

Corner Cuts are Close to Optimal

From Solid Grids to Polygons and Back

Report**Author(s):**

Feldmann, Andreas Emil; Das, Shantanu; Widmayer, Peter

Publication date:

2012-06

Permanent link:

<https://doi.org/10.3929/ethz-a-009937811>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Technical Report / ETH Zurich, Department of Computer Science 731(b)

Corner Cuts are Close to Optimal: From Solid Grids to Polygons and Back[☆]

Andreas Emil Feldmann*

*University of Waterloo, Combinatorics and Optimization Department,
200 University Avenue West, Waterloo, ON N2L 3G1, Canada,
Phone: +1 (519) 888 4567, Fax: +1 519 725 5441*

Shantanu Das

*Technion - Israel Institute of Technology, Faculty of Industrial Engineering and Management,
Technion City, Haifa 32000, Israel,
Phone: +972 4 829 2856, Fax: +972 4 829 5688*

Peter Widmayer

*ETH Zurich, Institut für Theoretische Informatik,
CAB H 39.2, Universitätstrasse 6, 8092 Zürich, Switzerland,
Phone: +41 44 632 74 00, Fax: +41 44 632 13 99*

Abstract

We study the solution quality for min-cut problems on graphs when restricting the shapes of the allowed cuts. In particular we are interested in min-cut problems with additional size constraints on the parts being cut out from the graph. Such problems include the *bisection* problem, *separator* problem, or the *sparsest cut* problem. We therefore aim at cutting out a given number m of vertices from a graph using as few edges as possible. We consider this problem on *solid grid graphs*, which are finite, connected subgraphs of the infinite two-dimensional grid that do not have holes. Our interest is in the tradeoff between the simplicity of the cut shapes and their solution quality: we study *corner cuts* in which each cut has at most one right-angled bend. We prove that optimum corner cuts get us arbitrarily close to a cut-out part of size m , and that this limitation makes us lose only a constant factor in the quality of the solution. We obtain our result by a thorough study of cuts in polygons and the effect of limiting these to corner cuts.

Keywords: min-cut, m -cut, bisection, corner cuts, grid graphs, polygons

[☆]This work was supported through the Swiss National Science Foundation under Grant No. 200021_125201/1. An extended abstract [1] of this report appeared at the *37th International Workshop on Graph-Theoretic Concepts in Computer Science*.

*Corresponding author

Email addresses: andreas.feldmann@uwaterloo.ca (Andreas Emil Feldmann),
shantanu.das@acm.org (Shantanu Das), widmayer@inf.ethz.ch (Peter Widmayer)

1. Comparing Optimal with Simple Shaped Cuts

Many problems consider cutting a graph into two parts with an additional constraint on the sizes of the resulting parts. We study the minimisation version of these types of problems in which as few edges as possible are to be used to cut the graph. More formally, we wish to partition the n vertices of a graph into two parts of sizes m and $n - m$ while minimising the *cut length*, i.e. the number of edges connecting vertices from different parts. The constraint on the sizes of the parts is realised by giving a bound on m that needs to be fulfilled. Alternatively the optimisation function may also depend on m . Some examples of these types of problems include the *bisection* problem in which $\lfloor n/2 \rfloor \leq m \leq \lceil n/2 \rceil$, the *edge separator* problem in which $n/3 \leq m \leq 2n/3$, or the *sparsest cut* problem in which the function $\frac{C}{m(n-m)}$ is to be optimised. In the latter C denotes the cut length of the solution. Note that in any of these problems the respective optimal solution cuts out m vertices using a minimum number of edges for this particular value of m . Our interest is in *solid grid graphs*: a *grid graph* is a finite, connected subgraph of the infinite two-dimensional grid. An interior face surrounded by more than four edges is called a *hole*. If a grid graph does not contain holes it is *solid*. Solid grid graphs appear in finite element simulations [2] for which also the considered types of problems are relevant in devising data distribution algorithms for parallel computations [2, 3]. The graphs and problems are also relevant to VLSI circuit design [4].

We aim at understanding the intricacies of optimally cutting out a fixed number m of vertices in a graph from a novel point of view: we study simple cut shapes, for which it can be shown that they compare well to optimal unrestricted ones. In related problems on polygons similar ideas have led to interesting insights. For instance *guillotine cuts* have been considered, which are orthogonal straight line cuts. When an orthogonal polygon is to be partitioned into rectangles these cuts lead to good approximations [5]. Also if a rectangular polygon is to be partitioned into rectangles fulfilling given size constraints, good solutions can be achieved by using guillotine cuts [6].

For the graphs considered in this article it will be convenient to determine the cut shapes in relation to their embedding in the plane. We therefore assume that a grid graph is given together with its natural embedding. That is, the grid is a plane graph in which the vertices are coordinates in \mathbb{N}^2 and all edges have unit length. It is well-known that any cut in a planar graph $G = (V, E)$ corresponds to a set of cycles in its dual graph (i.e. the (multi-)graph whose vertices are the faces of G and whose edges represent a shared boundary between faces). We call a set of edges in a planar graph a *segment* if it corresponds to a simple cycle in the dual graph. Hence a *cut* can be defined as a set $S \subseteq 2^E$ of segments (Figure 1). In this view the segments are the building blocks of a cut. In our setting a guillotine cut in a solid grid graph corresponds to a set of straight lined segments without bends. However these kinds of cuts do not yield satisfactory solutions since they can be far away from optimum (Figure 2). On




Figure 1: An optimal (left) and a corner cut (right) in a solid grid, each cutting out $m = 110$ vertices. The bold edges indicate the elements of the segments.

the other hand, it was observed [7] that there always exists an optimum solution to cut out m vertices in which almost all segments have a simple shape. More precisely, at most one segment in an optimal solution has more than one bend.

The above two observations on the shapes of segments and their relation to optimal solutions naturally lead to the question of how well so-called *corner cuts* perform. These contain only segments having at most one bend (Figure 1). In this article we prove that optimum corner cuts do not need a lot more cut-edges than arbitrary cuts. We achieve our result by proving a number of theorems for polygons that we relate to the case of solid grid graphs. The reason for choosing this approach is that polygons are continuous objects, which is in contrast to the discrete nature of graphs. This fact makes certain tools available for our proofs that otherwise would not be applicable. The main part of this paper will therefore be concerned with thoroughly analysing corner cuts in polygons.

We measure the quality of a cut S using the *cut length* which is the number of edges $\sum_{s \in S} |s|$ in S . Notice that some edges may be counted several times in this sum. However (if the cut shapes are not restricted) edges that appear more than once can be removed. In particular the number of cut out vertices stays the same while the cut length decreases when eliminating redundant edges. Hence this generalisation of measuring the quality of a cut does not change the solution which minimises the cut length among those cuts cutting out m vertices. More formally, consider the set of connected components left after removing the edges belonging to a cut.

If there exists a subset of these components whose sizes add up to exactly m , then we call the cut an *m-cut*. An *m-cut* minimising the cut length among all *m-cuts* is *optimal*. We distinguish between the set containing m vertices, the *A-part*, from the other set, the *B-part*, of the *m-cut*. As mentioned before, we propose to use only segments that correspond to orthogonal curves in the dual graph with at most one right-angled bend, when disregarding the part of the cycle that connects to the exterior face (Figure 1). If the corresponding curve of




Figure 2: A grid in which the optimum cut length for $m = n/2$ is constant (bold edges) but any cut containing only straight segments has cut length $\Omega(\sqrt{n})$.




Figure 3: Converting a grid to a polygon.

a segment contains no bend we call it a *straight segment* and if it contains exactly one right-angled bend a *corner segment*. An m -cut containing only straight and corner segments is called a *corner m -cut*. The main result of this article is summarised in the following theorem.

Theorem 1. *Let l be the cut length of an optimal m -cut in G and $\varepsilon \in]0, 1]$. Then there exists a corner m' -cut, for some $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$, which has a cut length that is at most a factor of $\mathcal{O}(1/\sqrt{\varepsilon})$ larger than l .*

This theorem was used [2] in order to compute constant approximations to sparsest cuts in solid grid graphs in linear time, using a method developed by Leighton and Rao [8]. Subsequently these approximate sparsest cuts can be put to work [2] via known techniques [3, 8] in order to speed up the computation of approximate separators and bisections.

We will prove Theorem 1 by going through several steps, each of which is an interesting problem in itself. We start by comparing cuts in grid graphs to cuts in polygons in order to be able to use the continuous nature of the polygons in our proofs. For this we convert a given solid grid graph into a simple orthogonal polygon, and hence all polygons considered in this article are orthogonal and simple. We also assume that any polygon has a fixed orientation in the plane towards which the used cuts are oriented, as the grids are given together with their embedding. In this article, we define a polygon using its interior point set.

Definition 2 (polygon). A *polygon*¹ $\mathcal{P} \subset \mathbb{R}^2$ is an open bounded set of points in the plane. Let β be the boundary of \mathcal{P} . If β only contains orthogonal line segments, we refer to \mathcal{P} as *orthogonal*. We call \mathcal{P} *simple* if any closed curve in \mathcal{P} can be shrunk to a point without leaving \mathcal{P} .

Given a solid grid graph G , the conversion is done by replacing each vertex $(x, y) \in V$ by a unit square that has its centre at the coordinate $(x, y) \in \mathbb{N}$ (Figure 3). Notice that the squares of two neighbouring vertices of V will share a boundary, but the converse is not necessarily true. Ignoring those boundaries that correspond to an edge in G leaves a connected curve that is the boundary of the polygon. It may happen that this boundary is degenerate in the sense that it can have overlapping edges (Figure 3). The region enclosed by the boundary

¹We use calligraphic capital letters to denote areas in the plane such as polygons or areas inside polygons, and we use lower case Greek letters to denote curves such as boundaries or segment curves.




Figure 4: A straight, corner, staircase, and rectangular line in a polygon denoted by λ_1 through λ_4 respectively.

is the polygon \mathcal{P}_G and has area exactly n , equal to the number of vertices in G . We will refer to the area covered by a polygon, or any open set of points, as its *size*.

All the notions used for cuts in grids carry over naturally to the case of polygons. Intuitively, the building blocks of a cut in a polygon \mathcal{P} are curves that can be drawn between points on the boundary of \mathcal{P} (Figure 4). In accordance with the grid case, we call them *segment curves* and a *cut* is a set of segment curves. Formally these curves are defined as follows.

Definition 3 (curve, boundary point, segment curve). Given a polygon \mathcal{P} , a *curve* $\lambda \subset \mathcal{P}$ is the image of a continuous map from the open unit interval to \mathcal{P} . The *length* of a curve is measured using the l_1 -norm. Unless otherwise stated, all considered curves have finite length. If β denotes the boundary of \mathcal{P} , we call a point $p \in \beta$ a *boundary point* of a curve λ in \mathcal{P} if the distance from p to λ is 0. If λ has two boundary points we call it a *segment curve*.

Note that a segment curve has exactly two boundary points since a polygon is an open set of points. Consider the connected areas left after removing the segment curves from a polygon. An m -cut is a set L of such curves that leaves a subset of these areas with total size m . The *cut length* of L is the sum of the lengths of the curves in L , which are measured using the Manhattan distance. This ensures that an m -cut in a grid graph G has a corresponding m -cut in the polygon \mathcal{P}_G with the same cut length. The curves in the latter cut reside on the boundaries of the unit squares used to construct \mathcal{P}_G . Note that the m -cut in \mathcal{P}_G that corresponds to the optimal m -cut in G obviously has a cut length that is at least the cut length of the *optimal* m -cut in \mathcal{P}_G . The latter is defined as an m -cut having the smallest cut length among all m -cuts. Those segment curves that we will use to cut out areas from polygons are rectilinear and we therefore call them *lines*. A *corner m -cut* in a polygon is an m -cut containing only *straight* and *corner lines*. Analogous to the case of grids, the former are orthogonal segment curves without bends, and the latter are orthogonal and have exactly one right-angled bend, as seen in the following definition.

Definition 4 (bar, straight, corner line). We call a curve λ a *bar line* if all points in λ share either the same x - or the same y -coordinate. In the former case we say that the orientation of the bar line is *vertical* and it is *horizontal* in the latter case. A bar line that is also a segment curve is called a *straight line*,

and a segment curve that consists of a horizontal and a vertical bar line is called a *corner line*. We refer to these bar lines as the horizontal, respectively vertical, *bar line of the straight or corner line*. Analogous to the corner segments, we call the point at which the two bar lines of a corner line meet its *corner*, and say that it *points* in two of the directions *up*, *down*, *left*, and *right*, depending on whether its horizontal and vertical bar lines go up, down, left, or right from its corner, respectively.

We will first show the existence of corner cuts in simple polygons that cut out almost the required area and have small cut length (close to optimal). We will then convert such a cut in a polygon \mathcal{P}_G derived from a grid graph G to a corresponding cut in G having the properties described in Theorem 1. More precisely we prove the following results for polygons which together imply the theorem:

1. We show that there is an optimal m -cut in a polygon that is almost a corner cut, in the sense that the cut consists of only straight and corner lines except at most one other curve (Section 2). This curve may be shaped like a staircase (a so-called *staircase line*; cf. Definition 5), or it may be a *rectangular line* (cf. Definition 6), which is a contiguous part of the boundary of an orthogonal rectangle (Figure 4). We call a cut a *1-rectangular*, respectively a *1-staircase* cut if it contains one rectangular, respectively staircase line, and only straight and corner lines otherwise.
2. We show how to remove a rectangular line from an optimal 1-rectangular m -cut (Section 3). We replace the rectangular line by a set of straight and corner lines, and at most one staircase line. Together these cut out the same area as the rectangular line. While doing this we need to take other curves from the cut into consideration so that the newly introduced curves do not interfere with these. The new cut will also be a m -cut but its cut length may not be optimal. However, we show that the cut length of the new cut is only a constant factor away from the optimal.
3. Given a (not necessarily optimal) 1-staircase m -cut of the polygon we next show how to replace the staircase line with a set of corner and straight lines, such that the new area that is cut out is close to m (Section 4). To be more precise, the new cut is an m' -cut where $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$ for any desired constant $\varepsilon \in]0, 1]$. Furthermore, the cut length of the new cut is only a constant factor (depending on ε) times the cut length of the original cut.
4. At last we show how to convert a cut containing only straight and corner lines in a polygon \mathcal{P}_G corresponding to a grid graph G into a cut in G (Section 5). Note that this step would be straightforward if all the curves in the cut were passing through exactly the midpoints of the edges of the grid. We call such curves *grid lines*. We show that all curves in the cut obtained in the previous steps can be moved to grid lines in such a way that the cut length remains the same, but we lose a small area a from the cut out area. Since a is small we can cut this area from the polygon using

a recursive method using only grid lines so that the cut length grows by only a small factor.

The following sections explain these techniques in more detail.

2. Cuts in Polygons

We will now show that for any polygon there is an optimal m -cut for which all but at most one curve are corner and straight lines. Curves with more bends include *staircase lines* and *rectangular lines*. The former have at least two bends and are monotone in x - and in y -direction. The latter have two or three bends and form part of the boundary of an orthogonal rectangle (Figure 4).

Definition 5 (staircase line). For any polygon \mathcal{P} a *staircase line* $\lambda \subseteq \mathcal{P}$ is a segment curve that consists of a sequence of bar lines such that of two adjacent bar lines one is horizontal and the other vertical. This sequence has length at least three and the resulting curve is monotonic in x - and y -direction. The orientation of λ is *up* if its left boundary point also is lower than the other, and *down* otherwise.

Definition 6 (rectangular line). Let $\mathcal{R} \subset \mathbb{R}^2$ be an axis-parallel rectangle in the plane and let γ be its boundary. Any segment curve $\lambda \subseteq \gamma \cap \mathcal{P}$ which contains either two or three corners of \mathcal{R} is called a *rectangular line*. These corners are called the *corners of λ* . We call \mathcal{R} the *defining rectangle of λ* if \mathcal{R} is the rectangle of smallest size among those from which λ can be constructed in this way.

Notice that a rectangular line contains either three or four bar lines between its corners and boundary points since \mathcal{P} is an open set. Notice also that a corner line and a staircase line that have the same boundary points have the same length.

In a first step, we convince ourselves that in any simple polygon there is an optimal m -cut that contains only straight, corner, staircase, and rectangular lines. Furthermore, none of these lines *cross* or *overlap*, which is defined as follows.

Definition 7 (\mathcal{A} - and \mathcal{B} -part, crossing, overlapping). Let $\mathcal{A}(L) \subseteq \mathcal{P} \setminus \{p \in \lambda \mid \lambda \in L\}$ be the open set of size m that is cut out by the m -cut L in \mathcal{P} and let $\mathcal{B}(L) = \mathcal{P} \setminus (\mathcal{A}(L) \cup \{p \in \lambda \mid \lambda \in L\})$ be the other cut out open set of size $n - m$. That is, the areas $\mathcal{A}(L)$ and $\mathcal{B}(L)$ do not include points that are contained in curves of L or the boundary of \mathcal{P} .

Let λ_1 and λ_2 be two segment curves in \mathcal{P} . We say that λ_1 and λ_2 *cross* if λ_2 contains points from both $\mathcal{A}(\{\lambda_1\})$ and $\mathcal{B}(\{\lambda_1\})$. A cut L is said to be *non-crossing* if no pair of curves in L cross. Two segment curves in \mathcal{P} *overlap* if they do not cross but share a curve of length greater than zero (i.e. the shared part is not just a point).

The following results are analogous to those obtained in [4] for grid graphs.




Figure 5: A rectangular line λ_1 with its defining rectangle \mathcal{R}_1 is replaced with the rectangular line λ'_1 (left). To compensate for the area a (shaded in grey), another curve λ_2 is replaced by λ'_2 (right). It can be a corner, staircase, or a rectangular line (with defining rectangle \mathcal{R}_2).

Lemma 8. *In any polygon \mathcal{P} there is an optimal m -cut L that is non-crossing and contains only straight, corner, staircase, and rectangular lines. Furthermore no curves in L overlap.*

Proof. Note that any pair of crossing segment curves can be seen as a (different) pair of segment curves that do not cross. Hence there always exists an optimal non-crossing m -cut. Additionally, removing overlapping parts of curves results in an m -cut of smaller cut length and thus no curves in L overlap. Also, as in the case of grids, it is easy to see that for any m -cut with cut length l there is an m -cut for which every curve is a segment curve and has a cut length of at most l . Thus, let λ be a curve from L and let $\mathcal{R} \subset \mathbb{R}^2$ be the smallest rectangle containing λ . Due to the well-known isoperimetric problem, using the l_1 -norm (see e.g. [9]) it follows that λ is a rectangular line if the boundary points of λ do not coincide with two of the opposing corners of \mathcal{R} . If λ 's boundary points coincide with two opposing corners of \mathcal{R} , it is easy to see that λ can be replaced with a straight, corner, or staircase line since these lines have minimum length between the boundary points using the l_1 -norm. \square

In a next step, we show that if an optimal m -cut contains a rectangular line, then all other curves are straight or corner lines. Generally speaking, the reason is that cuts can be modified so that the cut out area remains the same. This is easy to see for two rectangular lines where the \mathcal{A} -part of the cut out area is on the inside of one of the rectangles and on the outside of the other: we can simply make both rectangular lines smaller by the same area, thereby decreasing the length of the cut (Figure 5)—a contradiction to optimality. More generally, we call a corner line *convex* w.r.t. the area next to its 90 degree angle and *concave* w.r.t. the area next to its 270 degree angle (Figure 6). Similarly, a rectangular line is convex w.r.t. the area next to its 90 degree angles, and concave w.r.t. the area on its other side. Similar area exchange arguments as above show that for an optimal 1-rectangular m -cut, the area on the concave side of the rectangular line will belong to the same part of the cut as the area on the concave sides of all corner lines. This fact will become important later when a rectangular line is replaced with a staircase line.

Definition 9 (convex, concave). For any segment curve $\lambda \in L$ let $\mathcal{C} \subseteq \mathcal{P}$ be an open set of points such that λ is part of the boundary of \mathcal{C} . We define $\mathcal{Z}(\mathcal{C}) \subseteq \mathcal{C}$ as the set of points $p \in \mathcal{C}$ such that there exist a horizontal and a vertical bar line which both are contained in \mathcal{C} , and end in p and a point on λ . We call a corner or rectangular line *convex* w.r.t. \mathcal{C} if $\mathcal{Z}(\mathcal{C}) \neq \emptyset$ and *concave* w.r.t. \mathcal{C} otherwise.

Since in an optimal m -cut L no curves overlap, each set $\mathcal{A}(L)$ and $\mathcal{B}(L)$ can only lie to one side of a corner or rectangular line. Hence any such line in L is either concave or convex w.r.t. exactly one of the cut-out areas from \mathcal{P} .

For staircase lines, area exchange works by changing the staircase line while still keeping it monotonic between its end points. The potential area exchanged is the *deficit* or the *surplus*, which are areas with monotone boundaries contained in the \mathcal{B} - and \mathcal{A} -part respectively (Figure 7). These areas are used to prove that an optimal cut requires at most one staircase line: for more than one staircase line we trade the smaller deficit or surplus of one staircase with the larger of another one, turning the former into only straight and corner lines (Figure 9).

More formally, consider the simple case of an m -cut that contains only one curve λ which is a staircase line. In this case the set $\mathcal{Z}(\mathcal{C}) \cup \lambda$ is referred to as the surplus if $\mathcal{C} = \mathcal{A}(\{\lambda\})$ and as the deficit if $\mathcal{C} = \mathcal{B}(\{\lambda\})$. We are interested in the segment curves that are part of the boundary of the deficit and surplus. We need to add the points in λ to $\mathcal{Z}(\mathcal{C})$ so that the boundary of both the deficit and the surplus is made up of segment curves. (For instance the boundary of the surplus shown in Figure 7 would otherwise only contain λ as a segment curve.) If λ is the only curve in an m -cut, then the segment curves apart from λ in the surplus and deficit are all straight and corner lines by definition of $\mathcal{Z}(\mathcal{C})$. The surplus can be seen as the area that λ cuts out from the \mathcal{A} -part in addition to what these lines in the boundary of the surplus cut out. The deficit on the other hand can be seen as the area that λ does not cut out compared to the lines in the boundary of the deficit. If there are other curves apart from λ in an m -cut, then the definition has to be modified in the following way, in order to capture a similar notion. If there is a curve λ' that overlaps with λ (as shown in Figure 7), then it can happen that the intersection between $\mathcal{Z}(\mathcal{C})$ and a part cut out by λ' is non-empty. This would mean that a curve in $\mathcal{Z}(\mathcal{C}) \cup \lambda$ might cross λ' . Such a curve will later be used when transforming λ . Hence all the parts cut out by other curves that include λ are removed in the surplus and deficit.

Definition 10 (surplus, deficit). Let $\lambda \in L$ be a staircase line from a non-crossing m -cut L , $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$, and $L' = L \setminus \{\lambda\}$. For any curve $\lambda' \in L'$ let $\mathcal{D}_{\lambda'} \in \{\mathcal{A}(\{\lambda'\}), \mathcal{B}(\{\lambda'\})\}$ such that $\lambda \cap \mathcal{D}_{\lambda'} = \emptyset$. We call the set

$$(\mathcal{Z}(\mathcal{C}) \cup \lambda) \setminus \bigcup_{\lambda' \in L'} \mathcal{D}_{\lambda'}$$

the *surplus* of λ if $\mathcal{C} = \mathcal{A}(L)$ and we call it the *deficit* of λ if $\mathcal{C} = \mathcal{B}(L)$.

Using the above notions we are able to prove that if there is a rectangular line in an optimal m -cut it is the only curve that is not a corner or straight line. We proceed in two steps of which the following lemma is the first.




Figure 6: A corner, and rectangular line in a polygon denoted by λ_1, λ_2 , respectively. Both are concave with respect to the \mathcal{A} -part and convex with respect to the \mathcal{B} -part.




Figure 7: A staircase line λ together with its surplus (in light grey shading) and its deficit (in dark grey shading).

Lemma 11. *For any polygon \mathcal{P} , if an optimal m -cut L contains a rectangular line that is concave with respect to the area $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$, then it contains no staircase line and also no corner line that is convex with respect to \mathcal{C} .*

Proof. Let $\lambda_1 \in L$ be the rectangular line that w.l.o.g. is concave with respect to $\mathcal{C} = \mathcal{A}(L)$ (by Lemma 8 no curves overlap and hence any rectangular or corner line is either concave or convex w.r.t. \mathcal{C}). Let $\lambda_2 \in L$ be a staircase line. As we will show, there is a sufficiently small area of size $a > 0$ that can be locally “transferred” from λ_1 to λ_2 by making λ_1 shorter while transforming λ_2 such that the cut length is decreasing. Hence we get a contradiction to the optimality of L .

Any rectangular line has at least two adjacent corners, i.e. there is a bar line connecting them. For any rectangular line under consideration we can assume w.l.o.g. that these corners coincide with the lower right and upper right corners of its defining rectangle. Let \mathcal{R}_1 be the defining rectangle of λ_1 , let $\mathcal{Q}_1(x) = \{(x', y') \in \mathcal{R}_1 \mid x' > x\}$, and let $a_1(x)$ be the size of $\mathcal{Q}_1(x)$. For sufficiently small $a > 0$ there is a value x_a such that $a_1(x_a) = a$ and the rectangle $\mathcal{R}_1 \setminus \mathcal{Q}_1(x_a)$ defines a rectangular line λ'_1 which has the same boundary points as λ_1 and does not cross any curve in L . Observe that λ'_1 is shorter than λ_1 by twice the width of the area $\mathcal{Q}_1(x_a)$. When replacing λ_1 with λ'_1 in L we need to compensate for the area $\mathcal{Q}_1(x_a)$ in order to cut out an area of size m , by also replacing the staircase line λ_2 with some appropriate curve λ'_2 (see Figure 5). We show next how this is done.

Since λ_2 is a staircase line, for sufficiently small a we can find a staircase line λ'_2 that cuts out an area of size a from the surplus of λ_2 (remember that $\mathcal{C} = \mathcal{A}(L)$), such that λ_2 and λ'_2 have the same boundary points, and replacing λ_2 with λ'_2 will make the cut out area have size m (remember that λ_1 is concave w.r.t. \mathcal{C}). Notice that, by the definition of the surplus, λ'_2 does not cross any curve and therefore the new m -cut is non-crossing. The length of λ'_2 is equal to the length of λ_2 in the l_1 -norm and hence the cut length is decreasing when replacing λ_1 and λ_2 . This is a contradiction to the optimality of L and therefore λ_2 cannot be a staircase line.

A similar argument can be made when λ_2 is a corner line that is concave

w.r.t. \mathcal{C} . For sufficiently small a we can find a staircase line λ'_2 which has a deficit of size a in the m -cut that results from replacing λ_1 and λ_2 , and the boundary of the deficit is $\lambda_2 \cup \lambda'_2$. Also if a is small enough, λ'_2 does not cross any other curve since no curve overlaps with λ_2 by Lemma 8. Since this means that the boundary points of λ_2 and λ'_2 are the same, the length of these two lines are the same in the l_1 -norm, and therefore the cut length decreases when replacing λ_1 and λ_2 . This is a contradiction to the optimality of L and hence λ_2 cannot be a corner line that is concave w.r.t. \mathcal{C} . \square

Using the above lemma we can prove that if an optimal m -cut contains a rectangular line, then no other curve is a rectangular or staircase line, as the following lemma shows.

Lemma 12. *For any polygon \mathcal{P} , if an optimal m -cut L contains a rectangular line that is concave with respect to the area $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$, then all other curves are straight and corner lines, where the latter all are concave with respect to \mathcal{C} .*

Proof. By Lemma 8 we can assume that all curves in L are straight, corner, staircase, or rectangular lines. Additionally Lemma 11 shows that apart from the rectangular line $\lambda_1 \in L$, the only lines left that might violate the statement in this lemma are other rectangular lines. Let $\lambda_2 \in L$ be such a rectangular line. We first consider the case when λ_2 is convex w.r.t. \mathcal{C} . As in the proof of Lemma 11 let λ'_1 be the rectangular line that is defined by the rectangle $\mathcal{R}_1 \setminus \mathcal{Q}_1(x_a)$. Analogous to the definition of λ'_1 we can define a rectangular line λ'_2 such that the corresponding function $a_2(\cdot)$ equals a for an appropriate value x'_a if a is sufficiently small. The line λ'_2 is shorter than λ_2 by twice the width of the corresponding area $\mathcal{Q}_2(x'_a)$. But this means that replacing λ_1 with λ'_1 and λ_2 with λ'_2 results in an m -cut with smaller cut length than L . This contradicts the optimality of L and hence λ_2 cannot be a rectangular line that is convex w.r.t. \mathcal{C} .

Thus consider the case when λ_2 is a rectangular line that also is concave w.r.t. \mathcal{C} . For $i \in \{1, 2\}$ let h_i and w_i be the height and width of the defining rectangle \mathcal{R}_i of λ_i , respectively. Assume w.l.o.g. that $w_2 \geq h_2 \geq h_1$ (otherwise we can switch the identity of the width and height of \mathcal{R}_2 for the former, and the identity of λ_1 and λ_2 for the latter inequality). As noted before, the length l'_1 of λ'_1 is shorter than the length l_1 of λ_1 by twice the width of $\mathcal{Q}_1(x_a)$. Since the height of the latter equals the height of \mathcal{R}_1 this means that $l'_1 = l_1 - 2a/h_1$. If λ_2 has three corners, then let (x, y) be the corner that is adjacent to both the other two corners. In case λ_2 has two corners we can decompose it into three bar lines of which two are incident to exactly one corner. Let in this case (x, y) be the corner that is incident to the longer of these two bar lines. In all of these cases we can assume w.l.o.g. that (x, y) is the top right corner of \mathcal{R}_2 .

For sufficiently small a we can find a rectangular line λ'_2 that has the following properties (Figure 8). It is defined by a rectangle \mathcal{R}'_2 that has the same lower left corner as \mathcal{R}_2 and the top right corner $(x + z, y + z)$, for some $z > 0$, such that the area $(\mathcal{R}'_2 \setminus \mathcal{R}_2) \cap \mathcal{P}$ that is cut out between λ'_2 and λ_2 has size a . It




Figure 8: Constructing the line λ'_2 (dotted) when λ_2 (dashed) is a rectangular line that is concave w.r.t. \mathcal{C} . Both lines share the boundary point p since (x, y) is the upper right corner of λ_2 . For sufficiently small z the area $\mathcal{R}'_2 \setminus \mathcal{R}_2$ is entirely included in \mathcal{P} . If λ_2 has two corners, as depicted here, this is due to \mathcal{P} 's orthogonal boundary (solid black line).

also does not cross any other curve, and λ'_2 shares at least one boundary point p with λ_2 . By the assumption that (x, y) is the top right corner of \mathcal{R}_2 and the construction of \mathcal{R}'_2 , the boundary point p is the one that is incident to the lower horizontal bar lines of λ_2 and λ'_2 . Since by Lemma 8 no curves overlap in L , such that for sufficiently small a the constructed line does not cross any other curve, this means that λ'_2 always exists. Notice however that the two lines might differ in the other boundary point if λ_2 has two corners since the boundary of \mathcal{P} may overlap with the boundary of \mathcal{R}'_2 . Under the assumption that \mathcal{P} is orthogonal we can however always find some sufficiently small $z > 0$ such that the area $\mathcal{R}'_2 \setminus \mathcal{R}_2$ is entirely included in \mathcal{P} .

The area $\mathcal{R}'_2 \setminus \mathcal{R}_2$ can be decomposed into three rectangles of which one extends \mathcal{R}_2 to the right by z , one extends \mathcal{R}_2 to the top by z , and one which lies between these two extensions and has height and width z . By the assumption that $w_2 \geq h_2$ we can therefore conclude that $a = zw_2 + zh_2 + z^2 > 2zh_2$. It is easy to see that the length l'_2 of λ'_2 is at most $l_2 + 4z$. Solving the lower bound on a for z we can conclude that $l'_2 < l_2 + 4\frac{a}{2h_2} = l_2 + 2a/h_2$. Replacing λ_1 and λ_2 by λ'_1 and λ'_2 yields an m -cut that has a shorter cut length than L since we assumed that $h_2 \geq h_1$. But this contradicts the optimality of L which means that λ_2 cannot be a rectangular line that is concave w.r.t. \mathcal{C} . \square

After considering optimal m -cuts containing rectangular lines we turn to the case where they contain staircase lines. In this case we can show that there always exists an optimal m -cut in which at most one curve is a staircase line while all others are corner and straight lines.

Lemma 13. *For any polygon \mathcal{P} , if there is an optimal m -cut L that contains a staircase line, then there also is an optimal m -cut that contains at most one staircase line while all other curves are straight or corner lines.*




Figure 9: Two staircase lines λ_1 and λ_2 together with their respective surplus (or deficit) shaded in grey. The dotted lines indicate that the boundary of the surplus (or deficit) together with some parts of λ_i , $i \in \{1, 2\}$, form only corner and straight lines which are contained in L_i .

Proof. By Lemma 12 it can not happen that there is a rectangular line in L . Hence, by Lemma 8, the only case we have to consider is when there are two staircase lines λ_1 and λ_2 in L . It can happen that the boundary of the deficit of λ_2 contains parts of λ_1 , or that the boundary of the surplus of λ_1 contains parts of λ_2 . It is easy to see though that it can not happen that both boundaries contain parts of the respective other staircase line. Hence we can assume w.l.o.g. that the boundary of the surplus of λ_1 does not contain any parts of λ_2 . Let a_1 denote the size of the surplus of λ_1 . For any $a \in [0, a_1]$ we can find a set of curves $L_1(a)$ that cut out an area of size a from the surplus of λ_1 , such that removing λ_1 yields a $(m - a)$ -cut. If $a < a_1$ we can choose a single staircase line having the same boundary points as λ_1 for the set $L_1(a)$. If $a = a_1$ the curves in $L_1(a)$ are part of the boundary of λ_1 's surplus together with some parts of λ_1 (Figure 9).

If the boundary of the deficit of λ_2 contains parts of λ_1 , the deficit of λ_2 can grow when replacing λ_1 with $L_1(a)$. Hence let $d_2(a)$ denote the size of λ_2 's deficit in the constructed $(m - a)$ -cut. Similar as for the surplus of λ_1 , for a fixed a we can find a set of curves $L_2(d)$ for any $d \in [0, d_2(a)]$ cutting out an area of size d from the deficit of λ_2 . It either contains a single staircase line or curves that are part of the boundary of λ_2 's deficit. Let $b = \min\{a_1, d_2(a_1)\}$. Observe that it is possible to replace the line λ_2 with the set $L_2(b)$ after replacing λ_1 with $L_1(b)$. This yields a $(m - b + b)$ -cut, i.e. an m -cut which has a cut length that is at most the cut length of the original m -cut since distances are measured using the l_1 -norm. Assume that some curve in $L_1(b)$ or in $L_2(b)$ overlaps with some other curve in the m -cut (including those in $L_2(b)$ and $L_1(b)$ respectively). The overlapping part can be removed, again yielding an m -cut which now however has a shorter cut length. Since this is a contradiction to the optimality of L , we can conclude that no curve in $L_1(b)$ or $L_2(b)$ overlaps with any other curve. By the definition of b , at least one (both if $a_1 = d_2(a_1)$) of the sets $L_1(b)$ and $L_2(b)$ consists of curves that are part of the surplus or deficit, respectively, together with some parts of the respective staircase line. Since these curves do not overlap with any other curves they must either be straight or corner lines, by the definition of the surplus and deficit (cf. Figure 9).

Hence if L contains several staircase lines, then there is an optimal m -cut




Figure 10: A rectangular line ρ (dashed) together with the set of curves Ξ (dotted) with which it is replaced. The area of size a is shaded in grey.




Figure 11: A virtual corner line (black dashed) at (x, y) . The cut out area is shaded in grey.

which contains one staircase line less. By applying the argument repeatedly we can conclude that there is an optimal m -cut with at most one staircase line while all other curves are straight or corner lines. \square

To summarise the above results, the following observation immediately follows from Lemmas 8, 12, and 13.

Corollary 14. *For any simple polygon \mathcal{P} there is an optimal m -cut L such that all curves in L are corner or straight lines except at most one, which is either a staircase line or a rectangular line. If there is a rectangular line in L that is concave with respect to the cut out area $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$, then all corner lines in L are concave with respect to the same area \mathcal{C} .*

Because our interest is in cuts with only straight and corner lines, we need to study how to cope with a rectangular line, and how to cope with a staircase line. For a rectangular line we show how to convert the m -cut into a cut in which there is at most one staircase line while the cut length grows at most by some constant factor. With our observations on staircase lines we are then able to convert any optimal m -cut into one containing only straight and corner lines.

3. Removing Rectangular Lines

We now show how to convert an optimal 1-rectangular m -cut into an m -cut containing only straight and corner lines except at most one which is a staircase line. Consider the area inside the defining rectangle of the rectangular line (Figure 10). This region may contain a part of the boundary of the polygon (and possibly some other curves of the cut). We can replace the rectangular line with a set Ξ of straight and corner lines lying within the defining rectangle such that these curves have total length less than the length of the rectangular line. By doing this, we do not increase the length of our cut, but we now have to cut out an additional area of size a equal to the difference in sizes of the part cut out by the original cut and the part cut out by the new cut. We show how to find a set of curves that cut out the required area of size a and has total length not too large (compared to the cut length l of the optimal m -cut). Note that the length

of the rectangular line (and thus l) is at least \sqrt{a} . So, it is sufficient to show that the area of size a can be cut out using a set of curves of total length not much larger than \sqrt{a} .

In order to find this set of curves we need to abstract from the actual topology of the polygon. We achieve this by introducing the following notions (cf. Figure 11). In the proofs of this section we will restrict ourselves to the case of one specific orientation of the involved curves. Notice that in Definition 3 a segment curve can be defined for the plane by seeing \mathbb{R}^2 as the polygon with a boundary that lies infinitely far away. Hence we may define corner lines of infinite length in the plane due to Definition 4, and leverage the following definition.

Definition 15 (virtual corner line). Let μ be a corner line in the plane \mathbb{R}^2 . For any (open or closed) finite area $\mathcal{P} \subset \mathbb{R}^2$ the set Λ of corner and straight lines in \mathcal{P} for which $\lambda \in \Lambda$ if and only if $\lambda \subset \mu \cap \mathcal{P}$ is called a *virtual corner line*². The corner of μ is also referred to as the *corner of Λ* . The *length* of Λ is the sum of the lengths of the included straight and corner lines. The *horizontal length* of Λ is the total length of all horizontal bar lines covered by Λ , while the *vertical length* of Λ is the total length of all vertical bar lines covered by Λ . If Λ cuts out an area of size a on the upper right side of its corner, we say that it is a virtual corner line *for a* .

For a fixed value a let $\Lambda(x)$, if it exists, be a virtual corner line for a with corner (x, y) such that its underlying corner line in the plane points up and right. If there are several virtual corner lines that match the definition then $\Lambda(x)$ denotes the one having the largest y value for its corner. Let $l_h(x)$ be the horizontal, $l_v(x)$ the vertical, and $l(x) = l_h(x) + l_v(x)$ the total length of $\Lambda(x)$. Also let $\mathcal{P}(x) \subset \mathcal{P}$ be the cut out area of size a , i.e. $\Lambda(x)$ is the lower and left boundary of $\mathcal{P}(x)$.

Notice that if \mathcal{P} has size n , for any $a \in [0, n]$ there is a value x' such that $\Lambda(x) = \Lambda(x')$ for all $x \leq x'$ while $\Lambda(x) \neq \Lambda(x')$ for all $x > x'$. Also there is a value x'' such that $\Lambda(x)$ is defined for all $x \leq x''$ while $\Lambda(x)$ is not defined whenever $x > x''$. In this sense the points x' and x'' are extreme points for these virtual corner lines beyond which the function $\Lambda(x)$ is irrelevant for our purposes. Let $I_a = [x', x'']$ be the interval of relevant x values for the virtual corner lines for a . Note that the y values of the corners of these virtual corner lines are non-increasing with x in I_a .

The easy case is when the required area a can be cut out from the polygon using a single virtual corner line of short length (say, of length at most $c\sqrt{a}$ for some fixed constant c). However, depending on the shape of the polygon, it is not always possible to find such a virtual corner line. For example, in the polygon shown in Figure 12, any virtual corner line cutting out the required area has a long vertical or a long horizontal length.

Given any polygon we can search along the x -axis between the two extremities of the polygon, and for each value of x find a y such that the virtual corner line

²We use capital Greek letters to denote virtual lines.




Figure 12: A polygon in which every virtual corner line that cuts out an area of fixed size on the upper right side of its corner, is too long. At p_1 the vertical length switches from short to long and at p_2 the horizontal length switches from long to short.




Figure 13: The interval $[x_1, x_2]$ in a polygon and a virtual corner line (dashed black) for a whose horizontal and vertical lengths are both large. The diagonally shaded areas are $\mathcal{P}(x)$ and $\mathcal{P}(x_1)$, and the grey shaded area is \mathcal{Q}_y .

at (x, y) cuts out exactly an area of size a (Figure 13). We can show that if there does not exist any single virtual corner line for a having sufficiently small length, then there exist virtual corner lines for a at two points (x_1, y_1) and (x_2, y_2) such that the former has short (i.e. at most $c\sqrt{a}$) vertical length, the latter has short horizontal length, and for all virtual corner lines in between both lengths are large.

Lemma 16. *Let $\mathcal{P} \subset \mathbb{R}^2$ be an open set of points in the plane of size n , $a \in [0, n]$, and c be a constant. Suppose there is no virtual corner line for a with a length of at most $2c\sqrt{a}$. Then there is an interval $[x_1, x_2] \subseteq I_a$ such that*

- $l_v(x_1) \leq c\sqrt{a}$,
- $l_v(x) > c\sqrt{a}$ for all $x \in]x_1, x_2]$,

- $l_h(x_2) \leq c\sqrt{a}$, and
- $l_h(x) > c\sqrt{a}$ for all $x \in [x_1, x_2[$.

Proof. Let

$$x_2 = \inf \{x \in I_a \mid l_h(x) \leq c\sqrt{a}\} \text{ and}$$

$$x_1 = \sup \{x \in I_a \mid l_v(x) \leq c\sqrt{a} \wedge x \leq x_2\}.$$

We need to show that if the premise holds, i.e. if there is no $x \in I_a$ such that $l(x) \leq 2c\sqrt{a}$, then the interval $[x_1, x_2]$ fulfils the above listed properties. It is easy to see that $l_v(x') = 0$ and $l_h(x'') = 0$, where $I_a = [x', x'']$, and from the premise it then follows that $l_h(x') > 2c\sqrt{a}$ and $l_v(x'') > 2c\sqrt{a}$. Hence the points x_1 and x_2 must exist in I_a since the vertical length must switch from short to long and the horizontal length from long to short when traversing the interval. Assume that $l_v(x_1) > c\sqrt{a}$. Since \mathcal{P} is an open set of points there must then be some $z > 0$ such that $l_v(x) > c\sqrt{a}$ for all $x \in [x_1 - z, x_1]$. However this contradicts the definition of x_1 and we can hence conclude that $l_v(x_1) \leq c\sqrt{a}$. A similar argument can be given for $l_h(x_2)$ and thus also $l_h(x_2) \leq c\sqrt{a}$.

The premise states that $l_h(x) + l_v(x) > 2c\sqrt{a}$ for all $x \in I_a$. Thus we can conclude that $l_h(x) > c\sqrt{a}$ or $l_v(x) > c\sqrt{a}$ for any such x . By the definition of x_1 and x_2 it therefore holds that $x_1 < x_2$ and for all points $x \in]x_1, x_2[$ it holds that $l_v(x) > c\sqrt{a}$ and $l_h(x) > c\sqrt{a}$. Hence the properties on the horizontal and vertical lengths listed above are true. \square

Further we can show that the interval $[x_1, x_2]$ of the above lemma is also short.

Lemma 17. *Let \mathcal{P} be a polygon of size n , $a \in [0, n]$, and $c \geq 2$ be a constant. If (x_1, y_1) and (x_2, y_2) , where $x_1 < x_2$, are the corners of two virtual corner lines for a in \mathcal{P} such that the interval $[x_1, x_2]$ has the properties listed in Lemma 16, then*

$$x_2 - x_1 < \frac{2\sqrt{a}}{c} \text{ and } y_1 - y_2 < \frac{2\sqrt{a}}{c}.$$

Proof. Fix some $x \in]x_1, x_2]$ and let (x, y) be the corner of the virtual corner line $\Lambda(x)$ for a (see Figure 13). Let \mathcal{Q}_y be the area cut out by the virtual corner line (for some $a + b$ where $b > 0$) with corner (x_1, y) , i.e. both sets $\mathcal{P}(x_1)$ and $\mathcal{P}(x)$ are included in \mathcal{Q}_y . We can derive an upper bound on the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ by observing that this area can be split into two parts. Of these, one is contained in the area $\mathcal{P}(x_1)$ while the other is contained in the rectangle below this area. Hence we can conclude that the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ is at most $a + h(y) \cdot w(x)$, where $h(y) = y_1 - y$ and $w(x) = x - x_1$ are the height and width of the rectangle, respectively.

We can derive a lower bound on the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ by integrating along the vertical lengths of the virtual corner lines between x_1 and x , yielding

$$\lim_{z \rightarrow x_1} \int_z^x l_v(t) dt > w(x) \cdot c\sqrt{a}.$$

The area $\mathcal{Q}_y \setminus \mathcal{P}(x_1)$ is split into the part that is entirely contained in $\mathcal{P}(x)$ and the part that is contained in the rectangle to the left of that area. Note that the latter rectangle is the same as for the area $\mathcal{Q}_y \setminus \mathcal{P}(x)$. Since $\Lambda(x_1)$ and $\Lambda(x)$ both cut out an area of size a , we can conclude that the sizes of the areas $\mathcal{P}(x) \setminus \mathcal{P}(x_1)$ and $\mathcal{P}(x_1) \setminus \mathcal{P}(x)$ are equal. Therefore also the areas $\mathcal{Q}_y \setminus \mathcal{P}(x)$ and $\mathcal{Q}_y \setminus \mathcal{P}(x_1)$ have the same size. Hence we can derive a second lower bound of $h(x) \cdot c\sqrt{a}$ on the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ by integrating along the horizontal lengths of the virtual corner lines in a similar way as before. By defining $l(x) = \max\{w(x), h(x)\}$ we can thus conclude that the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ is greater than $l(x) \cdot c\sqrt{a}$. By using $l(x)$ as an estimate of $h(x)$ and $w(x)$ in the upper bound derived above, we get the following inequality:

$$a + l(x)^2 > l(x) \cdot c\sqrt{a}$$

This inequality is an invariant that is true for any $x \in]x_1, x_2]$. Using standard methods we can derive that the two terms in the invariant are equal if

$$l(x) = \frac{\sqrt{a}}{2} \left(c \pm \sqrt{c^2 - 4} \right).$$

Since $c \geq 2$ this means that the invariant amounts to one of the following terms:

$$l(x) < \frac{\sqrt{a}}{2} \left(c - \sqrt{c^2 - 4} \right) \quad \text{or} \quad (1)$$

$$l(x) > \frac{\sqrt{a}}{2} \left(c + \sqrt{c^2 - 4} \right). \quad (2)$$

Note that this means that there is an interval between these two bounds from which $l(x)$ cannot take a value. However, since the vertical and horizontal lengths of $\Lambda(x)$ are always greater than zero for $x \in]x_1, x_2[$, both $w(x)$ and $h(x)$ are continuous functions in the interval $]x_1, x_2]$. Therefore also $l(x)$ is a continuous function. Since $l(x)$ can be arbitrarily close to zero this means that Inequality (2) can never be fulfilled.

It is easily verifiable that the right-hand side of Inequality (1) can be upper-bounded by $\frac{2}{c}\sqrt{a}$ since $c \geq 2$. The proof is concluded by noticing that an upper bound on $l(x_2)$ is also one for $w(x_2)$ and $h(x_2)$. \square

Analogous to a virtual corner line we can define a *virtual staircase line* by considering any staircase line of infinite length in the plane and taking the parts of the line that lie inside some specific polygon.

Definition 18 (virtual staircase line). Let μ be a staircase line in the plane \mathbb{R}^2 of orientation down. For any finite area $\mathcal{P} \subset \mathbb{R}^2$ the set Λ of staircase, corner, and straight lines in \mathcal{P} for which $\lambda \in \Lambda$ if and only if $\lambda \subset \mu \cap \mathcal{P}$ is called a *virtual staircase line*. The *length* of Λ is the sum of the lengths of the included straight, corner, and staircase lines. If Λ cuts out an area of size a on its upper right side we say that it is a virtual corner line *for* a .

Notice that a virtual corner line is also a virtual staircase line. Using the above results we find a virtual staircase line which cuts out exactly the required

area a and has a short total length (Figure 14). The corresponding staircase line goes along the vertical section of the first virtual corner line, to some y^* and then turns to the right and goes to some x^* , turns again and then finally follows the horizontal part of the second virtual corner line.

Lemma 19. *Given a polygon \mathcal{P} of size n , for any $a \in [0, n]$ there is a virtual staircase line Λ for a that has a length of at most $7\sqrt{a}$.*

Proof. We attempt to cut out an area of size a from \mathcal{P} using a virtual corner line. Due to Lemmas 16 and 17 we can either find one with length at most $4\sqrt{a}$ (throughout this proof we set $c = 2$) or there is an interval $[x_1, x_2]$ with the properties listed in the lemmas. In the former case let Λ contain this set of lines. In the latter case we can find the desired set of curves as follows (Figure 14). We will use the same notation as in the proof of Lemma 17.

For any $x \geq x_2$ let $\Lambda'(x)$ be the virtual corner line with corner (x, y_2) and let $l'_v(x)$ be its vertical length. We attempt to find $x^* = \min\{x \geq x_2 \mid l'_v(x) \leq \sqrt{a}\}$ (which is well defined since \mathcal{P} is an open set). Notice that the vertical lines that are part of $\Lambda'(x)$ are contained in $\mathcal{P}(x_2)$, which has size a . Hence by the definition of x^* we can conclude that

$$a \geq \lim_{z \rightarrow x^*} \int_{x_2}^z l'_v(x) dx > (x^* - x_2)\sqrt{a},$$

which means that $x^* - x_2 < \sqrt{a}$.

Let $\mathcal{P}'(x^*) \subset \mathcal{P}(x_2)$ be the area that is cut out by $\Lambda'(x^*)$. We define y^* to be the coordinate where the virtual corner line with corner (x_1, y^*) cuts out an area \mathcal{Q} such that $\mathcal{Q} \cup \mathcal{P}'(x^*)$ has size a . Observe that $y^* \geq y_1$ since $\mathcal{P}(x_1)$ has size a and hence $\mathcal{Q} \subseteq \mathcal{P}(x_1)$. The desired set of curves Λ contains all curves λ




Figure 14: A virtual staircase line (black dotted) cutting out the area of size a shaded in grey. It is constructed using the two virtual corner lines at (x^*, y_2) and (x_1, y^*) .

that are segment curves and

$$\lambda \subseteq \{(x, y_2) \in \mathcal{P} \mid x \geq x^*\} \cup \quad (3)$$

$$\{(x^*, y) \in \mathcal{P} \mid y \in [y_2, y^*]\} \cup \quad (4)$$

$$\{(x, y^*) \in \mathcal{P} \mid x \in [x_1, x^*]\} \cup \quad (5)$$

$$\{(x_1, y) \in \mathcal{P} \mid y \geq y^*\}. \quad (6)$$

The points in the first set (in Row (3)) are contained in the horizontal parts of $\Lambda(x_2)$ and the points in the last set (in Row (6)) are contained in the vertical parts of $\Lambda(x_1)$, which each have a length of at most $2\sqrt{a}$ by Lemma 16. The points in the second set (in Row (4)) are contained in the vertical parts of $\Lambda'(x^*)$ which by definition has a length of at most \sqrt{a} . The length of the parts from the third set (in Row (5)) are at most the distance between x_1 and x^* . By Lemma 17 the distance between x_1 and x_2 is at most \sqrt{a} . By the observations above the distance from x_2 to x^* is also at most \sqrt{a} . In total this gives a length of at most $7\sqrt{a}$ for the curves in Λ . \square

After replacing the rectangular line we will see that we are left with a set L of straight and corner lines cutting out an area $m - a$ or $m + a$. We now know that there exists a virtual staircase line Λ that can be used to cut out the remaining area of size a from the \mathcal{A} - or \mathcal{B} -part. Notice that the underlying staircase line (of infinite length in the plane) may be intersecting with other curves in the cut (Figure 15). So the parts of the line included in Λ may not have endpoints on the boundary of the polygon. Thus, we need to convert Λ into a set M of staircase, corner, and straight lines, none of which ends at any other curve in L (however, the curves may partially overlap). This is done by adding those parts of curves in L to the curves in Λ that are monotonic extensions of the latter in x - and in y -direction. This is possible since the corner lines in L are all concave w.r.t. the same cut-out part, as shown in the previous section (Corollary 14). Thus the set M may contain several staircase lines, but its total length is at most that of L .

Lemma 20. *For any polygon \mathcal{P} , let L be a non-crossing corner m -cut with cut length l , such that all corner lines in L are concave with respect to $\mathcal{A}(L)$. For any $a \in [0, m]$ there is a set of segment curves M in \mathcal{P} that cuts out an area of size a from $\mathcal{A}(L)$ and has the following properties. The set $M \cup L$ is non-crossing, M contains only straight, corner, and staircase lines, and the cut length of M is at most $7\sqrt{a} + l$. Furthermore any staircase line in M is oriented down and its surplus, w.r.t. the $(m - a)$ -cut $M \cup L$, lies on its lower left side.*

Proof. By Lemma 19 we can find a virtual staircase line Λ in $\mathcal{A}(L)$ that cuts out an area of size a and has a length of at most $7\sqrt{a}$. The boundary points of a line $\lambda \in \Lambda$ with respect to $\mathcal{A}(L)$ are either boundary points with respect to \mathcal{P} or they are points on curves in L . If there is a $\lambda' \in L$ and a point $(x, y) \in \lambda'$ such that (x, y) is a boundary point of λ w.r.t. $\mathcal{A}(L)$, the assumption that all corner lines in L are concave w.r.t. $\mathcal{A}(L)$ lets us conclude that λ' lies on the opposite side of (x, y) than λ does. More formally, there is a relation $\preceq \in \{\leq, \geq\}$ such that for all $(x', y') \in \lambda'$ and all $(x'', y'') \in \lambda$ it either holds that $x' \preceq x$ while $x \preceq x''$ or




Figure 15: A virtual staircase line (dashed) that is converted to a set of staircase, corner, and straight lines (dashed and dotted). For this, parts of the corner lines (dotted) from the original cut are used. These are all concave w.r.t. the same part of the cut shaded in grey.

that $y' \leq y$ while $y \leq y''$. Let in the former case $\mu_{(x,y)} = \{(x', y') \in \lambda' \mid y \leq y'\}$ and in the latter case $\mu_{(x,y)} = \{(x', y') \in \lambda' \mid x \leq x'\}$. That is, if λ lies to the right or to the left of (x, y) the set $\mu_{(x,y)}$ contains the parts of λ' above or below (x, y) , respectively, and if λ lies above or below (x, y) the set $\mu_{(x,y)}$ contains the parts of λ' to the right or to the left of (x, y) , respectively.

To construct the desired a -cut M in \mathcal{P} we initially set $M = \Lambda$. If γ denotes the boundary of the area of size a that Λ cuts out, we add to M any curve in L that is contained in γ . Let P denote the set of boundary points of the curves in M w.r.t. $\mathcal{A}(L)$ which are contained in some curve from L . Since Λ , and hence initially also M , is a virtual staircase line in $\mathcal{A}(L)$, for any straight line $\lambda' \in L$ there can be at most one curve in M that has a boundary point on λ' . For any corner line $\lambda' \in L$ there can be at most two lines $\lambda_1, \lambda_2 \in M$ that have boundary points p and q on λ' . Of these points one must be on the vertical and one on the horizontal part of λ' . Hence the sets μ_p and μ_q intersect. In this case we replace the lines λ_1 and λ_2 by the line $\lambda_1 \cup \lambda_2 \cup (\mu_p \cap \mu_q)$ in M . At the same time we remove the points p and q from the set P . We repeat this process until no pair of points in P remain that both are part of some single line $\lambda' \in L$. For any remaining point in P we now know that if it is contained in some curve $\lambda' \in L$ then it is the only one. For any such point p we replace the line $\lambda \in M$, for which p is a boundary point, with the line $\lambda \cup \mu_p$ in M and remove p from P . This is repeated until no points remain in the set P .

Since the curves in L are straight lines or corner lines that are concave w.r.t. $\mathcal{A}(L)$, the added parts of the curves in L connect the curves in the original set Λ with the boundary of the polygon \mathcal{P} in such a way that in the end M contains only straight, corner, or staircase lines. Notice that the latter lines are all oriented down by the fact that Λ is oriented down and by the definition of the sets μ_p for $p \in P$. Furthermore the area $\mathcal{B}(M \cup L)$ of the $(m - a)$ -cut $M \cup L$ contains the parts of $\mathcal{A}(L)$ that were cut out by the virtual staircase line Λ on




Figure 16: A staircase line λ with its surplus shaded in grey. The curves on the boundary of the surplus can be replaced by a set of straight and corner lines (dotted). The corner line μ is also removed.

its upper right side. Hence the surplus, defined w.r.t. the $(m - a)$ -cut $M \cup L$, of a staircase line in M must lie on its lower left side. Finally each added part from a curve in L was only added once to a staircase line while constructing M . This means that the total length of the curves in M is at most $7\sqrt{a} + l$. \square

The next step is to convert the staircase lines from the set $M \cup L$ so that at most one of them remains but the cut length does not increase. Similar to the techniques seen before, we will use the curves contained in the boundary of the surplus or deficit of a staircase line for the transformation. Unfortunately some of the previous arguments can not be used here since $M \cup L$ is not an optimal cut. Instead we need some observations on the nature of the boundary of the deficit and surplus of a staircase line $\lambda \in M$: it turns out that any staircase line λ' different from λ at the boundary of the deficit or surplus of λ overlaps with exactly one corner line $\mu \in L$ (Figure 16). This corner line μ together with the staircase line λ' can be used to construct a pair of corner lines. These can be replaced with μ and λ' so that the same region is cut out by the new set of curves. The cut length decreases during this process.

Lemma 21. *For a polygon \mathcal{P} , let L be a set of non-crossing straight and corner lines and λ be a staircase line that does not cross any curve in L . Let Λ denote the set of segment curves in \mathcal{P} that are contained in the boundary of λ 's surplus (deficit), apart from λ itself, where the surplus (deficit) is defined w.r.t. the cut $L \cup \{\lambda\}$. If the set $L \cup \Lambda$ cuts out an area of size m , then there exists an m -cut that has a cut length at most that of $L \cup \Lambda$ and contains only straight and corner lines.*

Proof. We will prove the statement for the case when the curves in Λ are contained in the boundary of λ 's surplus. The other case is analogous. If Λ only contains straight and corner lines the lemma obviously holds. By the definition of the surplus the only problem that can arise is when Λ contains a staircase line λ' . Assume w.l.o.g. that λ is oriented down, which means that also λ' is. Assume furthermore that the surplus of λ lies on the lower left side of λ . If we partition λ' into a succession of bar lines that are alternating horizontally and

vertically oriented, λ' consists of at least three successive bar lines since it is a staircase line. Hence there must be a horizontal bar line σ_h that, to its right, is followed by a vertical bar line σ_v . Let (x, y) denote the point at which these two bar lines meet. This means that σ_h lies to the left of (x, y) and σ_v below (x, y) . In this sense (x, y) is a *concave corner* of the boundary of λ' 's surplus. We want to argue that there can be at most one such point and it is the corner of a corner line from L . These facts can then be used to convert λ' into a set of appropriate corner lines.

Since (x, y) is part of the boundary of λ' 's surplus, we know that for any $z > 0$ the point $(x - z, y - z)$ to the lower left of (x, y) is not part of the surplus. Since the point (x, y) is a concave corner of the surplus, for any sufficiently small z and $z_x, z_y \geq 0$ there must be two points $(x - z, y + z_y)$ and $(x + z_x, y - z)$, i.e. to the top left and the lower right of (x, y) , that are part of the surplus. (It holds that $z_x = 0$ or $z_y = 0$ if the respective point lies on λ . This may happen since λ is part of the surplus.) If z is small enough then there is no point $(x - z, y')$ or $(x', y - z)$, for any $y' \in [y - z, y + z_y]$ and $x' \in [x - z, x + z_x]$, that lies on the boundary of the polygon \mathcal{P} . Hence the only reason why $(x - z, y - z)$ is not part of the surplus can be that the point lies in $\mathcal{B}(L \cup \{\lambda\})$ and not in $\mathcal{A}(L \cup \{\lambda\})$. Letting z tend to zero it follows that (x, y) must be part of some curve $\mu \in L$ that cuts out the area to which the point $(x - z, y - z)$ belongs. Obviously μ is a corner line with corner (x, y) which includes the horizontal and vertical bar lines σ_h and σ_v .

Suppose there are more than one concave corner of λ' . Then there must be at least two of these that are adjacent in the sense that the vertical bar line σ_v^p of one of the corners p shares a point r with the horizontal bar line σ_h^q of the other corner q . By the arguments given above there must be two corner lines μ^p and μ^q in L such that $\sigma_v^p \subset \mu^p$ and $\sigma_h^q \subset \mu^q$. But since r is not part of the boundary of \mathcal{P} this means that μ^p and μ^q cross at this point, which is a contradiction. Hence there can only be one concave corner of λ' . In particular this means that σ_h is the only horizontal bar line of λ' that has an adjacent vertical bar line to its right while σ_v is the only vertical bar line that has an adjacent horizontal bar line to its left.

Consider the case when there is a horizontal bar line σ'_h to the right of the vertical bar line σ_v . As noted above, σ'_h must have a boundary point. Removing σ_v from the corner line μ that overlaps with λ' leaves the horizontal bar line of μ and some vertical bar line σ'_v that is the lower extension of σ_v in μ . Obviously σ'_v has a boundary point. Hence by removing σ_v from both μ and λ' we can construct a corner line $\sigma'_h \cup \sigma'_v$. Similarly the horizontal bar line σ_h can be removed from λ' and μ , leaving a corner line if there is a vertical bar line above σ_h in λ' . If λ' and μ share a boundary point then removing σ_v or σ_h as described above obviously leaves nothing to be taken care of. Hence any staircase line in the set Λ can, together with some corner line from L , be replaced with one or two corner lines. In a cut that includes the curves from Λ and L this transformation will not change the size of the cut out area and will decrease the cut length. \square

The above observations can now be used to convert the staircase lines

constructed in Lemma 20 in such a way that only one staircase line remains, as the next lemma shows.

Lemma 22. *In a polygon \mathcal{P} let L be a non-crossing corner m -cut with cut length l such that all corner lines in L are concave with respect to $\mathcal{A}(L)$. Then for any $a \in [0, m]$ there exists a $(m - a)$ -cut M in \mathcal{P} with cut length at most $7\sqrt{a} + 2l$ that contains only straight and corner lines except at most one which is a staircase line.*

Proof. By Lemma 20 we can find a set of curves M' such that $L \cup M'$ is a $(m - a)$ -cut that fulfils all properties of the statement except for the fact that $L \cup M'$ may contain more than one staircase line. Due to the additional properties that any staircase line is oriented down and its surplus lies on the lower left side, we can conclude that the boundaries of the surplus and the deficit can not contain any other staircase line. Hence we may use Lemma 21 to convert a staircase line into a set of straight and corner lines as follows.

We initially set $M = L \cup M'$. Let λ_1 and λ_2 be two distinct staircase lines from M' , and let b_1 be the size of the surplus of λ_1 and b_2 be the size of the deficit of λ_2 . Without loss of generality we can assume that $b_1 \leq b_2$. Similar to the proof of Lemma 13 we can find two sets of curves L_1 and L_2 that cut out an area of size b_1 from the surplus of λ_1 and the deficit of λ_2 , respectively, such that removing λ_i and adding L_i , for both $i \in \{1, 2\}$, in M' again yields a $(m - a)$ -cut M . The set L_2 can be chosen to consist of a single staircase line if $b_2 > b_1$ and it contains only curves that are part of the boundary of λ_2 's deficit if $b_2 = b_1$. The set L_1 always contains curves that are part of the boundary of λ_1 's surplus. The new $(m - a)$ -cut in which λ_1 and λ_2 were replaced has a cut length that is at most the cut length of the old m -cut since distances are measured in the l_1 -norm (it is decreasing if there are more than one curve in L_1 or L_2 since then parts of the boundary of \mathcal{P} act as a short cut for the curves).

Using Lemma 21 the staircase lines in L_1 can all be converted to corner and straight lines. If there are more than one staircase line in L_2 , i.e. L_2 is part of the boundary of λ_2 's deficit, using the same lemma all of them can be converted to straight and corner lines. Otherwise L_2 consists of only one staircase line. Hence repeating this procedure with any remaining pair of staircase lines in M will eventually yield a $(m - a)$ -cut in which there is at most one staircase line left. Since the cut length is non-increasing during each transformation step, the cut length of the final set M is at most $7\sqrt{a} + 2l$, which concludes the proof. \square

Using the above techniques we can find an m -cut containing at most one staircase line for any optimal m -cut containing a rectangular line, such that the cut length of the former m -cut is at most a constant times the cut length of the latter. The following theorem summarises these results.

Theorem 23. *For any polygon \mathcal{P} with an optimal m -cut L of \mathcal{P} containing a rectangular line, there exists a non-crossing m -cut M which contains only corner and straight lines except at most one which is a staircase line. Moreover M has a cut length of at most $9l$, where l is the cut length of L .*

Proof. Let $\rho \in L$ be a rectangular line in L which w.l.o.g. is concave w.r.t. $\mathcal{A}(L)$. Let \mathcal{R} be the defining rectangle of ρ and let $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$ be the boundary points of ρ . We consider \mathcal{R} to be a closed set, i.e. \mathcal{R} contains its boundary. Assume w.l.o.g. that ρ is oriented in a way such that p_1 is part of the left boundary of \mathcal{R} while p_2 is part of the lower boundary of \mathcal{R} . This in particular means that $x_1 \leq x_2$ and $y_1 \geq y_2$. Let β be the boundary of the polygon \mathcal{P} and

$$\begin{aligned} x_{\max} &= \max\{x \in \mathbb{R} \mid (x, y) \in \beta \cap \mathcal{R}\} \text{ and} \\ y_{\max} &= \max\{y \in \mathbb{R} \mid (x, y) \in \beta \cap \mathcal{R}\} \end{aligned}$$

be the extreme points of the boundary β in \mathcal{R} . Notice that $x_{\max} \geq x_2$ and $y_{\max} \geq y_1$ since p_1 and p_2 are boundary points and hence belong to β and (the boundary of) \mathcal{R} . We define the set of curves Ξ (Figure 10) such that $\lambda \in \Xi$ if and only if λ is a segment curve and

$$\begin{aligned} \lambda \subseteq & \{(x_{\max}, y) \in \mathcal{P} \mid y \in [y_2, y_{\max}]\} \cup \\ & \{(x_1, y) \in \mathcal{P} \mid y \in [y_1, y_{\max}]\} \cup \\ & \{(x, y_{\max}) \in \mathcal{P} \mid x \in [x_1, x_{\max}]\} \cup \\ & \{(x, y_2) \in \mathcal{P} \mid x \in [x_2, x_{\max}]\}. \end{aligned}$$

The set Ξ can be seen as a virtual rectangular line.

Let \mathcal{R}' be the “defining rectangle” of Ξ , i.e. the rectangle that is defined by the two opposing corners (x_1, y_2) and (x_{\max}, y_{\max}) . There are three corners of \mathcal{R}' that some curve in Ξ might include, namely (x_1, y_{\max}) , (x_{\max}, y_2) , and (x_{\max}, y_{\max}) . If (x_1, y_{\max}) is included in some curve $\lambda \in \Xi$ then this point cannot be a boundary point. Hence it must be the case that $y_{\max} > y_1$ and thus, by the definition of y_{\max} , there is some part of β that intersects with the upper boundary of \mathcal{R}' . But then there can be no single curve in Ξ that contains both (x_1, y_{\max}) and (x_{\max}, y_{\max}) since these are the endpoints of the upper boundary of \mathcal{R}' . A similar argument holds for (x_{\max}, y_2) and (x_{\max}, y_{\max}) . Therefore no curve in Ξ contains more than one corner, i.e. Ξ includes only straight and corner lines.

Let $\mathcal{D} \subseteq \mathcal{R}$ be the area that is cut out between ρ and Ξ in the a -cut $\Xi \cup \{\rho\}$, where a is the size of \mathcal{D} . (Remember that this means that \mathcal{D} is an open set.) Due to Lemma 12 and the fact that L is an optimal m -cut, apart from ρ the set L contains only straight and corner lines. Hence, since no curve crosses ρ and because by the construction of Ξ the set \mathcal{D} does not intersect with the boundary β , no curve from L crosses a curve in Ξ . Thus we can replace ρ with Ξ in L and yield a non-crossing $(m + a)$ -cut M' . By the construction of Ξ the corner lines in Ξ are concave w.r.t. $\mathcal{A}(M') = \mathcal{A}(M) \cup \mathcal{D}$, and by Lemma 12 the corner lines in L are concave w.r.t. $\mathcal{A}(M)$. Hence all corner lines in M' are concave w.r.t. $\mathcal{A}(M')$. By Lemma 22 we can thus find an m -cut M that has a cut length of at most $7\sqrt{a} + 2l'$, where l' is the cut length of M' , such that M contains only straight and corner lines except for at most one which is a staircase line.

By the construction of Ξ the total length of the curves in Ξ is at most the length of ρ and hence the cut length of M' is at most the cut length of L , i.e. $l \geq l'$. At the same time, since the size a of $\mathcal{D} \subseteq \mathcal{R}$ is smaller than the size of \mathcal{R} , and since the length l_ρ of ρ is greater than the height of \mathcal{R} plus the width of \mathcal{R} , it follows that $l \geq l_\rho \geq \sqrt{a}$. Hence we can upper bound the cut length of M by $7\sqrt{a} + 2l' \leq 9l$, which proves the claim. \square

4. Removing Staircase Lines

We now turn to the task of converting a (not necessarily optimal) 1-staircase cut L into a cut containing only straight and corner lines. Similar to the case of rectangular lines we will replace the staircase line with a set of appropriate corner and straight lines along the boundary of the deficit (or surplus). It is easy to see that if the deficit (or surplus) area of the staircase line λ has size a , then $\sqrt{a} < l$, where l is the cut length of L . Thus, if we can cut out the excess area a using straight and corner lines of total length in $\mathcal{O}(\sqrt{a})$, then our cut length will still be close to optimal. Given any simple polygon \mathcal{P} of area n , $a \in [0, n]$, and $\varepsilon \in]0, 1]$ we can find a set of at most three virtual corner lines that cut out an area whose size is in the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$ with a cut length that is a constant (depending on ε) times \sqrt{a} . Furthermore the corners of these virtual corner lines all have either the same x -coordinate or the same y -coordinate. They can be found using the short interval $[x_1, x_2]$ that was identified before (Figure 13). We use the virtual corner line with corner (x_1, y_2) which has short length but cuts out an area that is too large. To correct for the area we additionally find two virtual corner lines (of short length) with corners at either points (x', y_2) and (x'', y_2) , for some $x', x'' \geq x_2$, or points (x_1, y') and (x_1, y'') , for some $y', y'' \geq y_1$.

Lemma 24. *For any polygon \mathcal{P} of total area n , any $a \in [0, n]$, and any $\varepsilon \in]0, 1]$ there is a set L of straight and corner lines with the following properties. The lines in L cut out an area which has a size in the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$, and the cut length of L is at most $(6\sqrt{7/\varepsilon} + 2) \cdot \sqrt{a}$. Furthermore L is the union of at most three virtual corner lines with corners that either have the same x - or y -coordinate.*

Proof. If ε and a are chosen such that $(1 + \varepsilon)a \geq n$ the lemma trivially holds since L can be empty. Hence assume that $(1 + \varepsilon)a < n$ throughout this proof. Let $c = \sqrt{7/\varepsilon}$. If there is a virtual corner line for a that has a length of at most $2c\sqrt{a}$ the lemma obviously holds. Since $\sqrt{7/\varepsilon} > 2$ for $\varepsilon \in]0, 1]$, if there is no such virtual corner line then by Lemmas 16 and 17 there is an interval $[x_1, x_2] \subseteq I_a$ with the properties listed therein. We use the same notation as in the proof of Lemma 17, but for better readability let $\mathcal{Q} = \mathcal{Q}_{y_2}$, $w = w(x_2)$, and $h = h(y_2)$. The size of the area \mathcal{Q} is $a + d$ for some $d > 0$, i.e. the size of the area $\mathcal{Q} \setminus \mathcal{P}(x_2)$ is d (see Figure 17).

Our first goal in this proof is to establish an upper bound on the size b of the area $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ depending on d . For this we establish a lower bound




Figure 17: The interval $[x_1, x_2]$ of width w together with the virtual corner lines $\Lambda(x_1)$ and $\Lambda(x_2)$ (grey dashed lines). The dark grey area is $\mathcal{P}(x_1) \setminus \mathcal{P}(x_2)$ of size b' , while the light grey area is $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ of size b . The right-most point of $\mathcal{P}(x_2)$ is x_3 , and x and $\varphi(x)$ define points at which the virtual corner lines $\Lambda'(x)$ and $\Lambda'(\varphi(x))$ (black dotted lines) enclose an area of size d' between them (striped pattern).

on the size b' of $\mathcal{P}(x_1) \setminus \mathcal{P}(x_2)$ which we can then subtract from a , the size of $\mathcal{P}(x_1)$. One simple bound can be derived by subtracting from d the size of the area not in $\mathcal{P}(x_1) \setminus \mathcal{P}(x_2)$ but in $\mathcal{Q} \setminus \mathcal{P}(x_2)$. Since the size of this area can be upper-bounded by $h \cdot w$, we get $b' \geq d - hw$. We can derive an upper bound for w depending on d by integrating along the vertical lengths of the virtual corner lines for a between x_1 and x_2 . By Lemma 16 this gives

$$d \geq \lim_{z \rightarrow x_1} \int_z^{x_2} l_v(t) dt > w \cdot c\sqrt{a}.$$

Hence $w < \frac{d}{c\sqrt{a}}$, and using the upper bound on h given by Lemma 17 we can conclude that $b' > d(1 - \frac{2}{c^2})$. This directly translates into the upper bound on the size of the area $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ which is

$$b = a - b' < a - d \left(1 - \frac{2}{c^2}\right). \quad (7)$$

The next step is to find a lower bound on b (which also depends on c) under the assumption that no appropriate set L exists. We will show that for $c = \sqrt{7/\varepsilon}$ the upper and lower bounds contradict each other. Let Δ denote the virtual corner line for $a + d$ with corner (x_1, y_2) , i.e. Δ cuts out \mathcal{Q} . If $d \leq \varepsilon a$ we can cut out an area which has a size in the interval $[a, (1 + \varepsilon)a]$ by letting $L = \Delta$. By Lemmas 16 and 17 the virtual corner line Δ has a length of at most $2(c\sqrt{a} + 2\sqrt{a}/c) \leq (2\sqrt{7/\varepsilon} + 2)\sqrt{a}$ and hence in this case L satisfies the lemma. Therefore let $d > \varepsilon a$ in the remainder of the proof. We will attempt to find L by either including Δ in L and cutting out an area of size approximately d from \mathcal{Q} , or by cutting out an area of size approximately a from \mathcal{Q} directly. The decision on whether to include Δ in L is determined by distinguishing between small and big values for d . In case $d < a$ we include Δ in L and otherwise not.

Since we have the freedom to choose the size of the area that we cut out from the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$, we attempt to cut out the smallest possible area from \mathcal{Q} . Hence if d' denotes the size of this area, let $d' = \min\{d, a\} - \varepsilon a$. Notice that d' is well-defined since $d > \varepsilon a$, and that the size of the cut out area is $a - \varepsilon a$ if Δ is not included in L and it is $a + d - (d - \varepsilon a) = a + \varepsilon a$ otherwise. Hence the size of the cut out area lies in the given interval.

To cut out the area of size d' from \mathcal{Q} we use a pair of virtual corner lines such that for both lines either the horizontal parts overlap with $\Lambda(x_2)$ or the vertical parts overlap with $\Lambda(x_1)$. Notice that such a pair always exists, since $d' < a$ and the lines $\Lambda(x_1)$ and $\Lambda(x_2)$ each cut out an area of size a . Since by Lemma 16 both the horizontal length of $\Lambda(x_2)$ and the vertical length of $\Lambda(x_1)$ is short, we only have to guarantee that either the vertical or the horizontal lengths of the two desired virtual corner lines are short, respectively. We will concentrate on the case where we pick the virtual corner lines from those that overlap with $\Lambda(x_2)$, since the other case is analogous. Therefore, if x_3 is defined such that $x_3 - x_2$ is the width of $\mathcal{P}(x_2)$, let $\Lambda'(x)$ denote the virtual corner line with corner (x, y_2) for any $x \in]x_2, x_3]$ and let $l'_v(x)$ be its vertical length.

Let $\varphi(x)$ be a function that, for a given virtual corner line $\Lambda'(x)$, gives the x -coordinate of $\Lambda'(\varphi(x))$, such that $\Lambda'(x)$ and $\Lambda'(\varphi(x))$ enclose an area of size d' between them and $\varphi(x) > x$. This means that $\int_x^{\varphi(x)} l'_v(t) dt = d'$ and that the domain of φ is upper-bounded by $\varphi^{-1}(x_3)$ where φ^{-1} is the inverse function of φ (notice that the function φ is bijective). Assume that there is no pair of virtual corner lines $\Lambda'(x)$ and $\Lambda'(\varphi(x))$ for which both vertical lengths are shorter than $c\sqrt{a}$. From this assumption it follows that for all $x \in [x_2, \varphi^{-1}(x_3)]$ it holds that $l'_v(x) > c\sqrt{a}$ or $l'_v(\varphi(x)) > c\sqrt{a}$.

Let for any interval $I \subseteq [x_2, x_3]$ the function f be equal to the size of the area $\{(x', y') \in \mathcal{Q} \mid x' \in I\}$ in the vertical stripes in \mathcal{Q} defined by I , i.e.

$$f(I) = \int_I l'_v(t) dt.$$

Let $J = \{x \in [x_2, \varphi^{-1}(x_3)] \mid l'_v(x) > c\sqrt{a}\}$ be the subset of the domain of φ for which the vertical lengths of the virtual corner lines $\Lambda'(x)$ are long. Let \bar{J} be the set for which the vertical lengths are short, i.e. $\bar{J} = [x_2, \varphi^{-1}(x_3)] \setminus J$. Also let $K = \{x \in [\varphi(x_2), x_3] \mid l'_v(x) > c\sqrt{a}\}$ and $\bar{K} = [\varphi(x_2), x_3] \setminus K$ be the corresponding subsets from the domain of φ^{-1} . To establish the connection between the assumption on the lengths of the vertical lines and the lower bound on b we investigate $f(J \cup K)$ (see Figure 18).

Let $[l, r] \subseteq \bar{J}$ be a connected subset of \bar{J} . By the definitions of f and φ we get

$$\begin{aligned} f([l, r]) &= f([l, \varphi(l)]) + f([\varphi(l), \varphi(r)]) - f([r, \varphi(r)]) \\ &= f([\varphi(l), \varphi(r)]). \end{aligned}$$

Since \bar{J} is a union of connected subsets and φ is bijective we can conclude that $f(\bar{J}) = f(\varphi(\bar{J}))$, where $\varphi(\bar{J})$ is the image of \bar{J} . By the assumption that for all




Figure 18: The interval $[x_1, x_2]$ together with the virtual corner lines $\Lambda(x_1)$ and $\Lambda(x_2)$ (dashed lines). To estimate the size b of $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ we determine the size $f(J \cup K)$ of \mathcal{X} and the size of \mathcal{Y} .

$x \in [x_2, \varphi^{-1}(x_3)]$ the vertical length of $\Lambda'(x)$ or of $\Lambda'(\varphi(x))$ is long, $\varphi(\bar{J})$ must be a subset of K and hence $f(\bar{J}) \leq f(K)$. A similar observation can be made for \bar{K} , $\varphi^{-1}(\bar{K})$, and J so that $f(\bar{K}) \leq f(J)$.

By the definition of φ we know that $f([\varphi^{-1}(x_3), x_3]) = d'$ and $f([x_2, \varphi(x_2)]) = d'$, while the total area of $\mathcal{P}(x_2)$ has size a . Hence $f(J \cup \bar{J}) = f(K \cup \bar{K}) = a - d'$. From the bounds above and the fact that J and \bar{J} but also K and \bar{K} are disjoint we can conclude that

$$f(J) + f(K) \geq f(\bar{K}) + f(\bar{J}) = 2(a - d') - f(K) - f(J).$$

The sets J and K might not be disjoint but from the above inequality we get

$$f(J \cup K) + f(J \cap K) = f(J) + f(K) \geq a - d'.$$

By the pigeonhole principle and the fact that $(J \cap K) \subseteq (J \cup K)$ we can thus conclude that

$$f(J \cup K) \geq \frac{a - d'}{2}.$$

Let $\mathcal{X} = \{(x, y) \in \mathcal{P}(x_2) \mid x \in J \cup K\}$ and let $f_{\mathcal{X}}$ be the size of \mathcal{X} , i.e. $f_{\mathcal{X}} = f(J \cup K)$. We now want to also consider the assumption that there is no pair of virtual corner lines amongst those overlapping with $\Lambda(x_1)$ such that both their horizontal lengths are short while cutting out an area of d' between them. Let $\mathcal{Y} \subseteq \mathcal{P}(x_1)$ denote the area such that $(x, y) \in \mathcal{Y}$ if and only if there is a virtual corner line with corner (x_1, y) which has a horizontal length greater than $c\sqrt{a}$, analogous to the definition of \mathcal{X} . Using a similar argumentation as for the set \mathcal{X} , we can conclude that the size $f_{\mathcal{Y}}$ of \mathcal{Y} must also be at least $\frac{a-d'}{2}$ if no pair of virtual corner lines exists that cuts out an area of size d' such that both horizontal lengths are short.

To yield a lower bound on b we want to consider those parts of \mathcal{X} and \mathcal{Y} that are contained in $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ and determine their size. For this we need

to find an appropriate bound on the parts of \mathcal{X} and \mathcal{Y} that are not contained in $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$, but also a bound on the size of the intersection of \mathcal{X} and \mathcal{Y} . Therefore let $w_{\mathcal{X}}$ be the total width of \mathcal{X} , i.e. $w_{\mathcal{X}}$ is the total length of the interval $J \cup K$. Since \mathcal{X} is contained in $\mathcal{P}(x_2)$ and the latter has a size of a we can conclude that

$$a \geq \int_{J \cup K} l'_v(t) dt > w_{\mathcal{X}} \cdot c\sqrt{a},$$

and hence $w_{\mathcal{X}} < \frac{\sqrt{a}}{c}$. If $h_{\mathcal{Y}}$ denotes the total height of \mathcal{Y} , a similar argument yields that also $h_{\mathcal{Y}} < \frac{\sqrt{a}}{c}$.

Those parts of \mathcal{X} that are not contained in $\mathcal{P}(x_1)$ are confined to the area below $\mathcal{P}(x_1)$ in \mathcal{Q} which has height h . Hence the area $\mathcal{X} \setminus \mathcal{P}(x_1)$ has a size of at most $h \cdot w_{\mathcal{X}}$. Similarly the area $\mathcal{Y} \setminus \mathcal{P}(x_2)$ has a size of at most $w \cdot h_{\mathcal{Y}}$, since the area to the left of $\mathcal{P}(x_2)$ in \mathcal{Q} has width w . The size of the intersection of \mathcal{X} and \mathcal{Y} can be at most $w_{\mathcal{X}} \cdot h_{\mathcal{Y}}$. Thus, using the bounds on w and h given in Lemma 17 we can conclude that

$$b \geq f_{\mathcal{X}} - h \cdot w_{\mathcal{X}} + f_{\mathcal{Y}} - w \cdot h_{\mathcal{Y}} - w_{\mathcal{X}} \cdot h_{\mathcal{Y}} > a - d' - \frac{5}{c^2}a. \quad (8)$$

We make a case distinction on the value of d to compare the lower and upper bounds on b . If $d < a$ then $d' = d - \varepsilon a$ so that setting $c = \sqrt{7/\varepsilon}$ in Bounds (7) and (8) gives

$$\begin{aligned} b &< a - d \left(1 - \frac{2}{7}\varepsilon\right) < a - d + \frac{2}{7}\varepsilon a \text{ and} \\ b &> a - (d - \varepsilon a) - \frac{5}{7}\varepsilon a = a - d + \frac{2}{7}\varepsilon a, \end{aligned}$$

which is a contradiction. In the case when $d \geq a$ it holds that $d' = (1 - \varepsilon)a$ so that, using the fact that $\varepsilon \in]0, 1]$, Bounds (7) and (8) give

$$\begin{aligned} b &< a - d \left(1 - \frac{2}{7}\varepsilon\right) \leq a - a \left(1 - \frac{2}{7}\varepsilon\right) = \frac{2}{7}\varepsilon a \text{ and} \\ b &> a - (1 - \varepsilon)a - \frac{5}{7}\varepsilon a = \frac{2}{7}\varepsilon a, \end{aligned}$$

which again is a contradiction.

We can thus conclude that one of our assumptions must be wrong. Therefore there always exists a pair of virtual corner lines which cuts out an area of size d' and has a short total length. These can be found either amongst those overlapping with $\Lambda(x_1)$ or those overlapping with $\Lambda(x_2)$. Hence we can find the set L which is the union of these virtual corner lines and, depending on the value of d , also Δ in case we need it. The cut length of L is at most the length of the two corner lines cutting out the area d' between them, plus the length of Δ . Together these three virtual corner lines have a length of at most $4c\sqrt{a} + 2(c + 2/c)\sqrt{a} < (6\sqrt{7/\varepsilon} + 2)\sqrt{a}$, which concludes the proof. \square




Figure 19: A tail defined by the dotted line. Three (dashed) virtual corner lines cut out the area shaded in grey (the lines overlap on the bottom right).

Note that in the above proof the existence of $\varepsilon > 0$ guarantees that the intervals J and K have lengths greater than zero: it may happen that $d = a$ so that the size of $\mathcal{Q} \setminus \mathcal{P}(x_1)$ is a . In this case, whether Δ is included in L or not, $\varphi(x_2) = x_3$ if $\varepsilon = 0$. Thus for the proof technique used above we need to allow a deviation from cutting out an exact area size as given by the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$.

To apply the above result, we need to find a region of the polygon of size larger than a that does not contain any curves of the cut, so that we can cut out the excess area without interfering with the other curves. For this we define the concept of a *tail* of a polygon with respect to a cut: for any cut in a polygon \mathcal{P} , consider all the connected pieces of the polygon cut out by it. If there a connected piece \mathcal{T} that is defined by a single curve τ then we call \mathcal{T} a tail of the polygon.

Definition 25 (tail). For an m -cut L in a polygon \mathcal{P} , let $\mathcal{T} \subseteq \mathcal{P} \setminus \bigcup_{\mu \in L} \mu$ be a connected area that is cut out by L . We call \mathcal{T} a *tail* if there exists a single curve $\tau \in L$ that cuts out \mathcal{T} . We refer to τ as *the curve of \mathcal{T}* . In case L contains a staircase line λ , we call a tail $\mathcal{T} \subseteq \mathcal{A}(L)$ respectively $\mathcal{T} \subseteq \mathcal{B}(L)$ *small* if its area is strictly smaller than λ 's deficit respectively surplus.

Notice that there always exists a tail if \mathcal{P} is a simple polygon. Notice also that apart from the curve of a tail \mathcal{T} there might be other subsets of L that cut out \mathcal{T} if curves in L overlap.

To convert a cut containing a staircase line λ into one containing only straight and corner lines, we can shift λ in either direction, i.e. going into either the \mathcal{A} - or the \mathcal{B} -part. However all the tails in the polygon may belong to only one part. We need to consider two cases, one of which is when L contains only λ . This means that there are exactly two tails, one on each side of λ . If we assume w.l.o.g. that the size a of λ 's deficit is at most that of its surplus, we can replace the staircase line by the set of straight and corner lines on the boundary of its deficit. We then cut out the area $a' \in [(1 - \varepsilon)a, (1 + \varepsilon)a]$ from the original \mathcal{A} -part (containing the surplus) using the at most three virtual corner lines which were shown to exist above (Figure 19). The other case is when there is a tail contained

in, say, the \mathcal{A} -part whose curve μ is not λ . We can safely assume that the size of the tail is larger than the size a of the deficit of λ . If this was not the case then we could remove μ from the cut by using an area exchange with the staircase line λ , without increasing the cut length, as the next lemma shows.

Lemma 26. *Let L be a 1-staircase m -cut in \mathcal{P} with cut length l . There exists a 1-staircase m -cut L' in \mathcal{P} , has cut length of at most l , and the curve of any small tail cut out by L' equals the staircase line.*

Proof. If the curve of every small tail cut out by L is the staircase line $\lambda \in L$ there is nothing to prove. Hence let \mathcal{T} be a small tail cut out by L such that its curve is $\lambda' \neq \lambda$. Assume w.l.o.g. that $\mathcal{T} \subseteq \mathcal{A}(L)$, i.e. \mathcal{T} is strictly smaller than λ 's deficit. This means that we can find a staircase line λ'' that cuts out an area that has the same size as \mathcal{T} from λ 's deficit, such that removing λ' and replacing λ with λ'' in L yields an m -cut that has a cut length that is less than l by the length of λ' . The new cut has one straight or corner line less than the old one and it contains one staircase line. We repeat this process for any small tail cut out by the new set that does not conform with the desired property. This will eventually terminate in a state in which the resulting m -cut L' fulfils the property. \square

As this lemma shows we can replace the staircase line λ by the corner and straight lines on the boundary of its deficit and cut out an area a' from the tail, using the virtual corner lines of short length. It may be that some of the virtual corner lines end at the curve μ of the tail. If this happens we can find a set of straight and corner lines that overlap with parts of the virtual corner lines and μ , with which to replace the latter lines (in the same way as suggested by Figure 16). The cut out area is the same while the cut length only grows by a constant factor since there are at most three virtual corner lines. The result of the above described method is summarised in the following theorem.

Theorem 27. *Given a 1-staircase m -cut L of a polygon \mathcal{P} with cut length l , for any $\varepsilon \in]0, 1]$ there exists a corner m' -cut L' , where $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$, having a cut length of at most $(6\sqrt{7}/\varepsilon + 7) \cdot l$.*

Proof. Due to Lemma 26 we can assume that any tail cut out by L is not small or its curve is the staircase line $\lambda \in L$. Consider the case when there is a tail \mathcal{T} and its curve is $\lambda' \in L$ such that $\lambda' \neq \lambda$, i.e. \mathcal{T} is not small. In case $\mathcal{T} \subseteq \mathcal{A}(L)$ let a denote the size of λ 's deficit and in case $\mathcal{T} \subseteq \mathcal{B}(L)$ let a denote the size of λ 's surplus.

The curve λ' is either a straight or a corner line. We assume w.l.o.g. that the horizontal bar line σ'_h of λ' (if any) lies below \mathcal{T} and the vertical bar line σ'_v of λ' (if any) lies to the left of \mathcal{T} . That is, for all sufficiently small $z > 0$, $(x_h, y_h) \in \sigma'_h$, and $(x_v, y_v) \in \sigma'_v$ it holds that $(x_h, y_h + z) \in \mathcal{T}$, $(x_h, y_h - z) \notin \mathcal{T}$, $(x_h + z, y_h) \in \mathcal{T}$, and $(x_h - z, y_h) \notin \mathcal{T}$. Notice that this in particular means that if λ' is a corner line then it points up and right if λ' is convex w.r.t. \mathcal{T} , and it points down and left if λ' is concave w.r.t. \mathcal{T} . According to Lemma 24 there is an a' -cut L' , for some $a' \in [(1 - \varepsilon)a, (1 + \varepsilon)a]$, in \mathcal{T} such that L' is

the union of at most three virtual corner lines, i.e. L' contains only straight or corner lines where the latter point up and right. Let $\lambda'' \in L'$ be a curve that has a boundary point p with respect to \mathcal{T} such that $p \in \lambda'$. Assume that λ'' is a corner line. If $p \in \sigma'_v$, the corner of λ'' must lie to the right of p since $\lambda'' \subset \mathcal{T}$ and the assumption made on the location of \mathcal{T} with respect to λ' . However this contradicts the orientation of λ'' since its corner must lie to the left of or below its boundary point. A similar contradiction can be derived if $p \in \sigma'_h$. Hence it must be the case that λ'' is a straight line. If $p \in \sigma'_v$ then let $\sigma \subseteq \sigma'_v$ be the part of σ'_v that lies above p if λ' is a vertical straight line or λ' is a convex corner line w.r.t. \mathcal{T} , and let $\sigma \subseteq \sigma'_v$ be the part of σ'_v that lies below p if λ' is a concave corner line w.r.t. \mathcal{T} . If $p \in \sigma'_h$ then let $\sigma \subseteq \sigma'_h$ be the part of σ'_h that lies to the right of p if λ' is a horizontal straight line or λ' is a convex corner line w.r.t. \mathcal{T} , and let $\sigma \subseteq \sigma'_h$ be the part of σ'_h that lies to the left of p if λ' is a concave corner line w.r.t. \mathcal{T} . Notice that in all cases σ is a bar line between p and a boundary point of λ' . Hence we can convert λ'' into a corner line in \mathcal{P} by adding the point p and the line σ to it.

If there are at most two virtual corner lines that make up the set L' then there can be at most four straight lines that have to be converted to corner lines in \mathcal{P} : one for each horizontal and vertical part of the virtual corner lines. Lemma 24 states that the virtual corner lines in L' have corners that either have the same x - or y -coordinate. This means that the straight lines on either the vertical parts or the horizontal parts overlap. Hence if there are three virtual corner lines then two of each overlapping triple can be removed so that the resulting set of curves still is an a' -cut and the cut length decreases. Thus also in this case there are at most four straight lines in L' that have to be converted to corner lines in \mathcal{P} : three in either the horizontal or the vertical parts of the virtual corner lines and one in the other part. Therefore after converting L' and adding these curves to L , the resulting set of curves M' has a cut length of at most $5l + (6\sqrt{7/\varepsilon} + 2)\sqrt{a}$.

Notice that M' , apart from λ , only contains straight and corner lines. Hence using Lemma 21 we can replace the staircase line λ with a set of corner and straight lines, yielding the set M . It only contains straight and corner lines and has a cut length of at most that of M' . What remains to be shown, in case the boundary of \mathcal{T} does not contain λ , is that M cuts out an area of the desired size and that its cut length is of the desired length. The set M cuts out an area of size m' where $m' \in [m - \varepsilon a, m + \varepsilon a]$. Since \mathcal{T} is a tail that is not small, if $\mathcal{T} \subseteq \mathcal{A}(L)$ we can conclude that the size of \mathcal{T} is greater or equal to a and hence $m \geq a$. If $\mathcal{T} \subseteq \mathcal{B}(L)$, since the surplus is part of $\mathcal{A}(L)$ obviously $m \geq a$ also holds in this case. Thus $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$, which establishes the desired size for the area. Concerning the cut length, let \mathcal{R} be the rectangle that is defined by the boundary points of λ , let h be its height, w be its width, and let w.l.o.g. $h \geq w$. The length l_λ of λ is $l_\lambda = h + w > h$. Since both the deficit and the surplus of λ are contained in \mathcal{R} we know that $a < hw \leq h^2$. Hence we can conclude that $l \geq l_\lambda > \sqrt{a}$. This means that the cut length of M is at most $(6\sqrt{7/\varepsilon} + 7)l$, as claimed.

Now consider the case when there is no tail cut out by L such that its curve is

different from λ . This means that the only curve in L is λ . In this case we need to proceed differently than in the case before by reversing the transformation of the m -cut: we first remove λ and instead add the curves that, apart from λ , are contained in the boundary of λ 's deficit. This yields a $(m+a)$ -cut M' which contains only straight and corner lines, where a is the size of the deficit. Furthermore, if we assume w.l.o.g. that λ is oriented down and its deficit lies to the lower left side of λ , the corner lines all point up and right and are convex w.r.t. $\mathcal{A}(M')$. Since the deficit of λ is part of $\mathcal{A}(M')$, we can use Lemma 24 to find a set of straight and corner lines L' in $\mathcal{A}(M')$ that cuts out an area of size $a' \in [(1-\varepsilon)a, (1+\varepsilon)a]$. Again we need to convert those curves in L' that have a boundary point on one of the curves in M' into feasible curves in \mathcal{P} . Since the corner lines in both L' and M' point up and right, any curve in L' that has a boundary point in $\mathcal{A}(M')$ on one of the curves in M' can only have one such boundary point. Hence the same arguments as given above for the other case also apply for each such case here. We can thus make the necessary conversions of the curves in L' , add the curves in M' , and thereby yield the set of curves M which only contains straight and corner lines. As above it cuts out an area of size $m' \in [(1-\varepsilon)m, (1+\varepsilon)m]$, and has a cut length of at most $(6\sqrt{7/\varepsilon} + 7)l$, which concludes the proof. \square

5. Converting Curves in Polygons to Segments in Grids

We have learned that, for any desired area m to be cut out from a simple polygon, there exists a cut of only straight and corner lines that (1) cuts out at most a small amount $\varepsilon \cdot m$ more (or less) than the desired area, and (2) has a cut length that is close to the optimum (of arbitrary shape for area m). We summarise these results in the following corollary.

Corollary 28. *Let l be the cut length of an optimal m -cut L in some polygon \mathcal{P} . For any $\varepsilon \in]0, 1]$ there exists a non-crossing corner m' -cut for some $m' \in [(1-\varepsilon)m, (1+\varepsilon)m]$, which has a cut length of at most $(54\sqrt{7/\varepsilon} + 63) \cdot l$.*

Proof. According to Corollary 14 we can assume that L only contains straight and corner lines except at most one curve which can either be a staircase or a rectangular line. In case L only contains straight and corner lines there is nothing to prove. In case it contains a staircase line the claim holds according to Theorem 27. In case L contains a rectangular line we can use Theorem 23 to convert L into an m -cut L' . It has a cut length of at most $9l$ and contains only straight and corner lines except at most one staircase line. If L' does not contain a staircase line, the claim obviously holds. Otherwise, using Theorem 27 on L' , we can convert L into a non-crossing corner m' -cut, for some $m' \in [(1-\varepsilon)m, (1+\varepsilon)m]$, having a cut length of at most $(6\sqrt{7/\varepsilon} + 7) \cdot 9l$, which concludes the proof. \square

Because our real interest is in cuts in grids, we now face the task to find a cut in the grid G given a cut in the polygon \mathcal{P}_G constructed from G . Our transformation from a grid to a polygon implies that an optimal m -cut in G transforms into an m -cut in \mathcal{P}_G . But not necessarily into an optimal one, since




Figure 20: A grid line λ_1 and a non-grid line λ_2 . The corridor of λ_2 is shaded in grey. The boundary of the polygon is divided into lines of unit length.

the cut curves in the polygon are not limited to integer positions (these are integer positions in the dual of the grid, and thus halfway positions between grid points). In other words, a cut in the polygon does not generally translate directly into a cut in the grid (note that if we would just cut grid edges with polygon cut curves, that is, not cut them in the middle, this would not translate the cut out area into the same number of grid vertices). Whenever a curve in the cut of \mathcal{P}_G happens to lie in integer position however, we will just take the corresponding segment to cut the grid G (Figure 20).

Definition 29 (corridor, grid line). Given a grid $G = (V, E)$ let \mathcal{S}_v be the axis-parallel unit square that has $v \in V$ as its centre and let γ_v be the boundary of \mathcal{S}_v . We consider a unit square to be an open set, i.e. $\gamma_v \cap \mathcal{S}_v = \emptyset$ for all $v \in V$. For any curve λ in \mathcal{P}_G we refer to the set $\mathcal{K}_\lambda = \{p \in \mathcal{P}_G \mid \exists v \in V : \lambda \cap \mathcal{S}_v \neq \emptyset \wedge p \in \mathcal{S}_v \cup \gamma_v\}$ as the *corridor* of λ . It is the union over the unit squares that intersect with λ together with their boundaries that are not part of the boundary of \mathcal{P}_G . A curve λ in \mathcal{P}_G is called a *grid line* if $\mathcal{K}_\lambda = \emptyset$, i.e. λ lies on the boundaries of the unit squares.

For non-grid lines, we start with a clean-up phase that modifies a pair of these curves so that one of them becomes a grid line, and the other compensates for the area difference that this creates. We start the clean-up phase by first focussing on unit length open intervals on the polygon boundary between adjacent integer positions, as defined next.

Definition 30 (U_G). Given a grid G let β be the boundary of \mathcal{P}_G . We define $\mathbb{H} = \{x - \frac{1}{2} \mid x \in \mathbb{N}\}$ so that \mathbb{H}^2 denotes the points between integer positions in the plane. Let the set U_G contain all unit length curves in $\beta \setminus \mathbb{H}^2$.

Because a grid line λ does not hit any such open unit interval $\delta \in U_G$ we are only concerned with cut curves that do. For any open unit interval hit by more than one cut curve, we can shift one of these curves to the boundary and compensate for the area difference by also shifting one other of these curves accordingly. Repeating this leaves us with at most one cut curve per open unit interval in U_G on the boundary of \mathcal{P}_G (and ultimately G).




Figure 21: The considered curve δ and the lines λ_1 to $\lambda_{|K|}$ from left to right (dashed). The area shaded in grey is $\mathcal{D}(\lambda_2)$. The black contiguous lines in the picture are curves from U_G .

Lemma 31. *For any grid G and any non-crossing corner m -cut L of cut length l in \mathcal{P}_G , there is a non-crossing corner m -cut M of cut length at most l in \mathcal{P}_G such that there is no curve in U_G which includes more than one boundary point of curves in M .*

Proof. Consider the case when there is a curve $\delta \in U_G$ such that at least two curves in L have boundary points on δ . Without loss of generality let δ be the lower side of a unit square \mathcal{S}_v , i.e. the curves having a boundary point on δ lie above it. This means that any corner line having a boundary point on δ points down, while any such straight line is vertical. Let $K \subseteq L$ be the set of curves that have a boundary point on δ (Figure 21). We define (x_δ, y_δ) to be the lower left corner of the unit square \mathcal{S}_v to which δ is the lower side. Let for any curve $\lambda \in K$ the point (x_λ, y_λ) be either the corner of λ , if it is a corner line, or the boundary point of λ that does not lie on δ , if λ is a straight line. We define

$$\mathcal{D}(\lambda) = \{(x, y) \in \mathcal{K}_\lambda \mid x \in]x_\delta, x_\lambda[\wedge y \in]y_\delta, y_\lambda[\}$$

to be the open set of points in λ 's corridor that lie to the left of λ .

Since the curves in L are non-crossing, observe that if $\lambda \in K$ is a corner line pointing down and left any curve $\lambda' \in L$ that intersects $\mathcal{D}(\lambda)$, i.e. $\lambda' \cap \mathcal{D}(\lambda) \neq \emptyset$, must be a corner line and it must have the same orientation as λ . Thus λ' must also have a boundary point on δ since the lower boundary of $\mathcal{D}(\lambda)$ is part of δ . From this we can conclude that for a boundary point p on δ that belongs to a corner line pointing down and left, the boundary points to the left of p on δ all belong to corner lines of the same orientation. Furthermore they must all be of smaller vertical length since the height of $\mathcal{D}(\lambda)$ equals the vertical length of λ . An analogous observation can be made for a corner line $\lambda \in K$ pointing down and right, if we consider the open set of points in its corridor that lie to the right of λ . Hence we can order the curves in K by traversing their boundary points on δ from left to right such that we first encounter corner lines pointing down and left with increasing vertical length, then straight lines, and finally corner lines pointing down and right with decreasing vertical length. Obviously this




Figure 22: The line λ_1 (dashed), a point (x, y) on the boundary of \mathcal{P}_G , and the corresponding set of curves $\Lambda_1(x, y)$ (dotted). The black contiguous lines in the picture are curves from U_G .

is also possible if some of the curves in K share the same boundary point on δ . Let the indices of the curves in $K = \{\lambda_1, \dots, \lambda_{|K|}\}$ denote their position in this order (cf. Figure 21).

We will consider the curves λ_1 and λ_2 from K and in each case attempt to move the vertical part of λ_1 to the left until it intersects with the boundary of the unit squares, i.e. until the vertical line is a grid line. Thereafter we will find one or several curves that substitute λ_2 such that the resulting set of curves is an m -cut again. Towards this end, for any point (x, y) on the boundary of \mathcal{P}_G such that $x \leq x_{\lambda_1}$ and $y \in [y_{\delta}, y_{\lambda_1}]$, we define the set of curves $\Lambda_1(x, y)$ as those straight and corner lines between (x, y) and (x_{δ}, y_{δ}) (Figure 22). That is $\mu \in \Lambda_1(x, y)$ if and only if μ is a segment curve and

$$\mu \subseteq \{(x', y') \in \mathcal{P}_G \mid (x' = x_{\delta} \wedge y' \in [y_{\delta}, y]) \vee (y' = y \wedge x' \in [x, x_{\delta}])\}.$$

Note that $\Lambda_1(x, y)$ contains more than one curve if $\mathcal{D}(\lambda_1)$ touches the boundary of \mathcal{P}_G to its left.

Consider the case when λ_1 and λ_2 are corner lines pointing down and left (Figure 23). Let $(x, y) \notin \delta$ be the boundary point of λ_1 that does not lie on δ . We know that no curve from L intersects $\mathcal{D}(\lambda_1)$ by the observations made above. This means that removing λ_1 and adding the curves in $\Lambda_1(x, y)$ yields a non-crossing m' -cut for some m' . The difference between m and m' is equal to the size of $\mathcal{D}(\lambda_1)$ and is hence less than the size of $\mathcal{D}(\lambda_2)$. Also the only curve in L that intersects $\mathcal{D}(\lambda_2)$ is λ_1 . Hence, after replacing λ_1 , we can find a corner line μ pointing down and left that has its corner on λ_2 's horizontal bar line and a boundary point on δ , such that removing λ_2 and introducing μ will again result in a non-crossing m -cut. Notice that the total length of the curves in $\Lambda_1(x, y)$ is shorter than the length of λ_1 and also the length of μ is shorter than the length of λ_2 . Hence we obtain an m -cut of smaller cut length than l . Also the number of boundary points on δ is reduced by one.

In case the curves $\lambda_{|K|}$ and $\lambda_{|K|-1}$ are corner lines pointing down and right we can use an analogous argumentation as the one given above to obtain an




Figure 23: The case when λ_1 and λ_2 are corner lines (left). They are substituted with $\Lambda_1(x, y)$ and μ (right). The area shaded in dark grey is $\mathcal{D}(\lambda_1)$, and the area shaded in both light and dark grey is $\mathcal{D}(\lambda_2)$. The size of the cut out area (diagonally striped) remains the same. The black contiguous lines in the pictures are curves from U_G .




Figure 24: The case when λ_1 is a corner line and λ_2 is a straight line that overlap as λ_1 (right). They are substituted by the line pointing up and left that has the same corner q as λ_1 (right). The size of the cut out area (diagonally striped) remains the same. The black contiguous lines in the pictures are curves from U_G .

m -cut of cut length smaller than l . In the new cut the number of boundary points on δ is reduced by one. By repeating this procedure, we can transform L into an m -cut of smaller cut length. This can be done until there are at most two corner lines with boundary points on δ , such that they point down and left, and down and right, respectively. We thus assume in the remainder of the proof that K contains at most one such corner line each, while all others are straight lines.

Consider the case when λ_1 is a corner line pointing down and left and λ_2 is a straight line. In case the boundary points of λ_1 and λ_2 are the same on δ these two curves overlap (Figure 24). Let q denote the corner of λ_1 . In this case we can introduce the corner line pointing up and left that has q as its corner. Clearly then removing λ_1 and λ_2 , we obtain an m -cut of smaller cut length than l . The number of boundary points on δ will be reduced by two. In case λ_1 and λ_2 do not share the same boundary point on δ (Figure 25), we can again replace λ_1




Figure 25: The case when λ_1 is a corner line and λ_2 is a straight line that do not overlap (left). Together with λ' (dotted) they are substituted with $\Lambda_1(x, y)$ and μ (right). The area shaded in grey is \mathcal{L} . The size of the cut out area (diagonally striped) remains the same. The black contiguous lines in the pictures are curves from U_G .

with the curves in $\Lambda_1(x, y)$, exactly as above, yielding a non-crossing m' -cut L' . Let $\mathcal{C} \in \{\mathcal{A}(L'), \mathcal{B}(L')\}$ be the part of the m' -cut for which $\mathcal{D}(\lambda_1) \subseteq \mathcal{C}$. We define $\mathcal{L} \subseteq \mathcal{D}(\lambda_2) \cap \mathcal{C}$ to be the connected area for which $\mathcal{D}(\lambda_1) \subseteq \mathcal{L}$. That is, \mathcal{L} lies to the left of λ_2 in L' . Since λ_1 and λ_2 do not overlap, the size of $\mathcal{D}(\lambda_1)$ is smaller than the size of \mathcal{L} , i.e. $\mathcal{D}(\lambda_1) \subset \mathcal{L}$. Hence we can find a vertical straight line σ in \mathcal{L} that cuts out an area the size of $\mathcal{D}(\lambda_1)$ on its right-hand side. One of the boundary points of σ (w.r.t. \mathcal{L}) lies on δ and the other boundary point can either lie on the boundary of \mathcal{P}_G or on a curve $\lambda' \in L \setminus \{\lambda_1, \lambda_2\}$. In the former case we can replace λ_2 with σ and again yield an m -cut which has a smaller cut length and one boundary point less on δ . Otherwise, note that λ' lies in $\mathcal{D}(\lambda_2)$ and hence must be a corner line pointing up and left, since any other straight or corner line would either cross λ_1 or λ_2 , or would have a boundary point on δ . This is not possible due to the choice of λ_2 in the ordering of K . This means that we can extend σ by a horizontal bar line σ' to a corner line $\mu = \sigma \cup \sigma'$ pointing down and left that has a corner on the horizontal bar line of λ' . Removing λ_2 and λ' and introducing μ instead will yield an m -cut with a smaller cut length and one boundary point less on δ .

Since we assumed that there is at most one corner line pointing down and right in K , if λ_1 is a straight line and λ_2 is such a corner line then $|K| = 2$. Hence this case is analogous to the case just covered.

Now consider the case when both λ_1 and λ_2 are straight lines (Figure 26). In case they overlap we can simply remove both lines. Otherwise it holds that $\mathcal{D}(\lambda_1) \subset \mathcal{D}(\lambda_2)$. As above, any curve from L that intersects $\mathcal{D}(\lambda_1)$ must be a corner line pointing up and left. Let $L' \subseteq L$ be the curves that intersect $\mathcal{D}(\lambda_1)$, and if $L' \neq \emptyset$ let $\lambda' \in L'$ be the one with the lowest and right-most corner among these. Notice that λ' is well-defined since the curves in L' are non-crossing. In this case we replace both λ_1 and λ' with $\Lambda_1(x, y)$, where (x, y) is the boundary point of the horizontal bar line of λ' . In case L' is empty we replace λ_1 with $\Lambda_1(x_\delta, y_{\lambda_1})$. In both cases we obtain an m' -cut for some m' . Analogous to the




Figure 26: The case when λ_1 and λ_2 are straight lines (left). Together with some curves in L' (dotted) they are substituted with $\Lambda_1(x, y)$ and μ (right). The size of the cut out area (diagonally striped) remains the same. The black contiguous lines in the pictures are curves from U_G .

case when λ_1 is a corner line pointing down and left, and λ_2 is a straight line, we can find a curve with which to replace λ_2 . As above we possibly also need to remove some other curve in L' to yield an m -cut of smaller cut length than l , and in which there is one boundary point less on δ .

The only case left is the one where both λ_1 and λ_2 are corner lines, i.e. the former points down and left, the latter down and right, and $|K| = 2$. We assume w.l.o.g. that the vertical length of λ_1 is at most that of λ_2 . If λ_1 and λ_2 have the same boundary point on δ and they have the same vertical length, obviously we can remove these two curves and introduce a straight line that consists of the horizontal bars of λ_1 and λ_2 instead, and thereby obtain an m -cut with smaller cut length than l and with two boundary points less on δ . Consider the case when the two curves share the same boundary point on δ , they have different vertical lengths, and there is a corner line $\lambda' \in L$ pointing up and left having the same corner as λ_1 . Then we can remove λ_1 , λ_2 , and λ' and introduce the corner line pointing up and right that has the same corner as λ_2 . We thereby yield an m -cut of smaller cut length in which two boundary points on δ are removed. All remaining cases are handled in the following.

Let $(x, y) \notin \delta$ be the boundary point of λ_1 that does not lie on δ . Replacing λ_1 with $\Lambda_1(x, y)$ results in an m' -cut for some m' , as in the case when λ_2 is a straight line (Figure 27). We define $\mathcal{L} \subseteq \mathcal{D}(\lambda_2)$ analogous to that case, i.e. \mathcal{L} is the connected area to the left of λ_2 in the m' -cut. Furthermore let σ be the vertical straight line in \mathcal{L} that cuts out an area the size of $\mathcal{D}(\lambda_1)$ on its right-hand side. Notice that $\sigma \subset \mathcal{L}$ is well-defined since above we excluded all cases where $\mathcal{D}(\lambda_1) = \mathcal{L}$. In case there is a corner line μ pointing down and right that has σ as its vertical bar line and overlaps with λ_2 , we can replace λ_2 with μ and obtain an m -cut that has a cut length of at most l since the vertical length of λ_1 is at most that of λ_2 . Also the number of boundary points on δ is reduced by one. Otherwise, similar to the case when λ_2 is a straight line, the boundary point of σ (w.r.t. \mathcal{L}) that does not lie on δ is either part of the boundary of




Figure 27: The case when λ_1 and λ_2 are corner lines pointing in different directions (left). The area shaded in grey is \mathcal{L} . If σ (dashed and dotted) is the vertical bar line of μ , λ_1 and λ_2 can be substituted with $\Lambda_1(x, y)$ and μ (top). If σ ends at a line λ' (dotted), the lines λ_1 , λ_2 , and λ' can be substituted with $\Lambda_1(x, y)$, μ_1 , and μ_2 (bottom). The size of the cut out area (diagonally striped) remains the same. The black contiguous lines in the pictures are curves from U_G .

\mathcal{P}_G , or it lies on a curve $\lambda' \in L$ that must be a corner line pointing up and left. In the former case we replace λ_2 with σ and yield an m -cut with the desired properties. In the latter case there are two corner lines μ_1 and μ_2 with the following properties. The first line points down and left, it has σ as its vertical bar line, and its horizontal bar line overlaps with λ' . The second line points up and right, has a vertical bar line that is part of λ' , and a horizontal bar line that includes the horizontal bar line of λ_2 . We can then replace λ_2 and λ' with μ_1 and μ_2 and yield an m -cut with cut length at most l , since the vertical length of λ_1 is at most that of λ_2 . Also the number of boundary points on δ is reduced by one.

Notice that in all transformations above the number of boundary points on δ is reduced and at the same time the number of boundary points on other curves in U_G is never increased. We can hence repeat the above procedure for the curves in K and then in the same manner for all curves in U_G that include more than one boundary point of curves in the resulting m -cut. We yield an m -cut that has a smaller cut length than l , and for which any curve in U_G includes at most one boundary point of a curve in the m -cut. \square

As long as there is more than one non-grid line (now in different open unit intervals on the boundary), we can shift one of them to become a grid line, and shift the other one accordingly to compensate for the area difference. This results in a situation with at most one non-grid line in the cut. During the whole process, the cut length does not increase, as the next lemma shows.

Lemma 32. *For any grid G and any non-crossing corner m -cut L of cut length l in \mathcal{P}_G , there is a non-crossing corner m -cut M of cut length at most l in \mathcal{P}_G such that all curves in M except at most one are grid lines.*

Proof. According to Lemma 31 we can assume that L contains no two curves that have boundary points that lie on the same curve from U_G . Let $K \subseteq L$ be the set of curves that are not grid lines and assume that $|K| \geq 2$. For a straight line $\lambda \in K$, any curve from K that intersects the corridor \mathcal{K}_λ must have a boundary point on the same curve from U_G as λ . Hence no set of curves in L that have intersecting corridors include straight lines. For a corner line $\lambda \in K$, let $\mathcal{C} \in \{\mathcal{A}(\{\lambda\}), \mathcal{B}(\{\lambda\})\}$ be the area to which λ is convex and let $\bar{\mathcal{C}} = \{\mathcal{A}(\{\lambda\}), \mathcal{B}(\{\lambda\})\} \setminus \{\mathcal{C}\}$ be the area to which λ is concave. We define $\mathcal{X}(\lambda) = \mathcal{K}_\lambda \cap \mathcal{C}$ to be the part of λ 's corridor to which λ is convex and $\bar{\mathcal{X}}(\lambda) = \mathcal{K}_\lambda \cap \bar{\mathcal{C}}$ to be the part to which λ is concave. Any curve in K that intersects $\mathcal{X}(\lambda)$ must have a boundary point on the same curve in U_G as λ . Therefore no curve in L intersects with $\mathcal{X}(\lambda)$.

We can thus conclude that if $\lambda_1, \lambda_2 \in K$ is a pair of curves with intersecting corridors, then λ_1 and λ_2 must be corner lines and λ_i , for $i \in \{1, 2\}$, must intersect $\bar{\mathcal{X}}(\lambda_j)$, where $j \in \{1, 2\} \setminus \{i\}$. Assume w.l.o.g. that λ_1 points down and left. Observe that this means that λ_2 points up and right and no other curve in K can intersect the corridors of λ_1 or λ_2 , since otherwise there would be curves in L that have boundary points on the same curve from U_G . Notice that the corridors of λ_1 and λ_2 intersecting means that the corridors of the horizontal bar lines of λ_1 and λ_2 or the corridors of the corresponding vertical bar lines




Figure 28: A virtual pseudo-corner line with its corner at (x, y) and its unit sized step at \tilde{x} .

intersect. Assume w.l.o.g. that the vertical bar lines σ_1 and σ_2 of λ_1 and λ_2 , respectively, are not grid lines and that the length l_1 of σ_1 is at most the length l_2 of σ_2 . As in the proof of Lemma 31 we define $\mathcal{D}(\lambda_i)$, for both $i \in \{1, 2\}$, to be the open set of points to the left of λ_i , i.e. the height of $\mathcal{D}(\lambda_i)$ equals l_i and $\mathcal{D}(\lambda_1) \subseteq \mathcal{X}(\lambda_1)$ but $\mathcal{D}(\lambda_2) \subseteq \overline{\mathcal{X}(\lambda_2)}$. In this setting we assume w.l.o.g. that the size of $\mathcal{D}(\lambda_1)$ is at most the size of $\mathcal{D}(\lambda_2)$. We thus replace λ_1 with $\Lambda_1(x, y)$, as defined in Lemma 31, where (x, y) is the boundary point of the horizontal bar line of λ_1 , and yield a non-crossing m' -cut for some m' . If $l_1 < l_2$ there exists a corner line λ'_2 pointing up and right that contains the horizontal bar line of λ_2 and intersects $\mathcal{D}(\lambda_2)$, such that replacing λ_2 with λ'_2 yields a non-crossing m -cut. If $l_1 = l_2$ we can find an according virtual corner line that contains the horizontal bar line of λ_2 and intersects the boundary of $\mathcal{D}(\lambda_2)$. Notice that the new m -cut has a cut length of at most l since $l_1 \leq l_2$ and hence σ_1 was “moved” farther to the left than (or equally far as) σ_2 . Also note that the new cut has at least one curve less in U_G containing a boundary point since the vertical bar lines in $\Lambda_1(x, y)$ are grid lines.

For any two curves in K with non-intersecting corridors we can use an analogous transformation as above. Since each transformation yields an m -cut in which there is at least one curve less in U_G with a boundary point, by repeating the above procedure we can transform L into the m -cut M with the desired properties. \square

From now on, we can limit ourselves to the situation with only one non-grid line in the polygon cut. We shift this line to the nearest integer position (Figure 20), creating the need to compensate for the area difference. We do this by introducing more grid lines. But since this increases the cut length, we need to prove that the extra grid lines we introduce are short. In the end, this will preserve the property that the cut out area lies in the interval defined by m and ε , but will increase the cut length only by a small factor. Next, we will look at a way to cut out for compensation, and then argue that there is a place from which to cut out in this way.

We manage to compensate in a recursive manner. We compensate for an area difference a by first finding a particular way to cut out an area guaranteed to

be between a and $3a/2$, with the exact value not under our control. This leaves us with the problem to compensate for at most half the previous area (since we are at most $a/2$ away from a). A recursive repetition of this compensation step ends after at most $\log(a)$ steps. The particular way to cut out the area between a and $3a/2$ makes use of a staircase grid line of three consecutive bends, with a step of unit height at the middle bend (Figure 28). Furthermore, the middle bend is guaranteed to lie outside or on the boundary of the polygon, so that the intersection of the staircase with the polygon results in a set of corner and straight lines in the cut. We call this a *virtual pseudo-corner line*. The analysis of the recursion reveals that the total length of the additional curves to cut out area a is limited to $3a$.

Definition 33 (virtual pseudo-corner line). For any polygon \mathcal{P}_G of a grid G a *virtual pseudo-corner line* is a set of grid lines Λ in \mathcal{P}_G containing only straight and corner lines for which there are two points (x, y) and $(\tilde{x}, y - 1)$, where $\tilde{x} \geq x$ and $(\tilde{x}, y) \notin \mathcal{P}_G$, such that $\lambda \in \Lambda$ if and only if

$$\begin{aligned} \lambda \subseteq & \{(x, y') \in \mathcal{P}_G \mid y' \geq y\} \cup \\ & \{(x', y) \in \mathcal{P}_G \mid x' \in [x, \tilde{x}]\} \cup \\ & \{(\tilde{x}, y') \in \mathcal{P}_G \mid y' \in [y - 1, y]\} \cup \\ & \{(x', y - 1) \in \mathcal{P}_G \mid x' \geq \tilde{x}\}. \end{aligned}$$

We call the unit step $\{(\tilde{x}, y') \in \mathbb{R}^2 \mid y' \in [y - 1, y]\}$ the *break*, and (x, y) the *corner* of Λ . The length of Λ is the sum of the lengths of the included straight and corner lines. If Λ cuts out an area of size a on the upper right side of its corner, we say that it is a virtual pseudo-corner line *for* a .

A virtual pseudo-corner line is a special kind of virtual staircase line containing only grid lines. Notice that a virtual corner line containing only grid lines is a virtual pseudo-corner line. This is because the break of a virtual pseudo-corner line can entirely lie outside of the polygon.

We first convince ourselves that the needed virtual pseudo-corner line exists. In case there is a virtual corner line that cuts out the required area and contains only grid lines we are done. In the other case a suitable set of curves can be constructed using three virtual corner lines at some integer points (x^*, y^*) , $(x^* + 1, y^*)$, and $(x^*, y^* + 1)$ (Figure 29). These three virtual corner lines are chosen such that the first one cuts out an area larger than $3a/2$, while the other two each cut out at most $a - 1$. Using these properties it is then possible to show that there must be a unit sized step, i.e. a break, between the virtual corner lines at (x^*, y^*) and $(x^*, y^* + 1)$ with which a suitable virtual pseudo corner line can be constructed. That is, the corresponding set of curves cuts out an area between a and $3a/2$, and the upper most point of the break is on the boundary or outside of the polygon.

Lemma 34. *For any grid G with n vertices and any $b \in \{0, \dots, n\}$, there is a value $a \in [b, \frac{3}{2}b]$ for which there exists a virtual pseudo-corner line Λ for a in \mathcal{P}_G .*




Figure 29: The virtual corner line at (x^*, y^*) cuts out an area of size more than $\frac{3}{2}b$ while those at $(x^* + 1, y^*)$ and $(x^*, y^* + 1)$ cut out at most $b - 1$. If all points $(x^* + i, y^* + 1)$, where $i \in \{l, \dots, r\}$, are inside the polygon, then all unit squares from A are also inside. Since there are $r - l + 1 > b/2$ of them, the area $\mathcal{P}(x^* + 1, y^* + 1)$ has size greater than $b/2$.

Proof. Since the vertices of the grid G are points with integer coordinates, i.e. $V \subset \mathbb{N}^2$, a virtual corner line contains only grid lines if its corner is a point in the set \mathbb{H}^2 . If there is a virtual corner line for some $a \in [b, \frac{3}{2}b]$ with a corner in \mathbb{H}^2 then the lemma holds. Assume no such virtual corner line exists. Since any set of grid lines cuts out an area of integer size, this means that any virtual corner line with a corner from \mathbb{H}^2 either cuts out an area of size at least $\lceil \frac{3}{2} \rceil b$ or at most $b - 1$.

Let $\Lambda(p)$ denote the virtual corner line with corner $p \in \mathbb{H}^2$ and let $\mathcal{P}(p)$ denote the area cut out by $\Lambda(p)$ on the upper right side of p . Under the above assumption, clearly there must be a point $(x, y) \in \mathbb{H}^2$ such that the size of $\mathcal{P}(x, y)$ is greater than $\frac{3}{2}b$ since $b \leq n$, and obviously there is a point $(x', y') \in \mathbb{H}^2$ with $x' \geq x$ and $y' \geq y$ such that $\mathcal{P}(x', y') = \emptyset$. Because the area $\mathcal{P}(p)$ for any $p \in \mathbb{H}^2$ includes any area $\mathcal{P}(q)$ of a corner q above or to the right of p , the size of $\mathcal{P}(p)$ is monotonically decreasing in both coordinates of p . Hence we can find a point $(x^*, y^*) \in \mathbb{H}^2$ with $x^* \in [x, x']$ and $y^* \in [y, y']$ such that the size of $\mathcal{P}(x^*, y^*)$ is at least $\lceil \frac{3}{2} \rceil b$ while the size of both $\mathcal{P}(x^* + 1, y^*)$ and $\mathcal{P}(x^*, y^* + 1)$ are at most $b - 1$.

Let $\mathcal{P}_{ij} = \mathcal{P}(x^* + i, y^* + j)$ and $\Lambda_{ij} = \Lambda(x^* + i, y^* + j)$ for $i, j \in \mathbb{N}_0$. The area $\mathcal{P}_{00} \setminus \mathcal{P}_{01}$ has height 1 and contains a series of unit squares. For any $x \in \mathbb{N}$ the difference between the area \mathcal{P}_{00} and $\mathcal{P}_{01} \cup \mathcal{P}_{x0}$ includes only unit squares from $\mathcal{P}_{00} \setminus \mathcal{P}_{01}$. Hence the above bounds on the sizes of \mathcal{P}_{00} and \mathcal{P}_{01} mean that we can find two integers $l, r \in \mathbb{N}$ such that $l \leq r$ and the size of $\mathcal{P}_{01} \cup \mathcal{P}_{l0}$ equals $\lceil \frac{3}{2} \rceil b$ and the size of $\mathcal{P}_{01} \cup \mathcal{P}_{r0}$ equals b . If for a value $i \in \{l, \dots, r\}$ there is a pair of crossing curves in $\Lambda_{01} \cup \Lambda_{i0}$, their crossing point is $p_{i1} = (x^* + i, y^* + 1)$. If however there exists a corresponding value for i such that there are no curves in $\Lambda_{01} \cup \Lambda_{i0}$ that cross, then let Λ include the curves to the left of p_{i1} from the first set together with the curves below p_{i1} from the second set, i.e.

$$\Lambda = \{\lambda \in \Lambda_{01} \mid \forall (x, y) \in \lambda : x < x^* + i\} \cup \{\lambda \in \Lambda_{i0} \mid \forall (x, y) \in \lambda : y < y^* + 1\}.$$

Clearly the set Λ fulfils the lemma. Hence it remains to show that we can always

find a corresponding value i such that p_{i1} is not in \mathcal{P}_G .

Assume this is not the case, i.e. for any value $i \in \{l, \dots, r\}$ it holds that $p_{i1} \in \mathcal{P}_G$. This means that any unit square that has one of these points as a corner must be included in \mathcal{P}_G . Let A be the set of unit squares in \mathcal{P}_{11} that have such a point p_{i1} as their lower left corner. There are $r - l + 1 \geq \lfloor \frac{3}{2}b \rfloor - b + 1 > \frac{1}{2}b$ many points p_{i1} . We can conclude that there are at least $\frac{1}{2}b$ many unit squares in A . Since the squares have unit size and are included in \mathcal{P}_{11} the size of \mathcal{P}_{11} is at least $\frac{1}{2}b$.

Let us derive an upper bound on the size of the area $\mathcal{P}_{11} = \mathcal{P}_{10} \cap \mathcal{P}_{01}$. Since $\mathcal{P}_{10} \subseteq \mathcal{P}_{00}$ the size of $\mathcal{P}_{00} \setminus \mathcal{P}_{10}$ is at least $\frac{1}{2}b + 1$. The difference between the area $\mathcal{P}_{00} \setminus \mathcal{P}_{10}$ and $\mathcal{P}_{01} \setminus \mathcal{P}_{10}$ can at most include the unit square \mathcal{S}_v where $v = (x^* + \frac{1}{2}, y^* + \frac{1}{2})$. Whether $v \in V$ or not, this means that the size of the area $\mathcal{P}_{01} \setminus \mathcal{P}_{10}$ is at least $\frac{1}{2}b$. Since $\mathcal{P}_{01} \cap \mathcal{P}_{10} = \mathcal{P}_{01} \setminus (\mathcal{P}_{01} \setminus \mathcal{P}_{10})$, we can conclude that the size of \mathcal{P}_{11} is at most $\frac{1}{2}b - 1$. However this contradicts the lower bound derived above and hence the lemma holds. \square

Using the above lemma we can show that an area of arbitrary size can be cut out recursively as described before.

Lemma 35. *For a grid G let Λ be a virtual corner line for b in \mathcal{P}_G that contains only grid lines. If \mathcal{P} denotes the area cut out by Λ on the upper right side of its corner, then for any $a \in \{0, \dots, b\}$ there exists a set of non-crossing corner grid lines L in \mathcal{P}_G cutting out an area of size a from \mathcal{P} . Furthermore, the curves in $\Lambda \cup L$ are non-crossing and the cut length of L is at most $3a$.*

Proof. Let $a_1 = a$ and G_1 be the grid corresponding to the area \mathcal{P} . Consider the following recursive procedure. In each step $i \geq 1$ we attempt to cut out an area of size a_i from \mathcal{P}_{G_i} using only grid lines. According to Lemma 34 we can find a virtual pseudo-corner line Λ'_i in \mathcal{P}_{G_i} for some $a'_i \in [a_i, \frac{3}{2}a_i]$ with the properties listed therein. We need to transform the curves in Λ'_i into a valid virtual pseudo-corner line Λ_i in \mathcal{P}_G that cuts out the same area as Λ'_i . Assume for now that this can be done. We will describe the transformation later. If $a'_i = a_i$ the recursion terminates. Otherwise let $a_{i+1} = a'_i - a_i$ and let G_{i+1} be the grid that corresponds to the area $\mathcal{A}(\Lambda_i)$ of the a'_i -cut Λ_i , i.e. $\mathcal{P}_{G_{i+1}} = \mathcal{A}(\Lambda_i)$. From $a'_{i+1} \leq \frac{3}{2}a_{i+1} = \frac{3}{2}(a'_i - a_i)$ and $a_i \geq \frac{2}{3}a'_i$ we can conclude that $a'_{i+1} \leq \frac{1}{2}a'_i$, i.e. the area $\mathcal{A}(\Lambda_i)$ that is cut out from \mathcal{P}_{G_i} is smaller than \mathcal{P}_{G_i} . If this procedure terminates the set $L = \bigcup_{i \geq 1} \Lambda_i$ clearly cuts out an area of size exactly a . Since any set Λ_i contains only grid lines, the size a'_i of the cut out area must be integer. By the fact that the cut out area in step $i + 1$ has a size at most half the size of the cut out area in step i , this means that the procedure terminates after at most $\lceil \log_2(a) \rceil$ steps.

Since Λ_i contains only grid lines, the area \mathcal{P}_{G_i} that is cut out by Λ_i can be decomposed into a'_i unit squares. The set Λ_i contains at most two corner lines and therefore each, except at most two of the a'_i unit squares, has at most one side of its boundary coinciding with a curve in Λ_i . There may be two unit squares that each have two sides of their boundaries coincide with a corner line in Λ_i . Hence the length l_i of Λ_i is at most $a'_i + 2$. From $a'_{i+1} \leq \frac{1}{2}a'_i$ we can




Figure 30: A tail with its corner line at p (black dotted). The excess area in grey is cut out using the four virtual corner lines at p (thin dashed) together with the recursive method that uses virtual pseudo-corner lines (dashed and dotted).

conclude that $a'_i \leq a/2^i$ which means that $l_i \leq a/2^i + 2$. Therefore the cut length l of L is

$$l = \sum_{i \geq 0} l_i \leq \sum_{i=1}^{\lfloor \log_2(a) \rfloor} \left(\frac{a}{2^i} + 2 \right) \leq a \left(2 - \frac{2}{a} \right) + 2 \log_2(a) \leq 3a,$$

where the last inequality holds since $2 \log_2(a) - 2 < a$ for any $a > 0$. If $a = 0$ then $l = 0$ and the claimed bound still holds.

What remains to be shown is that we can convert the curve sets Λ'_i into valid virtual pseudo-corner lines Λ_i in \mathcal{P}_G . Since Λ is a virtual corner line containing only grid lines, it is also a virtual pseudo-corner line. We let $\Lambda_0 = \Lambda$ and then show by induction that each Λ'_i can be transformed into an appropriate Λ_i for $i \geq 1$. Assume that Λ_i is a virtual pseudo-corner line that cuts out the same area $\mathcal{A}(\Lambda_i)$ in \mathcal{P}_G as Λ'_i does in \mathcal{P}_{G_i} . The set Λ'_{i+1} is an a'_{i+1} -cut in $\mathcal{P}_{G_{i+1}} = \mathcal{A}(\Lambda_i)$ that cuts out the area $\mathcal{A}(\Lambda'_{i+1}) \subseteq \mathcal{P}_{G_{i+1}}$. If β denotes the boundary of \mathcal{P}_G and γ denotes the boundary of $\mathcal{A}(\Lambda'_{i+1})$, then we include all segment curves $\lambda \subseteq \gamma \setminus \beta$ in Λ_{i+1} and claim that it is a virtual pseudo-corner line. If it is then it clearly cuts out the same area as Λ'_{i+1} . The point set $\gamma \setminus \beta$ may contain parts of curves from Λ_i and Λ'_{i+1} . However, since these sets contain only grid lines and the length of a break is 1, $\gamma \setminus \beta$ can contain at most one break from Λ_i and Λ'_{i+1} . It is easy to see that this means that Λ_{i+1} is a virtual pseudo-corner line. \square

It remains to be shown that there is a place in the polygon to cut out from using the recursive method above. For this we use a tail of the cut (Figure 30), similar to the staircase line argument in the previous section. We have to make sure that there is a tail that is big enough to support an area of size a . For a non-crossing corner cut L containing only one curve λ that is not a grid line we call a tail $\mathcal{T} \subseteq \mathcal{A}(L)$ (respectively $\mathcal{T} \subseteq \mathcal{B}(L)$) *tiny* if the size of \mathcal{T} is strictly smaller than the size of $\mathcal{K}_\lambda \cap \mathcal{B}(L)$ (respectively $\mathcal{K}_\lambda \cap \mathcal{A}(L)$). In the following we give a similar observation on such tails as was given for the case when they are small.

Lemma 36. *For a grid G , let L be a non-crossing corner m -cut in \mathcal{P}_G with cut length l containing exactly one curve λ that is not a grid line. There exists a*

non-crossing corner m -cut M in \mathcal{P}_G which contains exactly one curve that is not a grid line, has cut length of at most $l + 1$, and the curve of any tiny tail cut out by M equals the curve that is not a grid line.

Proof. The proof of this lemma is analogous to the proof of Lemma 26. However the non-grid line λ may get longer when the area of a tail is transferred to the corridor of λ . This can only happen if λ is a corner line and then its length grows by at most 2. Since the curve of any tail is a grid line it has a length of at least 1. Hence the total cut length grows by at most 1. \square

We need to make sure that no additional curves are produced while cutting out the area of size a from a tail which would increase the cut length by some non-constant factor. For this we break the tail into four sectors using four virtual corner lines having the same corner as the curve of the tail. We then greedily assign these virtual corner lines to the cut as long as the cut out area does not exceed a . The remaining difference to reach the desired area a is finally cut out using the recursive method presented above from one of the four sectors that was not yet used.

Lemma 37. *For a grid G , let L be a set of grid lines in the polygon \mathcal{P}_G and let \mathcal{T} be a tail cut out by L . If b denotes the size of \mathcal{T} , then for any $a \in \{0, \dots, b\}$ there exists a set of non-crossing corner grid lines M in \mathcal{P}_G cutting out an area of size a from \mathcal{T} such that the curves in $M \cup L$ are non-crossing. Furthermore, the cut length of M is at most $3a$.*

Proof. Let $\lambda \in L$ be the curve of \mathcal{T} . If λ is a straight line then let p be one of its boundary points, and if λ is a corner line let p be its corner. There are four virtual corner lines in \mathcal{T} having p as their corner, one for each possible orientation. These virtual corner lines Λ_1 to Λ_4 partition \mathcal{T} into four (possibly empty) areas \mathcal{T}_1 to \mathcal{T}_4 such that \mathcal{T}_i is cut out by Λ_i , where $i \in \{1, 2, 3, 4\}$, on the “convex side” of its corner. Let $I \subseteq \{1, 2, 3, 4\}$ be the set for which $i \in I$ if and only if $\mathcal{T}_i \neq \emptyset$. If a equals 0 or b then the lemma obviously holds. Assume that $0 < a < b$. There exists a (possibly empty) subset $J \subset I$ such that the size b_J of the union area $\mathcal{T}_J = \bigcup_{i \in J} \mathcal{T}_i$ is at most a while for any $j \in I \setminus J$ the size of the area $\mathcal{T}_J \cup \mathcal{T}_j$ is greater than a . For each $i \in J$ the set M contains the curves in Λ_i . Notice that λ can not be included in any of the sets Λ_i since the latter are virtual corner lines in the open set of points \mathcal{T} . Hence, in case the boundary of \mathcal{T}_J includes λ we also include λ in M . So far these curves cut out an area of size b_J from \mathcal{P}_G .

Since all involved curves are grid lines, if b_i , for some $i \in I$, denotes the size of the area \mathcal{T}_i , we can decompose \mathcal{T}_i into b_i many unit squares. The set Λ_i contains at most one corner line and therefore each except at most one of the b_i unit squares has at most one side of its boundary coinciding with a curve in Λ_i . There may be one unit square that has two coinciding sides of its boundary with the corner line in Λ_i . Hence the length l_i of Λ_i is at most $b_i + 1$ and therefore the cut length of $\bigcup_{i \in J} \Lambda_i$ is at most $b_J + |J|$. Note that the same bound holds for the curves included in M so far, even if $\lambda \in M$.

Let $j \in I \setminus J$. According to Lemmas 35 and 34 we can find a set M' of non-crossing corner grid lines in \mathcal{P}_G that cut out an area of size $a - b_j$ from \mathcal{T}_j such that the cut length of M' is at most $3(a - b_j)$. If we also include M' in M , we cut out an area of size a from \mathcal{T} without crossing a curve in L . Furthermore the cut length of M is at most

$$b_j + |J| + 3(a - b_j) = 3a + |J| - 2b_j \leq 3a,$$

where the inequality holds since $\mathcal{T}_i \neq \emptyset$ and hence $b_i \geq 1$ for each $i \in J$. Thus the set M fulfils the required properties. \square

The main result as stated in Theorem 1, follows from the next theorem which summarises the results of this section.

Theorem 38. *Let l be the cut length of an optimal m -cut L , for some $m \in \mathbb{N}$, in the polygon \mathcal{P}_G of a grid G . For any $\varepsilon \in]0, 1]$ there exists a non-crossing corner m' -cut L' for some $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$, such that all curves in L' are grid lines and the cut length is at most $(216\sqrt{7/\varepsilon} + 261) \cdot l$.*

Proof. We can apply Corollary 28 and Lemma 32 to L , i.e. we know that there exists a non-crossing corner m' -cut M , for some $m' \in [(1 - \varepsilon)m, (1 + \varepsilon)m]$, with cut length at most $(54\sqrt{7/\varepsilon} + 63) \cdot l$ in \mathcal{P}_G such that M contains at most one curve that is not a grid line. If all curves in M are grid lines, we are done. If not then let $\lambda \in M$ be the curve that is not a grid line. In case there exists a tail cut out by M such that λ is not its curve, let \mathcal{T} be this tail and assume w.l.o.g. that $\mathcal{T} \subseteq \mathcal{A}(M)$. By Lemma 36 we can assume that the size of \mathcal{T} is at least the size of $\mathcal{K}_\lambda \cap \mathcal{B}(M)$ if we allow the cut length of M to be at most $(54\sqrt{7/\varepsilon} + 63) \cdot l + 1$. Notice that, if a denotes the size of $\mathcal{K}_\lambda \cap \mathcal{B}(M)$, a is not necessarily an integer since m' might not be a natural number. However we can conclude that the size of \mathcal{T} is at least $\lceil a \rceil$ since the curve of \mathcal{T} is a grid line and hence \mathcal{T} is of integer size. Let β denote the boundary of \mathcal{P}_G and γ the boundary of $\mathcal{K}_\lambda \cap \mathcal{B}(M)$. In case λ is a corner line and contains a bar line σ that is a grid line, let Λ contain all straight and corner lines that are contained in the set $(\gamma \setminus (\beta \cup \lambda)) \cup \sigma$. In any other case let Λ contain the straight and corner lines in the set $\gamma \setminus (\beta \cup \lambda)$. By replacing the curve λ with the curves in Λ we yield an m'' -cut M' where $m'' = m' + a$.

We attempt to cut out the excess area of size a in \mathcal{T} using only grid lines. Notice that m'' must be an integer since M' contains only grid lines. If we assume w.l.o.g. that $m' \geq m$, since m is also an integer this means that $m'' - \lceil a \rceil$ is a natural number in the interval $[m, m']$. The latter is contained in $[(1 - \varepsilon)m, (1 + \varepsilon)m]$. Using Lemma 37 we can find a set of grid lines M'' that cut out an area of size $\lceil a \rceil$ from \mathcal{T} . The union $M' \cup M''$ forms a non-crossing set of grid lines cutting out an area from the interval $[(1 - \varepsilon)m, (1 + \varepsilon)m]$. Hence it remains to show (for the case when λ is not the curve of \mathcal{T}) that the cut length of $L' = M' \cup M''$ is bounded from above as claimed in the theorem.

Since λ is a straight or corner line, the size of the corridor of λ is at most the length of λ plus 1. Since the length of λ is upper-bounded by the cut length

l' of M we can conclude that $\lceil a \rceil \leq a + 1 \leq l' + 2$. By Lemma 37 this means that the cut length of M'' is upper-bounded by $3(l' + 2)$. Clearly the length of Λ can be at most the length of λ plus 2. Hence also the cut length of M' is at most the cut length of M plus 2. Therefore the cut length of L' is at most $l' + 2 + 3(l' + 2) \leq (216\sqrt{7/\varepsilon} + 252) \cdot l + 9$. Cutting out an integer sized area greater than zero (and smaller than n) from the polygon \mathcal{P}_G , i.e. a polygon constructed from unit squares, will need a cut length l of at least 1. In this case the latter bound on the cut length of L' can be upper-bounded by the claimed bound of the theorem. If none (or all) of the area is to be cut out from \mathcal{P}_G , the trivial empty cut obviously also fulfils the requirements of this theorem.

Now consider the case when there is no tail such that λ is not its curve. This can only mean that there are two tails which both have λ as their curve and λ is the only curve in M . Let \mathcal{T} be the tail that corresponds to the area $\mathcal{A}(M)$. Replacing λ with Λ as before, we obtain an m'' -cut M' for which $\mathcal{A}(M') = \mathcal{A}(M) \cup (\mathcal{K}_\lambda \cap \mathcal{B}(M))$. Hence the size of $\mathcal{A}(M')$ is at least a . Since Λ may contain more than one curve, $\mathcal{A}(M')$ might not be a tail. Nevertheless, proving an analogous statement as Lemma 37 for this case we can come to the same conclusions as above. This is due to the fact that Λ is a virtual corner line, which conclude the proof. \square

6. Conclusions

We have seen that when restricting ourselves to simple shaped cuts in solid grid graphs, it is possible to cut out a number of vertices close to the desired number m , while not losing a lot in terms of the quality of the cut length. We proved this fact by considering polygons for which a similar result is true. The corresponding result for polygons (Corollary 28) is of independent interest and might be considered for further research on polygons in the future. For solid grid graphs it was already possible to use the obtained results [2] in order to speed up an algorithm computing sparsest cuts [8]. The latter insight can subsequently be put to work [2] in order to gain faster approximation algorithms for separators and bisections for solid grid graphs, using known techniques [3, 8].

One remaining question is whether the approximation guarantee given in Theorem 38 for the corner cuts can be improved. In particular it is not clear whether the factor ε , by which the size of the cut-out part deviates from the given value m , is necessary. Also the final constant given by Theorem 38, which has a value of at least 832, seems very large. The reason for this large value is that in many of the lemmas leading to the theorem, the cut length of the involved curves grow by a constant factor. This means that the resulting constant grows exponentially with the number of intermediate steps used by the proof. Hence an improvement on the guaranteed approximation ratio may be achievable with a more direct approach than the one chosen here. In particular the best lower bound we can provide to compare optimal corner cuts with optimum m -cuts is $1 + 1/\sqrt{2}$ (Figure 31). Interestingly the lower bound example is a very simple one. There also exist more complicated examples based on the insights gained in this article. For instance it is possible to construct examples where the optimum




Figure 31: A polygon in which the optimum corner $n/2$ -cut (dotted lines) has a cut length that is a factor of $1 + 1/\sqrt{2}$ larger than the optimum $n/2$ -cut (dashed line). Obviously this gives a lower bound for the corresponding grid graph.

corner m -cut needs three segments. For this, topologies such as the one shown in Figure 12 can be used. However in all found examples the corner m -cut with minimum cut length was also at most a factor $1 + 1/\sqrt{2}$ away from optimum.

References

- [1] A. E. Feldmann, S. Das, P. Widmayer, Restricted cuts for bisections in solid grids: A proof via polygons, in: Proceedings of the 37th International Workshop on Graph-Theoretic Concepts in Computer Science (WG), 2011, pp. 143–154.
- [2] A. E. Feldmann, Balanced partitioning of grids and related graphs: A theoretical study of data distribution in parallel finite element model simulations, Ph.D. thesis, ETH Zurich, diss.-Nr. ETH: 20371 (April 2012).
- [3] H. D. Simon, S. H. Teng, How good is recursive bisection?, SIAM Journal on Scientific Computing 18 (5) (1997) 1436–1445. doi:<http://dx.doi.org/10.1137/S1064827593255135>.
- [4] C. Papadimitriou, M. Sideri, [The bisection width of grid graphs](#), Theory of Computing Systems 29 (1996) 97–110. URL <http://dx.doi.org/10.1007/BF01305310>
- [5] X. Cheng, D.-Z. Du, J.-M. Kim, L. Ruan, [Optimal rectangular partitions](#), in: Handbook of Combinatorial Optimization, Springer US, 2005, pp. 313–327. URL http://dx.doi.org/10.1007/0-387-23830-1_7
- [6] C. Kenyon, E. Rémila, [A near-optimal solution to a two-dimensional cutting stock problem](#), Mathematics of Operations Research 25 (2000) 645–656. doi:[10.1287/moor.25.4.645.12118](https://doi.org/10.1287/moor.25.4.645.12118). URL <http://dl.acm.org/citation.cfm?id=366078.366094>

- [7] A. E. Feldmann, P. Widmayer, An $O(n^4)$ time algorithm to compute the bisection width of solid grid graphs, in: Proceedings of the 19th Annual European Symposium on Algorithms (ESA), 2011, pp. 143–154.
- [8] T. Leighton, S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, Journal of the ACM 46 (6) (1999) 787–832. doi:<http://doi.acm.org/10.1145/331524.331526>.
- [9] G. Strang, Maximum area with Minkowski measures of perimeter, in: Proceedings of the Royal Society of Edinburgh: Section A Mathematics, Vol. 138, 2008, pp. 189–199.