

Framework for $\exists\mathbb{R}$ -Completeness of Two-Dimensional Packing Problems

Received Apr 21, 2022
Accepted Feb 15, 2024
Published Apr 26, 2024

Key words and phrases
Packing, Polygons,
 $\exists\mathbb{R}$ -completeness

Mikkel Abrahamsen^a ✉ 

Tillmann Miltzow^b ✉ 

Nadja Seiferth^c ✉ 

^a Basic Algorithms Research
Copenhagen (BARC), University of
Copenhagen.

^b Utrecht University.

^c Freie Universität Berlin.

ABSTRACT. The aim in packing problems is to decide if a given set of pieces can be placed inside a given container. A packing problem is defined by the types of pieces and containers to be handled, and the motions that are allowed to move the pieces. The pieces must be placed so that in the resulting placement, they are pairwise interior-disjoint. We establish a framework which enables us to show that for many combinations of allowed pieces, containers and motions, the resulting problem is $\exists\mathbb{R}$ -complete. This means that the problem is equivalent (under polynomial time reductions) to deciding whether a given system of polynomial equations and inequalities with integer coefficients has a real solution.

We consider packing problems where only translations are allowed as the motions, and problems where arbitrary rigid motions are allowed, i.e., both translations and rotations. When rotations are allowed, we show that it is an $\exists\mathbb{R}$ -complete problem to decide if a set of convex polygons, each of which has at most 7 corners, can be packed into a square.

Restricted to translations, we show that the following problems are $\exists\mathbb{R}$ -complete: (i) pieces bounded by segments and hyperbolic curves to be packed in a square, and (ii) convex polygons to be packed in a container bounded by segments and hyperbolic curves.

The paper was presented at the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2020) [6]. Mikkel Abrahamsen is supported by 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. Tillmann Miltzow is supported by the Netherlands Organisation for Scientific Research (NWO) under project no. 016.Veni.192.25. Nadja Seiferth is supported by the German Science Foundation within the collaborative DACH project *Arrangements and Drawings* as DFG Project MU3501/3-1.




Figure 1. Real examples of nesting on a leather hide (left) and a piece of fabric (right), kindly provided by MIRISYS and produced by their software for automatic nesting, <https://www.mirisys.com/>.

1. Introduction

Packing problems are everywhere in our daily lives. To give a few examples, you solve packing problems when finding room for your tupperware in the kitchen, the manufacturer of your clothing arranges cutting patterns on a large piece of fabric in order to minimize waste, and at Christmas time you are trying to cut out as many cookies from a dough as you can. In a large number of industries, it is crucial to solve complicated instances of packing problems efficiently. In addition to clothing manufacturing, we mention leather, glass, wood and sheet metal cutting, selective laser sintering, shipping (packing goods in containers) and 3D printing (arranging the parts to be printed in the printing volume); see Figure 1.

Packing problems can be easily and precisely defined in a mathematical manner, but many important questions are still completely elusive. In this work, we uncover a fundamental aspect of many versions of geometric packing by settling their computational difficulty.

We denote $\text{PACK}[\mathcal{P} \rightarrow C, \mathcal{M}]$ as the packing problem with pieces of the type \mathcal{P} , containers of type C and motions of type \mathcal{M} . In an instance of $\text{PACK}[\mathcal{P} \rightarrow C, \mathcal{M}]$, we are given pieces p_1, \dots, p_n of type \mathcal{P} and a container C of type C . We want to decide if there is a motion of type \mathcal{M} for each piece such that after moving the pieces by these motions, each piece is in C and the pieces are pairwise interior-disjoint. Such a placement of the pieces is called a *valid* placement.

As the allowed motions, we consider *translations* ($+$) and *rigid motions* ($\circ+$), where a rigid motion is a combination of a translation and a rotation. As containers and pieces, we consider squares (\square), convex polygons (\triangleleft), simple polygons (\triangleleft) and curved polygons (\triangleleft), where a *curved polygon* is a region enclosed by a simple closed curve consisting of a finite number of line segments and arcs contained in hyperbolae (such as the graph of $y = 1/x$). Note that hyperbolae, like all conic sections, can be represented as rational quadratic Bézier curves [32], which are extensively used for computer-aided design and manufacturing. It would therefore not be uncommon that curved pieces as the ones used here could appear in practical settings.

The problems with only translations allowed are relevant to some industries; for instance when arranging cutting patterns on a roll of fabric for clothing production, where the orientation

of each piece has to follow the orientation of the weaving or a pattern printed on the fabric. In other contexts such as leather, glass, or sheet metal cutting, there are usually no such restrictions, so rotations can be used to minimize waste. As can be seen from Figure 1, it is relevant to study packing problems where the pieces as well as the containers may be non-convex and have boundaries consisting of many types of curves (not just straight line segments).

We show that many of the above mentioned variants of packing are $\exists\mathbb{R}$ -complete. The complexity class $\exists\mathbb{R}$ will be defined below. We call the techniques we developed a *framework*, since the same techniques turn out to be applicable to prove hardness for many versions of packing. With adjustments or additions, they can likely be used for other versions or proofs of other types of hardness as well.

The Existential Theory of the Reals

The term *Existential Theory of the Reals* refers ambiguously to a formal language, a corresponding algorithmic problem (ETR) and a complexity class ($\exists\mathbb{R}$). Let us start with the formal logic. Let

$$\Sigma := \{\forall, \exists, 0, 1, x_1, \dots, x_n, +, \cdot, =, \leq, <, \wedge, \vee, \neg\}$$

be an alphabet for some $n \geq 1$. A *sentence* over Σ is a well-formed formula with no free variables, i.e., so that every variable is bound to a quantifier. The *Existential Theory of the Reals* is the set of true sentences of the form

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

where Φ is a quantifier-free formula. The algorithmic problem ETR is to decide whether a sentence of this form is true or not. At last, this leads us to the complexity class *Existential Theory of the Reals* ($\exists\mathbb{R}$), which consists of all those languages that are many-one reducible to ETR in polynomial time. Given a quantifier-free formula Φ , we define the solution space of Φ as $V(\Phi) := \{\mathbf{x} \in \mathbb{R}^n : \Phi(\mathbf{x})\}$. Thus in other words, ETR is to decide if $V(\Phi)$ is empty or not. It is currently known that

$$\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}. \quad (1)$$

To show the first inclusion is an easy exercise, whereas the second is non-trivial and was first proven by Canny [17]. A problem P is $\exists\mathbb{R}$ -hard if ETR is many-one reducible to P in polynomial time, and P is $\exists\mathbb{R}$ -complete if P is $\exists\mathbb{R}$ -hard and in $\exists\mathbb{R}$. None of the inclusions (1) are known to be strict, but the first is widely believed to be [15], implying that the $\exists\mathbb{R}$ -hard problems are not in NP. As examples of $\exists\mathbb{R}$ -complete problems, we mention problems related to realization of order-types [52, 51, 58], graph drawing [16, 27, 40, 43], recognition of geometric graphs [19, 20, 38, 45], straightening of curves [30], the art gallery problem [3], minimum convex covers [2], Nash-equilibria [13, 34], linkages [1, 54, 55], matrix-decompositions [26, 56, 57], polytope theory [52],

motions	\uparrow (translations)		\uparrow (rigid motions)	
pieces container	\square (convex)	\curvearrowright (curved)	\square (convex)	\curvearrowright (curved)
\curvearrowright (curved)	$\exists\mathbb{R}^*$	$\exists\mathbb{R}$	$\exists\mathbb{R}$	$\exists\mathbb{R}$
\square (polygon)	NP	$\exists\mathbb{R}$	$\exists\mathbb{R}$	$\exists\mathbb{R}$
\square (square)	NP	$\exists\mathbb{R}^*$	$\exists\mathbb{R}^*$	$\exists\mathbb{R}$

Table 1. This table displays 6 variants of the packing problem with rotations and translations, and 6 with translations only. $\exists\mathbb{R}$ means $\exists\mathbb{R}$ -complete and NP means NP-complete. We show that 10 of the problems are $\exists\mathbb{R}$ -complete, and the remaining 2 are known to be NP-complete. The problems marked with * are the *basic* problems. The $\exists\mathbb{R}$ -completeness of the remaining problems follows since there is a basic problem in the table which is more restricted.

embedding of simplicial complexes [4] and training neural networks [5, 14]. See also the surveys [18, 44, 53].

$\exists\mathbb{R}$ -membership

Showing that the packing problems we are dealing with in this paper are contained in $\exists\mathbb{R}$ is easy using the following recent result.

THEOREM 1.1 (Erickson, Hoog, Miltzow [31]). *For any decision problem P , there is a real verification algorithm for P if and only if $P \in \exists\mathbb{R}$.*

A *real verification algorithm* is like a verification algorithm for a problem in NP with the additional feature that it accepts real inputs for the witness and runs on the real RAM. (We refer to [31] for the full definition, as it is too long to include here.)

Thus in order to show that our packing problems lie in $\exists\mathbb{R}$, we have to specify a witness and a real verification algorithm. The witness is simply the motions that move the pieces to a valid placement. The verification algorithm checks that the pieces are pairwise interior-disjoint and contained in the container. Note that without the theorem above, we would need to describe an ETR-formula equivalent to a given packing instance in order to show $\exists\mathbb{R}$ -membership. Although this is not difficult for packing, it would still require some work.

Results

We show that a wide range of two-dimensional packing problems are $\exists\mathbb{R}$ -complete. A compact overview of our results is displayed in Table 1. In the table, the second row (with problems $\text{PACK}[\mathcal{P} \rightarrow \square, \mathcal{M}]$) is in some sense redundant, since the $\exists\mathbb{R}$ -completeness results can be deduced from the more restricted third row (the problems $\text{PACK}[\mathcal{P} \rightarrow \square, \mathcal{M}]$). We anyway include the row since a majority of our reduction is to establish hardness of problems with

polygonal containers, and only later we reduce these problems to the case where the container is a square.

A strength of our reductions is that in the resulting constructions, all corners can be described with rational coordinates that require a number of bits only logarithmic in the total number of bits used to represent the instance. Therefore, we show that the problems are *strongly* $\exists\mathbb{R}$ -hard. Another strength is that all the pieces have constant complexity, i.e., each piece can be described by its boundary as a union of $O(1)$ straight line segments and arcs contained in hyperbolae.

In the following we sketch an argument why $\text{PACK}[\sqsupset \rightarrow \sqsupset, +]$ is in NP. This may be folklore; the second author learned the proof from Günter Rote. We show that a valid placement can be specified as the translations of the pieces represented by a number of bits polynomial in the input size. Consider a valid placement of the pieces. For each pair of a segment s and a corner c (of a piece or the container), we consider the line $\ell(s)$ containing s and note which of the closed half-planes bounded by $\ell(s)$ contains c . Then we build a linear program (LP) using that information in the natural way. Here, the translation of each piece is described by two variables and for each pair (s, c) , we have one constraint involving at most four variables, enforcing c to be on the correct side of $\ell(s)$. It is easy to verify that the constraint is linear. The solution of the LP gives a valid placement of every piece and as the LP is polynomial in the input, so is the number of the bits of the solution to the LP. Note that if rotations are included, the corresponding constraints become non-linear.

Our results show that some packing problems with rotations or non-polygonal features allowed are $\exists\mathbb{R}$ -hard and thus likely not in NP. This gives a confirmation to the operations research community that most likely, they cannot employ standard algorithm techniques (solvers for ILP and SAT, etc.) that work well for many NP-complete problems like scheduling and TSP. The main message for the theory community is that efficient algorithms dealing with rotations or non-polygonal shapes can probably only be found if we relax the problems considerably.

Basis problem

In the first version of this paper [6], we could not show $\exists\mathbb{R}$ -hardness of the problem $\text{PACK}[\square \rightarrow \square, \circ +]$. We introduced the problem RANGE-ETR-INV, which in turn was a restricted version of the problem ETR-INV used to prove $\exists\mathbb{R}$ -hardness of the art gallery problem [3]. When proving $\exists\mathbb{R}$ -hardness by reducing from these problems, one must create gadgets for simulating *inversion* constraints of the form $xy = 1$, or, equivalently, $xy - 1 = 0$. This is usually obtained by making gadgets for each of the inequalities $xy - 1 \geq 0$ and $-xy + 1 \geq 0$. We managed to make a gadget for the constraint $xy - 1 \geq 0$ using convex pieces only, but despite much effort, we were unable to realize $-xy + 1 \geq 0$ unless we introduced non-convex pieces. Note that the equation $xy - 1 = 0$ defines a convex curve while $-xy + 1 = 0$ defines a concave curve. In the recent paper [50], Miltzow and Schmiermann proved that it is not important to make a gadget realizing

the inversion constraint exactly, but that it suffices to make gadgets realizing constraints of the forms $f(x, y) \geq 0$ and $g(x, y) \geq 0$ for any sufficiently well-behaved functions f and g , where one is convex and the other is concave. This is the basis for the reduction in this paper. We now specify the details of the problem we are reducing from.

DEFINITION 1.2 (CURVE-ETR[f, g] formula). Let $\delta > 0$, $U := [-\delta, \delta]$ and $f, g : U^2 \rightarrow \mathbb{R}$. An CURVE-ETR[f, g] formula $\Phi = \Phi(x_1, \dots, x_n)$ is a conjunction

$$\bigwedge_{i=1}^m C_i,$$

where each constraint C_i has one of the forms

$$x \geq 0, \quad x = \delta, \quad x + y = z, \quad f(x, y) \geq 0, \quad g(x, y) \geq 0,$$

for $x, y, z \in \{x_1, \dots, x_n\}$. Each constraint of the form $f(x, y) \geq 0$ or $g(x, y) \geq 0$ is called a *curved* constraint.

DEFINITION 1.3 (Well-behaved and convexly/concavely curved function). Let $\delta > 0$ and $U := [-\delta, \delta]$. We say a function $f : U^2 \rightarrow \mathbb{R}$ is *well-behaved* if the following conditions are met.

- f is a C^3 -function, i.e., three times continuously differentiable,
- $f(0, 0) = 0$, and all partial derivatives f_x, f_y, f_{xx}, f_{xy} and f_{yy} are rational in $(0, 0)$, and
- $f_x(0, 0) \neq 0$ or $f_y(0, 0) \neq 0$.

We write the curvature of a well-behaved function f at $(0, 0)$ as

$$\kappa = \kappa(f) = \left(\frac{f_y^2 f_{xx} - 2f_x f_y f_{xy} + f_x^2 f_{yy}}{(f_x^2 + f_y^2)^{\frac{3}{2}}} \right) (0, 0).$$

We say f is *convexly curved* if $\kappa(f) < 0$, and *concavely curved* if $\kappa(f) > 0$.

THEOREM 1.4 (Miltzow and Schmiermann [50]). Let $\delta > 0$, $U := [-\delta, \delta]$ and $f, g : U^2 \rightarrow \mathbb{R}$ be two well-behaved functions, one being convexly curved, and the other being concavely curved. Then deciding whether a CURVE-ETR[f, g] formula Φ is satisfiable is an $\exists\mathbb{R}$ -hard problem, even when $\delta = n^{-c}$ for any constant $c > 0$ and when we are promised that $V(\Phi) \subset U^n$.

The promise that we only need to look for solutions to a CURVE-ETR[f, g] formula in the tiny hyper-cube U^n will be crucial in our reduction. Note that the paper [50] is also focused on proving $\exists\mathbb{R}$ -membership of problems like CURVE-ETR[f, g]. We will only use the problem to prove $\exists\mathbb{R}$ -hardness, which is why we present the theorem in a slightly simpler form.

Related work

Milenkovich [49, 47] described exact algorithms for PACK[$\sqcup \rightarrow \sqcup, +$] with running times exponential in the number of pieces; see also the papers by Milenkovich and Daniels [24, 23].




Figure 2. Left: The optimal packing of five unit squares already requires rotations. Right: The currently best known packing of eleven unit squares into a larger square [35].

In another paper, Milenkovich [48] gave an algorithm and described a robust floating point implementation for the problem $\text{PACK}[\square \rightarrow \square, \circ +]$ using a combination of computational geometry and mathematical programming.

Alt [8] provides a survey of the literature on packing problems from a theoretical point of view. A lot of work has been done on bin packing, strip packing and knapsack, usually with rectangular pieces that can be translated or rotated by 90° . We refer to the survey of Christensen, Khan, Pokutta and Tetali [21] for an overview.

Recently, Merino and Wiese studied a version of 2-dimensional knapsack where the pieces are convex polygons and arbitrary rotations are allowed and presented a QPTAS for the problem [46]. Note that the problem $\text{PACK}[\square \rightarrow \square, \circ +]$ is a special case of the knapsack problem, so it is also $\exists\mathbb{R}$ -hard by our result.

Several packing variants are known to be NP-hard. Here we mention the problem of packing squares into a square by translation [42], packing segments into a simple polygon by translation [39], packing disks into a square [25], packing identical simple polygons into a simple polygon by translation [7], and packing unit squares into a polygon with holes by translation [12, 33]. Alt [8] proves by a simple reduction from the partition problem that packing rectangles into a rectangle is NP-hard, and this reduction works with and without rotations allowed (note that *a priori*, it is not clear that rotations make the problems more difficult, and it is straightforward to define (artificial) problems that even get easier with rotations). It is easy to modify the reduction to the problem of packing rectangles into a square, so this implies NP-hardness of all problems in Table 1.

A fundamental problem related to packing is to find the smallest square containing a given number of *unit* squares, with rotations allowed. A long line of mathematical research has been devoted to this problem, initiated by Erdős and Graham [29] in 1975, and it is still an active research area [22]. Even for *eleven* unit squares, the exact answer is unknown [35]; see Figure 2. Other packing problems have much older roots, for instance Kepler’s conjecture on the densest packings of spheres from 1611, famously proven by Hales in 2005 [36]. The 2D analog, i.e., finding the densest packings of unit disks, was solved already in 1773 by Lagrange under the assumption that the disk configurations are lattices, and the general case was solved

by Fejes Tóth in 1940 [61] (Thue already published a proof in 1910 [60] which is considered incomplete by some experts [62]).

There is a staggering amount of papers in operations research on packing problems. The research is mainly experimental and focuses on the development of heuristics to solve benchmark instances efficiently. We refer to some surveys for an overview [11, 10, 28, 37, 41, 59]. In contrary to theoretical work, there is a lot of experimental work on packing pieces with irregular shapes and with arbitrary rotations allowed.

Open problems

A natural continuation of our research is to study even more restricted packing variants. The problem $\text{PACK}[\square \rightarrow \square, \circ +]$ (packing arbitrary rectangles) is particularly interesting, as rectangles are very simple and widely studied. The problem $\text{PACK}[\circ \rightarrow \square, +]$ (packing disks of arbitrary sizes, which curiously has some relevance to origami [25]) is interesting for similar reasons. Our techniques seem not to extend to these special cases. If both problems are $\exists\mathbb{R}$ -complete, we expect that the proof techniques must be very different, since the non-linearity stems from rotations in the rectangle case, but from the shape of the pieces in the disk case. A first step to show $\exists\mathbb{R}$ -completeness of $\text{PACK}[\circ \rightarrow \square, +]$ could be to show $\exists\mathbb{R}$ -completeness of $\text{PACK}[\diamond \rightarrow \sqsupset, +]$. (Here, \diamond denotes convex curved polygons, or another reasonable generalization of convex polygons to pieces with curved boundaries.) The access to both convex and concave constraints is key in the currently known techniques for proving $\exists\mathbb{R}$ -hardness [50]. If all the pieces are convex and only translations are allowed, it seems that we can only encode concave constraints. It is therefore conceivable that $\text{PACK}[\diamond \rightarrow \sqsupset, +]$ and thus also $\text{PACK}[\circ \rightarrow \square, +]$ are contained in NP, by an argument similar as to why $\text{PACK}[\sqsupset \rightarrow \sqsupset, +]$ is in NP.

Another direction of research would be to consider the parameterized complexity of geometric packing problems. A natural parameter is the number of pieces, k . Then an instance of the problem $\text{PACK}[\sqsupset \rightarrow \sqsupset, \circ +]$ can be formulated as an ETR formula with $3k$ real variables, since the placement of each piece can be described by a two-dimensional translation and a rotation. Using general purpose algorithms from real algebraic geometry [9], this implies that the instance can be decided in $L^{O(k)}$ time, where L is the total length of the formula. It is interesting to find out if this is best possible under widely believed hypotheses. As a first step, one could investigate whether packing is W[1]-hard, or to show that $L^{O(k)}$ is the best possible running time assuming the exponential time hypothesis (ETH). In a second step, it would be interesting to see if it is possible to solve the packing problem using an ETR formula with fewer real variables. We note that W[1]-hardness or ETH-based lower bounds are not enough to give lower bounds on the number of real variables that are needed, as there could be a two-phase algorithm: The first phase runs in $L^{O(k)}$ time, without using tools from real algebraic

geometry. The second phase solves an ETR formula with, say, only $O(\sqrt{k})$ real variables. Neither $W[1]$ -hardness nor an ETH lower bound of $L^{\Omega(k)}$ would exclude this scenario.

Acknowledgments

We thank Reinier Schmiermann for useful discussions related to the use of tools from [50]. We would also like to thank anonymous reviewers for comments on earlier versions of this article.

2. Reduction skeleton

In this section, we give an overview of the steps and concepts needed in our reductions. The rest of the paper will then fill out the details.

2.1 The problem WIRED-CURVE-ETR[f, g]

We will reduce from an auxiliary problem called WIRED-CURVE-ETR[f, g]. An instance of this problem is a graphical representation of a CURVE-ETR[f, g] formula Φ , i.e., a drawing of Φ of a specific form, which we call a *wiring diagram*; see Figure 3. We denote by n the number of variables of Φ .

The term “wiring diagram” is often used for drawings of a similar appearance used to represent electrical circuits or pseudoline arrangements. We define equidistant horizontal *diagram lines* ℓ_1, \dots, ℓ_{2n} so that ℓ_1 is the topmost one and ℓ_{2n} is bottommost. The distance between consecutive lines is 10. In a wiring diagram, each variable x_i in Φ is represented by two x -monotone polygonal curves \vec{x}_i and \overleftarrow{x}_i , which we call *wires*. We think of \vec{x}_i as oriented to the right and \overleftarrow{x}_i as oriented to the left. The wire \vec{x}_i starts and ends on ℓ_{2i-1} , and \overleftarrow{x}_i starts and ends on ℓ_{2i} . Each wire consists of horizontal segments contained in the diagram lines and *jump segments*, which are line segments connecting one diagram line ℓ_j to a neighbouring diagram line $\ell_{j\pm 1}$. The wires are disjoint except that each jump segment must cross exactly one other




Figure 3. A wiring diagram corresponding to the CURVE-ETR[f, g] formula

$$x_2 + x_3 = x_1 \wedge f(x_2, x_1) \geq 0 \wedge g(x_2, x_1) \geq 0.$$

jump segment. Thus, the jump segments are used when two wires following neighbouring diagram lines swap lines.

The left and right endpoints of \vec{x}_i and \overleftarrow{x}_i are vertically aligned, and the wires appear and disappear in the order $(\vec{x}_1, \overleftarrow{x}_1), \dots, (\vec{x}_n, \overleftarrow{x}_n)$ from left to right in a staircase-like fashion.

In the wiring diagram, we represent each addition constraint of Φ as two inequalities, i.e., $x_i + x_j = x_k$ becomes $x_i + x_j \leq x_k$ and $x_i + x_j \geq x_k$. Each addition inequality and each curved constraint ($f(x_i, x_j) \geq 0$ and $g(x_i, x_j) \geq 0$) is represented by an axis-parallel *constraint box* intersecting the three or two topmost diagram lines; three for addition constraints and two for curved constraints. These boxes are pairwise disjoint. For a constraint $x_i + x_j \leq x_k$, the right-oriented wires $\vec{x}_i, \vec{x}_j, \vec{x}_k$ must inside the box occupy the lines ℓ_1, ℓ_2, ℓ_3 , respectively. For $x_i + x_j \geq x_k$, we need the left-oriented wires $\overleftarrow{x}_i, \overleftarrow{x}_j, \overleftarrow{x}_k$ instead. For a curved constraint $f(x_i, x_j) \geq 0$ or $g(x_i, x_j) \geq 0$, we need one of the wires of x_i and one of the wires of x_j to occupy ℓ_1 and ℓ_2 . Which combination and which order depends on the particular variant of packing that we are reducing to.

As we define our packing instance using a vertical line sweeping over the wiring diagram from left to right, we require that each vertical line is allowed to cross either zero or two jump segments and in the latter case, these two must cross each other. A vertical line crossing a constraint box must not cross any jump segment. This ensures that we only make one new feature in each step of the construction.

DEFINITION 2.1. An instance $\mathcal{I} := [\Phi, D]$ of the WIRED-CURVE-ETR[f, g] *problem* consists of a CURVE-ETR[f, g] formula Φ together with a wiring diagram D of Φ .

LEMMA 2.2. Given a CURVE-ETR[f, g] formula Φ with variables x_1, \dots, x_n , we can in $O(n^4)$ time construct a wiring diagram of Φ .

PROOF. We may assume that Φ has $O(n^3)$ constraints, since there will otherwise be duplicates of some constraints. We construct a wiring diagram as follows; refer to Figure 3. We construct all the curves simultaneously from left to right. We handle the constraints in order and define the curves as we go along. For instance, for a constraint such as $x_i + x_j \leq x_k$, we route \vec{x}_i to the line ℓ_1 using jump segments. This defines how all other curves should behave in the same range of x -coordinates where we have routed \vec{x}_i . We then route \vec{x}_j to the line ℓ_2 , and then route \vec{x}_k to ℓ_3 . Each time we route a curve to a specific line, we introduce $O(n)$ crossings. Therefore, we make $O(n^4)$ crossings in total. ■

2.2 Constructing a packing instance

Let an instance \mathcal{I} of WIRED-CURVE-ETR[f, g] be given with CURVE-ETR[f, g] formula Φ of variables x_1, \dots, x_n and $\delta := n^{-300}$. We are going to construct an instance of a packing problem with $N = O(n^4)$ pieces, since this will be the complexity of the size of the wiring diagram of \mathcal{I} . The

general idea is to build a packing instance on top of the wiring diagram. We define a polygonal container $C := C(\mathcal{I})$ containing the wiring diagram in the interior, and a set of pieces \mathbf{p} to be placed in C . The container C is bounded from below by a line segment, from left and right by y -monotone chains, and from above by an x -monotone chain. See Figure 4 for a sketch of a complete example.

In Sections 4 and 5, we present reductions to the packing problems $\text{PACK}[\square \rightarrow \sqsupset, \circlearrowright+]$, $\text{PACK}[\square \rightarrow \triangleleft, +]$, and $\text{PACK}[\triangleleft \rightarrow \sqsupset, +]$, i.e., where the container is a polygon or a curved polygon. In Section 6, we show how for fixed polygon and motion type, packing into a polygonal container reduces to packing into a square container. Together, these results imply hardness for $\text{PACK}[\square \rightarrow \square, \circlearrowright+]$ and $\text{PACK}[\triangleleft \rightarrow \square, +]$.

Defining the construction in steps from left to right

We define the packing instance as we sweep over the wiring diagram of \mathcal{I} with a vertical sweep line from left to right. Each step corresponds to one of the following events:

- the introduction of a pair of wires $(\vec{x}_i, \overleftarrow{x}_i)$,
- a crossing of two wires,
- an addition or curved constraint,
- the termination of a pair of wires $(\vec{x}_i, \overleftarrow{x}_i)$.

In each step, we add one or more *gadgets*, each involving a constant number of pieces and possibly a constant number of edges to the boundary of the container C . When the sweep line passes over the right endpoints of the last wires $(\vec{x}_n, \overleftarrow{x}_n)$, the construction of the container C and all the pieces \mathbf{p} is complete.

The overall goal of the construction is to prove the following theorem.

THEOREM 2.3. *Let \mathcal{I} be an instance of $\text{WIRED-CURVE-ETR}[f, g]$. For each of the problems $\text{PACK}[\square \rightarrow \sqsupset, \circlearrowright+]$, $\text{PACK}[\square \rightarrow \triangleleft, +]$, and $\text{PACK}[\triangleleft \rightarrow \sqsupset, +]$, we can in polynomial time construct an instance of the problem consisting of a container C and a set of pieces \mathbf{p} such that \mathcal{I} has a solution if and only if there is a valid placement of \mathbf{p} in C .*

From Theorem 2.3 and the $\exists\mathbb{R}$ -hardness of $\text{WIRED-CURVE-ETR}[f, g]$, we now immediately get the claimed results of row one and two in Table 1.

COROLLARY 2.4. *The problems $\text{PACK}[\square \rightarrow \sqsupset, \circlearrowright+]$, $\text{PACK}[\square \rightarrow \triangleleft, +]$ and $\text{PACK}[\triangleleft \rightarrow \sqsupset, +]$ are $\exists\mathbb{R}$ -hard.*

The third row in Table 1 follows from Section 6, as described later in this section. We now describe how values of variables can be encoded by the placement of certain pieces in our constructed packing instances.




Figure 4. A sketch of the instance of $\text{PACK}[\boxplus \rightarrow \boxplus, \oplus \oplus]$ we get from the wiring diagram in Figure 3, broken over six lines. The adders and curvers (swings) are marked with gray boxes. The (light and dark) red, blue and green pieces are the variable pieces and the pieces of each nuance form a lane.




Figure 5. A wire \vec{x} and a variable piece representing x placed on top, showing the leftmost and rightmost canonical placements of the piece. The large arrow in the piece indicates that the piece is right-oriented.




Figure 6. (a): anchor, (b): swap, (c): adder, (d): swing, (e): gramophone.

Variable pieces

Each variable x of Φ will be represented by a number of *variable pieces* in our construction, each of which is a convex polygon. Each variable piece represents exactly one variable x , and we make a correspondence between certain placements of the piece and the values of x . When adding a variable piece to our construction, we also specify the *zero placement* of the piece, which is a specific placement where it encodes the value 0 of x . In the zero placement, the piece will have a pair of (long) horizontal edges which have distance 10. By sliding the piece to the left or to the right from the zero placement, we obtain placements of the piece that encode all real values of x , even values outside the range $[-\delta, \delta]$. Each variable piece will be defined to be either right- or left-oriented. By sliding a right-oriented (resp. left-oriented) variable piece to the right by some amount $t \geq 0$, we obtain a placement that encodes the value t (resp. $-t$), while sliding it to the left by t results in a placement encoding $-t$ (resp. t). If the piece is rotated differently or placed higher or lower than the zero placement, we do not define any value of x to be encoded by the placement (we will in fact prove that in any solution to the constructed packing instance, no piece can have such an undesirable placement). We define the *canonical placements* of a variable piece to be the placements that encode values in the interval $[-\delta, \delta]$; see Figure 5.

On each wire \vec{x}_i or \overleftarrow{x}_i in the wiring diagram of \mathcal{I} , we will place several variable pieces representing x_i . The pieces placed on one wire are called a *lane*. The variable pieces on \vec{x}_i and \overleftarrow{x}_i are oriented to the right and left, respectively. We also introduce some variable pieces which will be placed at other places than on the wires, namely above the topmost wire where they will be introduced in the steps of the construction corresponding to addition or curved constraints.

Gadgets

Sketches of some of the gadgets can be seen in Figure 6. In the wiring diagram, each variable x_i is represented by two wires \vec{x}_i and \overleftarrow{x}_i such that the left endpoints of \vec{x}_i and \overleftarrow{x}_i are vertically aligned at distance 10, as are the right endpoints. In both ends of the wires, we build an *anchor* (Section 4.1) which ensures that the pieces placed on \vec{x}_i and those placed on \overleftarrow{x}_i encode the value of x_i consistently. Furthermore, the anchor will ensure that the encoded value of x_i is in the *range* $I(x_i)$, which we define as

$$I(x) = \begin{cases} \{\delta\} & \text{if } \Phi \text{ has a constraint } x = \delta, \\ [0, \delta] & \text{else, if } \Phi \text{ has a constraint } x \geq 0, \\ [-\delta, \delta] & \text{otherwise.} \end{cases}$$

Whenever two wires cross, we build a *swap* (Section 4.2). The swap employs a central piece that can translate in all directions, so that when it is pushed by a variable piece, the push will propagate to the neighbouring variable piece on the other side of the crossing. We describe *adders* (Section 4.4) to implement the addition constraints and *curvers* (Section 5) for the curved constraints. We describe two curvers (see Figure 6 (d–e)), both of which exist in a convex and a concave variant. Which version we use depends on the variant of packing we are reducing to.

Every time we add a gadget to the construction, we also introduce a constant number of new pieces. Each variable piece is introduced in one gadget where the left end of the piece is defined. The piece then extends outside the gadget to the right. The right end of the piece will be defined in another gadget added later to the construction. The piece is *exiting* the former gadget and *entering* the latter. In between the left and right end of the piece, defined in these two gadgets, the piece is bounded by a pair of horizontal edges. All pieces that are not variable pieces are contained within a single gadget.

Canonical placements

Recall that we define canonical placements of each variable piece. We do not define individual canonical placements of pieces that are not variable pieces, but instead we define canonical placements of all pieces of one or more gadgets: A placement of a set of pieces is canonical if (1) the placement is valid (i.e., the pieces are in C and are non-overlapping), (2) all variable pieces have a canonical placement, and (3) the pieces have certain relationships such as edge-edge contacts between each other. Part (3) will be specified for each gadget individually.

Preservation of solutions

The following lemma will be used to prove that for every solution to the CURVE-ETR[f, g] formula Φ , there is a canonical placement where the pieces encode that solution, meaning

that for each variable x , all variable pieces representing x encode the same value of x as in the solution. Define g to be the total number of gadgets, and let \mathbf{p}_i denote the set of all pieces introduced in the first i gadgets, where $i \in \{0, \dots, g\}$, so that $\mathbf{p}_0 := \emptyset$. The proof will be given in the sections describing the individual types of gadgets.

LEMMA 2.5 (Solution preservation). *Consider any $i \in \{1, \dots, g\}$ and suppose that for every solution to Φ , there is a canonical placement of the pieces \mathbf{p}_{i-1} that encodes that solution. Then the same holds for \mathbf{p}_i .*

Soundness of the reduction

As mentioned, each variable x will be represented by many variable pieces in the complete construction. A difficulty is that conceivably, such a piece may not be placed in a way that encodes a value of x . Even if all the pieces happen to be placed such that they *do* encode values of x , these values could be different and therefore not represent a solution to the formula Φ .

For a small number $\Delta \geq \delta$, we are going to introduce a class of placements called *aligned Δ -placements*. These are defined from the canonical placements by relaxing the requirements a bit. In an aligned Δ -placement, each variable piece must be placed so that it encodes a value, but it may slide Δ sideways from the placement encoding the value 0 instead of at most δ as for the canonical placements. The requirements to the other pieces are likewise relaxed and will be given later. The values encoded by the variable pieces in an aligned Δ -placement may therefore conceivably be outside the required range $[-\delta, \delta]$. The following lemma tells us that this is not the case for Δ sufficiently small. In fact, the existence of such a placement is enough to ensure that \mathcal{I} has a solution. The number μ is the *slack* of the construction defined as the area of the container C minus the total area of the pieces \mathbf{p} , and as will be explained later, $\mu = O(n^{-296})$ in our construction. The number $g = O(n^4)$ is the number of gadgets.

LEMMA 2.6 (Soundness). *Consider an aligned $g\mu$ -placement and any variable x . There is a specific non-empty subset of the pieces representing x that encode the value of x consistently (i.e., they all encode the same value of x) and the value is in the range $I(x)$. Furthermore, these values of the variables satisfy the constraints of Φ .*

In fact, it will follow from Lemma 2.6 that every aligned $g\mu$ -placement is canonical, but this is not important for our proof of Theorem 2.3. The remaining work in proving the theorem will be to prove the following lemma.

LEMMA 2.7. *Every valid placement is an aligned $g\mu$ -placement.*

The proof of Theorem 2.3 is now straight-forward:

PROOF OF THEOREM 2.3. If \mathcal{I} has a solution, then it follows directly from repeated use of Lemma 2.5 that there is a valid placement. Suppose now that there is a valid placement. By




Figure 7. (a): A pocket and an augmentation that fit perfectly together as in a jigsaw puzzle. (b): A wedge of the empty space and a piece which fit together. (c): The corner of the piece we are fingerprinting is marked with a dot. (d) and (e): Two examples where space is wasted because a wedge is not occupied by a piece with a matching angle.

Lemma 2.7, the placement is an aligned $g\mu$ -placement. By Lemma 2.6, some variable pieces encode a solution to the formula Φ . Therefore, \mathcal{I} has a solution. ■

2.3 Basic tools: Slack, fingerprinting and unique angles

In the following we will describe some tools needed to prove Lemma 2.7.

The slack of the construction

The *slack* of an instance of a packing problem is the area of the container C minus the total area of the pieces, and we denote the slack of our construction by μ . We need the slack to be very small in order to use the fingerprinting technique which will be described later.

We now give an upper bound on the slack of the complete construction. Our construction will be described as depending on the number $\delta := n^{-300}$. We place each variable piece so that it encodes the value 0, and we place the remaining pieces as shown in the sections that describe the individual gadgets. We now define μ' to be the area of the container C that is not covered by pieces in this placement and thus trivially have $\mu \leq \mu'$. By checking each type of gadget, it is straightforward to verify that the placement can be realized as a canonical (and thus valid) placement of the pieces in all gadgets except for the anchors, where some pieces are not completely contained in C . The uncovered area in each gadget will appear as a thin layer along some of the edges of the pieces in the gadget. This layer has thickness $O(\delta)$, and the edges along which it appears have total length $O(1)$, so the area is $O(\delta)$ in each gadget. There will be no empty space outside the gadgets because that space will be completely covered by variable pieces. Since the final construction has $g = O(N) = O(n^4)$ gadgets, it follows that $\mu \leq \mu' = O(n^4\delta) = O(n^{-296})$. It may seem a little odd to measure the slack in this indirect way, but we found it to be the easiest way to get the bound since we do not explicitly specify the area of the container or the pieces in our construction. The bound $\mu = O(n^{-296})$ ensures that μ is sufficiently small that we can use of the fingerprinting technique, which is described in the following.

Fingerprinting

In order to prove Lemma 2.7, we first show that every piece must be placed very close to a canonical placement using a technique we call *fingerprinting*. To grasp the idea of this technique, we first present another simpler technique that only works for non-convex pieces, and which we call the *jigsaw puzzle* technique for obvious reasons; see Figure 7 (a). The idea behind this technique is to force each piece p_1 to be at a specific position by creating a pocket of the container and a corresponding augmentation of the piece p_1 intended to be placed there. This is done in a way that only the piece p_1 has an augmentation that fits into the pocket, just as the principle behind a jigsaw puzzle, and it can be done in a way that gives the piece freedom to slide back and forth or rotate by a slight amount, etc. The pocket can also be created in another piece p_2 if p_1 is intended to be placed next to p_2 . Making enough of these pairs of pockets and extensions, we can therefore deduce where all the pieces are placed in all valid placements.

In fact, the jigsaw puzzle technique can be used to prove $\exists\mathbb{R}$ -hardness of packing problems with non-convex pieces in a much simpler way than the proofs of this paper, but unfortunately, the technique is not directly realizable with convex pieces. In *fingerprinting*, instead of making complicated augmentations of the pieces, we only work with a piece p_1 with a convex corner c of a specific angle α_1 . In the canonical placements, the empty space left by the other pieces forms a wedge with apex corner d of angle α_1 which can thus be covered very efficiently by p_1 by placing the corner c at or very close to d , as in Figure 7 (b). We make sure that every corner of every other piece has an angle α_2 significantly different from α_1 , in the sense that $|\alpha_2 - \alpha_1| = \Omega(N^{-2})$. It should likewise hold that the total angle of any combination of corners of other pieces is different from α_1 in that sense. Furthermore, the slack μ is tiny, as described above. As a result, we can show that if c is *not* placed very close to d , this will result in the empty space in a neighbourhood around d with an area exceeding μ , because no other piece (or combination of pieces) can cover that neighbourhood efficiently, illustrated in Figure 7 (d-e). In our constructions, fingerprinted corners are marked with a dot; see Figure 7 (c). For technical reasons, the fingerprinted corners must have angles in the range from $5\pi/180$ to $\pi/2$.

We add a few remarks about the use of fingerprinting. First of all, the situation shown in Figure 7 (b) is simplified. In our applications, the fat segments (bounding the empty space left by the other pieces) do not need to meet at the apex corner d , since a short portion (of length $O(\delta)$) close to the corner can be missing. Furthermore, the angle between the two fat segments does not have to be exactly α_1 before the technique can be used; just very close to α_1 (this will be important when we are fingerprinting more than one piece in a row).

Second, the fingerprinting does not imply that the piece p_1 *must* be placed with the corner c coincident with d , but only that the distance $\|cd\|$ has to be small. This is used deliberately in our constructions, since it allows for the piece p_1 to move slightly.

Third, when we introduce a new gadget and its k new pieces p_i, \dots, p_{i+k-1} , we use fingerprinting iteratively to argue where the new pieces must be placed. Here, the j 'th piece p_{i+j-1} , $j \in \{1, \dots, k\}$, can be fingerprinted in a wedge of the empty space formed by the preceding pieces p_i, \dots, p_{i+j-2} . However, the bound on the uncertainty of where p_{i+j-1} is placed increases with j . Slightly simplified, the bound grows as $O(n^{O(j)}\sqrt{\mu})$, and we need the bound to be at most some small constant to be of any use. We prefer to create an instance where we need only a logarithmic number of bits to represent the coordinates of the container and the pieces, since this will prove that the packing problems are *strongly* $\exists\mathbb{R}$ -hard. It is therefore important that we only apply fingerprinting iteratively a constant number of times, i.e., that $k = O(1)$, as we will otherwise need to choose the slack μ to smaller than n^{-q} for every constant $q > 0$, and then it will require a superlogarithmic number of bits to represent the coordinates of our instance. In the construction, we will always have $k \leq 7$ and we can do with choosing $\delta := n^{-300}$ so that $\mu = O(n^{-296})$.

In Section 3, we will develop the fingerprinting technique in detail. We consider this part the technically most challenging of the paper. The technique is versatile and can likely be used in other reductions to packing. Let us for instance mention that fingerprinting also works for packing problems where we are allowed to reflect the pieces when placing them. This gives two possibilities for each piece, and one must be excluded by other reasons, such as overlap with other pieces. As the conditions for the fingerprinting technique are technical, we will list them in Section 3 in detail. We will show for each gadget that those conditions are met.

After using fingerprinting iteratively a constant number of times for the new pieces, we use other techniques, such as the *alignment* (to be described in the sequel), to argue about their placement.

Choosing unique angles

As described in the previous paragraph, whenever we apply the fingerprinting technique to argue that a corner with angle α_1 of some piece is placed close to some specific point in the container, we need that every combination of corners of the other pieces have angles that sum to an angle α_2 such that $|\alpha_2 - \alpha_1| = \Omega(N^{-2})$. This will be called the *unique angle property*. In order to obtain this property, the construction will be designed so that each piece has a special corner where the angle can be chosen freely (within some interval of angles of size $\Omega(1)$). Likewise, the wedge (where the special corner is intended to be placed) formed by the boundary of the container or the other pieces is flexible, so that the angle of the wedge can match the chosen angle of the corner. In our figures, the special fingerprinted corners are marked with a dot, as in Figure 7 (c). The following lemma is used to choose these free angles such that we get the unique angle property.

LEMMA 2.8. Let $S_k := \{\frac{i}{k} + \frac{1}{2k^2} \mid i \in \{1, \dots, k\}\}$, consider a subset $R \subseteq S_k$, and let $x := \sum_{r \in R} r$. If $x \in S_k$, then $R = \{x\}$.

PROOF. If R consists of m elements, then x has the form $\frac{j}{k} + \frac{m}{2k^2}$ for some $j \in \mathbb{N}$. A number of this form, for $m \leq k$, can only be in S_k if $m = 1$. ■

The lemma provides a set of k numbers in a range of size $O(1)$ and any number is $\Omega(k^{-2})$ away from the sum of any combination of other numbers. We multiply the numbers in S_k by π to get a set of rational angles and choose the free angles from such a set πS_k , for $k = O(N)$. The free angles are restricted to various subintervals of $[0, \pi]$, so we choose k so large that πS_k contains enough angles from each of these subintervals. However, as each subinterval has size $\Omega(1)$, we can do with $k = O(N)$.

2.4 Proof structure of Lemma 2.7

In the proof of Lemma 2.7, we use the fingerprinting technique to prove that in every valid placement, the pieces are placed almost as in a canonical placement. To explain the structure of the argument in more detail, we need some notions of placements that are close to being canonical, which will be defined in the following paragraph.

Almost-canonical placements and aligned placements

We say that a valid placement of the pieces of a gadget is *almost-canonical* if there exists rigid motions that move the pieces to a canonical placement such that every point in each piece is moved a distance of at most n^{-1} (in other words, the *displacement* between the actual placement and the canonical placement of each piece is n^{-1}).

We say that a placement of the pieces of a gadget is an *aligned Δ -placement* for $\Delta \geq \delta$ if (i) the placement is almost-canonical, and (ii) for each variable x , each variable piece representing x encodes a value in the range $[-\Delta, \Delta]$. Note that since the placement is almost-canonical, we can always assume $\Delta \leq n^{-1} + \delta = n^{-1} + n^{-300}$.

The following lemma says that the pieces of every new gadget can be assumed to be almost-canonical if the preceding pieces have an aligned Δ -placement, for Δ sufficiently small. Recall that \mathbf{p}_i is the set of all pieces introduced in the first i gadgets, where $i \in \{0, \dots, g\}$, so that $\mathbf{p}_0 := \emptyset$. The lemma will follow from the use of fingerprinting, and the proof will be given in the sections describing the individual types of gadgets.

LEMMA 2.9 (Almost-canonical Placement). For any $i \in \{1, \dots, g\}$, consider a valid placement P (of all the pieces) for which the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement. It then holds for P that the pieces \mathbf{p}_i have an almost-canonical placement.




Figure 8. The alignment segment ℓ makes us conclude that the pieces must be horizontally aligned: otherwise, they would overlap, cross the container boundary, or cover more of the alignment segment than what is available.

Aligning pieces

Once we know that the pieces \mathbf{p}_i have an almost-canonical placement, provided by the previous lemma, we can use so-called *alignment segments* to further restrict where the new pieces $\mathbf{p}_i \setminus \mathbf{p}_{i-1}$ introduced in gadget i can be placed. In particular, we will be able to fix the rotations of some pieces to be as in the canonical placements. The idea is sketched in Figure 8. From the rough placements we get from fingerprinting, we know that a set of the pieces each has a pair of parallel edges that are both cut through by a vertical alignment segment ℓ . If we sum the distance between the two parallel edges over all the pieces, we get exactly the length of ℓ . Since the portions of ℓ covered by the pieces must be pairwise disjoint in a valid placement, we can conclude that the pieces have to be rotated so that the parallel edges are perpendicular to ℓ . This technique will be used to prove the following lemma for each gadget individually. Using the lemma repeatedly together with Lemma 2.9, we get that every valid placement is also an aligned $g\mu$ -placement, proving Lemma 2.7.

LEMMA 2.10 (Aligned placement). *For any $i \in \{1, \dots, g\}$, consider a valid placement P (of all the pieces) for which the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and the pieces \mathbf{p}_i have an almost-canonical placement. It then holds for P that the pieces \mathbf{p}_i have an aligned $i\mu$ -placement.*

It now remains to prove Lemma 2.6.

2.5 Proof structure of Lemma 2.6

The proof of Lemma 2.6 goes along the following lines. In an aligned $g\mu$ -placement, each variable piece p_x encodes a value for the variable x it is representing, which we will denote by $\langle p_x \rangle$. The problem is that different pieces representing the same variable x may conceivably not encode the value consistently. However, recall that we build lanes of pieces on top of the two wires \vec{x} , \overleftarrow{x} , and these meet at the left and right endpoints of the wires. We prove that the values encoded by these pieces p_1, \dots, p_m make a cycle of inequalities: $\langle p_1 \rangle \leq \langle p_2 \rangle \leq \dots \leq \langle p_m \rangle \leq \langle p_1 \rangle$. It thus follows that all these pieces encode a value of x consistently. Furthermore, the anchors, which




Figure 9. An abstract drawing of the dependency graphs of the instance we get from the wiring diagram in Figure 3 and how the graphs connect to the gadgets for addition and curved constraints. The number of vertices on the cycles and paths are neither important nor correct.

are the gadgets that we place at the left and right endpoints of the wires \vec{x} , \overleftarrow{x} , will ensure that $\langle p_1 \rangle \in I(x)$, so that the encoded values are in the correct range.

In our construction, we also make additional lanes of pieces going to the adders and curvers. The functionality of the specific gadgets imply that the addition and curved constraints of Φ are all satisfied. In order to describe the structure of the argument, we introduce a graph G_x for each variable x as described in the next paragraph.

Dependency graph of variable pieces

For each variable x , we introduce a directed *dependency graph* G_x . The vertices of G_x are the variable pieces representing x . Consider a gadget and two variable pieces p_1, p_2 appearing in the gadget and both representing x . We add an edge from p_1 to p_2 in G_x if p_1 is an entering right-oriented piece or an exiting left-oriented piece and p_2 is an exiting right-oriented piece or an entering left-oriented piece. In crossings between the two wires \vec{x} , \overleftarrow{x} representing x , there will be a swap where this rule introduces unintended edges, so we make one exception described in Section 4.2 where the swap is described in detail.

The following lemma is going to follow trivially from the way we make the lanes on top of the wires \vec{x} , \overleftarrow{x} for each variable x , and the way we connect the gadgets representing addition and curved constraints to these lanes. See Figure 9 for an illustration.

LEMMA 2.11. *For each variable x , the graph G_x consists of a directed cycle K_x with some directed paths attached to it (oriented towards or away from K_x). The vertices of the cycle K_x are the variable pieces appearing on the wire \vec{x} from left to right and the wire \overleftarrow{x} from right to left in this order. For each path attached to K_x , the vertex farthest from K_x is a piece entering or leaving a gadget representing an addition or curved constraint.*

In the following, we consider a given aligned $g\mu$ -placement of all the pieces. Since all edges of G_x are between pieces appearing in the same gadget, the following lemma will be proven for each gadget individually.

LEMMA 2.12 (Edge inequality). *Consider a variable x and an edge (p_1, p_2) of G_x . Then $\langle p_1 \rangle \leq \langle p_2 \rangle$.*

From Lemma 2.11 and Lemma 2.12, we now get the following (except that the part about the anchor gadget will be proven in Section 4.1).

LEMMA 2.13. *For each variable x , all the pieces of the cycle K_x encode the value of x consistently. Furthermore, due to the design of the anchor gadget, the value is in $I(x)$.*

By the above lemma, we may write $\langle K_x \rangle$ to denote the value represented by all pieces of K_x .

Adders and curvers work

We will show in Section 4.4 and Section 5 that the adders and curvers actually enforce addition and curved constraints as they are supposed to. This entails showing that the gadgets implement the addition constraint or various convexly or concavely curved constraints in a geometric sense and also that the variable pieces of the gadgets are correctly connected to the cycles in the respective dependency graphs. In particular, we will show the following two lemmas.

LEMMA 2.14 (Adders work). *For each constraint $x + y = z$ in the formula Φ , we have $\langle K_x \rangle + \langle K_y \rangle = \langle K_z \rangle$.*

LEMMA 2.15 (Curvers work). *For each of the problems $\text{PACK}[\square \rightarrow \sqcup, \odot +]$, $\text{PACK}[\square \rightarrow \triangleleft, +]$, and $\text{PACK}[\triangleleft \rightarrow \sqcup, +]$, there exists well-behaved functions f and g that are convexly and concavely curved, respectively, such that for every constraint of the form $f(x, y) \geq 0$ in the $\text{CURVE-ETR}[f, g]$ formula Φ , we have $f(\langle K_x \rangle, \langle K_y \rangle) \geq 0$, and for every constraint $g(x, y) \geq 0$, we have $g(\langle K_x \rangle, \langle K_y \rangle) \geq 0$.*

Combining Lemmas 2.13 to 2.15, we then have a proof of Lemma 2.6.

2.6 Square container

In Section 6, we describe a reduction from problems of type $\text{PACK}[\mathcal{P} \rightarrow \sqcup, \mathcal{M}]$ to $\text{PACK}[\mathcal{P} \rightarrow \square, \mathcal{M}]$. It will be crucial that the container C is 4-monotone, as defined below.

DEFINITION 2.16. A simple closed curve γ is 4-monotone if γ can be partitioned into four parts $\gamma_1, \dots, \gamma_4$ in counterclockwise order that move monotonically down, to the right, up, and to the left, respectively. A polygon Q is 4-monotone if the boundary of Q is a 4-monotone curve.

LEMMA 2.17. *In the reductions resulting from using the gadgets described in Sections 4 and 5, the resulting container is 4-monotone.*

PROOF. The boundary of the resulting container has a left and a right staircase, γ_1 and γ_3 , created by the left and right anchors, respectively, and these staircases are y -monotone, and their upper and lower endpoints are horizontally aligned. The lower endpoints of the staircases




Figure 10. Construction used in the reduction to packing problems with a square container. The space left by the exterior pieces (blue, green, orange, and turquoise) is exactly the 4-monotone container C of the instance we are reducing from.

γ_1 and γ_3 are connected by a single horizontal line segment γ_2 bounding the bottom lane from below. The upper endpoints of the staircases are connected by a curve γ_4 which bound the topmost lane and the adders and curvers from above. The curve γ_4 is x -monotone, as can easily be verified by inspecting the boundary added due to the adders and curvers. Hence, the container is 4-monotone. ■

We get from the lemma that the packing problems are even $\exists\mathbb{R}$ -hard for 4-monotone containers. Let \mathcal{I}_1 be an instance of a packing problem where the container $C := C(\mathcal{I}_1)$ is 4-monotone. We place C in the middle of a larger square and fill out the area around C with pieces carefully; the details are given in Section 6 and Figure 10 shows an example of the construction. We call these new pieces the *exterior* pieces, whereas we call the pieces of \mathcal{I}_1 the *inner* pieces. Using fingerprinting and other arguments, we are able to prove that there is essentially only one way to fit the exterior pieces in the square, and the space left for the inner pieces is exactly the container C . Therefore, there exists a valid placement of the pieces in the resulting instance if and only if there is one of the inner pieces in C . We get the results in the third row of Table 1 as expressed by the following corollary.

COROLLARY 2.18. *The problems $\text{PACK}[\square \rightarrow \square, \circlearrowleft +]$ and $\text{PACK}[\triangle \rightarrow \square, +]$ are $\exists\mathbb{R}$ -hard.*

3. Fingerprinting

In this section, we develop a technique to argue that pieces are roughly at the position where we intend them to be. The high level idea is based on a few properties. First, the slack μ , i.e., the difference between the area of the container and the total area of the pieces, is very small. Second, every piece p has a specific corner v with a unique angle that fits precisely at one position. If a piece is placed at a different location than the intended one, the empty space would exceed μ . In order to make such arguments, we first need to carefully define a few concepts.

Motion and Placement

We encode a rotation by a *rotation matrix*, which is a matrix M of the form

$$M = \begin{pmatrix} a & -b \\ b & a \end{pmatrix},$$

with $\det M = 1$. From a translation $t \in \mathbb{R}^2$ and a rotation M , we get a *motion* $m := (t, M)$. If only translations are allowed, we require that M is the identity.

Given a piece p and a motion $m = (M, t)$, then we denote by p^m the piece p after moving p according to m , i.e.,

$$p^m := \{Mx + t : x \in p\}.$$

The set p^m is the *placement* of p by m . Given a tuple of $\mathbf{p} = (p_1, \dots, p_k)$ of pieces and a tuple $\mathbf{m} = (m_1, \dots, m_k)$ of motions, then we denote by

$$\mathbf{p}^{\mathbf{m}} = (p_1^{m_1}, \dots, p_k^{m_k})$$

the *placement* of \mathbf{p} by \mathbf{m} . We may write $p_i^{\mathbf{m}}$ instead of $p_i^{m_i}$.

Given a container C , pieces $\mathbf{p} = (p_1, \dots, p_k)$ and a motion \mathbf{m} , we say that \mathbf{m} (resp. $\mathbf{p}^{\mathbf{m}}$) is a *valid* motion (resp. placement), if (i) $p_i^{\mathbf{m}} \subset C$ for all i , and (ii) $p_i^{\mathbf{m}}$ and $p_j^{\mathbf{m}}$ are interior-disjoint for all $i \neq j$.

Other geometric definitions

Let ab and cd be two (oriented) line segments. The *angle between* ab and cd is the minimum angle that ab can be turned such that ab and cd become parallel and point in the same direction, i.e., after turning, we should have $(b - a)^\perp \cdot (d - c) = 0$ and $(b - a) \cdot (d - c) > 0$.

Consider two motions m_1 and m_2 of a piece p . The *displacement* between m_1 and m_2 is $\sup_{x \in p} \|x^{m_1} - x^{m_2}\|$. The *displacement angle* is the absolute difference in how much m_1 and m_2 rotate p in the interval $[0, \pi]$.

Given a compact set S in the plane, we denote by $\text{area}(S)$ the area of S , and we define the *diameter* of S as $\text{diam}(S) := \max_{a, b \in S} \|ab\|$.

Given a container C and pieces \mathbf{p} , we define the *slack* as $\mu = \text{area}(C) - \sum_{p \in \mathbf{p}} \text{area}(p)$.

Let $A \subset \mathbb{R}^2$. Then $A^c = \{x \in \mathbb{R}^2 \mid x \notin A\}$. Let $A, B \subset \mathbb{R}^2$. Then we denote the Minkowski sum as $A \oplus B := \{a + b \mid a \in A, b \in B\}$ and the Minkowski difference as $A \ominus B := (A^c \oplus B)^c$. For $\lambda > 0$, define $\text{disk}(\lambda) := \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq \lambda^2\}$.

3.1 Fingerprinting a single piece

Now let us go one level deeper into the details of the fingerprinting technique; see also Figure 11. We consider a case where we already know the position of some pieces (possibly with some




Figure 11. We are considering the placement of the pieces p_1, \dots, p_{i-1} according to a valid motion \mathbf{m} . The white area is the empty space E available for the remaining pieces p_i, \dots, p_N . The radius of the gray circles centered at x, y, z is the uncertainty value λ ; the first circle must contain a , the second b and c , and the third d , where ab and cd are segments on the boundary of E .

uncertainty), and we consider the empty space E where the remaining pieces must be placed. Ideally, we could identify a corner w of the empty space and a corner v of a remaining piece p which has exactly the same angle as w and deduce that p must be placed with v at w . Unfortunately, this is not the case, for two reasons. First, we want to give most pieces some tiny but non-zero amount of wiggle room. This is important as pieces are meant to represent variables. Second, we do not know the precise position of the other pieces as previous fingerprinting steps could only infer approximate and not exact positions of those pieces. Thus, we will identify for each piece p a triangle T , which has a corner y with the same angle as v . The edges adjacent to y will be very close to but not exactly on the boundary of the empty space E . The triangle T will be our main protagonist in the forthcoming proofs and formal definitions. It may partially overlap existing pieces or have some distance to already placed pieces. Another key player is the uncertainty value $\lambda \geq 0$, which is a measure of how much T is off from the ideal. We are now ready to go into the full details of the fingerprinting.

Setup

We are given a container C and pieces $\mathbf{p} = (p_1, \dots, p_N)$. Each piece $p \in \mathbf{p}$ is a simple polygon with the following properties.

- Each segment of p has length at least 1.
- The diameter of p is at most some number d_{\max} .
- The polygon p is *fat* in the following sense. For any two points v, w on different and non-neighbouring segments of p , we have $\|vw\| \geq \tau := 1/100$.

In Section 3.4, we will show that the results developed in the following for polygonal pieces also hold when the pieces are allowed to be curved polygons (provided that the curvature is sufficiently small and the segments do not curve within distance 1 from the corners).

The empty space E

We consider an arbitrary valid motion \mathbf{m} and analyze how we can infer something about the placement of the pieces p_i, \dots, p_N from the placement of the first $i - 1$ pieces p_1, \dots, p_{i-1} . We can think of this situation as if we have already decided where to place the first $i - 1$ pieces p_1, \dots, p_{i-1} in C so that they are interior-disjoint, and we are now reasoning about where to place the next piece.

Let

$$E := C \setminus \overline{\bigcup_{j=1}^{i-1} p_j^{\mathbf{m}}}$$

be (the closure of) the uncovered space available for the remaining pieces p_i, \dots, p_N , see Figure 11. Then E is a subset of C bounded by a finite number of line segments, and each of these segments is contained in edges of the pieces $p_1^{\mathbf{m}}, \dots, p_{i-1}^{\mathbf{m}}$ or C .

Covering a wedge of E

Assume that there is a special triangle $T \subset C$ with corners x, y, z and with the following properties. We have $\|xy\| = \|yz\| = 1$. Let $\lambda \geq 0$ be a (small) number that will be defined whenever we are going to apply the fingerprinting. The value λ can be thought of as a measure of uncertainty of the already placed pieces and the distance from the boundary of T to the boundary of E . Let $T^{\text{in}} = x^{\text{in}}y^{\text{in}}z^{\text{in}}$ denote the triangle $T \ominus \text{disk}(\lambda)$ (where \ominus is the Minkowski difference and $\text{disk}(\lambda)$ the disk of radius λ , as defined in the beginning of this section), such that $x^{\text{in}}, y^{\text{in}}, z^{\text{in}}$ are on the angular bisectors from x, y, z , respectively.

DEFINITION 3.1. We say that E is λ -bounding T at y if the following two conditions hold:

- There are segments ab and cd on the boundary of E such that each of the distances $\|ax\|, \|by\|, \|cy\|, \|dz\|$ is at most λ .
- The interior of T^{in} is a subset of E .
- The angle β of T at y satisfies $\beta \in [\alpha_{\min}, \alpha_{\max}]$, where $\alpha_{\min} := 5\pi/180$ and $\alpha_{\max} := \pi/2$.

The second requirement means that no piece among p_1, \dots, p_{i-1} covers anything of T^{in} when placed according to \mathbf{m} . The triangle T^{in} will have an area much larger than μ , which implies that almost all of T^{in} must be covered by the pieces $p_i^{\mathbf{m}}, \dots, p_N^{\mathbf{m}}$.

Unique angle property

We assume that the pieces p_i, \dots, p_N have the following property, which we denote as the *unique angle property* with respect to the angle β of y and a (small) number $\sigma > 0$:

Consider any set $S := \{v_1, \dots, v_m\}$ such that it contains at most one corner from each piece p_i, \dots, p_N . If the sum of angles of corners in S is in the interval $[\beta - \sigma, \beta + \sigma]$, then S consists of only one corner v , i.e., $S = \{v\}$.

In most applications of the fingerprinting technique, there will be only one such set $S = \{v\}$. In other words, the angle range $[\beta - \sigma, \beta + \sigma]$ uniquely identifies a specific piece and a specific corner v of the piece. However, in Section 6, we are going to consider a special case where the container is a square where there will be more such sets.

We will argue that almost all of T^{in} must be covered by a piece with a corner v with an angle in the range $[\beta - \sigma, \beta + \sigma]$, and the corner v must be placed close to y . Informally, since μ is much smaller than the area of T^{in} , almost all of T^{in} must be covered by the pieces p_i, \dots, p_N . Because of the unique angle property, it is only possible to cover a sufficient amount of T^{in} by placing a piece with such a corner v close to y and with the adjacent edges close to parallel to yx and yz , since the edges ab and cd of ∂E are preventing T^{in} from being covered in another way.

Main lemma

To sum up, we have made these assumptions:

- We consider a valid placement \mathbf{m} of the pieces \mathbf{p} .
- The empty space E is λ -bounding the triangle T at the corner y of T .
- The pieces p_i, \dots, p_N have the unique angle property with respect to the angle β of the corner y and the small number σ .

In Section 3.3, we are going to prove the following lemma in the setting described above.

LEMMA 3.2 (Single fingerprint). *There is a piece $p \in \{p_i, \dots, p_N\}$ with a corner v such that the angle of v is in $[\beta - \sigma, \beta + \sigma]$ and $\|yv^{\mathbf{m}}\| = O\left(\lambda/\sigma + \sqrt{\mu/\sigma}\right)$.*

Furthermore, let u, w be the corners preceding and succeeding v in counterclockwise direction, respectively. Then the angle between $v^{\mathbf{m}}u^{\mathbf{m}}$ and yx is $O\left(\lambda/\sigma + \sqrt{\mu/\sigma}\right)$, as is the angle between $v^{\mathbf{m}}w^{\mathbf{m}}$ and yz .

3.2 Fingerprinting more pieces at once

In this section, we consider the iterated use of the fingerprinting technique (in particular Lemma 3.2) for some number of times. This describes the situation whenever we have introduced the pieces of a new gadget to the construction. More precisely, we consider the situation where we know how the pieces p_1, \dots, p_{i-1} must be placed, and we want to deduce how the following k pieces p_i, \dots, p_{i+k-1} , for some $k \geq 1$, must then be placed. To this end, consider an arbitrary valid motion \mathbf{m} . Consider a set \mathbf{s} of *intended* motions s_i, \dots, s_{i+k-1} of the pieces p_i, \dots, p_{i+k-1} . We are going to define what it means for the intended motions \mathbf{s} to be *sound*, and

then we prove that if they are sound, then the valid motion \mathbf{m} must place the pieces p_i, \dots, p_{i+k-1} in a way similar to the intended motions \mathbf{s} .

To define soundness of the intended motions, we first define the empty space $E_j^{\mathbf{s}}$, for $j \in \{i, \dots, i+k-1\}$, as

$$E_j^{\mathbf{s}} := C \setminus \overline{\left(\bigcup_{l=1}^{i-1} p_l^{m_l} \cup \bigcup_{l=i}^{j-1} p_l^{s_l} \right)}.$$

Thus, $E_j^{\mathbf{s}}$ is the free space where the piece p_j can be placed if the pieces p_1, \dots, p_{i-1} are placed according to \mathbf{m} while the pieces p_i, \dots, p_{j-1} are placed according to the intended motions \mathbf{s} .

DEFINITION 3.3. We say that the intended motion s_j , $j \in \{i, \dots, i+k-1\}$, is λ -sound, for a value $\lambda \geq 0$, if there exists a triangle $T_j = x_j y_j z_j$ and a corner v_j of p_j such that the following holds,

- the angle β_j of T_j at y_j is in the range $[\alpha_{\min}, \alpha_{\max}]$,
- $\|x_j y_j\| = 1$ and $\|y_j z_j\| = 1$,
- $E_j^{\mathbf{s}}$ is λ -bounding T_j at y_j (recall Definition 3.1),
- if a set S of at most one corner from each of the pieces p_j, \dots, p_N has a sum of angles in the range $[\beta_j - \sigma, \beta_j + \sigma]$, then $S = \{v_j\}$ (note that this is a stronger version of the unique angle property since the original definition just requires S to be a singleton, while here S must contain a specific corner v_j),
- $T_j \subset p_j^{s_j}$, and
- $v_j^{s_j} = y_j$ and $x_j y_j, y_j z_j \subset \partial p_j^{s_j}$.

We likewise define the placement $p_j^{s_j}$ to be λ -sound if the motion s_j is λ -sound.

LEMMA 3.4. *There exists an absolute constant $c > 0$ such that the following holds. Define*

$$\begin{aligned} \Lambda_i &:= 0, \quad \text{and} \\ \Lambda_j &:= cd_{\max}/\sigma \cdot \Lambda_{j-1} + cd_{\max}(\lambda/\sigma + \sqrt{\mu/\sigma}), \end{aligned}$$

for $j > i$. If the motions s_i, \dots, s_{i+k-1} are λ -sound, then for each $j \in \{i, \dots, i+k-1\}$, the displacement between the motions m_j and s_j of the piece p_j is at most Λ_{j+1} . It holds that

$$\Lambda_{k+1} \leq (k+1)(cd_{\max}/\sigma)^{k+1}(\lambda/\sigma + \sqrt{\mu/\sigma}),$$

which is a bound on all the mentioned displacements.

PROOF. We proceed by induction on j . For $j = i$, we apply Lemma 3.2. We get that $\|y_i v_i^{\mathbf{m}}\| \leq O(\lambda/\sigma + \sqrt{\mu/\sigma})$. Furthermore, the second half of Lemma 3.2 implies that the displacement angle between the motions m_i and s_i is likewise at most $O(\lambda/\sigma + \sqrt{\mu/\sigma})$. We therefore get that the displacement between m_i and s_i is

$$cd_{\max}(\lambda/\sigma + \sqrt{\mu/\sigma}) = \Lambda_{i+1}$$

for some constant c .

Suppose now that the statement holds for indices $i, i+1, \dots, j-1$. Define

$$E_j^{\mathbf{m}} := C \setminus \overline{\bigcup_{l=1}^{j-1} p_l^{m_l}}.$$

Since E_j^s is λ -bounding T_j at y_j and the displacement between m_{j-1} and s_{j-1} is at most Λ_j , we get that $E_j^{\mathbf{m}}$ is $(\Lambda_j + \lambda)$ -bounding T_j at y_j . Therefore, Lemma 3.2 gives that the displacement between m_j and s_j is at most

$$cd_{\max}((\Lambda_j + \lambda)/\sigma + \sqrt{\mu/\sigma}) = cd_{\max}/\sigma \cdot \Lambda_j + cd_{\max}(\lambda/\sigma + \sqrt{\mu/\sigma}) = \Lambda_{j+1},$$

for the constant c introduced above. Unfolding the expression, we get

$$\begin{aligned} \Lambda_{k+1} &= \sum_{j=0}^k (cd_{\max}/\sigma)^j \cdot cd_{\max}(\lambda/\sigma + \sqrt{\mu/\sigma}) \\ &\leq \sum_{j=0}^k (cd_{\max}/\sigma)^{j+1} \cdot (\lambda/\sigma + \sqrt{\mu/\sigma}) \\ &\leq (k+1)(cd_{\max}/\sigma)^{k+1}(\lambda/\sigma + \sqrt{\mu/\sigma}). \end{aligned}$$

■

The following lemma will be used to fingerprint the pieces in each gadget individually.

LEMMA 3.5 (Multiple fingerprints). *Consider a gadget together with its pieces p_i, \dots, p_{i+k-1} , for $k \leq 7$, which are introduced in some step of the construction. Suppose that there is a valid motion \mathbf{m} of the complete construction. Furthermore, suppose that there exists intended motions s_i, \dots, s_{i+k-1} which are $g\mu$ -sound. Then the displacement between the intended motion s_j and the actual motion m_j is of the order $O(n^{-48})$.*

PROOF. In our construction, we have $d_{\max} = O(n^4)$, $\mu = O(n^{-296})$, and $g = O(n^4)$. We use the method described in the proof of Lemma 2.8 to choose unique angles. As our reduction results in a packing instance of $N = O(g) = O(n^4)$ pieces, we get the unique angle condition satisfied for a value of σ of the order $\Omega(N^{-2}) = \Omega(n^{-8})$. We now get from Lemma 3.4 with $\lambda = g\mu = O(n^{-292})$ that the displacement is at most

$$\begin{aligned} 8(cd_{\max}/\sigma)^8(\lambda/\sigma + \sqrt{\mu/\sigma}) &= O((n^4 n^8)^8 (n^{-292} n^8 + \sqrt{n^{-296} n^8})) \\ &= O(n^{96} n^{-144}) = O(n^{-48}). \end{aligned}$$

■




Figure 12. Setup in the proof of Lemma 3.2. The circles centered at x, y, z bound the disks of radius λ in which the points a, b, c, d are known to be. The figure is not to scale. In practice, λ is very small so that the triangles $T^{\text{out}}, T, T^{\text{in}}$ are almost equally large. Furthermore, the radius of A_2 is much larger than the radius of A_1 , which in turn is much larger than the radius of A_0 . The bottom figure shows the regions D_1 and D_2 , that together make up D . The segments ab and cd are drawn with thick lines to indicate that these act as a restriction to where the next piece can be placed.

3.3 Proof of Single fingerprint (Lemma 3.2)

Proof setup

We prove auxiliary Lemmas 3.6 to 3.11 and then show Lemma 3.2 (restated as Lemma 3.12). See Figure 12. Let $\zeta(\theta) := \frac{1}{\sin(\theta/2)}$. The function ζ is important when computing the distances between corresponding corners of offset versions of the same triangle, as the following lemma makes clear.

LEMMA 3.6. (1) Consider a triangle $U = efg$ and define for some $s > 0$ the triangle $U^{\text{in}} := U \ominus \text{disk}(s) = e^{\text{in}}f^{\text{in}}g^{\text{in}}$, so that $e^{\text{in}}, f^{\text{in}}, g^{\text{in}}$ are on the angular bisectors of e, f, g , respectively. Then $\|ee^{\text{in}}\| = s\zeta(\theta)$, where θ is the angle of U at e .

(2) Let e, f, g be points such that the distances from a point e^{in} to each of the segments ef and eg is at most s . Then $\|ee^{\text{in}}\| \leq s\zeta(\theta)$, where $\theta \in [0, \pi)$ is the angle between ef and eg .

PROOF. Proof of (1): Let p be the projection of e^{in} on ef . Then $\|ee^{\text{in}}\| = \|pe^{\text{in}}\|/\sin(\theta/2) = s\zeta(\theta)$.

Proof of (2): For a given angle θ , the distance $\|ee^{\text{in}}\|$ is maximum if the distances from e^{in} to each segment ef and eg are both s , so that, in particular, e^{in} is on the angular bisector between the segments ef and eg . We then proceed as in the proof for (1). ■

Lemma 3.6 gives that $\|yy^{\text{in}}\| = \lambda\zeta(\beta)$. Let $T^{\text{out}} = x^{\text{out}}y^{\text{out}}z^{\text{out}}$ be the triangle we get by offsetting the edges of T outwards in a parallel fashion by distance $\lambda \cdot \zeta(\alpha_{\min})$, i.e., T^{out} is the triangle such that $T^{\text{out}} \ominus \text{disk}(\lambda \cdot \zeta(\alpha_{\min})) = T$. The corners $x^{\text{out}}, y^{\text{out}}, z^{\text{out}}$ are on the angular bisectors of x, y, z , respectively.

Define $\psi := \zeta(\alpha_{\min}) + \zeta(\alpha_{\min})^2 = O(1)$. By Lemma 3.6, we have $\|yy^{\text{out}}\| = \lambda\zeta(\alpha_{\min})\zeta(\beta) \leq \lambda\zeta(\alpha_{\min})^2$, and $\|y^{\text{out}}y^{\text{in}}\| = \lambda(\zeta(\beta) + \zeta(\alpha_{\min})\zeta(\beta)) \leq \lambda\psi$.

Subdividing T^{out} by arcs

For some small constants $c_1, c_2 > 0$, we define three radii as

$$\begin{aligned} r_0 &:= \lambda\psi, \\ r_1 &:= c_1 \left(\lambda/\sigma + \sqrt{\mu/\sigma} \right), \\ r_2 &:= c_2 (r_1/\sigma). \end{aligned}$$

We require that r_2 is much smaller than 1, say $r_2 < 1/10$ (as it turns out, by choosing δ small enough, we can make μ and λ so small that r_2 is below any desired constant). For the ease of presentation, we will not explicitly specify the constants c_1, c_2 , but it will follow from the analysis that constants exist that will make the arguments work. Note that in our application of Lemma 3.5, we will have $\sigma = \Theta(n^{-8})$. Furthermore, one should think of r_1 as much larger than r_0 and r_2 as much larger than r_1 .

Let A_i be the arc with center y^{out} and radius r_i from the point s_i^{out} on segment $y^{\text{out}}z^{\text{out}}$ counterclockwise to the point t_i^{out} on segment $x^{\text{out}}y^{\text{out}}$. Let D be the region bounded by segments $t_2^{\text{out}}y^{\text{out}}$ and $y^{\text{out}}s_2^{\text{out}}$ and the arc A_2 . The arc A_1 separates D into two regions D_1 and D_2 , where A_2 appears on the boundary of D_2 .

Geometric core lemma

The following lemma is the geometric core of our argument and the setup is shown in Figure 13. We use this lemma to conclude that if a set Q of pieces cover most of D_1 , then they have corners whose angles sum to a number close to β . It then follows from the unique angle property that Q consists of just one piece.

LEMMA 3.7. *Let $\{W_i = x_i y_i z_i \mid i = 1, \dots, m\}$ be a collection of triangles where each y_i is a corner in D_1 and the edge $x_i z_i$ is disjoint from D . Let β_i be the angle of W_i at y_i and suppose that $\beta_i \in [\alpha_{\min}, \alpha_{\max}]$. Suppose that W_1, \dots, W_m are pairwise interior disjoint and that the interior of each W_i is disjoint from $t_2^{\text{out}}y^{\text{out}}$ and $y^{\text{out}}s_2^{\text{out}}$. Let $\mathcal{W} := \bigcup_{i=1}^m W_i$, and let $\rho \in [0, 1]$ be the fraction*




Figure 13. The setup of Lemma 3.7.

of A_2 covered by \mathcal{W} . Then

$$\frac{\text{area}(D_1 \cap \mathcal{W})}{\text{area}(D_1)} \leq \rho + O(r_1/r_2), \text{ and} \quad (2)$$

$$\sum_{i=1}^m \beta_i \in [\beta\rho - O(r_1/r_2), \beta\rho + O(r_1/r_2)]. \quad (3)$$

PROOF. We first analyze just a single triangle W_i and then generalize to all of W_1, \dots, W_m . For $j \in \{1, 2\}$, let $A_j^{(i)} := A_j \cap W_i$ be the arc on A_j contained in W_i , and let $\gamma_i \in (0, \beta]$ be the angle spanned by $A_2^{(i)}$. We claim that

$$\beta_i \in [\gamma_i - O(r_1/r_2), \gamma_i + O(r_1/r_2)], \text{ and} \quad (4)$$

$$\text{area}(D_1 \cap W_i) \leq r_1^2 \gamma_i / 2 + O(r_1^3/r_2). \quad (5)$$

Before proving (4) and (5), we show how (2) and (3) follow. We get from (5) that

$$\frac{\text{area}(D_1 \cap W_i)}{\text{area}(D_1)} \leq \frac{r_1^2 \gamma_i / 2 + O(r_1^3/r_2)}{r_1^2 \beta / 2} = \gamma_i / \beta + O(r_1/r_2).$$

Note that by (4) and since $\beta_i \in [\alpha_{\min}, \alpha_{\max}]$ for all i , we know that the number of triangles is $m = O(1)$. We now have that

$$\frac{\text{area}(D_1 \cap \mathcal{W})}{\text{area}(D_1)} = \sum_{i=1}^m \frac{\text{area}(D_1 \cap W_i)}{\text{area}(D_1)} \leq \sum_{i=1}^m (\gamma_i / \beta + O(r_1/r_2)) = \rho + O(r_1/r_2),$$

which proves (2). Likewise, (3) follows from (4) using that $m = O(1)$ as

$$\sum_{i=1}^m \beta_i \in \left[\sum_{i=1}^m (\gamma_i - O(r_1/r_2)), \sum_{i=1}^m (\gamma_i + O(r_1/r_2)) \right] = [\rho\beta - O(r_1/r_2), \rho\beta + O(r_1/r_2)].$$

For the following proof of (4), we refer to Figure 14 (left). Let $s'_1 \in y_i z_i$ and $t'_1 \in y_i x_i$ be the endpoints of $A_1^{(i)}$ and define s'_2 and t'_2 similarly as the endpoints of $A_2^{(i)}$. Note first that if $y_i = y^{\text{out}}$, we have $\beta_i = \gamma_i$. However, in general y_i is just a point within distance r_1 from y^{out} . Therefore, the angle α_s between the segments $y^{\text{out}} s'_2$ and $y_i s'_2$ is at most $\arcsin(r_1/r_2)$, which




Figure 14. Left: In the shown example, $\beta_i = \gamma_i - \alpha_s + \alpha_t$. Since $\alpha_s = O(r_1/r_2)$ and $\alpha_t = O(r_1/r_2)$, we always have $\beta_i = [\gamma_i - O(r_1/r_2), \gamma_i + O(r_1/r_2)]$. Right: The angle θ is an argument of the segment $y_i w$. The argument is bounded by $\beta + O(r_1/r_2)$.

is obtained when $y^{\text{out}}y_i s'_2$ is a triangle with $\|y^{\text{out}}y_i\| = r_1$ and a right angle at y_i . Since r_2 is much larger than r_1 , we have $\arcsin(r_1/r_2) = O(r_1/r_2)$. Similarly, the angle α_t between the segments $y^{\text{out}}t'_2$ and $y_i t'_2$ is at most $\arcsin(r_1/r_2) = O(r_1/r_2)$. Note that $\beta_i = \gamma_i \pm \alpha_s \pm \alpha_t$, so we get $\beta_i \in [\gamma_i - O(r_1/r_2), \gamma_i + O(r_1/r_2)]$.

By the *argument* of a line ℓ , we mean the counterclockwise angle from the x -axis to ℓ . The argument of a line segment s is the argument of the line containing s . Assume without loss of generality that $y^{\text{out}}s_2^{\text{out}}$ is horizontal with s_2^{out} to the right of y^{out} , so that the argument of any line through y^{out} and a point on A_2 is in the range $[0, \beta]$. We claim that then the argument of every segment $y_i w$, where $w \in W_i$, is in the range $[-O(r_1/r_2), \beta + O(r_1/r_2)]$. To verify the upper bound, note that the argument of $y_i w$ is maximum if $w = x_i$ and $t'_2 = t_2^{\text{out}}$, see Figure 14 (right). By an argument as the one used in the previous paragraph, we get that the argument can be at most $O(r_1/r_2)$ larger than β . The lower bound follows similarly.

We now observe that each of the segments $y_i s'_1$ and $y_i t'_1$ has length at most $r_1 + O(r_1^2/r_2)$, as follows. See Figure 15 (left). Since $\beta \leq \pi/2$, the longest segment ℓ_1 in D_1 with an argument in $[0, \beta]$ connects y^{out} to a point on A_1 and has length r_1 . The longest segment ℓ_2 in D_1 with an argument in $(\beta, \beta + O(r_1/r_2)]$ connects t_1^{out} to a point p on $y^{\text{out}}s_1^{\text{out}}$. We then get

$$\begin{aligned} \|\ell_2\| &= \|t_1^{\text{out}}p\| \leq \|t_1^{\text{out}}y^{\text{out}}\| + \|y^{\text{out}}p\| \\ &\leq r_1 + r_1 \tan(O(r_1/r_2)) = r_1 + r_1 \frac{\sin(O(r_1/r_2))}{\cos(O(r_1/r_2))} \\ &\leq r_1 + r_1 \frac{O(r_1/r_2)}{1 - O(r_1/r_2)} = r_1 + O\left(\frac{r_1^2}{r_2 - r_1}\right) = r_1 + O(r_1^2/r_2), \end{aligned}$$

where the last equality follows since r_2 is much larger than r_1 . Similarly, the longest segment in D_1 with an argument in $[-r_1/r_2, 0)$ connects s_1^{out} to a point on $y^{\text{out}}t_1^{\text{out}}$ and has length $r_1 + O(r_1^2/r_2)$. Hence, $r_1 + O(r_1^2/r_2)$ is also an upper bound on the length of $y_i s'_1$ and $y_i t'_1$.




Figure 15. Left: Longest segments in D_1 . The segment ℓ_1 shows the case that the argument is in $[0, \beta]$, and then the longest segment has length r_1 . The segment ℓ_2 shows the case where the argument is in $(\beta, \beta + O(r_1/r_2)]$, and then the segment has length $r_1 + O(r_1^2/r_2)$. Right: Figure to show that $\text{area}(D_1 \cap W_i) \leq r_1^2 \gamma_i / 2 + O(r_1^3/r_2)$. The segments $y_i s_1'$ and $y_i t_1'$ are assumed to have the same length $r_1 + O(r_1^2/r_2)$, and the angle β_i at y_i is $\gamma_i + O(r_1/r_2)$. We consider the point y_i' such that $\|y_i' s_1'\| = \|y_i' t_1'\| = r_1$. Then the angle at y_i' is likewise $\gamma_i + O(r_1/r_2)$, so the area of the blue triangle is $r_1^2 \gamma_i / 2 + O(r_1^3/r_2)$. The white triangles have area at most $O(r_1^3/r_2)$, so the total area of $D_1 \cap W_i$ is at most $r_1^2 \gamma_i / 2 + O(r_1^3/r_2)$.

We get an upper bound on $\text{area}(D_1 \cap W_i)$ in the case that β_i and the edges $y_i s_1'$ and $y_i t_1'$ all reach the upper bounds. This might not be realizable, but still provides an upper bound. See Figure 15 (right). If the edges have length $\|y_i s_1'\| = \|y_i t_1'\| = r_1$, the area is $r_1^2 (\gamma_i + O(r_1/r_2)) / 2 = r_1^2 \gamma_i / 2 + O(r_1^3/r_2)$. Extending the edges to $\|y_i s_1'\| = \|y_i t_1'\| = r_1 + O(r_1^2/r_2)$, we are adding two triangles each of which has area at most $r_1 \cdot O(r_1^2/r_2) = O(r_1^3/r_2)$, and the desired bound (5) follows. ■

We want to apply Lemma 3.7 to a set of pieces covering parts of D_1 ; see Figure 16. Let s_i^{in} and t_i^{in} be the intersection points of A_i with $y^{\text{in}} z^{\text{in}}$ and $x^{\text{in}} y^{\text{in}}$, respectively. Let $B \subset D_1 \cap T^{\text{in}}$ be the region bounded by segments $s_0^{\text{in}} s_1^{\text{in}}$, $t_0^{\text{in}} t_1^{\text{in}}$, and the arcs $A_0 \cap T^{\text{in}}$ and $A_1 \cap T^{\text{in}}$. We consider only pieces that cover a part of B . The reason we do not consider all pieces covering a part of D_1 is that a piece covering a part of D_1 but not B might violate the assumptions of Lemma 3.7. In particular, such a piece might not have an interior disjoint from $t_2^{\text{out}} y_2^{\text{out}}$ and $y_2^{\text{out}} s_2^{\text{out}}$, as is seen in case (ii.a) in Figure 18. Cases (i.b) and (ii.b) in Figures 17 and 18, respectively, show the possibilities of a piece covering a part of B , and here the piece fits the assumptions of Lemma 3.7. The following lemma makes this intuition precise. Conceivably, there may be some pieces that fulfill the conditions of Lemma 3.7 but do not cover a part of B . However, even considering only pieces covering a part of B , we will be able to arrive at our desired conclusion. Recall that the segments ab and cd are bounding some pieces (or the container C) in the valid motion \mathbf{m} , so these segments act as obstacles that restrict the placement of a piece $p_j^{\mathbf{m}}$ covering a part of B .

LEMMA 3.8. *A piece $p_j^{\mathbf{m}}$ covering a part of the interior of B has the following properties:*

- *There is a corner $v^{\mathbf{m}}$ of $p_j^{\mathbf{m}}$ contained in D_1 .*




Figure 16. Regions B and D' . The region D' is drawn with a pattern of falling gray lines. The fat part of A_2 is the arc A'_2 that bounds D' . A piece covering part of B has edges crossing one or both of A'_2 and bc . The figure is not to scale—in practice B will cover almost all of the region D_1 below the arc A_1 .

— The edges of p_j^m adjacent to v^m cross A_2 .

PROOF. Let s_2 and t_2 be the intersection points of A_2 with segment cd and ab , respectively, as shown in Figure 16. Let A'_2 be the part of A_2 from s_2 counterclockwise to t_2 . Let D' be the region bounded by segments s_2c , bc , t_2b , and the arc A'_2 . Then $B \subset D' \subset D$. Since $r_2 < 1/2$, the diameter of D is less than 1. Since p_j^m covers a part of the interior of B , there is one or more edges of p_j^m that cross the boundary of D' . An edge of p_j^m can only cross the boundary of D' at a point on the segment bc or the arc A'_2 , since the segments cs_2 and bt_2 are bounding some other pieces. We divide into the following cases, which are also shown in Figures 17 and 18, respectively:

- Case (i): No edge of p_j^m crosses bc . Then there is an edge ef crossing A'_2 . We have the following two cases:
 - Case (i.a): The edge ef crosses A'_2 twice. Since $\beta \leq \pi/2$, we get that the distance from y^{out} to ef is at least $r_2/\sqrt{2}$. Since r_1 is much smaller, this edge cannot contribute to covering a part of B , so there must be some edges of p_j^m crossing the boundary of D' that do not belong to this case.
 - Case (i.b): One of the endpoints e and f is inside D' while the other is outside. Assume without loss of generality that f is inside. It follows that the succeeding edge fg likewise intersects A'_2 due to the minimum length of the edges, and the claim holds.
- Case (ii): An edge ef of p_j^m crosses bc . Suppose that as we follow ef from e to f , we enter D' as we cross bc . In particular $e \notin D'$. There must likewise be another edge gh




Figure 17. Cases from the proof of Lemma 3.8. Left: Case (i.a). In this case, the piece p_j^m is too far from y^{out} to cover any of B . Right: Case (i.b). This case agrees with the statement of the lemma.

of p_j^m crossing bc , since otherwise, the interior of p_i^m would intersect ab or cd . We have the following cases depending on whether an endpoint of ef coincides with one of gh :

- Case (ii.a): f coincides with an endpoint of gh . Assume without loss of generality that $f = g$. By Lemma 3.6 part (1), we have $\|y^{\text{out}}y\| \leq \lambda\zeta(\alpha_{\min})^2$, and by part (2), we have $\|yf\| \leq \lambda\zeta(\alpha_{\min})$. Therefore $\|y^{\text{out}}f\| \leq \|y^{\text{out}}y\| + \|yf\| \leq r_0$. But then the edges ef and gh do not get far enough into D' so that the wedge they form can cover a part of B , as every point in B has distance at least r_0 to y^{out} . Therefore, there must be some edges of p_j^m crossing the boundary of D' that do not belong to this case.
- Case (ii.b): e coincides with an endpoint of gh . Assume without loss of generality that $e = g$. Because the angle at e is at least α_{\min} , it follows from Lemma 3.6 part (2) that $\|ye\| \leq \lambda\zeta(\alpha_{\min})$. Therefore, e is contained in D . It then follows that ef and gh both cross A_2 , since they must exit D , and the claim of the lemma thus holds.
- Case (ii.c): No endpoint of ef coincides with one of gh . Since $\|bc\| \leq 2\lambda < \tau$ and both segments ef and gh cross $\|bc\|$, we conclude that the fatness condition is violated in this case.

■

Here, we give an informal description of the following three lemmas. Lemma 3.9 states that the area of B grows quadratically in r_1 while $D_1 \setminus B$ grows only linearly. Lemma 3.10 says that almost all of B must be covered by pieces, as the uncovered area will otherwise be larger




Figure 18. Cases from the proof of Lemma 3.8. Top left: Case (ii.a). In this case, p_j^m cannot cover any of B . Top right: Case (ii.b). This case agrees with the statement of the lemma. Bottom left: Case (ii.c). This case violates the fatness assumption of the pieces.

than μ . We are then able to conclude in Lemma 3.11 that almost all of D_1 must be covered by the pieces covering B , as B has asymptotically the same area as D_1 . This eventually makes it possible to apply Lemma 3.7 in the proof of Lemma 3.2.

LEMMA 3.9. *We have*

- $\text{area}(D_1 \setminus B) = O(r_0^2 + \lambda r_1)$.
- $\text{area}(B) = \Omega(r_1^2 - r_0^2)$.

PROOF. The points in $D_1 \setminus B$ are either within distance r_0 from y^{out} or within distance $\lambda(1 + \zeta(\alpha_{\min})) = O(\lambda)$ from one of the line segments $y^{\text{out}}s_1^{\text{out}}$ or $y^{\text{out}}t_1^{\text{out}}$, each of length r_1 . It then follows that $\text{area}(D_1 \setminus B) = O(r_0^2 + \lambda r_1)$.

We thus have $\text{area}(B) = \text{area}(D_1) - \text{area}(D_1 \setminus B) = \Omega(r_1^2) - O(r_0^2 + \lambda r_1) = \Omega(r_1^2 - r_0^2)$. ■

Let $Q := \{p_j \mid j \in \{1, \dots, n\} \text{ and } p_j^{\text{m}} \cap B \neq \emptyset\}$, and let $Q := \bigcup_{p_j \in Q} p_j^{\text{m}}$.

LEMMA 3.10. *By choosing $r_1 := \Omega(\lambda/\sigma + \sqrt{\mu/\sigma})$, where Ω hides a sufficiently large constant, we get $\text{area}(B \cap Q) \geq (1 - \sigma/4)\text{area}(B)$.*

PROOF. Recall that the pieces $p_1^{\text{m}}, \dots, p_{i-1}^{\text{m}}$ are interior disjoint from T^{in} , as $T^{\text{in}} \subset E_i^{\text{m}}$. Since

$$r_1 = \Omega\left(\lambda/\sigma + \sqrt{\mu/\sigma}\right) = \Omega\left(r_0 + \sqrt{\frac{\mu}{\sigma}}\right) = \Omega\left(\sqrt{r_0^2 + \frac{\mu}{\sigma}}\right),$$

we get from Lemma 3.9 that

$$\text{area}(B) = \Omega(r_1^2 - r_0^2) = \Omega(\mu/\sigma).$$

Now, if the constant hidden in the Ω -notation is large enough, we have $\text{area}(B) \geq \frac{\mu}{\sigma/4}$, or equivalently, $\sigma/4 \cdot \text{area}(B) \geq \mu$. This means that the area of B covered by the pieces in Q is at least $(1 - \sigma/4)\text{area}(B)$, as otherwise the uncovered part would be larger than μ . ■

LEMMA 3.11. *By choosing $r_1 := \Omega(\lambda/\sigma + \sqrt{\mu/\sigma})$, where Ω hides a sufficiently large constant, we get $\text{area}(D_1 \cap Q) \geq (1 - \sigma/2)\text{area}(D_1)$.*

PROOF. Note that $r_1 = \Omega(\lambda/\sigma) = \Omega(r_0/\sqrt{\sigma} + \lambda/\sigma)$. We get $r_1^2 = \Omega\left(\frac{r_0^2 + \lambda r_1}{\sigma}\right)$, where we can choose the constant hidden in the Ω -notation as big as needed. Since $r_1 = \Omega(r_0)$, we then get $r_1^2 - r_0^2 = \Omega\left(\frac{r_0^2 + \lambda r_1}{\sigma}\right)$. Now, if we choose the constant big enough, we get from Lemma 3.9 that $\sigma/4 \cdot \text{area}(B) \geq \text{area}(D_1 \setminus B)$ and thus

$$\frac{\text{area}(D_1)}{\text{area}(B)} = \frac{\text{area}(B) + \text{area}(D_1 \setminus B)}{\text{area}(B)} \leq 1 + \sigma/4.$$

It follows that

$$\frac{\text{area}(B)}{\text{area}(D_1)} = \frac{\text{area}(B)}{\text{area}(B) + \text{area}(D_1 \setminus B)} \geq \frac{1}{1 + \sigma/4} \geq \frac{1 - \sigma/2}{1 - \sigma/4}.$$

Hence we have that

$$(1 - \sigma/4)\text{area}(B) \geq (1 - \sigma/2)\text{area}(D_1).$$

The claim now follows from Lemma 3.10 as

$$\text{area}(D_1 \cap Q) \geq \text{area}(B \cap Q) \geq (1 - \sigma/4)\text{area}(B) \geq (1 - \sigma/2)\text{area}(D_1).$$

■

We are now ready to prove Lemma 3.2. We rephrase the lemma as follows.

LEMMA 3.12. *By choosing $r_2 := \Omega(r_1/\sigma)$, where Ω hides a sufficiently large constant, we get that Q consists of just one piece p_j , and p_j has a corner y_j such that the angle of y_j is in $[\beta - \sigma, \beta + \sigma]$ and $y_j^{\mathbf{m}} \in D_1$. In particular, $\|y_j y_j^{\mathbf{m}}\| \leq 2r_1 = O(\lambda/\sigma + \sqrt{\mu/\sigma})$.*

Furthermore, let x_j, z_j be the corners preceding and succeeding y_j , respectively. Then the angle between $y_j^{\mathbf{m}} x_j^{\mathbf{m}}$ and $y_j x_j$ is $O(\lambda/\sigma + \sqrt{\mu/\sigma})$, as is the angle between $y_j^{\mathbf{m}} z_j^{\mathbf{m}}$ and $y_j z_j$.

PROOF. By Lemma 3.8, each piece $p_j \in Q$ has a corner y_j such that $y_j^{\mathbf{m}}$ is contained in D_1 , and the two adjacent edges cross A_2 . For each piece $p_j \in Q$, we consider the triangle $W_j := x_j y_j z_j$, such that $x_j y_j$ and $y_j z_j$ are the edges adjacent to y_j . The triangles W_j now fit in the setup of Lemma 3.7. Let ρ be the fraction of A_2 covered by Q . We then get by Lemma 3.11 and Lemma 3.7 that

$$1 - \sigma/2 \leq \frac{\text{area}(D_1 \cap Q)}{\text{area}(D_1)} \leq \rho + O(r_1/r_2).$$

Lemma 3.7 furthermore yields that the sum S of angles of the corners y_j is in the range $[\beta\rho - O(r_1/r_2), \beta\rho + O(r_1/r_2)]$. Since $1 - \sigma/2 - O(r_1/r_2) \leq \rho \leq 1$, we get

$$\beta\rho - O(r_1/r_2) \geq (1 - \sigma/2 - O(r_1/r_2))\beta - O(r_1/r_2) \geq \beta - \sigma\pi/4 - O(r_1/r_2), \text{ and} \quad (6)$$

$$\beta\rho + O(r_1/r_2) \leq \beta + O(r_1/r_2). \quad (7)$$

If we now choose $r_2 := \Omega(r_1/\sigma)$, where Ω hides a sufficiently large constant, we get $S \in [\beta - \sigma, \beta + \sigma]$. We then get from the unique angle property that Q consists of just one piece p_j .

Since $\|y_j y_j^{\mathbf{m}}\| = O(r_1)$ and $\|x_j a\| \leq \lambda = O(r_1)$ and r_1 is much smaller than $\|x_j y_j\| = 1$, we get that the angle between $y_j^{\mathbf{m}} x_j^{\mathbf{m}}$ and $y_j x_j$ is $O(r_1/\|y_j x_j\|) = O(r_1)$, and likewise for the angle between $y_j^{\mathbf{m}} z_j^{\mathbf{m}}$ and $y_j z_j$. ■

3.4 Generalization to curved polygons

Let $\gamma: [0, L] \rightarrow \mathbb{R}^2$ be a simple curve parameterized by arc-length and of length $L \geq 1$. We say that γ is a *curved segment* if

- the prefix $\gamma([0, 1])$ and suffix $\gamma([L - 1, L])$ are line segments (each of length 1),
- γ is differentiable,




Figure 19. Case (i.a) from the proof of Lemma 3.8 when the pieces can be curved. Since the curvature is bounded by some small constant κ , the distance from y^{out} to ef is at least $x \cdot r_2$ for some constant x .

- the *mean curvature* of γ at most some small constant $\kappa = O(1)$, i.e., for all $s_1, s_2 \in (0, L)$, we have

$$\|\gamma'(s_1) - \gamma'(s_2)\| \leq \kappa|s_1 - s_2|.$$

As an example, consider a simple curve γ satisfying the first condition and which is the concatenation of line segments and circular arcs. Then γ is a curved segment as long as each circular arc has radius at least κ and each transition from one line segment or circular arc to the next is tangential.

We claim that the above results on fingerprinting also hold when the pieces are curved polygons, each of which has a boundary which is a finite union of curved segments. The first requirement for a curved segment (i.e., that it has a straight prefix and suffix of length 1) ensures that near every corner of a curved polygon, the curved polygon behaves as a normal polygon. Because of this, the setup described in Sections 3.1 and 3.2 also makes sense for curved polygons.

We first check that Lemma 3.8 still holds. Because we require a curved segment to have a straight prefix and suffix of length 1, the Case (ii.a) is excluded for the same reason when using curved polygons as when using normal polygons. Likewise, Case (ii.c) is excluded because we still have the fatness assumption. Hence, Case (i.a) is the only case that should be excluded where the curved segments make a difference, namely the case where a curved segment ef crosses A'_2 twice. Here, the piece can cover a bit more of D' because the segment can curve; see Figure 19. However, the mean curvature is at most $\kappa = O(1)$ and we can choose r_2 to be an arbitrarily small constant. We therefore still get that the distance from y^{out} to ef is at least xr_2 for some constant $x < 1/\sqrt{2}$ (where, in the version for normal polygons, we had $x = 1/\sqrt{2}$). But

we have $r_1 = O(\sigma \cdot r^2) = O(n^{-8}r^2)$, since $\sigma = \Theta(n^{-8})$, so it follows that it is still impossible that the piece can cover anything of B in this case when n is large enough.

Lemmas 3.9 and 3.11 do not use any assumption on the geometry of the pieces, and can thus still be used. Finally, in the proof of Lemma 3.12, we can now apply Lemma 3.8 as before. Since the prefix and suffix of each curved segment are line segments of length 1, we can again define triangles W_j and apply Lemma 3.7 (the geometric core lemma), so the proof goes through unaltered.

4. Linear gadgets

In this section we are describing four types of gadgets called *anchor*, *swap*, *split*, and *adder*. They all work with convex polygonal pieces, a polygonal container, and translations. They also work when rotations are allowed and can thus be used for all packing variants studied in the paper.

For each gadget, we will define canonical placements and verify the four required lemmas of Section 2. Here we repeat the properties we need to verify for each gadget.

- For every solution to Φ , if the previously added pieces can be placed so that they encode the solution, then the same holds when the pieces of this gadget are added (Lemma 2.5).
- In a valid placement of all the pieces, if the earlier introduced pieces have an aligned placement, then the pieces of this gadget must have an almost-canonical placement (Lemma 2.9).
- In a valid placement of all the pieces, if the earlier introduced pieces have an aligned placement and the pieces of this gadget have an almost-canonical placement, then the pieces of this gadget must also have an aligned placement (Lemma 2.10).
- For each edge (p_1, p_2) of the dependency graph G_x , where p_1 and p_2 are pieces of this gadget, we have $\langle p_1 \rangle \leq \langle p_2 \rangle$, i.e., the value encoded by p_1 is at most that encoded by p_2 (Lemma 2.12).

Some of the steps will be very similar for all the gadgets. In order to avoid unnecessary repetition, we will handle the first two gadgets, the anchor and the swap, in greater detail than the subsequent gadgets.

4.1 Anchor

Recall that each variable x is represented by two wires \vec{x} and \overleftarrow{x} in the wiring diagram of the instance \mathcal{I} of WIRED-CURVE-ETR[f, g] which we reduce to a packing instance. Furthermore, the left endpoints of the wires are vertically aligned and occupy neighbouring diagram lines ℓ and ℓ' , as do the right endpoints. In our packing instance, we cover each wire with variable pieces that can slide back and forth and thus encode the value of x , and the pieces covering one wire are called a *lane*. In order to make the value represented by the lane on \vec{x} consistent with




Figure 20. Left: A simplified illustration of the anchor and how it is placed on top of the two wires representing a variable x . Right: If the blue piece is pushed to the left, the yellow piece must move by an equal amount to the right, and vice versa. Color codes: 1 blue, 2 yellow, 3 orange.

that of the lane on \overleftarrow{x} , we make an *anchor* at both ends, which will propagate a push from one lane to the other. Most of this section will be about the anchors at the left ends of the wires. The anchors at the right ends will be handled in the end of the section.

When we use an anchor in our construction, we also define part of the boundary of the container. Two of the three introduced pieces are variable pieces that will extend out through the right side of the gadget, and the remaining part of those will be defined as part of another gadget farther to the right, which will be described in other parts of the paper. It is a general convention in our figures of gadgets that if a part of the boundary of the gadget is drawn with thick full segments, it will be part of the container boundary. If part of the boundary is drawn with thick dashed segments, it means that the segments can be either part of the container boundary or part of the boundary of other pieces that have been introduced to the construction in earlier steps.

Simplified anchor

The anchor is meant to be a connection between the two lanes that represent a variable x ; see Figure 20 for an illustration. The gadget consists of part of the boundary of the container and three pieces: orange, yellow, and blue. The yellow and blue pieces are the two leftmost pieces on the lanes of \overrightarrow{x} and \overleftarrow{x} , respectively. The orange piece functions as a connection between the two lanes. The idea is that if we move the blue piece to the left by t , then we have to move the yellow piece to the right by at least t as well, and vice versa.

The segment bounding the gadget from below is part of the container boundary. The segment bounding the gadget from above is part of the boundary of a piece introduced in an earlier anchor, except for the very first anchor, which will be bounded from above by the container boundary.




Figure 21. Left: An illustration of the anchor gadget. The placement of the pieces corresponds to the value $x = 0$. The arrows show the orientation of the pieces, and the dots show the fingerprinted corners. The magnifying glass shows a scale of the correspondence between placements of the yellow piece and the encoded value of the variable x . The length a is $1 + O(\delta)$ and depends on α . Right: Another instance of the gadget with other angles chosen.

The actual anchor

See Figure 21 for an illustration of the following description. Recall that we need the slack added by each gadget to be only $O(\delta)$, where $\delta := n^{-300}$. We therefore design the boundary of the anchor to follow the pieces closely. The yellow and blue pieces are fingerprinted on the boundary, as indicated by the dots. The orange piece is fingerprinted in the wedge created by the yellow and blue pieces.

Lines that appear axis-parallel must be axis-parallel; those are important for the alignment. The height of the orange piece is 14. The angle α is in the range $[3\pi/4, 7\pi/8]$. The lower bound ensures that a range of size at most 2δ is needed for the orange piece, while the upper bound ensures that the length a is only $1 + O(\delta)$. The angle α and the angles where the yellow and blue pieces are fingerprinted are not fixed, and we therefore have flexibility to choose the angles of the fingerprinted corners freely to obtain the unique angle property; two different choices of angles are shown in Figure 21.

Canonical placements and solution preservation

As the next step, we define the set of canonical positions for the three pieces. The yellow and blue pieces are variable pieces, and in the placement shown in Figure 21, both pieces encode the value $x = 0$. By definition, the canonical placements of each are all placements obtained by sliding the piece to the left or right by distance at most δ from the placement shown. Recall that a placement of all pieces of a gadget is canonical if it is (1) valid, (2) canonical for each variable piece, and (3) have certain extra properties defined for each gadget individually. For the




Figure 22. The three triangles T_1, T_2, T_3 (shown a bit larger than they should be for clarity) are a witness that the intended placements (middle) are $O(\delta)$ -sound. We can therefore conclude that the pieces have an almost-canonical placement, as shown to the right (ignore here that the shown placement is not a valid placement).

anchor, we specify point (3) as having edge-edge contacts between the pieces and the container boundary as shown in the figure.

We are now ready to prove Lemma 2.5 for the anchor, namely that the reduction preserves solutions.

PROOF OF LEMMA 2.5 FOR THE ANCHOR. Suppose that for a given solution \mathbf{x} to the ETR-INV formula Φ , there exists a canonical placement of the previously introduced pieces \mathbf{p}_{i-1} that encodes that solution. To extend the placement to the pieces \mathbf{p}_i , i.e., with the yellow, blue, and orange piece of this anchor included, we place the yellow and blue pieces so that they represent the value of x in \mathbf{x} . This leaves room for the orange piece to be placed with the required edge-edge contacts to the blue and yellow pieces. ■

Fingerprinting and almost-canonical placement

We now prove Lemma 2.9. Recall that in this lemma, we assume that the previous pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and we want to conclude that the new pieces $\mathbf{p}_i \setminus \mathbf{p}_{i-1}$, i.e., the three pieces of this anchor, must have an almost-canonical placement.

PROOF OF LEMMA 2.9 FOR THE ANCHOR. Since the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement, the gadget is bounded from above by an earlier introduced piece (unless $i=1$, in which case it is bounded by the boundary of the container). We want to use Lemma 3.5 (Multiple Fingerprints) to prove that the three pieces have an almost-canonical placement. For this we need to point out intended placements that are $g\mu$ -sound; recall Definition 3.3. We choose the placement shown in Figure 21. To certify that the intended placements are $g\mu$ -sound, we need to point out three triangles T_1, T_2, T_3 , such that E_j^s is $g\mu$ -bounding T_j for $j \in \{1, 2, 3\}$. Here E_j^s is the empty space left by the pieces \mathbf{p}_{i-1} and the intended placements of the pieces q_1, \dots, q_{j-1} , where q_1, q_2, q_3 are the blue, yellow, and orange piece, respectively. We choose the triangles as the tips of fingerprinted corners, as shown in Figure 22. The placements of the blue and yellow pieces are δ -sound because the boundary of the container or the pieces of \mathbf{p}_{i-1} follow the relevant edges of the respective triangles T_1 and T_2 within distance δ . The placement of the orange piece



Figure 23. The parallel edges of the yellow and blue pieces must intersect the alignment segment ℓ . Therefore, the segments must be horizontal.

is 0-sound since the boundaries of the blue and yellow pieces contain the edges of T_3 . As the intended placements are δ -sound, we conclude using Lemma 3.5 (Multiple Fingerprints) that the displacement is $O(n^{-48})$, so the three pieces must have an almost-canonical placement. ■

Aligned placement

We show now Lemma 2.10, which assumes that the pieces in the anchor are almost-canonical and all previous pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement. The goal is to conclude that the pieces of this gadget have an aligned $i\mu$ -placement, which requires the variable pieces to be correctly aligned and encode values in the range $[-i\mu, i\mu]$.

PROOF OF LEMMA 2.10 FOR THE ANCHOR. Consider a valid placement where the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and the pieces $\mathbf{p}_i \setminus \mathbf{p}_{i-1}$, i.e., the yellow, blue, and orange pieces of this anchor, have almost-canonical placements; see Figure 23. We consider the alignment segment ℓ which has length 20. Since the placements are almost-canonical, it follows that ℓ crosses both of the parallel edges of the yellow and the blue pieces. Since the distance between the segments in each pair is 10, it follows that the segments must be horizontal.

In order to show that the pieces \mathbf{p}_i have an aligned $i\mu$ -placement, it remains to bound the horizontal displacement of the yellow and blue pieces as compared to the placements encoding the value 0. Consider the yellow piece p , the argument for the blue piece is similar. We need to prove $\langle p \rangle \in [-i\mu, i\mu]$. We will actually show the stronger statement that $\langle p \rangle \in [-\delta, \mu] \subset [-i\mu, i\mu]$.

Recall that p is right-oriented. The constraint that p must be inside C ensures that $\langle p \rangle \geq -\delta$. Since the pieces have an almost-canonical placement by assumption, we know that the displacement as compared to the situation in Figure 21 is at most n^{-1} . It therefore follows that no other pieces than the orange piece can fit in the region to the left of the yellow and blue pieces. We consider a canonical placement and analyze how much p can slide to the right before too much empty space has been made in the region. Observe that sliding p to the right by t creates empty space of area $10t$, since the height of the piece is 10. It therefore follows that we must have $t \leq \mu/10$. This translates to $\langle p \rangle \leq \delta + \mu/10 \leq \mu$, and we are done. ■




Figure 24. If the pieces do not have a canonical placement, there is room for the orange piece to be oriented in the canonical way and get a vertical edge-edge contact with the container. Then the blue and yellow pieces can be pushed to the left to obtain a canonical placement, which will decrease the value encoded by the yellow piece and increase the value encoded by the blue piece.




Figure 25. The figure shows how we adjust the left wall of an anchor to ensure that we respect the range constraint.

Edge inequalities

Recall that Lemma 2.12 states that for any edge (p_1, p_2) in the dependency graph G_x , we have the inequality $\langle p_1 \rangle \leq \langle p_2 \rangle$, where $\langle \cdot \rangle$ denotes the value encoded by a piece. We show this lemma now for the anchor.

PROOF OF LEMMA 2.12 FOR THE ANCHOR. Denote by p_1 and p_2 the blue and yellow piece, respectively. The pieces induce the edge (p_1, p_2) in G_x and we have to show $\langle p_1 \rangle \leq \langle p_2 \rangle$. We have that $\langle p_1 \rangle = \langle p_2 \rangle$ when the pieces have an edge-edge contact with the orange piece, as shown in Figure 21. We have to show that it is not possible that $\langle p_1 \rangle > \langle p_2 \rangle$. This could only potentially happen if the pieces do not have a canonical placement. However, some straightforward rotation arguments show that in this case we even have $\langle p_1 \rangle < \langle p_2 \rangle$; see Figure 24. It follows that $\langle p_1 \rangle \leq \langle p_2 \rangle$. ■

Range Constraints

Recall that together with the variable x is also given an interval $I(x)$ of one of the forms $[-\delta, \delta]$, $[-0, \delta]$, $\{\delta\}$. It is claimed in Lemma 2.13 that the pieces representing the variable x encode a consistent value $\langle K_x \rangle$, which is in the range $I(x)$ due to the design of the anchor. This is ensured by adjusting the left wall of the anchor as showed in Figure 25, so that the value




Figure 26. Anchors are placed one on top of the other to form two staircases. The alignment segments ℓ_1, \dots, ℓ_n are used to align the blue and yellow pieces of the anchors in this order. The dots in the orange pieces in the right staircase mark the fingerprinted corners (all other pieces are fingerprinted in other gadgets).

encoded by the yellow piece is bounded from below according to $I(x)$. Since the blue piece is restricted to encode a value of at most δ , it then follows that $\langle K_x \rangle \in I(x)$.

Staircases of anchors

As the next step, we describe how we organize all the anchors of the construction. Recall that the wires of the wiring diagram appear and disappear in the order $(\vec{x}_1, \overleftarrow{x}_1), \dots, (\vec{x}_n, \overleftarrow{x}_n)$ from left to right in a staircase-like fashion. Therefore, we also stack the anchors onto one another as displayed in Figure 26, so that the boundary of each stack appears similar to a staircase. This ensures that the container will be 4-monotone, which is used when proving hardness for packing into a square container in Section 6. The rest of the construction will be in between these two staircases. Anchors in the right staircase are treated in the following paragraph.

Anchors in the right staircase

For an anchor on the right side, the entering yellow and blue variable pieces have been started in other gadgets farther to the left, and we only add the orange piece. We define the canonical placements in an analogous way as for the left anchors. The proof of Lemma 2.5 (solution preservation) is trivial since no new variable pieces are introduced. The orange piece in a right anchor is fingerprinted as in the left anchors, and Lemma 2.9 follows. Lemma 2.10 is trivial since no new variable pieces are introduced. The edge inequality of Lemma 2.12 is proven as for the left anchors.

4.2 Swap

Idea

Recall that in the wiring diagram, the wires may cross each other (see Figure 3). On top of such a crossing, we build a swap. The purpose of the swap is thus to make a crossing of two neighbouring lanes of pieces. To get intuition about how the gadget works, consider Figure 27.




Figure 27. The simplified swap. The yellow pieces representing x are either both right- or both left-oriented. Likewise for the blue pieces representing y . For color codes, see Figure 28.




Figure 28. The actual swap. Color codes: 1 blue, 2 yellow, 3 turquoise, 4 orange, 5 yellow, 6 blue.

The yellow pieces encode a variable x and the blue pieces encode a variable y . It is possible that $x = y$, which will happen only when the two wires \vec{x} , \overleftarrow{x} cross each other. Therefore, the yellow pieces will have the opposite orientation of the blue pieces in this special case.

We want to show that when the pieces have edge-edge contacts to the orange piece, the variables are encoded consistently, so that the lanes have been swapped. The key observation is that if the left blue piece pushes to the right and the yellow pieces are fixed, then the orange piece will slide along the yellow pieces and push the right blue piece by an equal amount. Similarly, if the left yellow piece pushes to the right, then the orange piece will slide along the blue pieces, and will push the right yellow piece by an equal amount. For this to work the two opposite edges of the orange piece in contact with the blue pieces must be parallel and similarly the two other edges in contact with the yellow pieces must be parallel as well. The conclusion is that for all placements of the orange piece where it has edge-edge contacts to all yellow and blue pieces, the horizontal distance between the two yellow pieces is the same, as is the distance between the two blue pieces.




Figure 29. Left: the construction without the turquoise piece. Middle: We remove part of the pieces, leaving a triangular empty space with edges of length more than 1. Right: The turquoise piece is designed to fit in the triangle of empty space, but it is surrounded by empty space of thickness δ in every direction. This leaves enough wiggle room for the other pieces to encode all solutions to Φ .

The actual swap

See Figure 28 for the following description. The swap consists of six pieces. Those are the left and right yellow piece, the left and right blue piece, the orange piece and the turquoise piece.

The left yellow and blue pieces extend outside the gadget to the left, where they have been introduced in other gadgets added earlier to the construction. Likewise, the right yellow and blue pieces extend outside the gadget to the right, where their ends will be defined in other gadgets added later to the construction. The gadget is bounded by horizontal segments from above and below, and these are either part of the container boundary or the boundary of other pieces that have been added earlier. As is seen from the figure, the yellow pieces have corners with the same angle α , and the blue pieces have corners with the same angle β . The precise value of those angles is chosen freely for the fingerprinting. Similarly, the orange and turquoise piece have a corner with a flexible angle that can be chosen freely for fingerprinting. The orange piece has a horizontal top and bottom edge of length 2.

For the way we construct the turquoise piece, we refer the reader to Figure 29. The role of the turquoise piece is solely to be able to fingerprint the orange piece. It has no direct use in the functionality of the swap. If we avoid using the turquoise piece and instead use pieces as in the simplified Figure 27, the left corner of the orange piece, that we want to fingerprint in the wedge between the left yellow and blue pieces, has angle $\alpha + \beta$, and thus the unique angle property is violated; indeed, the right yellow and blue pieces can cover the wedge equally well. The left endpoints of the horizontal segments of the orange piece cannot be used for fingerprinting, because the angles are more than $\pi/2$. It may be tempting to believe that one could avoid the turquoise piece, but we could not find such a way using only convex pieces. Note that we have only two degrees of freedom without the turquoise piece, as certain edges must be parallel. This is not enough to choose three angles freely for fingerprinting.

Canonical placements and solution preservation

Recall that the yellow and blue pieces represent the variables x and y , respectively. In Figure 28, all variable pieces encode the value 0. For a placement of the six pieces to be canonical, we require that the yellow and blue pieces have edge-edge contacts to the orange piece as shown and that the turquoise piece is enclosed by the left yellow and blue pieces and the orange piece as shown.

We prove now Lemma 2.5, about solution preservation, for the swap.

PROOF OF LEMMA 2.5 FOR THE SWAP. Suppose that for a given solution \mathbf{x} to the ETR-INV formula Φ , there exists a canonical placement of the previously introduced pieces \mathbf{p}_{i-1} that encodes that solution. To extend the placement to the pieces \mathbf{p}_i , i.e., with the yellow, blue, turquoise and orange pieces of this swap included, we place the yellow and blue pieces so that they represent the value of x and y in \mathbf{x} . This leaves room for the orange piece to be placed with the required edge-edge contacts to the blue and yellow pieces. The turquoise piece has enough wiggle room to be placed correctly as well. ■

Fingerprinting and almost-canonical placement

Here we are less detailed in the application of Lemma 3.5 (Multiple Fingerprints) than in the section about the anchor in order to avoid unnecessary repetition.

Recall that in Lemma 2.9, we assume that the previous pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and we want to conclude that the new pieces \mathbf{p}_i have an almost canonical placement.

PROOF OF LEMMA 2.9 FOR THE SWAP. Recall that \mathbf{p}_{i-1} consists of all pieces introduced previously, including the left blue and yellow pieces, whereas \mathbf{p}_i additionally includes the orange, turquoise and the right blue and yellow pieces. We now fingerprint the turquoise, orange, right blue, and right yellow pieces in this order. We use the intended placements of those pieces shown in Figure 28 (corresponding to the case where the variables encode the value 0) and fingerprint the corners marked with dots. These intended placements are $O(\delta + (i-1)\mu)$ -sound for any aligned $(i-1)\mu$ -placement of the pieces \mathbf{p}_{i-1} , since the empty space has thickness $O(\delta)$ and the left blue and yellow pieces have displacement of at most $(i-1)\mu$ compared to the shown placement, by the assumption of Lemma 2.9. We conclude using Lemma 3.5 (Multiple Fingerprints) that the displacement is $O(n^{-48})$, so the pieces must have an almost-canonical placement. ■

Aligned placement

PROOF OF LEMMA 2.10 FOR THE SWAP. Consider a valid placement where the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and the pieces of this swap have an almost-canonical place-




Figure 30. Left: Due to the alignment segment ℓ_1 , the right yellow and blue pieces must be axis-parallel. Right: Due to the alignment segment ℓ_2 , the orange piece must also be in a canonical placement.

ment; see Figure 30 (left). Similarly as in the proof for the anchor, we consider the alignment segment ℓ_1 and get that the right yellow and blue edges are in horizontal position.

In order to show that the pieces p_i have an aligned $i\mu$ -placement, it remains to bound the horizontal displacement of the right yellow and blue pieces as compared to the placements encoding the value 0. Consider the yellow piece, the argument for the blue is similar. Let p_1 be the left yellow piece and p_2 the right one. By assumption, we have $\langle p_1 \rangle \in [-(i-1)\mu, (i-1)\mu]$, and we need to prove $\langle p_2 \rangle \in [-i\mu, i\mu]$. Suppose that the yellow pieces are right-oriented; the other case is similar. From the proof of Lemma 2.12 for the swap (given below), we have that $\langle p_1 \rangle \leq \langle p_2 \rangle$, so we just need to show $\langle p_2 \rangle \leq i\mu$. It is therefore sufficient to show $\langle p_2 \rangle \leq \langle p_1 \rangle + \mu$. This follows as in the proof for the anchor by considering how much the piece p_2 can be slid to the right before the empty space thus created gets larger than μ . ■

Edge inequalities

PROOF OF LEMMA 2.12 FOR THE SWAP. In the swap, the yellow pieces induce an edge in the graph G_x and the blue induce an edge in G_y . In the special case that $x = y$, the blue and yellow pieces have opposite orientations, so according to the rule about when to add edges to the dependency graph G_x , there will also be an edge between the left pieces and one between the right pieces. We make an exception to the rule and do not add these edges to the dependency graph.

We now prove the edge inequality for the edge between the yellow pieces; the argument for the blue pieces is analogous. Suppose that the yellow pieces are right-oriented and let p_1 be the left and p_2 be the right, so that they induce the edge (p_1, p_2) of G_x . The argument is analogous if they are left-oriented. Recall that $\langle p_1 \rangle = \langle p_2 \rangle$ exactly when the pieces have the horizontal distance shown in Figure 28, where the orange piece has edge-edge contacts with both yellow pieces. Since the pieces have an almost-canonical placement by assumption, the displacement of the orange piece is at most n^{-1} . It therefore follows that the same pair of parallel edges of the orange piece will prevent the yellow pieces from being closer than in the figure, so we have $\langle p_1 \rangle \leq \langle p_2 \rangle$. ■




Figure 31. The split. The yellow, blue, and green pieces represent the variable x and are either all right-oriented or all left-oriented. Color codes: 1 yellow, 2 orange, 3 blue, 4 green.

4.3 Split

The purpose of the split is to make an extra lane representing a variable x . This will be needed in order to lead lanes into the gadgets for the adders and curvers.

Description

See Figure 31 for an illustration of the split. We always split the topmost lane in the construction, so that there is room to expand above with one more lane. Therefore, the split is bounded from above by the boundary of the container and from below by pieces in the second-highest lane, which have been added to the construction earlier.

The yellow piece extends outside the gadget to the left, where it has been introduced to the construction in an earlier step. The yellow piece is in contact to the right with an orange piece with height 20, i.e., twice the height of a lane. The orange piece is in contact to the right with the blue and green pieces, which extend outside the gadget to the right, where they will enter other gadgets defined later in the construction. Each of the orange, blue, and green piece has a corner with an angle that can be freely chosen for fingerprinting.

Canonical placements and solution preservation

The yellow, blue, and green pieces are all variable pieces encoding a variable x . The position as indicated in Figure 31 show the situation where they all encode the value 0. The placement of all four pieces is canonical if they have the shown edge-edge contacts. Lemma 2.5 (solution preservation) follows trivially.

Fingerprinting and almost-canonical placement

The proof of Lemma 2.9 for the split is completely analogous to the one for the swap.




Figure 32. The left-split.

Aligned placement

All pieces must be aligned because of the container boundary and the edge of a piece from \mathbf{p}_{i-1} bounding the gadget from below. As for the swap, to get Lemma 2.10 for the split, we need to verify that the blue and green piece encodes a value that is at most μ larger than that encoded by the yellow piece. The proof is analogous as the one for the swap.

Edge inequalities

In the split, we get two edges in the dependency graph. If the pieces are right-oriented, we have edges from the yellow to green and to the blue pieces. Otherwise, we have edges from green to yellow and blue to yellow. That the edges satisfy the edge inequality (Lemma 2.12) follows by construction, since the orange piece restricts how close the yellow piece can get to the green and blue pieces.

Left-split

We explained above how to split an entering variable piece into two exiting pieces. We will also need the opposite, i.e., splitting an exiting piece into two entering pieces (or, in other words, merging two entering pieces into one exiting piece). It is straightforward to construct such a gadget by similar principles, as shown in Figure 32. In order to avoid disambiguation, we will occasionally denote these splits as *right-splits* and *left-splits*, respectively.

4.4 Adder

Idea

For the following description, see Figure 33. Here we explain the principle behind the adder for $x + y \leq z$. The adder for $x + y \geq z$ is identical, but has the entering variable pieces for x, y, z oriented to the left instead of to the right. The adder has three entering variable pieces (yellow, blue, left green), representing three variable (x, y, z) . There is also one exiting green variable piece representing z . In addition to this, there are three pieces that are not variable pieces (turquoise, orange, pink). The role of the turquoise piece is to transform horizontal




Figure 33. Left: The simplified adder. Right: The variable x is increased by two units. The variable y is decreased by one unit. Thus the variable z is increased by one unit. Color codes: 1 blue, 2 yellow, 3 green, 4 turquoise, 5 orange, 6 pink, 7 green.




Figure 34. The actual adder.

motion to the right of the blue piece to downwards vertical motion. Motion of the orange piece downwards or to the right both make the pink piece move to the right by an equal amount. Therefore, when the blue and yellow pieces push to the right, the pink piece will be pushed to the right by the sum of the two motions.

Actual description

The actual adder is shown in Figure 34. The figure shows the situation where all the variable pieces encode the value 0. The actual adder varies in several points from the simplified version in Figure 33, which is needed in order to use fingerprinting. The orange, turquoise, pink, and right green pieces must be fingerprinted, so they cannot have only nice angles as in the simplified gadget. The turquoise piece can easily be fingerprinted using the top right corner. The orange piece is fingerprinted in the upper left corner. The pink and right green pieces are fingerprinted at their lower left corners. In each case, the angle can be chosen freely by changing the slope of the edge-edge contact with the piece to the left.

As a consequence of changing the angle on the top left of the orange piece to something else than $\pi/4$, we also have to change the angle of the top right corner. See Figure 35 for an




Figure 35. In each of these three pictures, thick segments drawn with the same color are equally long. Left: The construction of the orange piece. Middle: Pushing the orange piece down by its height makes it push the pink piece by the same amount. Right: This also holds when pushing less.




Figure 36. The lines through corners and perpendicular to diagonals are disjoint from the interiors.

illustration of the following. First we describe how to construct the orange piece and then we explain why it actually works. The orange piece is a trapezoid with horizontal bottom and top edges and height $10 - 10\delta$. The left edge of the orange piece is parallel to the right edge of the yellow piece. The length of the top edge should be at least as long as the bottom edge of the turquoise piece. The length of the bottom piece is the length of the top edge minus the height $10 - 10\delta$. The right edge is determined by the description of all the other edges.

We need to explain why pushing the blue piece to the right by some amount $t > 0$ will push the pink piece to the right by t as well; see Figure 35. It is helpful to consider the case where t equals the height h of the orange piece (even though there is not room for pushing the pieces so much). This push of the blue piece will push the orange piece down by h . Since the length of the top edge of the orange piece equals the length of the bottom edge plus h , the pink piece will be pushed to the right by h as well. All the pieces move linearly, so it will also be the case for smaller values of t .

We furthermore want the property that for each of the top corners of the turquoise piece, the line through the corner and perpendicular to the diagonal of the corner is a tangent to the piece (i.e., the line intersects the piece only at the corner). The same must hold for the bottom corners of the orange piece; see Figure 36. It is easy to choose the fingerprinted angles so that the pieces have this property. We will use the property in the proof of Lemma 2.14 given later to conclude that if the turquoise or the orange piece is not aligned as in the canonical placements, they will take up too much space.




Figure 37. Incorporation of the adder. Splits are marked in turquoise, swaps are marked red.

Canonical placements and solution preservation

The canonical placements are defined as the placements with the edge-edge contacts as shown in Figure 34. That there is a canonical placement encoding any given solution to Φ (Lemma 2.5) follows by construction.

Incorporating the gadget

In order to incorporate the adder into the construction, we use two splits and eight swaps in order to organize the in-going and out-going lanes to the gadget; see Figure 37 for a schematic illustration.

Fingerprinting and almost-canonical placement

The proof of Lemma 2.9 for the adder is completely analogous to the one for the swap.

Aligned placement

The right green piece must be aligned because of the container boundary and the edge of a piece from \mathbf{p}_{i-1} bounding the gadget from below. As for the swap, to get Lemma 2.10 for the adder, we need to verify that the right green piece encodes a value that is at most μ larger than that encoded by the left green piece. The proof is analogous as the one for the swap.

Edge inequalities

In the adder, the left and right green pieces induce an edge in the dependency graph. That the edge satisfies the edge inequality (Lemma 2.12) is proven as for the swap.

The adder works

In this paragraph we prove Lemma 2.14. Here we are considering an aligned $g\mu$ -placement, and we need to prove that for every addition constraint $x + y = z$ of Φ , we have $\langle K_x \rangle + \langle K_y \rangle = \langle K_z \rangle$.




Figure 38. Top: If the turquoise piece does not have edge-edge contacts to the blue piece and the boundary, then all points of the lower edge have strictly smaller y -coordinates than otherwise. Bottom: If the orange piece does not have edge-edge contacts to the yellow and pink pieces, then the top edge is strictly higher than otherwise. The dashed segments are the edges in the canonical situations shown to the left.

PROOF OF LEMMA 2.14. We prove the inequality $\langle K_x \rangle + \langle K_y \rangle \leq \langle K_z \rangle$. The other inequality follows from analogous arguments about the gadget for $x + y \geq z$. Let p_x, p_y, p_{z1} be the yellow, blue, and left green pieces, and p_{z2} be the right green piece.

Note first that due to the way we incorporate the adder, there are paths P_x and P_y in G_x and G_y attached to and directed away from the cycles K_x and K_y , and the vertices farthest away from the cycles are the pieces p_x and p_y . Furthermore, the pieces p_{z1} and p_{z2} are two consecutive vertices on the cycle K_z , so by Lemma 2.13, we have $\langle p_{z1} \rangle = \langle p_{z2} \rangle = \langle K_z \rangle$.

We argue that if $\langle p_x \rangle + \langle p_y \rangle = \langle K_z \rangle$, the only way to place the pieces is the canonical way. This excludes the situation $\langle p_x \rangle + \langle p_y \rangle > \langle K_z \rangle$, since there the turquoise, orange, and pink pieces have strictly less space. Consider first the turquoise piece. It is straightforward to check that if it does not have edge-edge contacts to the blue piece and the container boundary, then the lower edge will be strictly below the placement of the edge where these contacts were present, in the sense that every point on the segment will have a placement with a smaller y -coordinate. This can be seen in Figure 38 and is due to the property that the lines through corners perpendicular to diagonals are tangents, as shown in Figure 36. Similarly for the orange piece, if it does not have edge-edge contacts with the yellow and pink pieces, the top edge is strictly above the edge in the placement where it does have these contacts. In the canonical placements, the bottom edge of the turquoise piece is contained in the top edge of the orange piece. Therefore, in any other placement of the turquoise and orange pieces, there will be points on the turquoise edge below the orange edge, which makes the placement invalid. We can now conclude that in general, we must have $\langle p_x \rangle + \langle p_y \rangle \leq \langle K_z \rangle$.

To finish the proof, note that since the paths P_x and P_y are directed away from the cycles K_x and K_y , we get from the edge inequalities (Lemma 2.12) that

$$\langle K_x \rangle + \langle K_y \rangle \leq \langle p_x \rangle + \langle p_y \rangle \leq \langle K_z \rangle. \quad \blacksquare$$

5. Curvers

We are going to construct two different types of curvers. Both curvers will be described in a convex and a concave version. To prove $\exists\mathbb{R}$ -hardness of the problem $\text{PACK}[\square \rightarrow \sqsupset, \circlearrowleft +]$, we will use the convex and the concave *swing*. For the problems $\text{PACK}[\square \rightarrow \triangleleft, +]$ and $\text{PACK}[\triangleleft \rightarrow \sqsupset, +]$, we will instead use variants of the convex and the concave *gramophone*.

5.1 Swing

The swing is reminiscent of the swings found on playgrounds, especially the twin version where two kids face each other and must collaborate to get the swing going; see Figure 39. In the preliminary version of this paper [6], we also used other playground inspired gadgets, namely the *teeter-totter* and the *seesaw*, which could encode a convex and a concave constraint, respectively. However, the seesaw needed non-convex pieces, and since our new swing can also encode a convex constraint, we need neither of them anymore. The orange piece in the swing plays the role of the swing itself and can rotate around the top corner o , while the blue and yellow pieces correspond to the children, pushing the swing from each side. Surprisingly, by adjusting the height of the contact points c_x and c_y of the orange piece with the other pieces, we can make versions of the swing that enforce convex and concave constraints, as we like.

Principle

For the following argument, see Figure 40. We assume for simplicity that $|oa| = 1$ and consider the situation where the orange piece is rotated so that the segment oa makes an angle of $\varphi \in (-\pi/4, \pi/4)$ with the vertical axis. The corners of the blue and yellow pieces that are in contact with the orange piece lie on the horizontal line ℓ_z which is $z > 0$ below the horizontal line ℓ_o through o . We consider the extensions of the edges of the orange piece incident at a and construct the intersections c_x and c_y with ℓ_z . We want to find formulas for horizontal distances




Figure 39. A simplified drawing of the swing. The corner a is a right angle. Color codes: 1 blue, 2 yellow, 3 orange.




Figure 40. The geometry of the swing.

x_1 and y_1 from o to c_x and c_y , respectively, since the rightmost corner of the blue piece will be at c_x or to the left, and the leftmost corner of the yellow piece will be at c_y or to the right, so x_1 and y_1 will lead to lower bounds for the values encoded by the pieces, as the pieces are oriented away from c_x and c_y , respectively. (In the figure, c_x is in the exterior of the orange piece and would thus not be a bound on the placement of the blue piece, but when using the gadget, the point will always be on the boundary of the orange piece, since it will only be slightly rotated.) The distances x_1 and y_1 are signed distances, so that x_1 is positive if and only if c_x is to the left of o_z and y_1 is positive if and only if c_y is to the right of o_z .

Note that $|af| = \cos \varphi$. Considering the right triangles afb and afc , we then obtain

$$|ab| = \frac{\cos \varphi}{\cos(\pi/4 + \varphi)};$$

$$|ac| = \frac{\cos \varphi}{\cos(\pi/4 - \varphi)}.$$

We then get

$$x_0 = |ob| = |fb| - |of| = \sin(\pi/4 + \varphi)|ab| - |of| = \tan(\pi/4 + \varphi) \cos \varphi - \sin \varphi;$$

$$y_0 = |oc| = |fc| + |of| = \sin(\pi/4 - \varphi)|ac| + |of| = \tan(\pi/4 - \varphi) \cos \varphi + \sin \varphi.$$

We observe that if l_z coincides with l_o , then $x_1 = x_0$, while if l_z passes through a , then $x_1 = -\sin \varphi$. Since the vertical distance from l_o down to a is $\cos \varphi$, we define $v := z/\cos \varphi$ and get x_1 by interpolating between x_0 and $-\sin \varphi$ using the factor v , and we can express y_1 in an analogous way. By standard trigonometric identities, we obtain the expressions

$$x_1 = x_0 \cdot (1 - v) - \sin \varphi \cdot v = \frac{1 + z(-\cos \varphi - \sin \varphi)}{\cos \varphi - \sin \varphi};$$

$$y_1 = y_0 \cdot (1 - v) + \sin \varphi \cdot v = \frac{1 + z(-\cos \varphi + \sin \varphi)}{\cos \varphi + \sin \varphi}.$$

We consider x_1 and y_1 as functions of φ . Note that $x_1(0) = y_1(0) = 1 - z$. The curve $\gamma_z : \varphi \mapsto (x_1(\varphi), y_1(\varphi))$ undergoes a remarkable metamorphosis as we increase z from 0 to 1;




Figure 41. The curve γ_z for each of the values $z \in \{i/20 \mid i = 0, \dots, 20\}$. The curves for $z = 0.7$ and $z = 0.9$, that we will use in our construction, have been highlighted.

see Figure 41. Note the sudden change from the curve being convex to concave at the diagonal point $(1 - z, 1 - z)$ as z changes from $z = 0.7$ to $z = 0.8$. This is what we exploit to make convex and concave variants of the swing.

To obtain well-behaved functions f and g , as required in Theorem 1.4, we will use an implicit function $F_z : U^2 \rightarrow \mathbb{R}$, for an interval $U \subset \mathbb{R}$ around 0, such that $F_z(x, y) = 0$ if and only if $(x, y) = \gamma_z(\varphi) + (z - 1, z - 1)$ for some $\varphi \in (-\pi/4, \pi/4)$. We add $(z - 1, z - 1)$ to obtain a curve that passes through $(0, 0)$ for $\varphi = 0$. In order to show that there exists such a function which is well-behaved, we use the following lemma from [50, Sec. 1.5], which has been slightly rephrased.

LEMMA 5.1 (Miltzow and Schmiermann [50]). *Let $\gamma = (\gamma_x, \gamma_y) : (-\pi/4, \pi/4) \rightarrow \mathbb{R}^2$ be a C^3 -function such that $\gamma(t) = (0, 0)$ if and only if $t = 0$. Assume that the derivatives $\gamma'_x, \gamma''_x, \gamma'_y, \gamma''_y$ are all rational in 0 and that $\gamma'_x(t) \neq 0$ for all $t \in (-\pi/4, \pi/4)$. Then there exists an interval around*




Figure 42. The convex swing (left) and the concave swing (right).

the origin $J \subset \mathbb{R}$ and a well-behaved function F such that

$$\{\gamma(t) \in \mathbb{R}^2 \mid t \in (-\pi/4, \pi/4)\} = \{(x, y) \in J \times \mathbb{R} \mid F(x, y) = 0\}.$$

In our case, it follows trivially from the definitions that the first and second derivatives of x_1 and y_1 in $\varphi = 0$ are rational and that $x'_1(0) \neq 0$, so the lemma ensures the existence of a function F_z as described.

We will use $f := F_{7/10}$ for our convex constraint and $g := F_{9/10}$ for our concave constraint. It is then straightforward to check that f is convex and g concave at $(0, 0)$ by evaluating the sign of the curvature of the curves $\gamma_{7/10}$ and $\gamma_{9/10}$ at $\varphi = 0$.

Actual swing

Figure 42 shows a more accurate illustration of the convex and the concave swing. The blue and yellow variable pieces represent variables x and y , respectively. The blue piece enters the gadget and is left-oriented, while the yellow piece exits the gadget and is right-oriented. The top corner o of the orange piece is fingerprinted at a wedge formed by the boundary of the container C . We need to ensure that the angle at the corner o can be changed continuously. By moving the orange piece up or down, we can also change the lengths of the edges meeting at the corner a , and then the angle at o also changes. We need to move the contact corners c_x and c_y of the blue and yellow pieces correspondingly, so that the z -value (described above) remains invariant. The yellow piece is fingerprinted at a wedge formed by the blue piece and the lower boundary of the gadget.

The figure shows the case where the blue and the yellow piece encode the value 0 of both x and y . In the convex swing, the vertical distance from o to the contact corners c_x, c_y is $7|oa|/10$. In the concave swing, the distance is $9|oa|/10$.

Incorporation of the swing

Figure 43 shows a schematic illustration of how the swing is incorporated in the construction. We use a right-split and a swap before the swing itself, and a left-split and a swap after.




Figure 43. Incorporation of the swing.

Canonical placements and solution preservation

A placement of the three pieces of the gadget is canonical if the orange piece separates the yellow and blue pieces in the sense that the segment $c_x c_y$ (connecting the contact corners of the blue and yellow pieces) intersects both edges of the orange piece incident at the corner a .

PROOF OF LEMMA 2.5 FOR THE SWING. Suppose that for a given solution to the CURVE-ETR[f, g] formula Φ , there exists a canonical placement of the previously introduced pieces \mathbf{p}_{i-1} that encodes that solution. We place the yellow piece so that it encodes the value of y as specified by Φ . We need to verify that there is enough empty room such that the orange piece can be placed in the gadget. Consider without loss of generality the case that we introduced the gadget because of a convex constraint of the form $f(x, y) \geq 0$ in Φ . If the blue and yellow pieces touch the orange piece, it corresponds to the case $f(x, y) = 0$. Since we consider a solution to Φ , we have $f(x, y) \geq 0$, so there is enough room for the orange piece. ■

Fingerprinting and almost-canonical placement

The proof of Lemma 2.9 for the swing follows exactly as for the swap. We first fingerprint the orange and then the yellow piece.

Aligned placement

PROOF OF LEMMA 2.10 FOR THE SWING. Consider a valid placement where the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement and the pieces of this swing have an almost-canonical placement. We need to verify that the exiting yellow piece has an aligned $i\mu$ -placement. It is clear that the yellow piece has an aligned placement because it is bounded from above and below by horizontal edges of the container C and a piece below. It is however not immediately clear that it has an $i\mu$ -placement, since we have no entering piece representing y , so we cannot use the assumption that the pieces \mathbf{p}_{i-1} have an aligned $(i-1)\mu$ -placement directly. Instead of considering the swing as an isolated gadget, we therefore take the two extra pieces of the left-split, which is to the right of the swing, into account; see Figure 43. We then get that if these pieces do not have an aligned $i\mu$ -placement, then the slack would exceed μ , as we argued for the swap. ■

Edge inequalities

The pieces of the swing induce no edges in the dependency graphs, so there are no edge inequalities to verify.

The swing works

In this paragraph we prove the part of Lemma 2.15 that the swing is responsible for, namely that in a given aligned $g\mu$ -placement, the convex swing used for the variables x and y implies that $f(\langle K_x \rangle, \langle K_y \rangle) \geq 0$ and the concave swing implies $g(\langle K_x \rangle, \langle K_y \rangle) \geq 0$.

PROOF OF LEMMA 2.15 FOR THE SWING. We prove that $f(\langle K_x \rangle, \langle K_y \rangle) \geq 0$ when we use a convex swing; the statement for the concave swing follows in an analogous way. Let p_x and p_y be the blue and yellow pieces, respectively. We note that there are paths P_x and P_y in G_x and G_y attached to and directed towards the cycles K_x and K_y , and the vertices farthest away from the cycles are the yellow and blue pieces p_x and p_y , so that we have $\langle K_x \rangle \geq \langle p_x \rangle$ and $\langle K_y \rangle \geq \langle p_y \rangle$. Since we consider an aligned $g\mu$ -placement, the orange piece separates the yellow and blue pieces, as defined for a canonical placement. Without loss of generality, we can consider a situation where the corners c_x and c_y of the blue and yellow pieces are in contact with the orange piece (otherwise, we can slide them towards the orange piece until they make such a contact). We then get from the principle behind the gadget (described in the beginning of this section) that $f(\langle p_x \rangle, \langle p_y \rangle) = 0$. It is straightforward to verify that we have derivatives $f_x(0,0) = 1$ and $f_y(0,0) = 1$. Since f_x and f_y are also continuous, we have $f_x > 0$ and $f_y > 0$ in a neighbourhood of constant size around the origin $(0,0)$. Hence, we have $f(\langle K_x \rangle, \langle K_y \rangle) \geq f(\langle p_x \rangle, \langle p_y \rangle) = 0$. ■

5.2 Gramophone

The gramophone works when only translations of the pieces are allowed. It works in fact also when rotations are allowed, but we will use the swing in that case since the gramophone requires non-polygonal features of the container or the pieces. The idea behind the gramophone is to have a special pink piece which is translated horizontally by motions of one lane and translated vertically by motions of another lane; see Figure 44 for an illustration. Therefore, the placement of the piece in a sense encodes two variables at once. The pink piece has a corner c which is bounded from above by a curve, and that induces a dependency between the two lanes, translating to an inequality of the variables represented by the lanes. The name is chosen since motion of one lane makes the pink piece trace and “read” the curve and thus act as the stylus of a gramophone.

The inequality can be changed by choosing another curve bounding the corner c . We can therefore make gramophones for inequalities $(x+1) \cdot (y+1) - 1 \geq 0$ and $-(x+1) \cdot (y+1) + 1 \geq 0$,




Figure 44. The concave and convex gramophone. Here, the gray pieces marked with k cannot move, so these can be thought of as representing a constant k . The purpose of these is to fix the horizontal placement of the orange piece such that horizontal motion of the blue pieces translates directly to vertical motion of the pink piece. In all cases the gramophone relies on translations only. For the resulting dependency between x and y to be non-linear, the curve restricting the corner c must be non-linear, requiring the boundary of the container to have curved parts. For color codes, see Figure 46.

which correspond to the regions on one side of a hyperbola. We could also have chosen a circular arc or another well-behaved curve. The curved arc is part of the boundary of the container C , so the reduction leads to an instance of the problem $\text{PACK}[\square \rightarrow \triangle, +]$, but as we will see, we can also obtain an instance of the problem $\text{PACK}[\triangle \rightarrow \square, +]$ by introducing a curved piece instead of using a curved container boundary.

In the following, we will first explain the principle in more detail. Then we describe how to turn the principle into a gadget that can be installed in our framework.

Principle

To understand the principle behind the gramophone, it helps to draw a coordinate system showing the correspondence between the position of the corner c and the variables x and y represented by the two lanes; see Figure 45. The curve restricting the corner c from above is the




Figure 45. All three graphs display the curve for the equation $(x+1) \cdot (y+1) - 1 = 0$. The drawings differ because of different y -axes. The left and middle diagrams correspond to the concave gramophone, while the right corresponds to the convex gramophone.




Figure 46. The actual concave gramophone (left) and a variant with polygonal container (right). The solid boundary bounding the gadgets from above are part of the boundary of the container, whereas the dashed segment bounding the gadgets from below is on the boundary of a piece that has been added to the construction earlier. The gray pieces marked with k cannot move, so these can be thought of as representing a constant k . For a detailed picture of the pink and green pieces, see Figure 47. Color codes: 1 yellow, 2 gray, 3 blue, 4 turquoise, 5 orange, 6 green, 7 pink, 8 yellow, 9 blue, 10 gray.

curve described by the equation $(x+1) \cdot (y+1) - 1 = 0$ in the respective coordinate system. Note that the region of possible positions of the corner c is below the curve. That region is concave in the left and middle diagram but convex in the right diagram. In order to get from concave to convex, we have flipped the y -axis, but that simply corresponds to changing the orientation of the blue pieces; see Figure 44. It then follows that gramophones can be made to encode both the concave constraint $-(x+1) \cdot (y+1) + 1 \geq 0$ and the convex constraint $(x+1) \cdot (y+1) - 1 \geq 0$.

Actual gadget

The actual design of the gramophone can be seen in Figure 46 (left) and many details differ from the simplified drawing in Figure 44. We only go through the concave gramophones. The convex gramophone differs only in the curved part and the orientation of the blue pieces. In order to incorporate the gadget into the complete construction, we need to add some swaps and a constant lane, as we will explain below.

Note first that since x and y are restricted to tiny ranges $I(x), I(y) \subseteq [-\delta, \delta]$, the curved part needed is likewise of length $O(\delta)$, i.e., extremely short. The gray, orange, blue and turquoise pieces form a swap. We refer to the details and correctness of the swap to Section 4.2. The gray pieces will have a fixed placement and can therefore be considered as pieces encoding a constant value. A lane for the constant will be started to the left of the gramophone and terminated to the right, which will be explained later.




Figure 47. Left: Detail of the pink and green pieces and their point of contact in the convex and the concave version of the polygonal gramophone. Right: Fingerprinting the pink piece. By moving the green piece up and down, the angle of the fingerprinted corner of the pink piece can be continuously altered.

Above the swap is the yellow lane and a pink piece. The pink piece has edge-edge contacts with the yellow and orange pieces. The top vertex of the pink piece is our special vertex c that is bounded from above by a curve. In the simplified description, the pink piece had a pair of parallel segments that were fixed at a vertical orientation by the yellow pieces. However, in order to fingerprint the right yellow piece, we rotate the pink piece a bit clockwise (and adjust the edges of the yellow pieces accordingly) so that the angle of the corner where the right yellow piece is fingerprinted can be chosen freely. The gadget is bounded from above by the container boundary and this part contains a tiny curved part, which corresponds to the curve $(x + 1) \cdot (y + 1) - 1 = 0$, marked red. However, because of the slanted orientation of the pink piece, the y -axis of the corner c is now likewise slanted as in Figure 45 (middle).

In the simplified construction, the angles were chosen so that horizontal motion of the blue pieces was translated to vertical motion of the orange piece with no scaling involved. In the actual gadget, since we want to fingerprint the right blue piece, we need to choose the slope of the edge-edge contact between the blue and orange pieces freely, and then a horizontal motion of the blue pieces is scaled when transformed to a vertical motion of the orange piece.

These two differences (slanted and scaled y -axis) result in a linear deformation of the red curve, as compared to the simplified situation. An example of such a deformation can be seen in Figure 45 (middle). The curve will however still be contained in a hyperbola. Moving the red curve up and down (where up and down is defined by the transformed y -axis) makes it possible to freely choose the fingerprinted angle of the pink piece.

Polygonal gramophone

We introduce special *polygonal* gramophones that work with a polygonal container C , but has curved pieces in order to show hardness of the problem $\text{PACK}[\triangleleft \rightarrow \sqsupset, +]$. The polygonal concave gramophone can be seen in Figure 46 (right). In this version, we have an additional green piece which is the only piece that is curved. If we encode the concave constraint $-(x+1) \cdot (y+1) + 1 \geq 0$,




Figure 48. Left: Installation manual for the gramophone. The gray lane consists of pieces with a fixed placement and can be interpreted as encoding a constant. This lane starts and ends just outside the gramophone. Right: The start of a constant lane consists of a single gray piece, which is fingerprinted at the left corner. The top of the gadget is the container boundary and we assume that the gadget is bounded from below by a yellow piece representing the variable x , which has been introduced to the construction earlier. The yellow piece is already known to be horizontally aligned, and it then follows that the gray is as well. The lane of the gray pieces is terminated in an analogous way.

then the green piece is convex. If we encode the constraint $(x + 1) \cdot (y + 1) - 1 \geq 0$, the green piece will not be convex. Recall from Section 3.4 that for pieces with curved segments, it is needed that the curvature is bounded and that the curved segments have a straight line segment of length 1 as prefix and suffix. In Figure 47 (left), the red curve and the blue segments together form a curved segment that satisfies the requirements. The pink piece is changed accordingly: In this version of the gramophone, the pink piece has one corner which can touch the curved part of the green piece and an extra corner on each side to fill out the space below the green piece. The left corner is fingerprinted. It is shown in Figure 47 (right) how the angle of the fingerprinted corner can be changed freely by moving up and down the green piece.

Installation of the gramophone

Figure 48 (left) shows how to install the gramophone in the complete construction. We make a new lane representing a constant which is started just to the left of the gramophone and terminated just to the right; see Figure 48 (right) on how a constant lane can start and end. In addition to this, three swaps are needed to organize the lanes in the right way.

Canonical placements and solution preservation

The variable pieces are the yellow and blue pieces. We define the canonical placements to be placements where the pieces have the edge-edge contacts as in Figure 46 and the turquoise piece is enclosed by the left gray and blue pieces and the orange piece. If the green piece is present, it should have edge-edge contacts with the container boundary as shown.

PROOF OF LEMMA 2.5 FOR THE GRAMOPHONE. Suppose that for a given solution to the ETR-INV formula Φ , there exists a canonical placement of the previously introduced pieces \mathbf{p}_{i-1} that encodes that solution. The placement of the left yellow, gray, and blue piece is then fixed. It

is now clear from the construction that the remaining pieces can be placed, since the inequality of the gramophone is satisfied by the solution to Φ . ■

Fingerprinting and almost-canonical placement

The proof of Lemma 2.9 for the gramophone follows exactly as for the swap. We fingerprint the pieces in the order shown in Figure 46 used for color codes.

Aligned placement

From the alignment line ℓ (shown in Figure 46), we get that the right yellow, right blue, and right gray pieces are correctly aligned. As for the swap, to get Lemma 2.10 for the gramophone, we can argue that the exiting blue and yellow pieces encode a value which is at most μ away from that encoded by the entering pieces. It therefore follows that the placement must be an aligned $i\mu$ -placement.

Edge inequalities

We have an edge from the left to the right yellow piece and from the left to the right blue piece. The edge inequalities (Lemma 2.12) follow similarly as for the split (yellow) and swap (blue).

Constant lane of gray pieces

As seen from the installation manual in Figure 48 (left), there are four swaps in the gray lane (including the one inside the gramophone itself), so the lane consists of 5 pieces p_1, \dots, p_5 in order from left to right. Here, p_2 and p_3 are the left and right gray piece in the gramophone itself. The pieces p_2, \dots, p_5 are introduced in swaps and in the gramophone, but p_1 is introduced in its own minuscule gadget shown in Figure 48 (right). The canonical placement of p_1 is when it has edge contacts with the boundary and the yellow piece as shown, and p_1 is fingerprinted as indicated by the dot in the figure. We now argue that these five pieces must have a fixed placement independent of the rest of the construction, so that they can indeed be considered as variable pieces encoding a constant. Let us define the placement of the first piece p_1 shown in Figure 48 to encode the value $\langle p_1 \rangle = 0$. We then know that in every placement, $0 \leq \langle p_1 \rangle$, since the piece may slide to the right, but not to the left due to the container boundary. Similarly as for the proof of the edge inequality in Lemma 2.12, we get that $0 \leq \langle p_1 \rangle \leq \dots \leq \langle p_5 \rangle \leq 0$, where the last inequality follows since p_5 is bounded from the right by the container boundary. We then have $\langle p_1 \rangle = \dots = \langle p_5 \rangle = 0$.

The gramophone works

In this paragraph we prove the part of Lemma 2.15 that the gramophone is responsible for, namely that in a given aligned $g\mu$ -placement, a concave gramophone used for the variables x and y implies that $-(\langle K_x \rangle + 1) \cdot (\langle K_y \rangle + 1) + 1 \geq 0$ while a convex gramophone implies that $(\langle K_x \rangle + 1) \cdot (\langle K_y \rangle + 1) - 1 \geq 0$.

PROOF OF LEMMA 2.15 FOR THE GRAMOPHONE. We show the statement for the concave gramophone; the proof for the convex gramophone is analogous. The yellow pieces are part of the cycle K_x . Since they encode values consistently by Lemma 2.13, the rotation of the pink piece is fixed and it has edge-edge contacts with the yellow pieces. Likewise, the blue pieces are part of K_y , so they also encode the same value of y . Since the gray pieces encode the same constant, we now get that the orange piece has edge-edge contacts to all blue and gray pieces. We can slide the pink piece down until it gets edge-edge contact to the orange piece. We now consider sliding the yellow pieces to the right until the pink piece gets in contact with the curved part of C . In this situation, we have $-(\langle K_x \rangle + 1) \cdot (\langle K_y \rangle + 1) + 1 = 0$ by construction. Since we have slid the yellow pieces to the right and thus increased the value of $\langle K_x \rangle$, we started with a placement where $-(\langle K_x \rangle + 1) \cdot (\langle K_y \rangle + 1) + 1 \geq 0$. ■

6. Square container

Recall from Lemma 2.17 that when we reduce to problems where the container is a polygon, the resulting container is 4-monotone (for the definition of 4-monotone, see Definition 2.16). Let \mathcal{I}_1 be an instance of a packing problem where the container $C := C(\mathcal{I}_1)$ is a 4-monotone polygon. We show how to make a reduction from $\text{PACK}[\mathcal{P} \rightarrow \sqsupset, \mathcal{M}]$ to $\text{PACK}[\mathcal{P} \rightarrow \square, \mathcal{M}]$. We will use the reduction for the problem $\text{PACK}[\square \rightarrow \sqsupset, \circ+]$ and obtain that $\text{PACK}[\square \rightarrow \square, \circ+] is $\exists\mathbb{R}$ -hard. Likewise, we will use the reduction for the problem $\text{PACK}[\triangle \rightarrow \sqsupset, +]$ (where we use polygonal gramophones as curvers) and get that $\text{PACK}[\triangle \rightarrow \square, +]$ is $\exists\mathbb{R}$ -hard. In conclusion, we aim at proving the following theorem.$

THEOREM 6.1. *Given an instance \mathcal{I}_1 of $\text{PACK}[\mathcal{P} \rightarrow \sqsupset, \mathcal{M}]$ where $\mathcal{P} \in \{\square, \triangle\}$ and $\mathcal{M} \in \{+, \circ+\}$, we can in polynomial time compute an instance \mathcal{I}_2 of $\text{PACK}[\mathcal{P} \rightarrow \square, \mathcal{M}]$ such that \mathcal{I}_1 is feasible if and only if \mathcal{I}_2 is feasible.*

Most of the remaining part of this section is devoted to the proof of this theorem. We show that the instance \mathcal{I}_1 can be reduced to an instance \mathcal{I}_2 where the container is a square S with corners $b_1b_2b_3b_4$. To this end, we introduce some auxiliary pieces that can essentially be placed in only one way in S . We call these new pieces the *exterior* pieces, whereas we call the pieces of \mathcal{I}_1 the *inner* pieces. The empty space left by the exterior pieces is the 4-monotone polygon C which act as the container C of \mathcal{I}_1 . We scale down the container C and the inner pieces so that C

fits in an $\varepsilon \times \varepsilon$ square in the middle of S , for a small value $\varepsilon = \Theta(1)$ to be defined shortly. Since the original container C has size $O(n^4) \times O(n)$, we scale it down by a factor of $O(n^4)$.

An example of the construction can be seen in Figure 49, where only the exterior pieces are shown. Our construction is parameterized by a number $\varepsilon > 0$, and the container S is the square $[0, 1 + \varepsilon] \times [0, 1 + \varepsilon]$. The blue pieces B_1, B_2, B_3, B_4 are independent of ε , but the green, turquoise, and orange pieces depend on ε . Only the exterior pieces are shown, but the full instance \mathcal{I}_2 does also include the inner pieces, which are intended to be packed in C . Each green piece G_i has a pair of corners c_i, d_i of interior angle α_i such that $\alpha_i \neq \alpha_j$ for $i \neq j$. In reality, these angles α_i are only slightly less than $\pi/2$, so that all the eight segments from a point c_i or d_i to the closest corner b_j are almost equally large. The angles α_i are independent of ε and the instance \mathcal{I}_1 .

Each orange piece is a trapezoid, and edges of the orange pieces form the boundary of C . At the opposite end of the pieces, they meet the green or turquoise pieces. To make sure that they are placed in the right order, they have unique angles so that fingerprinting can be used to argue about their placement. This will be explained in more detail in the end of this section.

We denote the placement of the exterior pieces shown in Figure 49 to be *canonical*. We also define the placements we get from the figure by sliding the orange and turquoise pieces towards C to be *canonical*. We aim at proving the following lemma.

LEMMA 6.2. *For a sufficiently small constant $\varepsilon > 0$, the following holds. For all valid placements of the pieces in \mathcal{I}_2 , the exterior pieces have a canonical placement.*

It follows from the lemma that \mathcal{I}_1 has a solution if and only if \mathcal{I}_2 has one, so that the problems are equivalent under polynomial time reductions.

The first step in proving Lemma 6.2 is to prove that the blue and green pieces, up to a rotation, can only be placed in the canonical way, i.e., even when we disregard the turquoise and orange pieces and the inner pieces. To prove this, we consider the situation when ε gets very small, as shown in Figure 50 (left), so that the blue and green pieces cover almost all of C . Then all the turquoise and orange pieces are skinny, and the 4-monotone polygon C is very small. When changing ε , we keep all angles constant and the blue pieces constant, and C is scaled appropriately so that it fits in the central square of size $\varepsilon \times \varepsilon$ (this square is drawn with dashed segments in Figure 49). We will eventually scale the pieces from \mathcal{I}_1 so that they fit in a square of size $\varepsilon \times \varepsilon$. Since they originally fit in a container of size $O(n^4) \times O(n)$ and ε is a constant, this scaling is polynomial and thus allowed in our reduction.

LEMMA 6.3. *If $\varepsilon > 0$ is a sufficiently small constant, the canonical placement is the only way to place the blue and green pieces into S , i.e., even when the turquoise and orange pieces and the inner pieces do not have to be placed. (When rotations are allowed, the three rotations of this packing by angles $\pi/2, \pi, 3\pi/2$ are also possible.)*




Figure 49. An example of the instance we get from the reduction using a square container for a given 4-monotone polygon C . In reality, the angles α_i are just slightly less than $\pi/2$, and all the eight segments from a point c_i or d_i to the closest corner b_j have length almost $1/2$. Color codes: B_1, \dots, B_4 are blue, G_1, \dots, G_4 are green, 1 is turquoise, 2 is orange.




Figure 50. Left: The situation from Figure 49 in the limit $\varepsilon = 0$. Right: The canonical placement of the green and blue pieces. For some sufficiently small ε , Q_m must be the shown placement. We have four alignment segments $\ell_1, \ell_2, \ell_3, \ell_4$.

PROOF. Suppose for the purpose of contradiction that for arbitrarily small $\varepsilon > 0$, there are other ways in which the eight pieces can be placed. For every $m \in \mathbb{N}$, let Q_m be such a placement for some $\varepsilon \in (0, 1/m)$. Recall that the blue pieces are independent of ε , and note that as $m \rightarrow \infty$, the shape of the green pieces converge to the pieces shown in Figure 50 (left). Note that the pieces are compact sets in the plane, and recall that the Hausdorff distance turns the set of non-empty compact sets into a compact metric space in its own right. By passing to a subsequence, we may therefore assume that for each piece p , the placement of p according to Q_m is likewise converging with respect to the Hausdorff distance.

We are going to apply the Single Fingerprint Lemma 3.2 to the corner b_1 of the square container S , and we want to prove that in the limit placement, a right corner of a blue piece coincides with b_1 . For the following consider the notation of Lemma 3.2. We define the triangle T so that $x \in b_1b_4$, $y = b_1$, and $z \in b_1b_2$, and we set $u := 0$. Choose σ so small that the eight blue and green pieces have the unique angle property. We get that a blue piece must be placed such that one of its right corners v is within distance $O(\sqrt{\mu/\sigma})$ from b_1 . Now, as $m \rightarrow \infty$, μ gets arbitrarily small, and therefore we get that v coincides with b_1 in the limit.

We likewise get that corners with right angles of the other blue segments are placed at the other corners b_2, b_3, b_4 of S .

Conceivably, the blue pieces can be placed incorrectly in two ways: (i) their cyclic order around the boundary of S can be different from B_1, B_2, B_3, B_4 and (ii) one of the right corners e_i which is supposed to be placed in the interior of S can be placed at a corner of S . Due to the difference in the angles $\alpha_1, \dots, \alpha_4$, it is clear that in each of these cases, the green pieces cannot be placed. Hence, the pieces must be placed as shown in Figure 50 (left) in the limit.

We conclude that by choosing m sufficiently large (and thus ε sufficiently small), the difference between each piece of Q_m and the packing of Figure 50 (left) can be made arbitrarily




Figure 51. Left: The red segments show the parts of the alignment segments that is occupied by the pieces B_1 and G_3 . Right: The fat red segments show the parts of the alignment segment ℓ_3 that is occupied by the piece G_3 in three different situations. The thin red segments are the parts occupied in the preceding situations. First is shown the canonical situation from Figure 50 (right). The middle situation shows that as G_3 is rotated, it occupies more. The last situation shows that when G_3 is translated, it occupies even more.

small. We now prove that if the difference is sufficiently small, the placement Q_m must be the canonical placement. Intuitively, this means that the canonical placement is “locked” in the sense that it is not possible to move the pieces just a little bit and obtain another valid placement.

In order to argue about the precise placement, we make an alignment argument using four alignment segments $\ell_1, \ell_2, \ell_3, \ell_4$ simultaneously. For each edge of S , we have an alignment segment ℓ_i parallel to and close to the edge, as shown in Figure 50 (right). For each of the eight blue and green pieces p , we measure how much p occupies of each alignment segment, and we take the sum of all these measures and show that it is strictly minimum in the canonical placement. In that placement, the segments are fully covered by the pieces, which means that there cannot be any other placement since they would occupy more of the segments than what is available.

Note that as we consider a placement that is close to the canonical placement, we get that the blue piece B_1 only intersects the segments ℓ_1 and ℓ_2 , and the other blue pieces likewise only intersect two segments each. Each green piece G_i only intersects ℓ_i . We now define the occupied parts of the alignment segments as follows, see Figure 51 for an illustration. Each green piece G_i occupies the part $G_i \cap \ell_i$ of the segment ℓ_i and nothing of the other segments. The piece blue piece B_1 occupies the part of ℓ_1 from the upper endpoint of ℓ_1 to the lowest point in $B_1 \cap \ell_1$. Similarly, B_1 occupies the part of ℓ_2 from the right endpoint to the leftmost point in $B_1 \cap \ell_2$. Each of the other pieces B_2, B_3, B_4 occupies parts of the two segments it intersects, defined in the analogous way. It follows that the parts of a segment ℓ_i occupied by two different pieces are interior-disjoint: This is trivial for the parts covered by the green pieces, but a blue piece B_j can also occupy parts that it does not cover close to the boundary of S . However, these parts of the




Figure 52. The fat red segments show the parts of the alignment segments ℓ_1 and ℓ_2 that are occupied by a piece in three different situations. The thin red segments are the parts occupied in the preceding situation.

alignment segments are close to the corner b_j , and since the constructed placement is close to canonical, the other pieces are far away, so they cannot cover anything close to the corner b_j .

Let L be the sum of lengths of the occupied parts. Since the occupied parts are interior-disjoint, we get that L is at most $4 \cdot (1 + \varepsilon)$, i.e., the sum of lengths of ℓ_1, \dots, ℓ_4 . We now show that L has a strict local minimum in the canonical placement. To this end, we observe that starting with the canonical placement, any small movement of one of the pieces makes L increase. This is seen for a green piece G_3 in Figure 51 (right). For a blue piece, consider the blue piece in Figure 52. First is shown the canonical placement. The middle figure shows a situation where the piece is rotated slightly while keeping it as far up and to the right as possible. We see that it occupies more of both alignment segments. The last situation shows that when the piece is translated to the left, the part occupied of ℓ_2 increases as much as the translation, while the part occupied of ℓ_1 slightly decreases (*slightly* since the piece has been only slightly rotated and the angles α_i are close to $\pi/2$). In total, the piece occupies more of ℓ_1 and ℓ_2 . The occupied amount likewise increases when the piece is translated down. It then follows that moving the piece from the canonical placement by any small movement (rotation and translation combined) will strictly increase the occupied amount of the alignment segments.

To sum up, choosing m sufficiently big, the packing Q_m can be made arbitrarily close to the canonical packing, but this implies that it must be identical to the packing, as it would otherwise occupy more of the alignment segments ℓ_1, \dots, ℓ_4 than what is possible. Hence, when ε is small enough, the eight blue and green pieces can only be placed in the canonical way. ■

To finish the proof of Theorem 6.1, it remains to argue about the fixed placement of the orange and turquoise pieces. The use of fingerprinting and aligning used here is similar to how we argued about the gadgets. We consider the piece G_3 as an example, and refer to Figure 53 for an illustration of the process. We fix the pieces in layers, so that there are a constant number of pieces in each layer. We first fingerprint the pieces touching G_3 in the canonical packing. We then align them. This may leave a bit of empty space to the right of G_3 and to the left of the orange and turquoise pieces, but that is allowed in the canonical placements, and the empty




Figure 53. Left: Fixing the alignment of the first layer of pieces. The three pieces are fingerprinted and then aligned. The empty space between G_3 and the three pieces is so small that no piece can fit in there. Right: The three layers which together fix all the orange and turquoise pieces. Note that the empty space between the exterior pieces leaves less room in C for the inner pieces (refer to Figure 49).

space will be too small for any piece to fit there. We then fingerprint the pieces touching the first turquoise piece, and repeat. This leads to a canonical placement of the exterior pieces, proving Lemma 6.2.

Recall that the (unscaled) instance \mathcal{I}_1 has $\mu = \Theta(n^{-296})$ and size $O(n^4) \times O(n)$. We scale it down to be contained in a square of size $\varepsilon \times \varepsilon$, where $\varepsilon = \Theta(1)$, so we get that the slack of our created instance \mathcal{I}_2 is $\mu = \Theta(n^{-304})$, which is polynomial. The reason that we use the turquoise pieces is that if instead all orange pieces were adjacent to the green pieces, we would need to fingerprint a superconstant number of pieces at once, and then we would need μ to be smaller than polynomial, as mentioned in Section 2.

As mentioned above, the empty space to the left of the orange and turquoise pieces sticking out from G_3 is too small that any pieces can fit there. We can therefore without loss of generality slide the pieces to the left so that they create all the edge-edge contacts shown in Figure 49. Similarly, those sticking out from G_4 can be pushed down, etc. After pushing all the orange and turquoise pieces towards their green piece, the empty area left for the inner pieces is exactly the container C . We conclude that there is a solution to the created instance \mathcal{I}_2 with a square container if and only if there is a solution to the given instance \mathcal{I}_1 , and Theorem 6.1 has been proven. We then have the following corollary.

COROLLARY 2.18. *The problems $\text{PACK}[\square \rightarrow \square, \circlearrowleft +]$ and $\text{PACK}[\triangleleft \rightarrow \square, +]$ are $\exists\mathbb{R}$ -hard.*

References

- [1] Zachary Abel, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Jayson Lynch, and Tao B. Schardl. Who needs crossings? Hardness of plane graph rigidity. *32nd International Symposium on Computational Geometry (SoCG 2016)*, 3:1–3:15, 2016. DOI (3)
- [2] Mikkel Abrahamsen. Covering polygons is even harder. *62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2021)*, pages 375–386, 2021. DOI (3)
- [3] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$ -complete. *Journal of the ACM*, 69(1):4:1–4:70, 2021. DOI (3, 5)
- [4] Mikkel Abrahamsen, Linda Kleist, and Tillmann Miltzow. Geometric embeddability of complexes is $\exists\mathbb{R}$ -complete, *39th International Symposium on Computational Geometry (SoCG 2023)*, 1:1–1:19, 2023. DOI (4)
- [5] Mikkel Abrahamsen, Linda Kleist, and Tillmann Miltzow. Training neural networks is $\exists\mathbb{R}$ -complete. *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, 2021. URL (4)
- [6] Mikkel Abrahamsen, Tillmann Miltzow, and Nadja Seiferth. Framework for $\exists\mathbb{R}$ -completeness of two-dimensional packing problems. *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS 2020)*, pages 1014–1021, 2020. DOI (1, 5, 58)
- [7] Sarah R. Allen and John Iacono. Packing identical simple polygons is NP-hard, 2012. Preprint. DOI (7)
- [8] Helmut Alt. Computational aspects of packing problems. *Bull. EATCS*, 118, 2016. URL (7)
- [9] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. Algorithms in real algebraic geometry. Springer-Verlag Berlin Heidelberg, 2006. (8)
- [10] Julia Bennell and Jose Oliveira. A tutorial in irregular shape packing problems. *Journal of the Operational Research Society*, 60:S93–S105, 2009. DOI (8)
- [11] Julia Bennell and Jose Oliveira. The geometry of nesting problems: a tutorial. *European Journal of Operational Research*, 184(2):397–415, 2008. DOI (8)
- [12] Francine Berman, David S. Johnson, Frank Thomson Leighton, Peter W. Shor, and Larry Snyder. Generalized planar matching. *J. Algorithms*, 11(2):153–184, 1990. DOI (7)
- [13] Marie Louisa Tølbøll Berthelsen and Kristoffer Arnsfelt Hansen. On the computational complexity of decision problems about multi-player nash equilibria. *Theory Comput. Syst.* 66(3):519–545, 2022. DOI (3)
- [14] Daniel Bertschinger, Christoph Hertrich, Paul Jungeblut, Tillmann Miltzow, and Simon Weber. Training fully connected neural networks is $\exists\mathbb{R}$ -complete. *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, 2023. URL (4)
- [15] Therese Biedl. Segment representations with small resolution. *Information Processing Letters*, 153:105868, 2020. DOI (3)
- [16] Daniel Bienstock. Some provably hard crossing number problems. *Discrete & Computational Geometry*, 6(3):443–459, 1991. DOI (3)
- [17] John Canny. Some algebraic and geometric computations in PSPACE. *20th Annual ACM Symposium on Theory of Computing (STOC 1988)*, pages 460–467, 1988. DOI (3)
- [18] Jean Cardinal. Computational geometry column 62. *SIGACT News*, 46(4):69–78, 2015. DOI (4)
- [19] Jean Cardinal, Stefan Felsner, Tillmann Miltzow, Casey Tompkins, and Birgit Vogtenhuber. Intersection graphs of rays and grounded segments. *J. Graph Algorithms Appl.* 22(2):273–295, 2018. DOI (3)
- [20] Jean Cardinal and Udo Hoffmann. Recognition and complexity of point visibility graphs. *Discrete & Computational Geometry*, 57(1):164–178, 2017. DOI (3)
- [21] Henrik I. Christensen, Arindam Khan, Sebastian Pokutta, and Prasad Tetali. Approximation and online algorithms for multidimensional bin packing: a survey. *Computer Science Review*, 24:63–79, 2017. DOI (7)
- [22] Fan Chung and Ron Graham. Efficient packings of unit squares in a large square. *Discret. Comput. Geom.* 64(3):690–699, 2020. DOI (7)
- [23] Karen L. Daniels and Victor Milenkovic. Multiple translational containment. part I: an approximate algorithm. *Algorithmica*, 19(1/2):148–182, 1997. DOI (6)
- [24] Karen L. Daniels and Victor Milenkovic. Translational polygon containment and minimal enclosure using mathematical programming. *International Transactions in Operational Research*, 6(5):525–554, 1999. DOI (6)
- [25] Erik D. Demaine, Sándor P. Fekete, and Robert J. Lang. Circle packing for origami design is hard, 2010. Preprint. DOI (7, 8)
- [26] Michael Gene Dobbins, Andreas Holmsen, and Tillmann Miltzow. A universality theorem for nested polytopes, 2019. Preprint. DOI (3)
- [27] Michael Gene Dobbins, Linda Kleist, Tillmann Miltzow, and Pawel Rżazewski. Completeness for the complexity class $\forall\exists\mathbb{R}$ and area-universality. *Discret. Comput. Geom.* 70(1):154–188, 2023. DOI (3)
- [28] Harald Dyckhoff and Ute Finke. Cutting and packing in production and distribution: A typology and bibliography. 1992. DOI (8)

- [29] Paul Erdős and Ron Graham. On packing squares with equal squares. *Journal of Combinatorial Theory, Series A*, 19(1):119–123, 1975. DOI (7)
- [30] Jeff Erickson. Optimal curve straightening is $\exists\mathbb{R}$ -complete, 2019. Preprint. DOI (3)
- [31] Jeff Erickson, Ivor van der Hoog, and Tillmann Miltzow. Smoothing the gap between NP and ER. *61st IEEE Annual Symposium on Foundations of Computer Science (FOCS 2020)*, pages 1022–1033, 2020. DOI (4)
- [32] Lian Fang. A rational quartic bézier representation for conics. *Computer Aided Geometric Design*, 19(5):297–312, 2002. DOI (2)
- [33] Robert J. Fowler, Michael S. Paterson, and Steven L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Information processing letters*, 12(3):133–137, 1981. DOI (7)
- [34] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. $\exists\mathbb{R}$ -completeness for decision versions of multi-player (symmetric) nash equilibria. *ACM Trans. Economics and Comput.* 6(1):1:1–1:23, 2018. DOI (3)
- [35] Thierry Gensane and Philippe Ryckelynck. Improved dense packings of congruent squares in a square. *Discret. Comput. Geom.* 34(1):97–109, 2005. DOI (7)
- [36] Thomas C. Hales. A proof of the Kepler conjecture. *Annals of mathematics*, (3):1065–1185, 2005. DOI (7)
- [37] Eva Hopper and Brian Turton. A review of the application of meta-heuristic algorithms to 2D strip packing problems. *Artificial Intelligence Review*, 16(4):257–300, 2001. DOI (8)
- [38] Ross J. Kang and Tobias Müller. Sphere and dot product representations of graphs. *Discret. Comput. Geom.* 47(3):548–568, 2012. DOI (3)
- [39] Heuna Kim and Tillmann Miltzow. Packing segments in a simple polygon is APX-hard. *European Conference on Computational Geometry (EuroCG 2015)*, pages 24–27, 2015. URL (7)
- [40] Linda Kleist. Planar graphs and faces areas – Area-Universality. PhD thesis, 2018. PhD thesis, Technische Universität Berlin. DOI (3)
- [41] Aline A.S. Leao, Franklina M.B. Toledo, José Fernando Oliveira, Maria Antónia Carravilla, and Ramón Alvarez-Valdés. Irregular packing problems: a review of mathematical models. *European Journal of Operational Research*, 282(3):803–822, 2020. DOI (8)
- [42] Joseph Leung, Tommy Tam, Chin Wong, Gilbert Young, and Francis Chin. Packing squares into a square. *Journal of Parallel and Distributed Computing*, 10(3):271–275, 1990. DOI (7)
- [43] Anna Lubiw, Tillmann Miltzow, and Debajyoti Mondal. The complexity of drawing a graph in a polygonal region. *J. Graph Algorithms Appl.* 26(4):421–446, 2022. DOI (3)
- [44] Jiří Matoušek. Intersection graphs of segments and $\exists\mathbb{R}$, 2014. Survey. DOI (4)
- [45] Colin McDiarmid and Tobias Müller. Integer realizations of disk and segment graphs. *Journal of Combinatorial Theory, Series B*, 103(1):114–143, 2013. DOI (3)
- [46] Arturo I. Merino and Andreas Wiese. On the two-dimensional knapsack problem for convex polygons. *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, 84:1–84:16, 2020. DOI (7)
- [47] Victor Milenkovic. Multiple translational containment. part II: exact algorithms. *Algorithmica*, 19(1/2):183–218, 1997. DOI (6)
- [48] Victor Milenkovic. Rotational polygon containment and minimum enclosure using only robust 2D constructions. *Comput. Geom.* 13(1):3–19, 1999. DOI (7)
- [49] Victor Milenkovic. Translational polygon containment and minimal enclosure using linear programming based restriction. *28th Annual ACM Symposium on the Theory of Computing (STOC 1996)*, pages 109–118, 1996. DOI (6)
- [50] Tillmann Miltzow and Reinier F. Schmiermann. On classifying continuous constraint satisfaction problems. *TheoretCS*, 2024. DOI. To appear. (5, 6, 8, 9, 60)
- [51] Nicolai E Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. *Topology and geometry – Rohlin seminar*, pages 527–543, 1988. DOI (3)
- [52] Jürgen Richter-Gebert and Günter M. Ziegler. Realization spaces of 4-polytopes are universal. *Bulletin of the American Mathematical Society*, 32(4):403–412, 1995. DOI (3)
- [53] Marcus Schaefer. Complexity of some geometric and topological problems. *17th International Symposium on Graph Drawing (GD 2009)*, pages 334–344, 2009. DOI (4)
- [54] Marcus Schaefer. Realizability of graphs and linkages. János Pach, editor, *Thirty Essays on Geometric Graph Theory*, chapter 23, pages 461–482. 2013. DOI (3)
- [55] Marcus Schaefer and Daniel Štefankovič. Fixed points, Nash equilibria, and the existential theory of the reals. *Theory of Computing Systems*, 60(2):172–193, 2017. DOI (3)
- [56] Yaroslav Shitov. A universality theorem for nonnegative matrix factorizations, 2016. Preprint. DOI (3)
- [57] Yaroslav Shitov. The complexity of positive semidefinite matrix factorization. *SIAM Journal on Optimization*, 27(3):1898–1909, 2017. DOI (3)
- [58] Peter W. Shor. Stretchability of pseudolines is NP-hard. Peter Gritzmann and Bernd Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*. Volume 4, DIMACS – Series in Discrete Mathematics and Theoretical Computer Science, pages 531–554. 1991. DOI (3)

- [59] Paul Sweeney and Elizabeth Ridenour Paternoster. Cutting and packing problems: a categorized, application-orientated research bibliography. *Journal of the Operational Research Society*, 43(7):691–706, 1992. DOI (8)
- [60] A. Thue. Über die dichteste Zusammenstellung von kongruenten Kreisen in einer Ebene. *Skifter udgivne af Videnskabselskabet i Christiania. I. Matematisk-naturvidenskabelig klasse*, (1), 1910. URL (8)
- [61] László Fejes Tóth. Über die dichteste Kugellagerung. *Mathematische Zeitschrift*, 48:676–684, 1942. URL (8)
- [62] Maryna S. Viazovska. The sphere packing problem in dimension 8. *Annals of Mathematics*, 185(3):991–1015, 2017. DOI (8)