Da Wei Zheng. Polynomial-Time Algorithms for Contiguous Art Gallery and Related Problems. In Oswin Aichholzer and Haitao Wang, editors, *41st International Symposium on Computational Geometry (SoCG 2025)*, volume 332 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:21, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SoCG.2025.20>, doi:10.4230/LIPIcs.SoCG.2025.20.
- [9] Janusz A. Brzozowski and Edward Joseph McCluskey. Signal flow graph techniques for sequential circuit state diagrams. *IEEE Trans. Electron. Comput.*, 12:67–76, 1963. URL: <https://api.semanticscholar.org/CorpusID:708459>.
- [10] John Canny. Some algebraic and geometric computations in pspace. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing, STOC '88*, page 460–467, New York, NY, USA, 1988. Association for Computing Machinery. doi:10.1145/62212.62257.
- [11] Witold Charatonik and Michal Wrona. 2-sat problems in some multi-valued logics based on finite lattices. In *37th International Symposium on Multiple-Valued Logic (ISMVL'07)*, pages 21–21, 2007. doi:10.1109/ISMVL.2007.1.
- [12] Daniel Dadush, Zhuan Khye Koh, Bento Natura, Neil Olver, and László A. Végh. A strongly polynomial algorithm for linear programs with at most two nonzero entries per row or column. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024*, page 1561–1572, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649764.
- [13] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, July 1960. doi:10.1145/321033.321034.
- [14] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM Journal on Discrete Mathematics*, 3(2):255–265, 1990. arXiv:<https://doi.org/10.1137/0403021>, doi:10.1137/0403021.
- [15] M. R. Krom. The decision problem for a class of first-order formulas in which all disjunctions are binary. *Mathematical Logic Quarterly*, 13(1-2):15–20, 1967. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/malq.19670130104>, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/malq.19670130104>, doi:10.1002/malq.19670130104.
- [16] Aldo Laurentini. Guarding the walls of an art gallery. *The Visual Computer*, 15:265–278, 10 1999. doi:10.1007/s003710050177.
- [17] D. T. Lee and Arthur K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986. doi:10.1109/TIT.1986.1057165.
- [18] Nimrod Megiddo. Towards a genuinely polynomial algorithm for linear programming. *SIAM Journal on Computing*, 12(2):347–353, 1983. arXiv:<https://doi.org/10.1137/0212022>, doi:10.1137/0212022.

- [19] Tillmann Miltzow and Reinier F. Schmiermann. On classifying continuous constraint satisfaction problems. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 781–791, 2022. doi:10.1109/FOCS52979.2021.00081.
- [20] B. Mishra and P. Pedersen. Arithmetic with real algebraic numbers is in nc. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, ISSAC '90, page 120–126, New York, NY, USA, 1990. Association for Computing Machinery. doi:10.1145/96877.96909.
- [21] N. E. Mnev. *The universality theorems on the classification problem of configuration varieties and convex polytopes varieties*, pages 527–543. Springer Berlin Heidelberg, Berlin, Heidelberg, 1988. doi:10.1007/BFb0082792.
- [22] Joseph O'Rourke. *Art gallery theorems and algorithms*, volume 57. Oxford New York, NY, USA, 1987.
- [23] Marcus Schaefer. *Realizability of Graphs and Linkages*, pages 461–482. Springer New York, New York, NY, 2013. doi:10.1007/978-1-4614-0110-0_24.
- [24] Marcus Schaefer, Jean Cardinal, and Tillmann Miltzow. The existential theory of the reals as a complexity class: A compendium, 2024. URL: <https://arxiv.org/abs/2407.18006>, arXiv:2407.18006.
- [25] Marcus Schaefer and Danie Stefankovič. Fixed points, nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60, 02 2017. doi:10.1007/s00224-015-9662-0.
- [26] Micha Sharir and Pankaj K. Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge University Press, 2010.
- [27] Peter W. Shor. Stretchability of pseudolines is np-hard. In *Applied Geometry And Discrete Mathematics*, 1990. URL: <https://api.semanticscholar.org/CorpusID:43637487>.
- [28] Jack Stade. The Point-Boundary Art Gallery Problem Is $\exists\mathbb{R}$ -Hard. In Oswin Aichholzer and Haitao Wang, editors, *41st International Symposium on Computational Geometry (SoCG 2025)*, volume 332 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 74:1–74:23, Dagstuhl, Germany, 2025. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.SocG.2025.74>, doi:10.4230/LIPIcs.SocG.2025.74.
- [29] Alfred Tarski and J. C. C. McKinsey. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, dgo - digital original, 1 edition, 1951. URL: <http://www.jstor.org/stable/jj.8501420>.