Covering Polygons is Even Harder*

Mikkel Abrahamsen
BARC, University of Copenhagen
Email: miab@di.ku.dk

***Abstract*—In the MINIMUM CONVEX COVER (MCC) problem, we are given a simple polygon $\mathcal{P}$ and an integer k, and the question is if there exist k convex polygons whose union is $\mathcal{P}$. It is known that MCC is NP-hard [Culberson & Reckhow: Covering polygons is hard, FOCS 1988/Journal of Algorithms 1994] and in $\exists\mathbb{R}$ [O'Rourke: The complexity of computing minimum convex covers for polygons, Allerton 1982]. We prove that MCC is $\exists\mathbb{R}$-hard, and the problem is thus $\exists\mathbb{R}$-complete. In other words, the problem is equivalent to deciding whether a system of polynomial equations and inequalities with integer coefficients has a real solution.**

If a cover for our constructed polygon exists, then so does a cover consisting entirely of triangles. As a byproduct, we therefore also establish that it is $\exists\mathbb{R}$-complete to decide whether k triangles cover a given polygon.

The issue that it was not known if finding a minimum cover is in NP has repeatedly been raised in the literature, and it was mentioned as a "long-standing open question" already in 2001 [Eidenbenz & Widmayer: An approximation algorithm for minimum convex cover with logarithmic performance guarantee, ESA 2001/SIAM Journal on Computing 2003]. We prove that assuming the widespread belief that NP $\neq \exists\mathbb{R}$, the problem is not in NP.

An implication of the result is that many natural approaches to finding small covers are bound to give suboptimal solutions in some cases, since irrational coordinates of arbitrarily high algebraic degree can be needed for the corners of the pieces in an optimal solution.

***Keywords*-Polygon decomposition, Minimum Convex Cover, Existential Theory of the Reals**

I. INTRODUCTION

Polygons are among the geometric structures that are most frequently used to model physical objects, as they are suitable for representing a wide variety of shapes and figures in computer graphics and vision, pattern recognition, robotics, computer-aided design and manufacturing, and other computational fields. Polygons may have very complicated shapes that make it difficult to find algorithms to process them directly. A natural first step in designing algorithms is to decompose the given polygon $\mathcal{P}$ into more basic *pieces* of a restricted type that permits very efficient processing. Here, the union of the pieces must be exactly the given polygon $\mathcal{P}$. When such a decomposition has been obtained, the partial solutions to the individual pieces can be combined to obtain a solution for the complete polygon $\mathcal{P}$. By "more basic pieces," we mean pieces that belong to a more restricted class of polygons. Ideally, we find a decomposition of the polygon $\mathcal{P}$ into the minimum number of pieces.

Many different decomposition problems arise this way, depending on restrictions on the polygon $\mathcal{P}$, the type of basic pieces, and whether the decomposition is a cover or a partition. In covering problems, we just require the union of the pieces to equal $\mathcal{P}$, whereas in partition problems, we have the further requirement that the pieces be interior-disjoint. Another important distinction is whether or not *Steiner points* are allowed. A Steiner point is a corner of a piece in a decomposition which is not also a corner of $\mathcal{P}$. Furthermore, it often makes a big difference whether or not $\mathcal{P}$ is allowed to have holes.

Because of the many variants of decomposition problems, their applicability within many practical domains, and the very appealing fundamental nature and the creativity and technical skills required to solve them, numerous papers have been written about these problems. The vast literature is documented in several highly-cited books and survey papers that give an overview of the state-of-the-art at the time of publication [2]–[8].

One of the first decomposition problems to be studied was that of covering a polygon with convex polygons. Pavlidis studied this problem from a practical angle in relation to shape analysis and pattern recognition in a series of papers and a book, the first paper from 1968 [9]–[14]. The related NP-complete problem of covering an orthogonal polygon with a given number of rectangles was mentioned by Garey and Johnson's famous book [15, p. 232], who denoted the problem RECTILINEAR PICTURE COMPRESSION, since a collection of pixels can be compactly represented as their minimum rectangle cover. For more recent practical work involving versions of convex covering problems, see the surveys [16], [17] and the paper [18].

O'Rourke and Supowit [19] proved that for polygons with holes, it is NP-hard to find minimum covers using convex, star-shaped, and spiral polygons as pieces. This was established with and without Steiner points allowed. When Steiner points are not allowed, the covering problems are clearly in NP and are thus NP-complete, whereas the ones with Steiner points allowed are not immediately seen to be decidable at all. In the paper [20], O'Rourke proved the convex covering problem with Steiner points to be decidable by showing that any given instance can be expressed as a

∗ This is an extended abstract of the paper [1].

2575-8454/21/$31.00 ©2021 IEEE
DOI 10.1109/FOCS52979.2021.00045

Tarski formula. In modern terms, we can say that he proved $\exists\mathbb{R}$-membership. He gave an example of a polygon $\mathcal{P}$ where Steiner points are needed and edge extensions of $\mathcal{P}$ are not sufficient to form the pieces of a minimum cover. The figure is now the logo of The Society of Computational Geometry and often also of Symposium on Computational Geometry.

Since the NP-hardness reductions in [19] relied on polygons with holes, it was still not known if the covering problems for polygons without holes could be solved efficiently. Chazelle and Dobkin [21] showed already in 1979 that a polygon without holes can be *partitioned* into a minimum number of convex pieces in polynomial time; a problem that had also been believed to be NP-hard, so this might have given hope for the covering problems as well. However, it was soon proved that some covering problems are NP-hard even for polygons without holes. The first one was apparently the problem of covering a polygon with a minimum number of star-shaped polygons, which is usually known as the ART GALLERY problem. The first proof is often attributed to Alok Aggarwal's PhD thesis [22] (see for instance [23]), which is unfortunately practically unavailable. The first published proof appears to be in a paper by Lee and Lin [24]. Then followed the proof by Culberson and Reckhow [23] that it is likewise NP-hard to cover polygons without holes with a minimum number of convex pieces. The authors added the comment: *"We are unable to show that general convex covering is in* NP*"*. The issue that the problem is not known to be in NP has been raised in many other papers and books [3], [4], [6], [7], [20], [25]–[27], and was mentioned as a "long-standing open question" already in 2001 [26]. Christ [27] proved that deciding if a polygon can be covered by a minimum number of *triangles* is also NP-hard.

In this paper, we prove that it is $\exists\mathbb{R}$-complete to decide if a polygon can be covered by a given number of convex pieces. The problem is thus not in NP, assuming the widespread belief that NP $\neq \exists\mathbb{R}$. Our reduction uses techniques that were developed for proving that ART GALLERY and some versions of geometric packing are $\exists\mathbb{R}$-complete [28], [29]. As a biproduct, we show that it is even $\exists\mathbb{R}$-complete to find a minimum triangle cover. The hardness holds also when the corners of $\mathcal{P}$ are in general position, i.e., no three are collinear.

A. Existential theory of the reals

In order to define the complexity class $\exists\mathbb{R}$, we first define the problem ETR in the style of Garey and Johnson [15].

Instance: A well-formed formula $\Phi(x_1, x_2, \ldots, x_n)$ using symbols from the set

$$\{x_1, x_2, \ldots, x_n, \wedge, \vee, \neg, 0, 1, +, -, \cdot, \ (\ , \ ) \ , =, <, \leq\}.$$

Question: Is the expression

$$\exists x_1, x_2, \ldots, x_n \in \mathbb{R} \colon \Phi(x_1, x_2, \ldots, x_n)$$

true?

The complexity class $\exists\mathbb{R}$ consists of all problems that are many-one reducible to ETR in polynomial time, and a problem is $\exists\mathbb{R}$-hard if there is a reduction in the other direction. It is currently known that

$$\mathsf{NP} \subseteq \exists\mathbb{R} \subseteq \mathsf{PSPACE}.$$

It is not hard see that the problem ETR is NP-hard, yielding the first inclusion. The containment $\exists\mathbb{R} \subseteq$ PSPACE is highly non-trivial, and it was first established by Canny [30].

As examples of $\exists\mathbb{R}$-complete problems, we mention problems related to realization of order-types [31]–[33], graph drawing [34]–[36], recognition of geometric graphs [37]–[40], straightening of curves [41], guarding polygons [28], Nash-equilibria [42], [43], linkages [44]–[46], matrix-decompositions [47]–[49], polytope theory [31], and geometric packing [29]. See also the surveys [50]–[52].

B. Results

Before stating our result, let us define the covering problem in more detail. We define a *polygon* $\mathcal{P}$ to be a compact region in the plane such that the boundary $\partial\mathcal{P}$ is a closed, simple curve consisting of finitely many line segments. We now define the problem MINIMUM CONVEX COVER (MCC for brevity) as follows.

Instance: A polygon $\mathcal{P}$ represented as an array of the coordinates of the corners in cyclic order, and a positive integer k. The corners have rational coordinates.

Question: Do there exist k convex polygons $Q_1, \ldots, Q_k$ such that $\bigcup_{i=1}^{k} Q_i = \mathcal{P}$?

We get the problem MINIMUM TRIANGLE COVER by requiring that each piece Q_i is a triangle. We can now state the main result of the paper.

Theorem 1: MINIMUM CONVEX COVER and MINIMUM TRIANGLE COVER are $\exists\mathbb{R}$-complete.

Recall that O'Rourke [20] proved $\exists\mathbb{R}$-membership already in 1982. Alternatively, $\exists\mathbb{R}$-membership can be easily proven using the recent framework by Erickson, van der Hoog, and Miltzow [53]. This paper is therefore only concerned with proving $\exists\mathbb{R}$-hardness.

An implication of the result is that many natural approaches to finding small covers are bound to give suboptimal solutions in some cases. One might attempt to make covers consisting of pieces with corners chosen from some discrete set of points inside the given polygon. For instance, Eidenbenz and Widmayer [26] gave a $O(\log n)$-factor approximation algorithm for finding a minimum convex cover for a polygon $\mathcal{P}$ by choosing pieces with corners from a set of $O(n^{16})$ points. These points are obtained by first making the arrangement of lines through all pairs of corners of $\mathcal{P}$. We then construct all lines through pairs of intersection points in the first arrangement. The intersection points of this final arrangement, of which there are $O(n^{16})$, are the candidate set of corners of the pieces from which a cover is constructed.

Eidenbenz and Widmayer showed that for this particular set, there exists a cover consisting of at most 3 times more pieces than in the unrestricted optimum. Our result shows that the optimal solution cannot always be found by choosing the corners from such a set of points, even if the process of making new lines and intersection points was repeated any finite number of times. It follows from our construction that there are polygons where the corners of some pieces in any optimal cover have irrational coordinates with arbitrarily high algebraic degree. As an example, we can consider an equation such as $x^5 - 4x + 2 = 0$ [54], which has one real solution and that solution cannot be expressed by radicals (that is, the solution cannot be written using integers and basic arithmetic operations including powers and roots). We can then transform the equation to an instance of MCC, so that in the optimal cover, some pieces have coordinates that cannot be expressible by radicals. We state this as a corollary to our construction.

Corollary 2: There exists a polygon $\mathcal{P}$ such that in any minimum convex cover for $\mathcal{P}$, a piece has a corner with a coordinate that is not expressible by radicals.

C. Structure of the paper and basic techniques

We prove that ETR reduces to MCC using two intermediate problems: RANGE-ETR-INV and MRCC, which we define in Section II. Essentially, an instance of RANGE-ETR-INV is a conjunction of addition constrains of the form $x + y = z$ and inversion constraints of the form $x \cdot y = 1$. Each variable is restricted to a tiny subinterval of $[\frac{1}{2}, 2]$, and the goal is to decide if there exist values of the variables that satisfy all the addition and inversion constraints. This problem was developed recently in order to prove $\exists\mathbb{R}$-hardness of geometric packing problems [29].

The problem MRCC is a technical version of MCC where only some crucial parts of the polygon have to be covered, namely a specific set of *marked corners* and a set of *marked rectangles*. The introduction of the problem MRCC is crucial in order to keep our reductions manageable. In the proof, we show that everything else than the marked corners and rectangles can be covered in a generic way by adding some spikes to the boundary of the polygon. We therefore get an instance of MCC which is equivalent to the instance of MRCC that we are reducing from.

In Section III, we give the reduction from RANGE-ETR-INV to MRCC, which is the major part of the work.

One of our basic tools in the reduction to MRCC is to represent variables using specific horizontal *variable segments* contained in the constructed polygon; see Figure 1. A segment s representing a variable x corresponds to the interval $[1/2, 2]$ of values that x can attain, with the endpoints representing $x = 1/2$ and $x = 2$. The values of points in between the endpoints are defined by linear interpolation. A part of a segment s must be covered by a piece Q that also covers a particular marked corner of the polygon. The right-most (or left-most) point in the intersection $s \cap Q$ defines the value that Q represents on s. The rest of s must then be covered by another piece (which also covers another marked corner).

In the bottom of the polygon, we have some *base pockets* containing variable segments representing all the variables. We describe gadgets for addition inequalities, $x+y \geq z$ and $x+y \leq z$, and inversion inequalities, $x \cdot y \geq 1$ and $x \cdot y \leq 1$, and these gadgets are placed far to the right and above the base pockets. The gadgets also contain variable segments. We create some corridors that connect the gadgets and the base pockets, and they ensure that if a cover exists, then appropriate inequalities hold between the values represented in the base pockets and those in the gadgets. We are therefore able to conclude that if a cover exists, then so does a solution to the instance of RANGE-ETR-INV that we are reducing from.

Our reduction shares some resemblance with that from [28], and many of the geometric tools underlying our construction are similar to the ones used in that paper. In [28], variables are also represented by segments in the polygon, but the actual value is defined by the position of a guard standing on the segment (as shown in Figure 1), whereas in our case, it is the left- or right-most point covered by a convex piece. The geometric principles underlying the addition and inversion gadgets are the same in the two papers, but again the actual realizations of the gadgets are very different.

A key insight is that one convex piece in a base pocket can partially cover any number of variable segments, and thus represent numerous copies of the same variable, as sketched in Figure 1 (right). Each of these segments will have its own *propagation corner*, which is a marked corner. The propagation corner can then be covered by a *propagation piece*, which must cover the remaining part of the segment. Each propagation piece is very long and thin and sticks into a corridor far away. This is needed so that a variable that appears in many addition and inversion constraints can be copied into all the corresponding gadgets. In contrast to this, the reduction to ART GALLERY in [28] relied on the fact that a single guard can *look* into several different corridors. That technique is not possible to realize with convex pieces, because points in different corridors cannot see each other and hence no convex piece covers points in more than one corridor.

II. AUXILLIARY PROBLEMS: RANGE-ETR-INV AND MRCC

In this section we introduce the two problems RANGE-ETR-INV and MRCC that will be used as intermediate problems in our reduction from ETR to MCC.

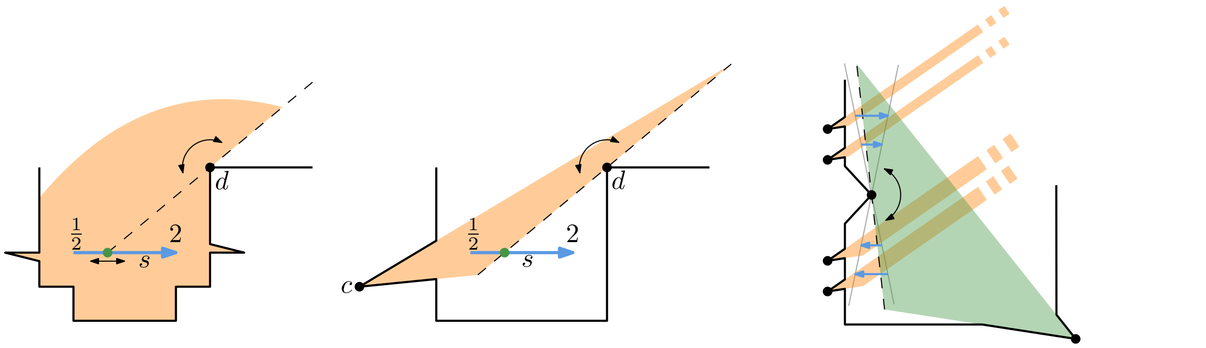

Figure 1. Representing variables in the paper [28] and this paper. Left: A guard segment s with a guard representing the value 1. When the guard moves to the right, more visibility is blocked due to the corner d. Middle: A piece covering the marked corner c represents the value 1 on the variable segment s. By rotating the dashed edge around the corner d, more or less of s will be covered. Right: Several variable segments can be partially covered by a single piece and then "copied" into different gadgets by the orange propagation pieces.

A. Range-Etr-Inv

Definition 3: An Etr-Inv *formula* $\Phi = \Phi(x_1, \ldots, x_n)$ is a conjunction

$$\left(\bigwedge_{i=1}^{n} 1/2 \leq x_i \leq 2\right) \wedge \left(\bigwedge_{i=1}^{m} C_i\right),$$

where $m \geq 0$ and each C_i is of one of the forms

$$x + y = z, \quad x \cdot y = 1$$

for $x, y, z \in \{x_1, \ldots, x_n\}$.

We can now define the problem Range-Etr-Inv as follows.

Instance: A tuple $\mathcal{I} = [\Phi, \delta, (I(x_1), \ldots, I(x_n))]$, where Φ is an Etr-Inv formula, $\delta > 0$ is a (small) number, and, for each variable $x \in \{x_1, \ldots, x_n\}$, $I(x)$ is an interval $I(x) \subseteq [1/2, 2]$ such that $|I(x)| \leq \delta$. Define $V(\Phi) := \{\mathbf{x} \in \mathbb{R}^n : \Phi(\mathbf{x})\}$. We are promised that $V(\Phi) \subset I(x_1) \times \cdots \times I(x_n)$.
Question: Is $V(\Phi) \neq \emptyset$?

The following theorem establishes that it suffices to make a reduction from Range-Etr-Inv in order to prove that MCC is $\exists\mathbb{R}$-hard.

Theorem 4 ([29]): Range-Etr-Inv is $\exists\mathbb{R}$-complete, even when $\delta = O(n^{-c})$ for any constant $c > 0$.

B. Minimum Restricted Convex Cover (MRCC)

We now turn our attention to the problem Minimum Restricted Convex Cover (MRCC). As we will see later in this section, MRCC can be reduced to MCC. Therefore, $\exists\mathbb{R}$-hardness for MRCC implies the same hardness for MCC. We define MRCC as follows.

Instance: A tuple $\langle \mathcal{P}, \mathcal{C}, \mathcal{R} \rangle$ consisting of the following parts:

- A simple polygon $\mathcal{P}$.
- A subset $\mathcal{C} = \{c_1, c_2, \ldots, c_k\}$ of the corners of $\mathcal{P}$ called *marked corners*.
- A set of pairwise disjoint axis-parallel rectangles $\mathcal{R} = \{R_1, \ldots, R_m\}$ contained in $\mathcal{P}$, called *marked rectangles.*

Question: Do there exist $k = |\mathcal{C}|$ convex polygons $Q_1, \ldots, Q_k$ contained in $\mathcal{P}$ such that $\mathcal{C} \subset \bigcup Q_i$ and $R \subset \bigcup Q_i$ for all marked rectangles $R \in \mathcal{R}$?

An instance of MRCC comes with a promise which we will explain below. In order to state the promise, we need some notation. For a corner c of a polygon $\mathcal{P}$, define $\Delta(c)$ to be the triangle with corners at c and the two neighbouring corners of $\mathcal{P}$. Note that $\Delta(c)$ may not be contained in $\mathcal{P}$, but that will always be the case when we use the notation.

For a marked rectangle $R \in \mathcal{R}$, define the *vertical bar* $\Pi_v(R)$ to be the region of points that are vertically visible from R, i.e., the points $p \in \mathcal{P}$ such that there is a point $q \in R$ where pq is a vertical line segment contained in $\mathcal{P}$. In a similar way, we define the *horizontal bar* $\Pi_h(R)$ as the points that are horizontally visible from R; see Figure 2.

In order to define the *vertical trapezoidation* of a polygon $\mathcal{P}$, we consider for each corner c of $\mathcal{P}$ the vertical segment that (i) contains c, (ii) is contained in $\mathcal{P}$, and (iii) is maximal with respect to inclusion. These segments partition $\mathcal{P}$ into trapezoids, some of which are only triangles. The *horizontal trapezoidation* is defined similarly, using horizontal segments.

Promise: For the instances $\mathcal{I} = \langle \mathcal{P}, \mathcal{C}, \mathcal{R} \rangle$ of MRCC that we will consider, we promise that the following properties hold, which we call the MRCC *promise*. For convenience, we have named each individual point of the promise.

1) *Skew triangle promise:* For each marked corner c, the triangle $\Delta(c)$ intersects the vertical and the horizontal line containing c only at the point c.
2) *Trapezoid generality promise:* In any trapezoid T in the vertical trapezoidation of $\mathcal{P}$, at least one right corner of T is not a corner of $\mathcal{P}$, unless the right corners of T are coincident (so that T is a triangle). Similarly

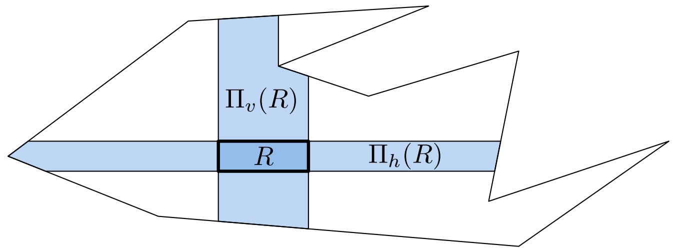

Figure 2. A polygon with a marked rectangle R the vertical and horizontal bars shown.

for the left corners. Likewise, in any trapezoid in the horizontal trapezoidation, at least one top corner is not a corner of $\mathcal{P}$ unless they are coincident, and similarly for the bottom corners.

3) *Bar intersection promise:* For any two marked rectangles R_i and R_j, the vertical bar of R_i and the horizontal bar of R_j are either disjoint or their intersection is contained in a marked rectangle.
4) *Broad cover promise:* If $\mathcal{I}$ has a cover, then there also exists a cover $\mathcal{Q}$ for $\mathcal{I}$ such that $\Delta(c) \subset \bigcup_{Q\in\mathcal{Q}} Q$ for all $c \in \mathcal{C}$.

Lemma 5: Suppose that MINIMUM RESTRICTED CONVEX COVER is $\exists\mathbb{R}$-hard, even when restricted to instances satisfying the MRCC promise. Then MINIMUM CONVEX COVER is also $\exists\mathbb{R}$-hard.

We here give a sketch of the proof of the lemma.

Proof (sketch): Let $\mathcal{I} := \langle \mathcal{P}, \mathcal{C}, \mathcal{R} \rangle$ be an instance of MRCC that satisfies the promise. We here describe the high-level idea behind our reduction to MCC. We construct a larger polygon $\mathcal{P}' \supset \mathcal{P}$ by adding some number s of spikes to $\mathcal{P}$. We add the spikes so that they can be covered by triangles that also cover everything of $\mathcal{P}$ except the marked rectangles $\mathcal{R}$ and regions close to the marked corners $\mathcal{C}$. In order to carry out this idea, we consider a vertical and a horizontal trapezoidation of $\mathcal{P}$ except for the vertical and horizontal bars of the marked rectangles. We add the spikes to some of the trapezoids and make sure that nothing outside these trapezoids can be covered when covering the spikes; see Figure 3.

We then get the instance $\mathcal{I}' := \langle \mathcal{P}', k' \rangle$ of MCC, where the parameter k' is defined as $k' := k + s$ with $k := |\mathcal{C}|$. In the full proof we argue that a cover for $\mathcal{I}'$ using $k' = k + s$ convex polygons contains s polygons covering the added spikes. The spikes have been chosen so that they cannot be covered while also covering $\mathcal{R}$ nor $\mathcal{C}$, so the remaining k polygons must cover $\mathcal{R}$ and $\mathcal{C}$ and thus constitute a cover for $\mathcal{I}$. On the other hand, given a cover $\mathcal{Q}$ for $\mathcal{I}$, one can add s convex polygons contained in $\mathcal{P}'$ that cover the spikes and everything else of $\mathcal{P}'$ not covered by $\mathcal{Q}$, and thus obtain a cover for $\mathcal{I}'$. Therefore, the instances are equivalent. In order to carry out the details, all the parts of the MRCC promise turn out to be needed. ■

III. REDUCTION FROM RANGE-ETR-INV TO MRCC

Let $\mathcal{I}_1$ be an instance of RANGE-ETR-INV with an ETR-INV formula Φ of n variables $X := \{x_1, \ldots, x_n\}$. We show that there exists an instance $\mathcal{I}_2 := \langle \mathcal{P}, \mathcal{C}, \mathcal{R} \rangle$ of MRCC which can be computed in polynomial time such that Φ has a solution if and only if $\mathcal{I}_2$ has a cover. Recall that a cover $\mathcal{Q}$ of $\mathcal{I}_2$ is a set of $|\mathcal{C}|$ convex polygons contained in $\mathcal{P}$ that together cover the marked corners $\mathcal{C}$ and the marked rectangles in $\mathcal{R}$. A high-level sketch of the polygon $\mathcal{P}$ is shown in Figure 4.

A. Basic tools

Before describing the actual geometry of the instance $\mathcal{I}_2$, we explain the basic tools that underlie our construction.

1) Invisibility property: We make sure that our constructed instance $\mathcal{I}_2$ satisfies the following *invisibility property*: No two marked corners can see each other.

For a marked corner $c \in \mathcal{C}$, let e_1 and e_2 be the two edges of $\mathcal{P}$ incident at c. Except for one marked corner in each $\leq$-inversion gadget, it holds that c is the lower endpoint of both e_1 and e_2. The invisibility thus follows for all other corners. Additionally, it holds without exception that c is either the left endpoint of both e_1 and e_2 or the right endpoint of both.

Note that since a cover $\mathcal{Q}$ for $\mathcal{I}_2$ consists of $|\mathcal{C}|$ pieces and no piece can cover more than one marked corner by the invisibility property, it follows that each piece $Q \in \mathcal{Q}$ contains exactly one marked corner $c \in \mathcal{C}$.

2) Critical segments and the bi-cover property: In the instance $\mathcal{I}_2$ that we construct, a marked rectangle $R \in \mathcal{R}$ often contains a special line segment s which we call a *critical segment*; see Figure 5. The following property turns out to be crucial, and we call it the *bi-cover property*: For each critical segment s, there are at least two and at most three marked corners that can see s. There are three if and only if s is not horizontal, and in that case, one of these three corners h is a *helper corner* (to be described later), and h is incident to an edge contained in the extension of s. In particular, s is on the boundary of the visibility polygon of h. Otherwise, if s is horizontal, then s is a variable segment, to be described in more detail below.

A critical segment will always be contained in a marked rectangle, but we will also make some marked rectangles

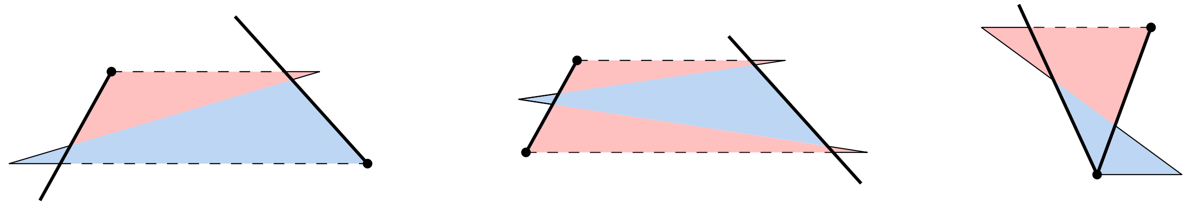

Figure 3. The figure shows how we add spikes to trapezoids in a trapezoidation of $\mathcal{P}$. The fat edges are edges of $\mathcal{P}$ and the dots show the corners of $\mathcal{P}$. The trapezoid generality promise is used to ensure that we can avoid making spikes where $\mathcal{P}$ has a corner, unless the trapezoid is a triangle. We can therefore prevent the spikes from overlapping with $\mathcal{P}$. Left: When the trapezoid has a left and a right corner that are not corners of $\mathcal{P}$, we add two spikes such that the trapezoid and the two spikes can be covered by two triangles. Middle: When both trapezoid corners in the same side are also corners of $\mathcal{P}$, we add three spikes. Right: When the trapezoid is a triangle, we add two spikes as shown.

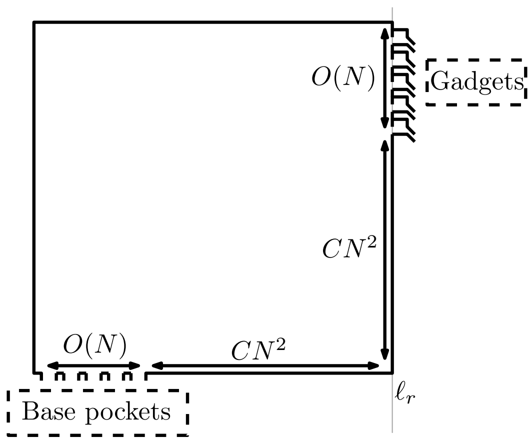

Figure 4. A high-level sketch of our construction. The polygon $\mathcal{P}$ will have dimensions quadratic in N, which is proportional to the size of the ETR-INV instance Φ.

that do not contain critical segments. Often these are introduced to satisfy the bar intersection promise of the MRCC instance, and they will be defined as the intersection of a vertical and a horizontal bar of two other marked rectangles.

Consider a critical segment s and the marked rectangle R containing s. Let $c_1, c_2 \in \mathcal{C}$ be the marked corners that see s and are not helper corners. These are called *lever corners*, and we say that c_1 and c_2 are *responsible* for s and for the marked rectangle R. Since a cover for $\mathcal{I}_2$ must cover all of the rectangle R, it follows from the invisibility property and the bi-cover property that c_1 and c_2 must be covered by two pieces Q_1 and Q_2 that also together cover s. All the marked corners that we make will be either helper corners or lever corners. A piece that covers a lever corner is called a *lever piece* and otherwise it is a *helper piece*.

3) Helper corners: As mentioned above, we will make two types of marked corners, namely lever corners and helper corners. Furthermore, we will make two kinds of helper corners. In general, we make a helper corner when a marked rectangle or a part of a marked rectangle cannot be covered by lever pieces. One situation is when we introduce a marked rectangle with the sole purpose of satisfying the bar intersection promise described in Section II-B; see Figure 6 for an example.

The other situation is when we have a marked rectangle R containing a critical segment s which is not horizontal; see Figure 5 (right). In this case, the two pieces Q_1 and Q_2 covering the lever corners c_1 and c_2 responsible for s can cover the part of R above or below s, but they cannot cover everything on the other side. We then make a helper corner h so that one of the edges incident at h is contained in the extension of s. Thus, s is on the boundary of the visibility polygon of h, and a piece covering h can cover the part of R on the other side of s.

If a critical segment s is horizontal, we make the marked rectangle R so that s is contained in the bottom or top edge of R, as seen in Figure 5 (left). Then the pieces covering the responsible corners c_1 and c_2 will be able to cover all of R without the use of a helper corner.

4) Representing variables: Each variable $x \in X$ is represented by a collection of *variable segments*, which are horizontal critical segments. Consider one variable segment $s := ab$ where a is to the left of b, and assume that s represents x. The segment s can be oriented either to the left or to the right. Each point p on s represents a value of x in the interval $[\frac{1}{2}, 2]$, which we denote as $\mathbf{x}(s,p)$. If s is right-oriented, then we define $\mathbf{x}(s,p) := 1/2 + \frac{3\|ap\|}{2\|ab\|}$, and otherwise we define $\mathbf{x}(s,p) := 1/2 + \frac{3\|bp\|}{2\|ab\|}$. In particular, $\mathbf{x}(s,\cdot)$ is a linear map from s to $[\frac{1}{2}, 2]$.

Consider a variable segment s and a lever corner c responsible for s. In a cover $\mathcal{Q}$ for $\mathcal{I}_2$, the piece Q covering c specifies a value at s, defined as follows. Each lever corner comes with a *pivot* which is also a corner of $\mathcal{P}$. Let d be the pivot of c. It is then always the case that d has another y-coordinate than s. For a point $q \in s$, consider the line $\overleftrightarrow{qd}$ through q and the pivot d, which is never horizontal by the previous remark, so the line has a well-defined left and right side. Furthermore, it will be the case that whether c is to the right or left of $\overleftrightarrow{qd}$ does not depend on the particular point q. We say that c is the *right-responsible* or *left-responsible* for s, depending on this side. For instance, in Figure 8, c_c is right-responsible and c_b is left-responsible for the segment s_3.

If c is a right-responsible, then we define $\mathbf{p}_{\mathcal{Q}}(s,c)$ as

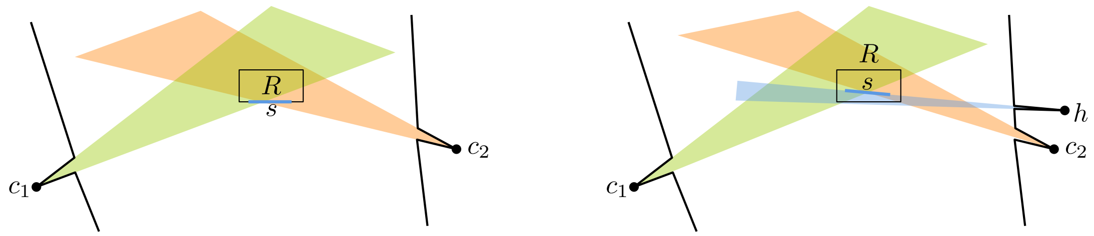

Figure 5. Left: A horizontal critical segment s contained in the bottom edge of its marked rectangle R. Two pieces cover R and the responsible lever corners c_1 and c_2. Right: A critical segment that is not horizontal. We use a helper corner h with an incident edge contained in the extension of s. The three pieces cover R and the responsible lever corners c_1, c_2 and the helper corner h.

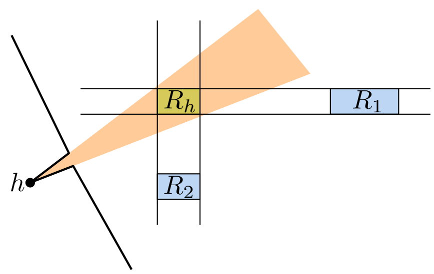

Figure 6. In order to satisfy the bar intersection promise, we make a marked rectangle R_h which is the intersection of the horizontal bar of R_1 and the vertical bar of R_2. We then make a helper corner h so that a piece covering h can also cover R_h.

the leftmost point in $Q \cap s$ (or the right endpoint of s if $Q \cap s = \emptyset$). Otherwise, we define $\mathbf{p}_{\mathcal{Q}}(s, c)$ as the rightmost point in $Q \cap s$ (or the left endpoint of s if $Q \cap s = \emptyset$). We now define the value represented by $\mathcal{Q}$ at s with respect to c as $\mathbf{v}_{\mathcal{Q}}(s, c) := \mathbf{x}(s, \mathbf{p}_{\mathcal{Q}}(s, c))$. We will usually use the simplified notation $\mathbf{p}(s, c)$ and $\mathbf{v}(s, c)$ if $\mathcal{Q}$ is clear from the context.

Recall that for each critical segment s, there are two responsible lever corners c_1, c_2, so that in a cover for $\mathcal{I}_2$, the pieces covering c_1 and c_2 must together cover s. It will always be the case in our construction that one is left-responsible and the other is right-responsible. We then have the following observation.

Observation 6: Consider a variable segment $s := ab$, where a is the left endpoint, and let c_1 be the left-responsible for s and c_2 the right-responsible. Consider any cover $\mathcal{Q}$. Since the pieces covering c_1 and c_2 also cover s, we have $\mathbf{p}(s, c_2) \in a\mathbf{p}(s, c_1)$. It follows that if s is right-oriented, then $\mathbf{v}(s, c_1) \geq \mathbf{v}(s, c_2)$. Otherwise, $\mathbf{v}(s, c_1) \leq \mathbf{v}(s, c_2)$.

We are now ready to state the theorem that expresses that our reduction works as desired.

Theorem 7: Suppose that there is a solution to Φ. Then our constructed instance $\mathcal{I}_2 := \langle \mathcal{P}, \mathcal{C}, \mathcal{R} \rangle$ has a cover, and any cover $\mathcal{Q}$ of $\mathcal{I}_2$ *specifies* a solution to Φ, as described by the following two properties.

- Each variable $x \in X$ is specified *consistently* by $\mathcal{Q}$: For each variable segment s representing x, consider the corners $c_1, c_2 \in \mathcal{C}$ that are responsible for s. Then $\mathbf{x}(s, c_1) = \mathbf{x}(s, c_2)$, and this value is the same for all segments s representing x.
- The cover $\mathcal{Q}$ is *feasible*, i.e., the values of X thus specified is a solution to Φ.

On the other hand, if there is no solution to Φ, then there is no cover for $\mathcal{I}_2$. The instance $\mathcal{I}_2$ can be computed in time polynomial in the input size $|\mathcal{I}_1|$. It thus follows that MRCC is $\exists\mathbb{R}$-complete.

5) Lever mechanism: The basic mechanism of our construction is similar to that of a lever (as known from mechanics) and works as follows; see Figure 7. As mentioned earlier, a marked corner will either be a helper corner or a *lever corner*, and each lever corner has a *pivot* d, which is also a corner of $\mathcal{P}$ but not a marked corner. The two corners c and d can always see each other. There is also a set of marked rectangles $\mathcal{S} := \{R_1, \ldots, R_l\}$, $l \geq 2$, for which c is one of the responsible corners. In particular, the rectangles $\mathcal{S}$ can be seen from c. One or more of the rectangles $\mathcal{S}_t := \{s_1, \ldots, s_m\}$, $m < l$, are above d and one or more $\mathcal{S}_b := \{s_{m+1}, \ldots, s_l\}$ are below d. Recall that some marked rectangles contain variable segments. The variable segments contained in the rectangles $\mathcal{S}$ all represent the same variable $x \in X$, and the variable segments in $\mathcal{S}_t$ have the same orientation, which is the opposite of that of the variable segments in $\mathcal{S}_l$.

A piece of a cover for $\mathcal{I}_2$ that covers a lever corner c is called a *lever piece*. A lever piece Q can cover a part of the rectangles $\mathcal{S}$, but the pivot d is preventing Q from covering all the rectangles entirely. Suppose that we want the piece Q to have non-empty intersection with all rectangles in $\mathcal{S}$ and cover as much as possible of each one. Such a piece must have an edge e that contains d, as we could otherwise cover more of the top rectangles $\mathcal{S}_t$ or the bottom rectangles $\mathcal{S}_b$. We say that e is a *lever edge*, and e has the property that it intersects all the rectangles $\mathcal{S}$. Furthermore, the lever piece Q contains the part of each rectangle R on the same side of e, either to the left (if c is the left-responsible for the rectangles $\mathcal{S}$) or to the right (if c is the right-responsible). We can now imagine that the edge e rotates around the pivot d. Doing so, the piece Q will cover more of the top rectangles $\mathcal{S}_t$ and less of the bottom rectangles $\mathcal{S}_b$, or vice versa.

Consider the subset $\mathcal{S}' \subseteq \mathcal{S}$ that contain critical segments

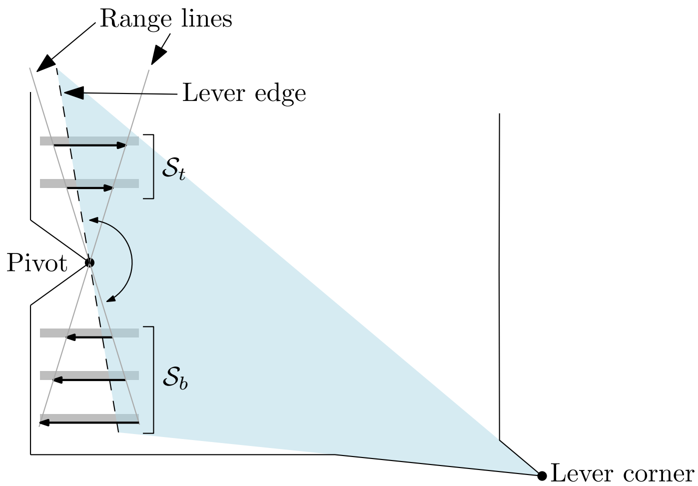

Figure 7. A lever piece partially covering the associated marked rectangles, each of which contains a variable segment. Two overlapping and maximal lever pieces are shown, and the lever edge of one of them is shown as a dashed segment.

(the reader can conveniently imagine the case $\mathcal{S}' = \mathcal{S}$; the only exception will be in the addition gadgets). Then there are two *range lines*, both of which pass through d, and each critical segment in $\mathcal{S}'$ has an endpoint on each range line. When the lever edge of Q coincides with a range line, then the critical segments in one of the sets $\mathcal{S}_b$ or $\mathcal{S}_t$ are completely covered, while the segments in the other set are not covered at all.

6) Chains of inequalities: The following observation follows immediately from the description above and is one of our main tools for copying values from one variable segment to another.

Observation 8: Let $R_t \in \mathcal{S}_t$ and $R_b \in \mathcal{S}_b$ be two marked rectangles and assume that they contain variable segments s_t and s_b, respectively. Suppose that the lever corner c is the right-responsible for both s_t and s_b. If s_t is right-oriented and s_b is left-oriented, then $\mathbf{v}(s_t, c) \geq \mathbf{v}(s_b, c)$. If s_t is left-oriented and s_b is right-oriented, then $\mathbf{v}(s_t, c) \leq \mathbf{v}(s_b, c)$. If on the other hand c is the left-responsible, then the inequalities are swapped. Equalities hold if and only if Q has a lever edge, and then Q represents the same value on all variable segments in $\mathcal{S}$.

Using chains of lever mechanisms, we can now make chains of inequalities of the values that the lever pieces represent of a variable $x \in X$, using Observations 6 and 8 alternatingly. See for instance Figure 8. The observations give $\mathbf{v}(s_1, c_a) \leq \mathbf{v}(s_2, c_a) \leq \mathbf{v}(s_2, c_b) \leq \mathbf{v}(s_3, c_b) \leq \mathbf{v}(s_3, c_c) \leq \mathbf{v}(s_4, c_c)$. The figure shows the case that Q_a, Q_b, Q_c all have lever edges (containing the pivots d_a, d_b, d_c, respectively) and cover s_2 and s_3 with no overlap, and then we have equalities $\mathbf{v}(s_1, c_a) = \mathbf{v}(s_2, c_a) = \mathbf{v}(s_2, c_b) = \mathbf{v}(s_3, c_b) = \mathbf{v}(s_3, c_c) = \mathbf{v}(s_4, c_c)$. Otherwise, one or more of the inequalities will be strict.

7) Infinitesimal range: Recall that each variable $x \in X$ comes with an interval $I(x) \subset [\frac{1}{2}, 2]$ of size $|I(x)| \leq \delta$, for some small δ, such that all solutions to Φ are guaranteed to be contained in these intervals. In this reduction, we use $\delta := n^{-7}$. Consider a variable segment s representing x. Let s' be the part of s such that the points of s' correspond to the range $I(x)$. We then call s' the *restricted range* of s, and since $\delta := n^{-7}$, we have that $\frac{\|s'\|}{\|s\|} = O(n^{-7})$. Since we can assume n to be arbitrarily large, we can usually think of s' as being just a single point on s. In our construction, we will always show the full segment s in the figures, and we will mark the position of the restricted range s' as a single point on s. We just need to ensure that the construction is generic in the sense that it works for every possible placement of the restricted range s' on the full range s. In particular, when we construct a lever corner $c \in \mathcal{C}$, it may be needed to move c and the edges of $\mathcal{P}$ incident at c depending on where the restricted range s' is located at s.

We need the lever corner c to be placed so that it can see all points on the restricted range s' of s. When that is the case for all segments s that c is responsible for, then a lever edge of a piece covering c will have the freedom to rotate over all of s', thus representing all values in the required range $I(x)$. See Figure 9 for examples with enough freedom and too little freedom.

In many parts of our construction, it is obvious that c can see all of the restricted ranges of the variable segments. The more delicate case happens for the propagation corners, to be described in more detail in the next section. These are responsible for variable segments in the base pockets, but also for critical segments in the corridors. Here, it is not clear that the propagation corners can see enough of the critical segments in the corridors, but as is shown in the full version, it works for our choice of $\delta := n^{-7}$.

We choose the marked rectangles to be similarly tiny, as they only need to contain the restricted ranges of the

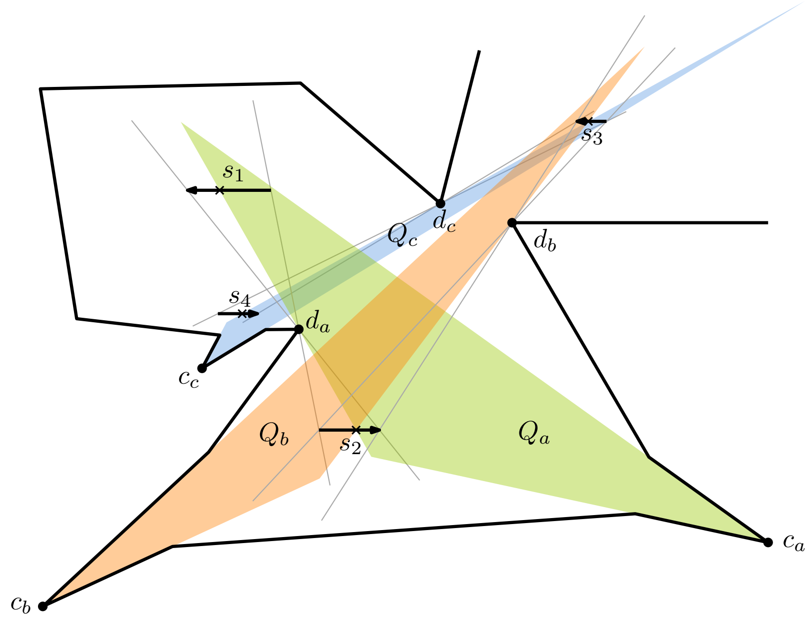

Figure 8. Three lever corners and lever pieces that together cover the variable segments s_2 and s_3.

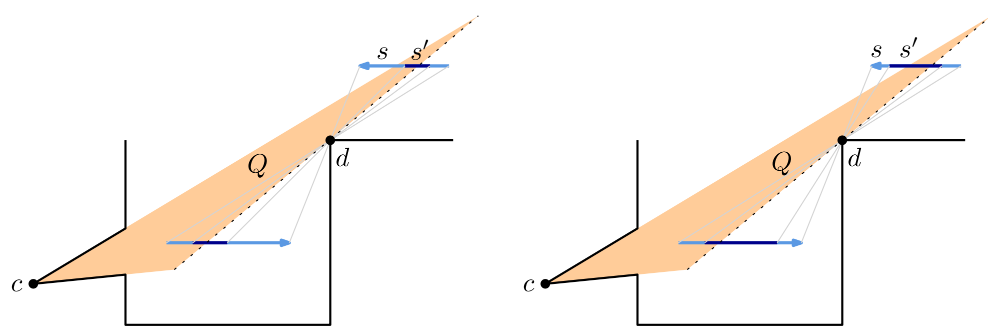

Figure 9. The (light and dark) blue arrows show the full variable segments representing a variable x and the dark blue parts are the parts corresponding to the restricted range $I(x)$. To the left, the dotted lever edge has enough freedom to rotate over the entire restricted range, since the corner c sees all points on the restricted ranges of the variable segments. To the right, c does not see all of the upper restricted range s', so not all values can be represented.

critical segments. Hence, the vertical and horizontal bars of the rectangles can be thought of as vertical or horizontal line segments contained in $\mathcal{P}$.

B. High-level description of $\mathcal{P}$

We are now ready to describe the actual design of the polygon $\mathcal{P}$ in more detail; see Figure 4 for a sketch of the entire polygon. In the bottom of $\mathcal{P}$, there are pockets with lever corners representing all the variables; see Figure 10. Each pocket is a polygon in $\mathcal{P}$ bounded by a chain of $\partial\mathcal{P}$ and from above by one horizontal chord of $\mathcal{P}$. A pocket can contain a large number of variable segments $\mathcal{S}$, all representing the same variable $x \in X$. In the lower right corner of the pocket, there is a lever corner, which we call a *base corner*, and which is the right-responsible lever corner for all the segments $\mathcal{S}$. A piece covering a base corner is called a *base piece*. For each variable segment, the left-responsible lever corner is called a *propagation corner*. The pivot of a propagation corner is very far away approximately in direction $(1,1)$, namely at the door of a corridor that leads into a gadget. A piece covering a propagation corner is called a *propagation piece*, and in any cover, the propagation pieces will be very skinny and long, sticking out of the base pockets and into the corridors, where each of them is responsible for covering another critical segment.

The bottom part of $\mathcal{P}$ consists of a collection of $3n$ base pockets. In order from left to right, we denote the pockets as $P_1,\dots,P_{3n}$. The pockets $P_1,\dots,P_n$ represent the variables $x_1,\dots,x_n$, respectively, as do the pockets $P_{n+1},\dots,P_{2n}$, and $P_{2n+1},\dots,P_{3n}$. At the right side of $\mathcal{P}$, there are some *corridors* attached, each of which leads into a *gadget*. The *doors* to the corridors are line segments contained in a vertical line ℓ_r (in fact, we will move the doors slightly away from ℓ_r in order to satisfy the trapezoid

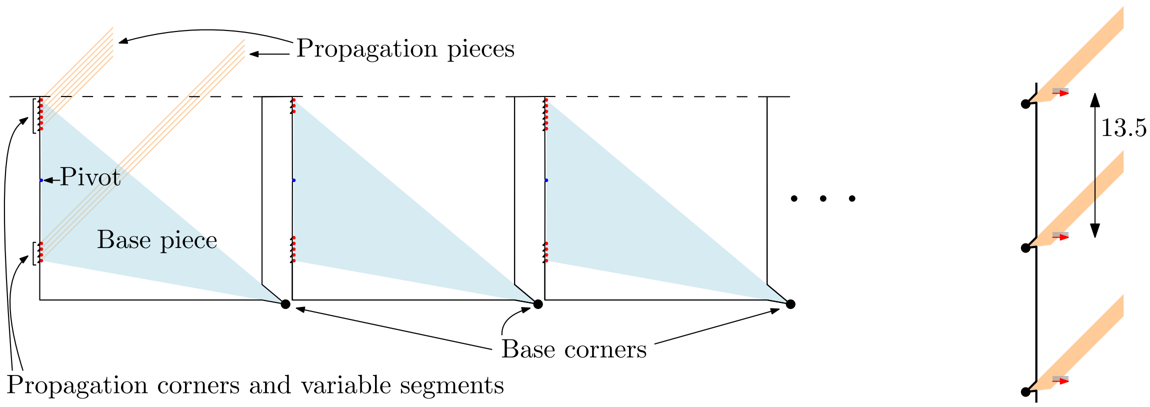

Figure 10. Left: A sketch of three consecutive base pockets with various numbers of variable segments. The variable segments are shown as red dots. Right: A closeup of three consecutive variable segments and the marked rectangles containing them.

generality promise, but that does not make any conceptual difference). The gadgets also contain marked rectangles and corners, and they are used to impose dependencies between the pieces covering variable segments in the base pockets. The corridors are used to make dependencies between the base pockets and the gadgets. Each gadget corresponds to a constraint of one of the types $x+y \geq z$, $x+y \leq z$, $x \cdot y \geq 1$, $x \cdot y \leq 1$, and $x \leq y$. The first four types of constraints are used to encode the dependencies between the variables in X as specified by Φ, whereas the latter constraint is used to make sure that the three base pockets representing a variable $x \in X$ specify the value of x consistently.

The reason that we need three pockets P_i, P_{i+n}, P_{i+2n} representing each variable x_i is that in the addition gadgets, such as the one encoding the restriction $x_i + x_j \geq x_l$, we need the base pockets representing x_l, x_j, x_i, that we connect to the gadget, to appear in the specific order $P_{i'}, P_{j'}, P_{l'}$ from left to right, respectively. This is obtained by choosing $i' := i$, $j' := j + n$, and $l' := l + 2n$.

Acknowledgements

I thank Aurélien Ooms for useful discussions in the initial phases of this work, and Joseph O'Rourke for sending me his paper [20]. This research is supported by the Starting Grant 1054-00032B from the Independent Research Fund Denmark under the Sapere Aude research career programme. BARC is supported by the VILLUM Foundation grant 16582.

References

[1] M. Abrahamsen, "Covering polygons is even harder," 2021, full version, https://arxiv.org/abs/2106.02335.

[2] B. Chazelle, "Approximation and decomposition of shapes," in *Algorithmic and Geometric Aspects of Robotics*, ser. Advances in Robotics, J. T. Schwartz and C.-K. Yap, Eds. Lawrence Erlbaum Associates, 1987, vol. 1, pp. 145–185.

[3] J. M. Keil and J.-R. Sack, "Minimum decompositions of polygonal objects," in *Computational Geometry*, ser. Machine Intelligence and Pattern Recognition, G. T. Toussaint, Ed. North-Holland, 1985, vol. 2, pp. 197–216.

[4] T. C. Shermer, "Recent results in art galleries," *Proceedings of the IEEE*, vol. 80, no. 9, pp. 1384–1399, 1992.

[5] B. Chazelle and L. Palios, "Decomposition algorithms in geometry," in *Algebraic Geometry and its applications*, C. L. Bajaj, Ed. Springer, 1994, pp. 419–447.

[6] J. M. Keil, "Polygon decomposition," in *Handbook of computational geometry*, J.-R. Sack and J. Urrutia, Eds., 1999, ch. 11, pp. 491–518.

[7] J. O'Rourke, *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

[8] J. O'Rourke, S. Suri, and C. D. Tóth, "Polygons," in *Handbook of discrete and computational geometry*, third edition ed., J. E. Goodman and J. O'Rourke, Eds., 2018, ch. 30, pp. 787–810.

[9] T. Pavlidis, "Analysis of set patterns," *Pattern Recognition*, vol. 1, no. 2, pp. 165–178, 1968.

[10] ——, "Structural pattern recognition: Primitives and juxtaposition relations," in *Frontiers of Pattern Recognition*, S. Watanabe, Ed. Academic Press, 1972, pp. 421–451.

[11] ——, "Representation of figures by labeled graphs," *Pattern Recognition*, vol. 4, no. 1, pp. 5–17, 1972.

[12] H.-Y. F. Feng and T. Pavlidis, "Decomposition of polygons into simpler components: Feature generation for syntactic pattern recognition," *IEEE Transactions on Computers*, vol. C-24, no. 6, pp. 636–650, 1975.

[13] T. Pavlidis, *Structural Pattern Recognition*, ser. Springer Series in Electrophysics, 1977, vol. 1.

[14] ——, "A review of algorithms for shape analysis," *Computer Graphics and Image Processing*, vol. 7, no. 2, pp. 243–258, 1978.

[15] M. R. Garey and D. S. Johnson, *Computers and intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., 1979.

[16] R. S. V. Rodrigues, J. F. M. Morgado, and A. J. P. Gomes, "Part-based mesh segmentation: A survey," *Computer Graphics Forum*, vol. 37, no. 6, pp. 235–274, 2018.

[17] T. Suk, C. H. IV, and J. Flusser, "Decomposition of binary images – A survey and comparison," *Pattern Recognition*, vol. 45, no. 12, pp. 4279–4291, 2012.

[18] S. Asafi, A. Goren, and D. Cohen-Or, "Weak convex decomposition by lines-of-sight," *Computer Graphics Forum*, vol. 32, no. 5, pp. 23–31, 2013.

[19] J. O'Rourke and K. J. Supowit, "Some np-hard polygon decomposition problems," *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 181–190, 1983.

[20] J. O'Rourke, "The complexity of computing minimum convex covers for polygons," in *20th Annual Allerton Conference on Communication, Control, and Computing*, 1982, pp. 75–84.

[21] B. Chazelle and D. P. Dobkin, "Optimal convex decompositions," in *Computational Geometry*, ser. Machine Intelligence and Pattern Recognition, G. T. Toussaint, Ed. North-Holland, 1985, vol. 2, pp. 63–133.

[22] A. Aggarwal, "The art gallery theorem: Its variations, applications and algorithmic aspects," Ph.D. dissertation, Department of Electrical Engineering and Computer Science, Johns Hopkins University, 1984.

[23] J. C. Culberson and R. A. Reckhow, "Covering polygons is hard," *Journal of Algorithms*, vol. 17, no. 1, pp. 2–44, 1994.

[24] D. T. Lee and A. K. Lin, "Computational complexity of art gallery problems," *IEEE Transactions on Information Theory*, vol. 32, no. 2, pp. 276–282, 1986.

[25] D. S. Johnson, "The NP-completeness column: An ongoing guide," *Journal of Algorithms*, vol. 3, no. 2, pp. 182–195, 1982.

[26] S. J. Eidenbenz and P. Widmayer, "An approximation algorithm for minimum convex cover with logarithmic performance guarantee," *SIAM Journal on Computing*, vol. 32, no. 3, pp. 654–670, 2003.

[27] T. Christ, "Beyond triangulation: Covering polygons with triangles," in *Algorithms and Data Structures – 12th International Symposium (WADS 2011)*, 2011, pp. 231–242.

[28] M. Abrahamsen, A. Adamaszek, and T. Miltzow, "The art gallery problem is $\exists\mathbb{R}$-complete," *Journal of the ACM*, 2021, to appear. Preliminary versions presented at STOC 2018 and SoCG 2017.

[29] M. Abrahamsen, T. Miltzow, and N. Seiferth, "Framework for $\exists\mathbb{R}$-completeness of two-dimensional packing problems," in *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS 2020)*, 2020.

[30] J. Canny, "Some algebraic and geometric computations in pspace," in *20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, 1988, pp. 460–467.

[31] J. Richter-Gebert and G. M. Ziegler, "Realization spaces of 4-polytopes are universal," *Bulletin of the American Mathematical Society*, vol. 32, no. 4, pp. 403–412, 1995.

[32] N. E. Mnëv, "The universality theorems on the classification problem of configuration varieties and convex polytopes varieties," in *Topology and geometry – Rohlin seminar*, O. Y. Viro, Ed., 1988, pp. 527–543.

[33] P. W. Shor, "Stretchability of pseudolines is NP-hard," in *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, ser. DIMACS – Series in Discrete Mathematics and Theoretical Computer Science, P. Gritzmann and B. Sturmfels, Eds., 1991, vol. 4, pp. 531–554.

[34] D. Bienstock, "Some provably hard crossing number problems," *Discrete & Computational Geometry*, vol. 6, no. 3, pp. 443–459, 1991.

[35] M. G. Dobbins, L. Kleist, T. Miltzow, and P. Rzazewski, "$\forall\exists\mathbb{R}$-completeness and area-universality," in *Graph-Theoretic Concepts in Computer Science – 44th International Workshop (WG 2018)*, 2018, pp. 164–175.

[36] A. Lubiw, T. Miltzow, and D. Mondal, "The complexity of drawing a graph in a polygonal region," in *Proceedings on the 26th International Symposium on Graph Drawing and Network Visualization (GD 2018)*, 2018, pp. 387–401.

[37] J. Cardinal, S. Felsner, T. Miltzow, C. Tompkins, and B. Vogtenhuber, "Intersection graphs of rays and grounded segments," vol. 22, no. 2, 2018, pp. 273–295.

[38] J. Cardinal and U. Hoffmann, "Recognition and complexity of point visibility graphs," *Discrete & Computational Geometry*, vol. 57, no. 1, pp. 164–178, 2017.

[39] R. J. Kang and T. Müller, "Sphere and dot product representations of graphs," *Discrete & Computational Geometry*, vol. 47, no. 3, pp. 548–568, 2012.

[40] C. McDiarmid and T. Müller, "Integer realizations of disk and segment graphs," *Journal of Combinatorial Theory, Series B*, vol. 103, no. 1, pp. 114–143, 2013.

[41] J. Erickson, "Optimal curve straightening is $\exists\mathbb{R}$-complete," 2019, preprint, https://arxiv.org/abs/1908.09400.

[42] M. L. T. Berthelsen and K. A. Hansen, "On the computational complexity of decision problems about multi-player Nash equilibria," in *International Symposium on Algorithmic Game Theory*. Springer, 2019, pp. 153–167.

[43] J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbod, "$\exists$r-completeness for decision versions of multi-player (symmetric) nash equilibria," *ACM Transactions on Economics and Computation*, vol. 6, no. 1, pp. 1:1–1:23, 2018.

[44] Z. Abel, E. D. Demaine, M. L. Demaine, S. Eisenstat, J. Lynch, and T. B. Schardl, "Who needs crossings? Hardness of plane graph rigidity," in *32nd International Symposium on Computational Geometry (SoCG 2016)*, 2016, pp. 3:1–3:15.

[45] M. Schaefer, "Realizability of graphs and linkages," in *Thirty Essays on Geometric Graph Theory*, J. Pach, Ed. Springer-Verlag New York, 2013, ch. 23, pp. 461–482.

[46] M. Schaefer and D. Štefankovič, "Fixed points, Nash equilibria, and the existential theory of the reals," *Theory of Computing Systems*, vol. 60, no. 2, pp. 172–193, 2017.

[47] M. G. Dobbins, A. Holmsen, and T. Miltzow, "A universality theorem for nested polytopes," 2019, preprint, arxiv.org/abs/1908.02213.

[48] Y. Shitov, "A universality theorem for nonnegative matrix factorizations," 2016, preprint, https://arxiv.org/abs/1606.09068.

[49] ——, "The complexity of positive semidefinite matrix factorization," *SIAM Journal on Optimization*, vol. 27, no. 3, pp. 1898–1909, 2017.

[50] J. Cardinal, "Computational geometry column 62," *SIGACT News*, vol. 46, no. 4, pp. 69–78, 2015.

[51] J. Matoušek, "Intersection graphs of segments and $\exists\mathbb{R}$," 2014, preprint, https://arxiv.org/abs/1406.2636.

[52] M. Schaefer, "Complexity of some geometric and topological problems," in *Proceedings of the 17th International Symposium on Graph Drawing (GD 2009)*. Springer, 2009, pp. 334–344.

[53] J. Erickson, I. van der Hoog, and T. Miltzow, "Smoothing the gap between NP and ER," in *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS 2020)*, 2020, pp. 1022–1033.

[54] J. R. Bastida and R. Lyndon, *Field Extensions and Galois Theory*, ser. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984.