

Restricted cuts for bisections in solid grids

A proof via polygons

Report

Author(s):

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

Publication date:

2011

Permanent link:

<https://doi.org/10.3929/ethz-a-006906296>

Rights / license:

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

Originally published in:

Technical report 731

Restricted Cuts for Bisections in Solid Grids: A Proof via Polygons^{*†}

Andreas Emil Feldmann¹, Shantanu Das², and Peter Widmayer¹

¹Institute of Theoretical Computer Science, ETH Zürich, Switzerland
`{feldmann,widmayer}@inf.ethz.ch`

²Laboratoire d'Informatique Fondamentale, Aix-Marseille University, France
`shantanu.das@lif.univ-mrs.fr`

July 18, 2011

Abstract

The *graph bisection problem* asks to partition the n vertices of a graph into two sets of equal size so that the number of edges across the cut is minimum. We study finite, connected subgraphs of the infinite two-dimensional grid that do not have holes. Since bisection is an intricate problem, our interest is in the tradeoff between runtime and solution quality that we get by limiting ourselves to a special type of cut, namely cuts with at most one bend each (corner cuts). We prove that optimum corner cuts get us arbitrarily close to equal sized parts, 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.

1 Comparing Optimal with Restricted Cuts

We consider the *bisection problem*: partition the vertex set of a given graph into two (almost) equal sized subsets such that the number of edges with an endpoint in each partition is minimised. The problem has been studied extensively, due to its utility in divide-and-conquer algorithms. It is NP-hard in general [5] and the best approximation algorithm known for general graphs [9] guarantees an approximation ratio of $\mathcal{O}(\log n)$. For planar graphs a PTAS [2] has been found, while for trees an optimum solution can be computed in $\mathcal{O}(n^2)$ time [6, 7]. Our motivation to study the bisection problem comes from the need to parallelise a finite element computation of a human bone structure model in order to diagnose osteoporosis [1]. In such an application the aim is to distribute the data, modelled by the vertices (of a graph G), evenly onto a given number p

^{*}We gratefully acknowledge discussions with Peter Arbenz who introduced the human bone simulation problem to us, and the support of this work through the Swiss National Science Foundation under Grant No. 200021_125201/1.

[†]An extended abstract of this report appeared at the *37th International Workshop on Graph-Theoretic Concepts in Computer Science* [4]

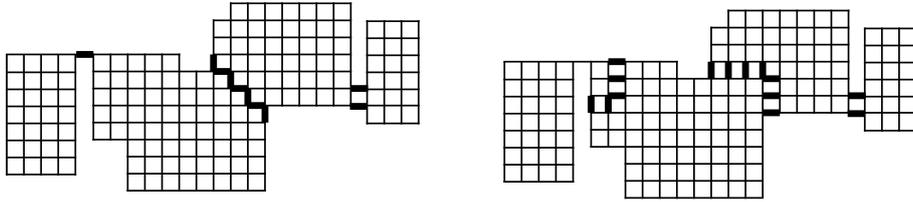


Figure 1: An optimal (left) and a corner cut (right) in a solid grid, each cutting out $k = 110$ many vertices. The bold edges indicate the segments.

of processors to achieve a balanced computation load. At the same time the communication between data points, modelled by the edges (of G), needs to be kept at a minimum between processors since this constitutes a bottleneck in parallel computing. One way to distribute the data is to recursively solve the bisection problem. The finite element models corresponding to the bones in our application constitute a porous 3D grid. In a first step towards understanding the problem for 3D grids, in this article we limit ourselves to 2D *grid graphs*, which are finite, connected subgraphs of the infinite two-dimensional grid. A grid graph is called *solid* if it has no hole, i.e. if it has no interior face surrounded by more than four edges. For solid grid graphs the fastest algorithm known [8] computes an optimum bisection in time $\mathcal{O}(n^5)$.

We aim at understanding the intricacies of bisection from a novel point of view: we study particular classes of cuts, for which it is possible to find optimum solutions faster. In this article we consider cuts that can have at most one bend, so-called *corner cuts* (Figure 1). We prove that optimum corner cuts do not need a lot more cut-edges than arbitrary cuts. Cuts without any bend (straight cuts) do not achieve this high quality: they can be a \sqrt{n} factor away from optimum. We achieve our result by proving a number of theorems for polygons that we relate to the case of grid graphs. The main part of this article will therefore be concerned with thoroughly analysing corner cuts in polygons.

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 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. Given a predefined *basic set* of segments Γ to choose from, we say that a cut S is Γ -*restricted* if $S \subseteq \Gamma$. Our overall goal is to identify such basic sets which (1) can be computed efficiently, (2) allow fast algorithms to compute an optimal restricted bisection, and (3) for which an optimal restricted solution also is a good solution for the non-restricted case. In this article we show that for solid grid graphs a basic set that consists of segments with at most one bend, i.e. constituting corner cuts, is a good approximation to the optimal bisection.

We call the number of edges $\sum_{s \in S} |s|$ in a cut S its *cut-size*. The cut minimising the cut-size is *optimal*. Notice that some edges may be counted several times in the sum. However in the non-restricted case, edges that appear more than once can be removed. Hence this generalisation does not change the optimal solution. A simple cycle in the dual of a planar graph corresponds to a closed curve in its embedding in the plane. Hence the cycle divides the plane into an interior and an exterior area. We say that a pair of cycles *cross* if the

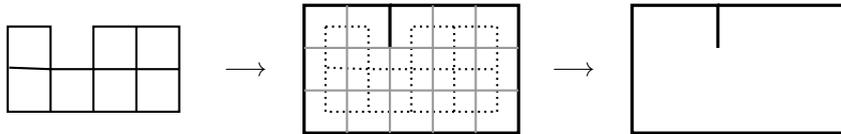


Figure 2: Converting a grid to a polygon.

corresponding closed curve of one of them both contains points belonging to the interior and to the exterior area into which the other cycle divides the plane. Note that any pair of crossing simple cycles can be seen as a (different) pair of simple cycles that do not cross. Hence for the non-restricted optimal cut we may consider only *non-crossing* cuts in which no corresponding cycles cross. In our application we can assume that the grid graph G is given together with its natural embedding in the plane, i.e. G is a plane graph in which the vertices are coordinates in \mathbb{N}^2 and all edges have unit length. We propose to use only segments that correspond to orthogonal lines with at most one right-angled bend in the dual graph, when disregarding the part of the cycle that connects to the exterior face (Figure 1). If the corresponding line of a segment contains no bend we call it a *straight segment* and if it contains exactly one right-angled bend a *corner segment*.

If removing the edges included in a cut leaves a set of connected components together containing exactly k vertices, we call it a k -cut. We say that it *cuts out* k vertices. We call the set containing k vertices the \mathcal{A} -part, and the other set the \mathcal{B} -part of the k -cut. A natural generalisation of the bisection problem is to find the optimal k -cut for a given integer $k \in \{0, \dots, n\}$. The main result of this article is summarised in the following theorem.

Theorem 1. *Let l be the cut-size of an optimal k -cut in G and $\varepsilon \in]0, 1]$. Then there exists a non-crossing corner k' -cut, for some $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, which has a cut-size of $\mathcal{O}(1/\sqrt{\varepsilon}) \cdot l$.*

This result was already used in [3] where we showed that the optimal corner k -cuts for all $k \in \{0, \dots, n\}$ can be computed in time $\mathcal{O}(n^4)$ for a given solid grid graph G . Thus the above theorem implies that one can find a good approximation to the optimal bisection of G faster than by using the $\mathcal{O}(n^5)$ time algorithm from [8]. Since the set sizes of the resulting k -cut correspond to the load on each machine in our application, the deviation by a factor of ε on the load balance is acceptable as it does not significantly affect the resulting running time of the parallel computation. However the speed-up of the bisection algorithm is a considerable improvement since a typical data set will have billions of vertices.

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. 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 (but can be degenerate in the sense that the polygonal chain corresponding to its border can have overlapping edges; see Figure 2). In this article we define a polygon using its interior point set.

Definition 2. A *polygon*¹ $\mathcal{P} \subset \mathbb{R}^2$ is an open bounded set of points in the plane, i.e. if β denotes its boundary then $\beta \cap \mathcal{P} = \emptyset$. If β only contains orthogonal lines

¹We will use calligraphic capital letters to denote areas in the plane such as polygons or

we refer to \mathcal{P} as *orthogonal*. We call \mathcal{P} *simple* if any two points in \mathcal{P} lie on a closed curve in \mathcal{P} which 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) . 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 line that is the boundary of the polygon. The region enclosed by the boundary is of size exactly n , the number of vertices in G . The k -cut in the polygon that corresponds to the optimal k -cut in the grid obviously has a cut-size that is at least the cut-size of the optimal k -cut in the polygon.

All the notions used for cuts in grids carry over naturally to the case of polygons. For a polygon \mathcal{P} with fixed orientation in the plane we call an orthogonal line² within \mathcal{P} that has no bend and starts and ends at the boundary of \mathcal{P} , a *straight line*. Accordingly an orthogonal line within \mathcal{P} that has exactly one right-angled bend and starts and ends at the boundary of \mathcal{P} is called a *corner line*. The cut-size of a set of lines L in a polygon is the sum of the lengths of the lines in L , which are measured using the Manhattan distance.

Definition 3. Given a polygon \mathcal{P} we call a line $\lambda \subset \mathcal{P}$ 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.

If β denotes the boundary of \mathcal{P} , we call a point $p \in \beta$ a *boundary point* of a line λ in \mathcal{P} if the distance from p to λ is 0. If λ has two boundary points we call it a *segment line*.

A bar line that is also a segment line is called a *straight line*, and a segment line 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. The point at which the two bar lines of a corner line meet is called its *corner*. The orientation of a corner line is either one of *up-right*, *up-left*, *down-right*, or *down-left*, depending on whether its horizontal and vertical bar lines go up, down, left, or right from its corner, respectively.

We say that two segment lines *cross* if the intersection as well as the symmetric difference of the respective cut out areas are non-empty. A set L of segment lines is called *non-crossing* if no pair of lines in L cross. If two segment lines share some points of \mathcal{P} but do not cross we say that they *touch*. If two lines touch and there is a connected part in their intersection which contains more than a single point, we refer to them as *overlapping*.

We will first show the existence of corner cuts in simple polygons that cut out almost the required area and have small cut-size (close to optimal). We will then convert such a cut in a polygon derived from a grid graph to a corresponding cut in the grid 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 k -cut in a polygon that is almost a corner cut, in the sense that the cut consists of only straight and corner lines

areas inside polygons, and we will use lower case Greek letters to denote lines such as borders or segment lines.

²All lines considered in this article have finite length unless otherwise stated.

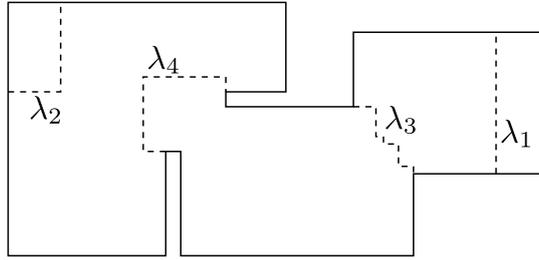


Figure 3: A straight, corner, staircase, and rectangular line in a polygon denoted by λ_1 through λ_4 respectively.

- except at most one other line. This line may be shaped like a staircase (a so called *staircase line*), or it may be a *rectangular line*, which is defined as a continuous part of the boundary of an orthogonal rectangle (Figure 3).
2. We show how to remove a rectangular line from a cut containing only straight and corner lines otherwise. 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 lines from the cut into consideration so that the newly introduced lines do not interfere with these. The new cut will also be a k -cut but its cut-size may not be optimal. However, we show that the cut-size of the new cut is only a constant factor away from the optimal.
 3. Given a k -cut of the polygon consisting of straight and corner lines, and one staircase line, 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 k . To be more precise, the new cut is a k' -cut where $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$ for any desired constant $\varepsilon \in]0, 1]$. Further, the cut-size of the new cut is only a constant factor (depending on ε) times the cut-size of the original cut.
 4. At last we show how to convert a cut containing only straight and corner lines in a polygon corresponding to a grid graph G into a cut in G . Note that this step would be straightforward if all the lines in the cut were passing through exactly the midpoints of the edges of the grid. We call such lines *grid lines*. We show that all lines in the cut obtained in the previous steps can be moved to grid lines in such a way that the cut-size 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-size grows by only a small factor.

The next sections explain these techniques in more detail.

2 Cuts in Polygons

We will now show that in an optimal k -cut of a polygon all but at most one line are corner and straight lines. Lines with more bends include *staircase lines* and *rectangular lines*. The former have at least two bends and are monotonic in x - and in y -direction. The latter have two or three bends and lie on the boundary of an orthogonal rectangle (Figure 3).

Definition 4. For any polygon \mathcal{P} a *staircase line* $\lambda \subseteq \mathcal{P}$ is a segment line that contains at least three bar lines and is monotonic in x - and y -direction. The orientation of λ is *up* if the left one of its boundary points also is lower than the other, and *down* otherwise.

Let $\mathcal{R} \subset \mathbb{R}^2$ be an axis parallel rectangle in the plane and let γ be its boundary. Any segment line $\lambda \subseteq \gamma \cap \mathcal{P}$ 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 according to the just stated definition.

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 having the same boundary points have the same length.

In a first step, we convince ourselves that in any simple polygon there is an optimal k -cut that contains only straight, corner, staircase, and rectangular lines. Furthermore, none of these lines cross or overlap. These results are analogous to those attained in [8] for grid graphs.

Lemma 5. *In any polygon \mathcal{P} there is an optimal k -cut L that is non-crossing and contains only straight, corner, staircase, and rectangular lines. Furthermore no lines in L overlap.*

Proof sketch. We will not present a full formal proof at this point since we feel that the ideas are straight forward.

Notice that, as in the case for grids, for any k -cut there exists a non-crossing k -cut with the same cut-size. Additionally, removing overlapping parts of lines results in a k -cut of smaller cut-size and thus no lines in L overlap. Also, as in the case for grids, it is easy to see that for any k -cut with cut-size l there is a k -cut for which every line is a segment line and has a cut-size of at most l . Thus let λ be a line from L and let $\mathcal{R} \subset \mathbb{R}^2$ be the smallest rectangle containing λ . Due to the well-known isoperimetrical problem using the L_1 -norm (see e.g. [10]) 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 exchanged 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 k -cut contains a rectangular line,

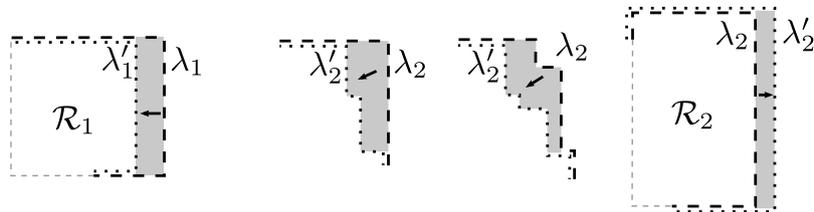


Figure 4: A rectangular line λ_1 with its defining rectangle \mathcal{R}_1 is exchanged with the rectangular line λ'_1 . To compensate for the area a (shaded grey), another line λ_2 is replaced by λ'_2 . It can be a corner, staircase, or a rectangular line (with defining rectangle \mathcal{R}_2).

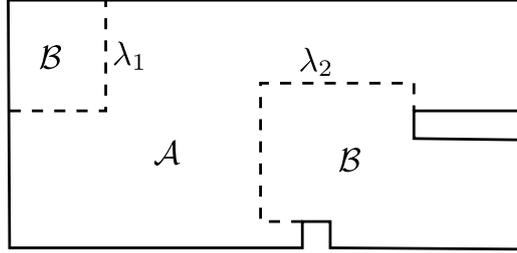


Figure 5: A corner, and rectangular line in a polygon denoted by λ_1 , λ_2 , respectively. Both are concave with respect to the \mathcal{A} -part and convex with respect to the \mathcal{B} -part.

then all other lines 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 4)—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 5). 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 show that for an optimal k -cut with a rectangular line, the area on its concave side 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.

Let in the definition below $[[a, b]]$ denote the interval $[a, b]$ if $a \leq b$ and $[b, a]$ otherwise.

Definition 6. Let $\mathcal{A}(L) \subseteq \mathcal{P} \setminus \{p \in \lambda \mid \lambda \in L\}$ be the open set of size k that is cut out by the k -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 $n - k$. That is, the areas $\mathcal{A}(L)$ and $\mathcal{B}(L)$ do not include points that are contained in lines of L or the boundary of \mathcal{P} .

For any line $\lambda \in L$ and $C \in \{\mathcal{A}(L), \mathcal{B}(L)\}$ let the sets $\mathcal{X}_\lambda(C)$ and $\mathcal{Y}_\lambda(C)$ contain the points in C for which there is a horizontal or vertical bar line, respectively, in C between such a point and λ , i.e.

$$\mathcal{X}_\lambda(C) := \{(x, y) \in C \cup \lambda \mid \exists x' : ((x', y) \in \lambda \wedge \forall x'' \in [[x, x']] : (x'', y) \in C \cup \lambda)\}$$

and

$$\mathcal{Y}_\lambda(C) := \{(x, y) \in C \cup \lambda \mid \exists y' : ((x, y') \in \lambda \wedge \forall y'' \in [[y, y']] : (x, y'') \in C \cup \lambda)\}.$$

Note that these sets also include λ itself.

We call a corner or rectangular line $\lambda \in L$ *concave* with respect to C if for every $(x', y') \in \mathcal{X}_\lambda(C)$ there is a relation $\leq \in \{\leq, \geq\}$ such that for all $(x, y) \in \lambda$ it holds that $x \leq x'$, and for every $(x', y') \in \mathcal{Y}_\lambda(C)$ there is a relation $\leq \in \{\leq, \geq\}$ such that for all $(x, y) \in \lambda$ it holds that $y \leq y'$. That is, there is no point p in $\mathcal{X}_\lambda(C)$ or in $\mathcal{Y}_\lambda(C)$ such that there are points on λ which lie to the right and the left of p or above and below p . We call λ *convex* with respect to C if it is concave with respect to the other area \bar{C} , i.e. $\bar{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\} \setminus \{C\}$.

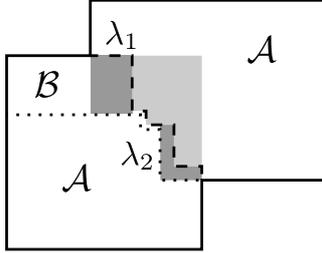


Figure 6: A staircase line λ_1 together with its surplus (in light grey shading) and its deficit (in dark grey shading).

Since in an optimal k -cut L no lines overlap and hence the set $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$ can only lie to one side of a corner or rectangular line, any such line in L is either concave or convex w.r.t. \mathcal{C} .

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 the areas of maximal size with monotone borders contained in the \mathcal{B} - and \mathcal{A} -part respectively (Figure 6). 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 7).

Consider the simple case of a k -cut that contains only one line λ which is a staircase line. In this case the set $\mathcal{Z}_\lambda(\mathcal{C}) = \mathcal{X}_\lambda(\mathcal{C}) \cap \mathcal{Y}_\lambda(\mathcal{C})$ is referred to as the surplus if $\mathcal{C} = \mathcal{A}(\{\lambda\})$ and as the deficit if $\mathcal{C} = \mathcal{B}(\{\lambda\})$. By the definition of $\mathcal{X}_\lambda(\mathcal{C})$ and $\mathcal{Y}_\lambda(\mathcal{C})$ the part of the boundary of $\mathcal{Z}(\mathcal{C})$ that does not intersect with λ forms a union of straight and corner lines in this simple case. Let $L_{\mathcal{C}}$ denote this set of lines for each $\mathcal{C} \in \{\mathcal{A}(\{\lambda\}), \mathcal{B}(\{\lambda\})\}$. The surplus can thus be seen as the area from $\mathcal{A}(\{\lambda\})$ that λ cuts out in addition to the area that the lines $L_{\mathcal{A}(\{\lambda\})}$ cut out. The deficit can be seen as the area from $\mathcal{A}(L_{\mathcal{B}(\{\lambda\})})$ that λ does not cut out. However if there are lines that overlap with λ in a k -cut then the definition has to be modified in the following way, in order to capture a similar notion.

Definition 7. Let $\lambda \in L$ be a staircase line from a non-crossing k -cut L , $\mathcal{C} \in \{\mathcal{A}(L), \mathcal{B}(L)\}$, and $L' := L \setminus \{\lambda\}$. For any line $\lambda' \in L'$ let $\mathcal{D}_{\lambda'} \in \{\mathcal{A}(\{\lambda'\}), \mathcal{B}(\{\lambda'\})\}$ such that $\lambda \cap \mathcal{D}_{\lambda'} = \emptyset$. We call the set

$$\mathcal{Z}_\lambda(\mathcal{C}) := (\mathcal{X}_\lambda(\mathcal{C}) \cap \mathcal{Y}_\lambda(\mathcal{C})) \setminus \bigcup_{\lambda' \in L'} \mathcal{D}_{\lambda'}$$

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

If there is a line λ' that overlaps with a staircase line λ then it can happen that the intersection between $\mathcal{X}_\lambda(\mathcal{C}) \cap \mathcal{Y}_\lambda(\mathcal{C})$ and $\mathcal{D}_{\lambda'}$ is non-empty. This would mean that a line in $\mathcal{X}_\lambda(\mathcal{C}) \cap \mathcal{Y}_\lambda(\mathcal{C})$, which will later be used when transforming λ , might cross λ' . Hence all the sets $\mathcal{D}_{\lambda'}$ are removed in the surplus and deficit.

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

Lemma 8. *For any polygon \mathcal{P} , if an optimal k -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 λ_1 be the rectangular line that w.l.o.g. is concave with respect to $\mathcal{C} = \mathcal{A}(L)$ (by Lemma 5 no lines overlap and hence any rectangular or corner line is either concave or convex w.r.t. \mathcal{C}). Let λ_2 be a staircase line. As we will show, there is a sufficiently small area of size $a > 0$ which can be locally “transferred” from λ_1 to λ_2 by making λ_1 shorter while transforming λ_2 such that the cut-size 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 line in L . Observe that λ'_1 is shorter than λ_1 by twice the width of the area $\mathcal{Q}(x_a)$. When exchanging λ_1 with λ'_1 in L we need to compensate for the area $\mathcal{Q}(x_a)$ in order to cut out an area of size k , by also exchanging the line λ_2 with some appropriate line λ'_2 (see Figure 4).

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 exchanging λ_2 with λ'_2 will make the cut out area have size k (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 line and therefore the new k -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-size 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 k -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 line since no line overlaps with λ_2 by Lemma 5. 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-size 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 k -cut contains a rectangular line then no other line is a rectangular or staircase line, as the following lemma shows.

Lemma 9. *For any polygon \mathcal{P} , if an optimal k -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 lines are straight and corner lines, where the latter all are concave with respect to \mathcal{C} .*

Proof. By Lemma 5 we can assume that all lines in L are straight, corner, staircase, or rectangular lines. Additionally Lemma 8 shows that apart from the rectangular line λ_1 there can be other rectangular lines that do not fit the

statement in this lemma. Let λ_2 be such a rectangular line, and we first consider the case when λ_2 is convex w.r.t. \mathcal{C} . As in the proof of Lemma 8 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 a k -cut with smaller cut-size 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} (see Figure 4). 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 l_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 line. In both 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. 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 does not cross any other line, 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 5 no lines overlap in L , such that for sufficiently small a the constructed line does not cross any other line, 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 and we will return to this case later.

We first consider the case when λ_2 and λ'_2 share both boundary points. In this case the area that resides between the two lines is $\mathcal{R}'_2 \setminus \mathcal{R}_2$ and 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 . Since $a > 0$ it must also hold that $z > 0$ and 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 equals $l_2 + 4z$. Solving the lower bound on a for z we can hence 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 would hence yield a k -cut that has a shorter cut-size than L since we assumed that $h_2 \geq h_1$. But this contradicts the optimality of L which means that λ_2 and λ'_2 can not share both boundary points.

If there exists some sufficiently small $a > 0$ such that λ_2 and λ'_2 share both boundary points then by the above argumentation we get a contradiction. Hence for any (sufficiently small) $a > 0$, respectively for any (sufficiently small) $z > 0$, it must be the case that the boundary points q and q' of λ_2 and λ'_2 , respectively, that differ from p are unequal. This means that the upper left corners of \mathcal{R}_2 and \mathcal{R}'_2 can never be connected by a bar line in \mathcal{P} for any $z > 0$. But this can only be the case if for some $z' > 0$ and all $z \in]0, z'[$ the boundary point q' lies to the

right of q . In this interval there must exist some value for z such that the right half of the rectangle that extends \mathcal{R}_2 to the top by z does not contain any parts of the boundary of the polygon, i.e.

$$\{(x', y') \in \mathbb{R}^2 \mid x' \in]x - w_2/2, x[\wedge y' \in]y, y + z[\} \subset \mathcal{P}.$$

By considering only the right half of the rectangle extending \mathcal{R}_2 to the top in addition to the rectangles that extend \mathcal{R}_2 to the right and the rectangle between those two, we can conclude that $a > z \frac{w_2}{2} + zh_2 + z^2 > \frac{3}{2}zh_2$. Since we chose z such that q' lies to the right of q we can conclude that $l'_2 < l_2 + 3z$ and solving the lower bound on a for z we yield $l'_2 < l_2 + 3 \frac{a}{3/2 \cdot h_2} = l_2 + 2a/h_2$. Hence, as in the case before, we get a contradiction which means that λ_2 cannot be a rectangular line that is concave w.r.t. \mathcal{C} . \square

After considering optimal k -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 k -cut in which at most one line is a staircase line while all others are corner and straight lines.

Lemma 10. *For any polygon \mathcal{P} without holes, if there is an optimal k -cut L that contains a staircase line then there also is an optimal k -cut that contains at most one staircase line while all other lines are straight or corner lines.*

Proof. By Lemma 9 it can not happen that there is a rectangular line in L . Hence, by Lemma 5, the only case we have to consider is when there are two staircase lines λ_1 and λ_2 . 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 . However since \mathcal{P} does not contain any holes, 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 lines $L_1(a)$ that cut out an area of size a from the surplus of λ_1 , such that removing λ_1 yields a $(k - a)$ -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 lines in $L_1(a)$ are part of the boundary of λ_1 's surplus together with some parts of λ_1 (see Figure 7).

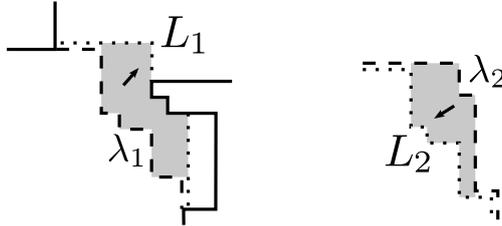


Figure 7: Two staircase lines λ_1 and λ_2 together with their respective surplus (or deficit) shaded in grey. The dotted lines indicate that the border of the surplus (or deficit) together with some parts of λ_i , $i \in \{1, 2\}$, form only corner and straight lines which are contained in L_i .

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 $(k - a)$ -cut. Similar as for the surplus of λ_1 , for a fixed a we can find a set of lines $L_2(d)$ for any $d \in [0, d_2(a)]$ cutting out an area of size d from the deficit of λ_2 , which either contains a single staircase line or lines 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 $(k - b + b)$ -cut, i.e. a k -cut, which has a cut-size that is at most the cut-size of the original k -cut since distances are measured using the L_1 -norm. Assume that some line in $L_1(b)$ or in $L_2(b)$ overlaps with some other line in the k -cut (including those in $L_2(b)$ and $L_1(b)$ respectively). The overlapping part can be removed, again yielding a k -cut which now however has a shorter cut-size. Since this is a contradiction to the optimality of L , we can conclude that no line in $L_1(b)$ or $L_2(b)$ overlaps with any other line. 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 lines that are part of the boundary of the surplus or deficit, respectively, together with some parts of the respective staircase line. Since these lines do not overlap with any other lines they must either be straight or corner lines, by the definition of the surplus and deficit (see also Figure 7).

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

To summarise the above results, the following observation immediately follows from Lemmas 5, 9, and 10.

Corollary 11. *For any polygon \mathcal{P} without holes there is an optimal k -cut L such that all lines 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 we can cope with a rectangular line, and how with a staircase line. For a rectangular line we show how to convert the k -cut into a cut in which there is at most one staircase line while the cut-size grows at most by some constant factor. With our observations on staircase lines we are then able to convert any optimal k -cut into one containing only straight and corner lines.

3 Removing Rectangular Lines

We now show how to convert an optimal k -cut containing straight and corner lines and one rectangular line into a k -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 8). This region may contain a part of the boundary of the polygon (and possibly some other lines 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 lines have total length less than the length of the rectangular line. By doing this, we do not increase the

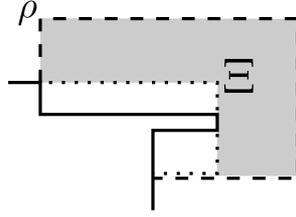


Figure 8: A rectangular line ρ (dashed) together with the set of lines Ξ (dotted) with which it is replaced. The area of size a is shaded in grey.

size 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 lines that cut out the required area of size a and has total length not too large (compared to the cut-size l of the optimal k -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 lines of total length not much larger than \sqrt{a} .

In order to find this set of lines we need to abstract away from the actual topology of the polygon. We achieve this by introducing the following notions (see also Figure 9), where we will restrict ourselves to the case of one specific orientation of the involved lines. Notice that in the definition below we can choose \mathcal{P} to be either a polygon, i.e. an open set, or a closed set of points, to our liking.

Definition 12. Let μ be a corner line of orientation up-right (and of infinite length) in the plane \mathbb{R}^2 . For any 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*³. 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* .

³We will use capital Greek letters to denote virtual lines.

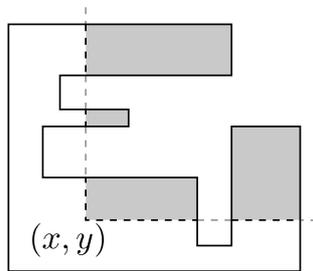


Figure 9: A virtual corner line (black dashed) at (x, y) . The cut out area is shaded in grey.

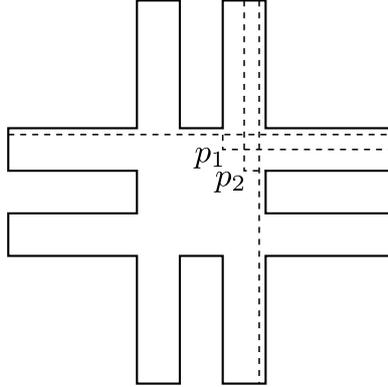


Figure 10: A polygon in which every virtual corner line for a is too long. At p_1 the vertical section switches from short to long and at p_2 the horizontal section switches from long to short.

For a fixed a let $\Lambda(x)$ be a virtual corner line for a with corner (x, y) , if it exists. 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 10, any virtual corner line cutting out the required area contains a long vertical section or a long horizontal section.

Given any simple polygon, we can search along the x -axis between the two extremities of the polygon, and for each value of x try to find a y such that the virtual corner line at (x, y) cuts out exactly an area of size a (Figure 11). 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 a short (i.e. at most $c\sqrt{a}$) vertical section, the latter has a short horizontal section, and for all virtual corner lines in between both sections are long.

Lemma 13. *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 by the pigeon-hole principle 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 14. *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 13, then*

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

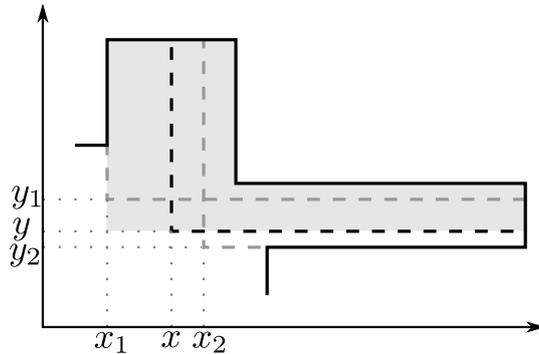


Figure 11: The interval $[x_1, x_2]$ in a polygon and a virtual corner line (dashed black) for k whose horizontal and vertical sections are both long.

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 11). 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 $m(x) := \max\{w(x), h(x)\}$ we can thus conclude that the size of $\mathcal{Q}_y \setminus \mathcal{P}(x)$ is greater than $m(x) \cdot c\sqrt{a}$. By using $m(x)$ as an estimate of $h(x)$ and $w(x)$ in the upper bound derived above, we get the following inequality:

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

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

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

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

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

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

Note that this means that there is an interval between these two bounds from which $m(x)$ 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 $m(x)$ is a continuous function. Since $m(x)$ 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 $m(x_2)$ is also one for $w(x_2)$ and $h(x_2)$. \square

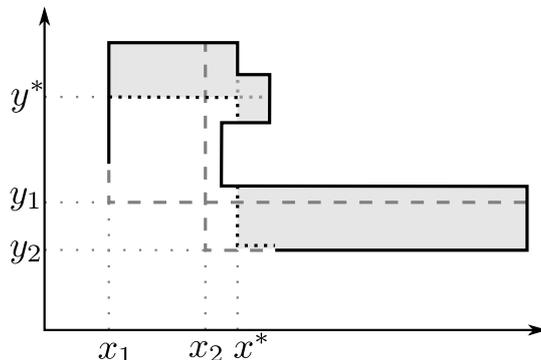


Figure 12: A virtual staircase line (dotted black) cutting out the area of size a shaded in grey. It is constructed using the two virtual corner lines at (x^*, y_2) and (x_1, y^*) .

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 15. Let μ be a staircase line of orientation down (and of infinite length) in the plane \mathbb{R}^2 . 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 k on its upper right side we say that it is a virtual corner line for k .

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. 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 (Figure 12).

Lemma 16. *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 13 and 14 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 lines as follows (see Figure 12 for a depiction). We will use the same notation as in the proof of Lemma 14.

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 lines Λ contains all lines λ that are segment lines 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 13. 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 14 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 lines in Λ . \square

Let L be the set of straight and corner lines in the cut after replacing the rectangular line, such that L cuts out an area $k - a$ or $k + a$. 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 lines in the cut (Figure 13). 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 line in L (however, the lines may partially overlap). This is done by adding those parts of lines in L to the lines 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 pointed in the previous section. Thus the set M may contain several staircase lines, but its total length is at most that of L .

Lemma 17. *For any polygon \mathcal{P} , let L be a non-crossing corner k -cut with cut-size l , such that all corner lines in L are concave with respect to $\mathcal{A}(L)$. For any $a \in [0, k]$ there is a set of segment lines 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-size 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 $(k - a)$ -cut $M \cup L$, lies on its lower left side.*

Proof. By Lemma 16 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 lines 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 λ' is either a straight line or a corner line that is concave w.r.t. $\mathcal{A}(L)$, the definition of a virtual staircase line, and the fact that $\lambda \subseteq \mathcal{A}(L)$ lets us conclude

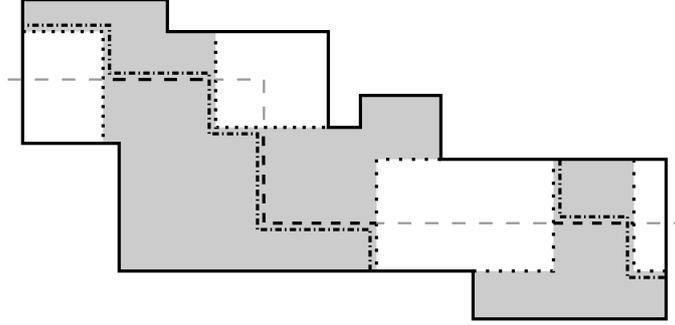


Figure 13: A virtual staircase lined (dashed) that is converted to a set of staircase, corner, and straight lines (dashed and dotted). For this parts of the corner lines (dotted) from the original cut are used. These are all concave w.r.t. the same part of the cut shaded in grey.

that 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 that $y' \preceq y$ while $y \preceq y''$. Let in the former case $\mu_{(x,y)} := \{(x', y') \in \lambda' \mid y \preceq y'\}$ and in the latter case $\mu_{(x,y)} := \{(x', y') \in \lambda' \mid x \preceq x'\}$. 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 border of the area of size a that Λ cuts out, we add to M any line in L that is contained in γ . Let P denote the set of boundary points of the lines in M w.r.t. $\mathcal{A}(L)$ which are contained in some line 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 line 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 exchange 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 line $\lambda' \in L$ then it is the only one. For any such point p we exchange 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 lines in L are straight lines or corner lines that are concave w.r.t. $\mathcal{A}(L)$, the added parts of the lines in L connect the lines 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 $(k - a)$ -cut $M \cup L$ contains the parts of $\mathcal{A}(L)$ that were cut out by the virtual staircase line Λ on its upper right side. Hence the surplus, defined w.r.t. the $(k - a)$ -cut $M \cup L$, of a staircase line

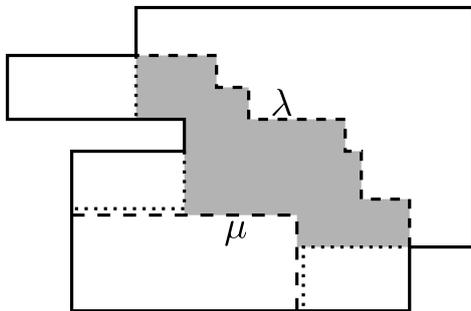


Figure 14: A staircase line λ with its surplus shaded in grey. The lines on the border of the surplus can be replaced by a set of straight and corner lines (dotted). The corner line μ is also removed.

in M must lie on its lower left side. Finally each added part from a line in L was only added once to a staircase line while constructing M , which means that the total length of the lines 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-size does not increase. Similar to the techniques seen before, we will use the lines 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 14). This corner line μ together with the staircase line λ' can be used to construct a pair of corner lines which can be replaced with μ and λ' so that the same region is cut out by the new set of lines. The cut-size decreases during this process.

Lemma 18. *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 line in L . Let Λ denote the set of segment lines in \mathcal{P} that are contained in the border 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 k then there exists a k -cut that has a cut-size 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 lines in Λ are contained in the border 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 which 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 $z_t = 0$ for some $t \in \{x, y\}$ if the respective point lies on λ which 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 line $\mu \in L$ 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 $\sigma'_h \cup \sigma'_v \cup \{p'\}$, where p' is the point at which σ'_h and σ'_v meet, i.e. it is the corner of the corner line. 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 exchanged with one or two corner lines. In a cut that includes the lines from Λ and L this transformation will not change the size of the cut out area and will decrease the cut-size. \square

The above observations can now be used to convert the staircase lines constructed in Lemma 17 in such a way that only one staircase line remains, as the next lemma shows.

Lemma 19. *In a polygon \mathcal{P} let L be a non-crossing corner k -cut with cut-size l such that all corner lines in L are concave with respect to $\mathcal{A}(L)$. Then for any $a \in [0, k]$ there exists a $(k - a)$ -cut M in \mathcal{P} with cut-size 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 17 we can find a set of lines M' such that $L \cup M'$ is a $(k - a)$ -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 18 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 10 we can find two sets of lines 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 $(k - a)$ -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 lines that are part of the boundary of λ_2 's deficit if $b_2 = b_1$. The set L_1 always contains lines that are part of the boundary of λ_1 's surplus. The new $(k - a)$ -cut in which λ_1 and λ_2 were replaced has a cut-size that is at most the cut-size of the old k -cut since distances are measured in the L_1 -norm (it is decreasing if there are more than one line in L_1 or L_2 since then parts of the boundary of \mathcal{P} act as a short cut for the lines).

Using Lemma 18 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 $(k - a)$ -cut in which there is at most one staircase line left. Since the cut-size is non-increasing during each transformation step, the cut-size of the final set M is at most $7\sqrt{a} + 2l$, which concludes the proof. \square

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

Theorem 20. *For any simple polygon \mathcal{P} with an optimal k -cut L of \mathcal{P} containing a rectangular line, there exists a non-crossing k -cut M which contains only corner and straight lines except at most one which is a staircase line and M has a cut-size of at most $9l$, where l is the cut-size 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} , which 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

$$x_{\max} := \max\{x \in \mathbb{R} \mid (x, y) \in \beta \cap \mathcal{R}\} \text{ and}$$

$$y_{\max} := \max\{y \in \mathbb{R} \mid (x, y) \in \beta \cap \mathcal{R}\}$$

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 lines Ξ such that $\lambda \in \Xi$ if and only if λ is a segment line 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 line 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 line $\lambda \in \Xi$ 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 line 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 therefore no line 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 9, apart from ρ the set L contains only straight and corner lines. Hence, since no line crosses ρ and by the construction of Ξ the set \mathcal{D} does not intersect with the boundary β , no line from L crosses a line in Ξ . Thus we can exchange ρ with Ξ in L and yield a non-crossing $(k + a)$ -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 9 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 19 we can thus find a k -cut M that has a cut-size of at most $7\sqrt{a} + 2l'$, where l' is the cut-size 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 lines in Ξ is at most the length of ρ and hence the cut-size of M' is at most the cut-size 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-size 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) cut L containing only straight and corner lines except one which is a staircase line, 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 having a short cut-size. 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-size of L . Thus, if we can cut out the excess area a using straight and corner lines of

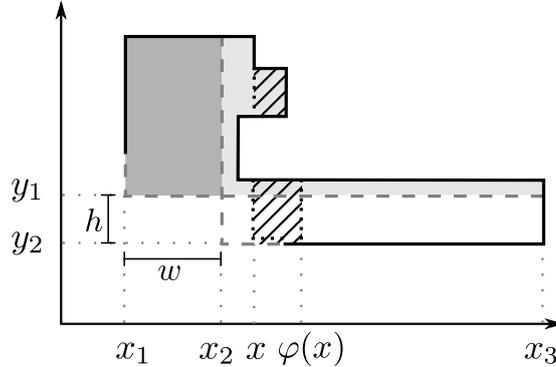


Figure 15: The interval $[x_1, x_2]$ of width w together with the virtual corner lines $\Lambda(x_1)$ and $\Lambda(x_2)$ (grey dashed lines). The dark grey area is $\mathcal{P}(x_1) \setminus \mathcal{P}(x_2)$ of size b' , while the light grey area is $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ of size 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 dashed lines) enclose an area of size d' between them (checked pattern).

total length in $\mathcal{O}(\sqrt{a})$, then our cut-size will still be close to optimal. Given any simple polygon \mathcal{P} of area n , $a \in [0, n]$, and constant $\varepsilon \in]0, 1]$ we can find a set of at most three virtual corner lines that cut out an area whose size is in the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$, and has a cut-size 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 11). 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 21. *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-size 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 13 and 14 there is an interval $[x_1, x_2] \subseteq I_k$ with the properties listed there. We use the same notation as in the proof of Lemma 14, 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 15).

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 on the size b' of $\mathcal{P}(x_1) \setminus \mathcal{P}(x_2)$ which we can then subtract from a since the size of

$\mathcal{P}(x_1)$ is a . 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 yield $b' \geq d - hw$. 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 13 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 14 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 13 and 14 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 13 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 lengths or the horizontal lengths of the two sought after 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

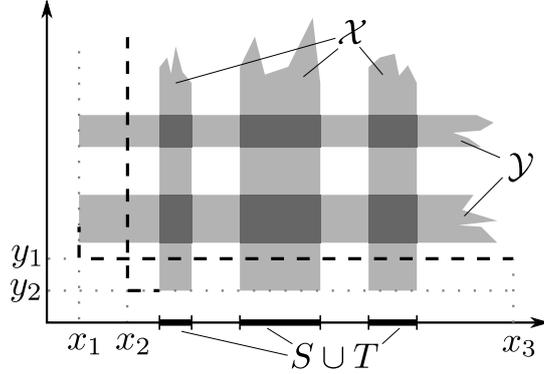


Figure 16: The interval $[x_1, x_2]$ together with the virtual corner lines $\Lambda(x_1)$ and $\Lambda(x_2)$ (dashed lines). To estimate the size b of $\mathcal{P}(x_1) \cap \mathcal{P}(x_2)$ we determine the size $f(S \cup T)$ of \mathcal{X} and the size of \mathcal{Y} .

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

$$f(U) := \int_U l'_v(t) dt.$$

Let $S := \{x \in [x_2, \varphi^{-1}(x_3)] \mid l'_v(x) > c\sqrt{a}\}$ be the subset of the domain of φ for which the vertical lengths of the virtual corner lines $\Lambda'(x)$ are long. Let \bar{S} be the set for which the vertical lengths are short, i.e. $\bar{S} := [x_2, \varphi^{-1}(x_3)] \setminus S$. Also let $T := \{x \in [\varphi(x_2), x_3] \mid l'_v(x) > c\sqrt{a}\}$ and $\bar{T} := [\varphi(x_2), x_3] \setminus T$ 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 $f(S \cup T)$ (see Figure 16).

Let $[u_1, u_2] \subseteq \bar{S}$ be a connected subset of \bar{S} . By the definitions of f and φ we get

$$\begin{aligned} f([u_1, u_2]) &= f([u_1, \varphi(u_1)]) + f([\varphi(u_1), \varphi(u_2)]) - f([u_2, \varphi(u_2)]) \\ &= f([\varphi(u_1), \varphi(u_2)]). \end{aligned}$$

Since \bar{S} is a union of connected subsets and φ is bijective we can conclude that $f(\bar{S}) = f(\varphi(\bar{S}))$, where $\varphi(\bar{S})$ is the image of \bar{S} . By the assumption that for all $x \in [x_2, \varphi^{-1}(x_3)]$ the vertical length of $\Lambda'(x)$ or of $\Lambda'(\varphi(x))$ is long, $\varphi(\bar{S})$ must be a subset of T and hence $f(\bar{S}) \leq f(T)$. A similar observation can be made for \bar{T} , $\varphi^{-1}(\bar{T})$, and S so that $f(\bar{T}) \leq f(S)$.

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 a size of a . Hence $f(S \cup \bar{S}) = f(T \cup \bar{T}) = a - d'$. From the bounds above and the fact that S and \bar{S} but also T and \bar{T} are disjoint we can conclude that

$$f(S) + f(T) \geq f(\bar{T}) + f(\bar{S}) = 2(a - d') - f(T) - f(S).$$

The sets S and T might not be disjoint but from the above inequality we can deduce that

$$f(S \cup T) + f(S \cap T) = f(S) + f(T) \geq a - d'.$$

By the pigeon-hole principle and the fact that $(S \cap T) \subseteq (S \cup T)$ we can thus conclude that

$$f(S \cup T) \geq \frac{a - d'}{2}.$$

Let $\mathcal{X} := \{(x, y) \in \mathcal{P}(x_2) \mid x \in S \cup T\}$ and let $f_{\mathcal{X}}$ be the size of \mathcal{X} , i.e. $f_{\mathcal{X}} = f(S \cup T)$. 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) and it 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 set $S \cup T$. Since \mathcal{X} is contained in $\mathcal{P}(x_2)$ and the latter has a size of a we can conclude that

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

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 14 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 that $d \geq a$ it holds that $d' = (1 - \varepsilon)a$ so

that, using the fact that $\varepsilon \in]0, 1]$, Bounds (7) and (8) give

$$b < a - d \left(1 - \frac{2}{7}\varepsilon\right) \leq a - a \left(1 - \frac{2}{7}\varepsilon\right) = \frac{2}{7}\varepsilon a \text{ and}$$

$$b > a - (1 - \varepsilon)a - \frac{5}{7}\varepsilon a = \frac{2}{7}\varepsilon a,$$

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, either amongst those overlapping with $\Lambda(x_1)$ or those overlapping with $\Lambda(x_2)$, which cut outs an area of size d' and has a short total length. 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-size 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

To apply the above result, we need to find a region of the polygon of size larger than a that does not contain any lines of the cut, so that we can cut out the excess area without interfering with the other lines. 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 line τ then we call \mathcal{T} a tail of the polygon (Figure 17).

Definition 22. For a k -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 line $\tau \in L$ that cuts out \mathcal{T} . We refer to τ as *the line 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 since \mathcal{P}_G is a simple polygon. Notice also that apart from the line of a tail \mathcal{T} there might be other subsets of L that cut out \mathcal{T} if lines in L overlap.

Note that we can shift the staircase line 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

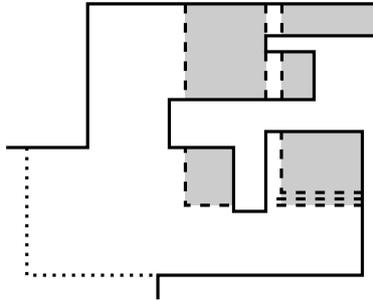


Figure 17: A tail defined by the dotted line. Three (dashed) virtual corner lines cut out the area shaded in grey (The lines overlap on the bottom right).

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 exchange the staircase line by the set of straight and corner lines on the border 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. The other case is when there is a tail contained in, say, the \mathcal{A} -part whose line 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-size.

Lemma 23. *Let L be a k -cut in \mathcal{P} with cut-size l containing one staircase line while all other lines in L are straight and corner lines. There exists a k -cut L' in \mathcal{P} which contains one staircase line and only straight and corner lines otherwise, has cut-size of at most l , and the line of any small tail cut out by L' equals the staircase line.*

Proof. If the line 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 line is $\lambda' \in L \setminus \{\lambda\}$. 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 exchanging λ with λ'' in L yields a k -cut that has a cut-size that is less than l by the length of λ' . The new cut has one straight or corner line less than the old and it contains one staircase line. Repeating this process for any small tail cut out by the new set that does not conform with the sought after property, will eventually terminate in a state in which the resulting k -cut L' fulfils the property. \square

Hence we can replace λ by the corner and straight lines on the border of its deficit and cut out the area a' from the tail, again using the virtual corner lines of short length. In both cases it may be that some of the virtual corner lines end at the line 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 exchange the latter lines (in the same way as suggested by Figure 14). The cut out area is the same while the cut-size 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 24. *Given a k -cut L of a simple polygon \mathcal{P} with cut-size l containing only straight and corner lines except one which is a staircase line, for any desired $\varepsilon \in]0, 1]$ there exists a corner k' -cut L' , where $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, having a cut-size of at most $(6\sqrt{7/\varepsilon} + 7) \cdot l$.*

Proof. Due to Lemma 23 we can assume that any tail cut out by L is not small or its line is the staircase line $\lambda \in L$. Consider the case when there is a tail \mathcal{T} and its line 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 line 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} , i.e. 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 has orientation up-right if λ' is convex w.r.t. \mathcal{T} , and it is down-left if λ' is concave w.r.t. \mathcal{T} . According to Lemma 21 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 have orientation up-right. Let $\lambda'' \in L'$ be a line 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. However Lemma 21 also states that the virtual corner lines 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 lines still is an a' -cut and the cut-size 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 lines to L , the resulting set of lines M' has a cut-size 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 18 we can replace the staircase line λ with a set of corner and straight lines, yielding the set M which only contains straight and corner lines and has a cut-size 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-size is of the desired length. Obviously the set M cuts out an area of size k' where $k' \in [k - \varepsilon a, k + \varepsilon a]$. 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 $k \geq a$. If $\mathcal{T} \subseteq \mathcal{B}(L)$, since the surplus is part of $\mathcal{A}(L)$ obviously $k \geq a$ also holds in this case. Thus $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, which establishes the desired size for the area. For the cut-size, 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 $d < hw \leq h^2$. Hence we can conclude that $l \geq l_\lambda > \sqrt{a}$ which means that the cut-size 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 line is different from λ . This means that the only line in L is λ . In this case we need to proceed differently than in the case before by reversing the transformation of the k -cut: We first remove λ and instead add the lines that, apart from λ , are contained in the boundary of λ 's deficit. This yields a $(k + a)$ -cut M' , where a is the size of the deficit, which contains only straight and corner lines. 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 have orientation up-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 21 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 lines in L' that have a boundary point on one of the lines in M' into feasible lines in \mathcal{P} . Since the orientation of the corner lines in both L' and M' is up-right, any line in L' that has a boundary point in $\mathcal{A}(M')$ on one of the lines 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 lines in L' , add the lines in M' , and thereby yield the set of lines M which only contains straight and corner lines, cuts out an area of size $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, and has a cut-size of at most $(6\sqrt{7/\varepsilon} + 7)l$, which concludes the proof. \square

5 Converting Lines in Polygons to Segments in Grids

We have learnt that for any desired area k to be cut out from an simple polygon there exists a cut of only straight and corner lines that is (1) at most a small amount $\varepsilon \cdot k$ away from the desired area, and (2) at most a constant factor times the optimal cut-size (of arbitrary shape for area k).

Corollary 25. *Let l be the cut-size of an optimal k -cut L in some polygon \mathcal{P} . For any $\varepsilon \in]0, 1]$ there exists a non-crossing corner k' -cut for some $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, which has a cut-size of at most $(54\sqrt{7/\varepsilon} + 63) \cdot l$.*

Proof. According to Corollary 11 we can assume that L only contains straight and corner lines except at most one which either is 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 24. In case L contains a rectangular line we can use Theorem 20 to convert L into a k -cut L' which has a cut-size of at most $9l$ and which contains only straight and corner lines except at most one which is a staircase line. If L' does not contain a staircase line the claim obviously holds. Otherwise, using Theorem 24 on L' , we can convert L into a non-crossing corner k' -cut, for some $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, having a cut-size of at most $(6\sqrt{7/\varepsilon} + 7) \cdot 9l$, which concludes the proof. \square

Because our real interest, however, 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 k -cut in G transforms into a k -cut in \mathcal{P}_G , but not necessarily into an optimal one, since the cut lines 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 in general translate directly into a cut in the grid (note that if we would just cut grid edges with polygon cut lines, 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 line in \mathcal{P}_G happens to lie in integer position however, we will just take the corresponding segment to cut the grid G (Figure 18). For non-grid lines, we start with a clean-up phase that modifies a pair of these lines 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 the unit length open intervals on the polygon boundary between adjacent integer positions, as shown in Figure 18.

Definition 26. Given a grid $G = (V, E)$ let \mathcal{S}_v be the orthogonal 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 line λ in \mathcal{P}_G we refer to the set $\mathcal{K}_\lambda := \{p \in \mathcal{P}_G \mid \exists v \in V : \lambda \in \mathcal{S}_v \wedge p \in \mathcal{S}_v \cup \gamma_v\}$, which is the union over the unit squares, together with their boundaries that are not part of the boundary of \mathcal{P}_G , that intersect with λ , as the *corridor* of λ . A line 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.

Because a grid line λ does not hit any such open unit interval \mathcal{K}_λ , we are concerned only with cut lines that do. For any open unit interval hit by more than one cut line, we can shift one of these cut lines to the boundary and compensate for the area difference by also shifting one other of these cut lines accordingly. Repeating this leaves us with at most one cut line per open unit interval on the boundary.

Lemma 27. *For any grid G and any non-crossing corner k -cut L of cut-size l in \mathcal{P}_G , there is a non-crossing corner k -cut M of cut-size at most l in \mathcal{P}_G such that there is no line in D which includes more than one boundary point of lines in M .*

Proof. Consider the case when there is a line $\delta \in D$ such that at least two lines 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 lines having a boundary point on δ lie above it. This means that any corner line having a boundary point on δ either has orientation down-left or down-right, while any such straight line is vertical. Let $K \subseteq L$ be the set of lines that have a boundary point on δ . Let (x_δ, y_δ) be the lower left corner of the unit square \mathcal{S}_v to which δ is the lower side, and let for

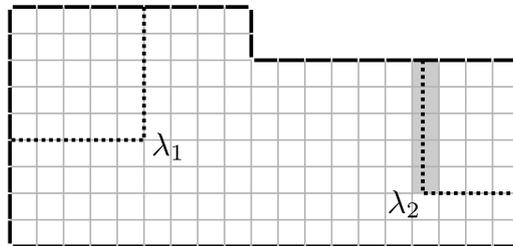


Figure 18: A grid line λ_1 and a non-grid line λ_2 . The corridor of λ_2 is shaded in grey. The boundary of the polygon is divided into the unit length lines.

any line $\lambda \in K$ the point (x_λ, y_λ) either be 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 lines in L are non-crossing, observe that if $\lambda \in K$ is a corner line of orientation down-left any line $\lambda' \in L$ 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 of orientation down-left, the boundary points to the left of p 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 $\lambda \in K$ of orientation down-right, if we consider the open set of points in its corridor that lie to the right of λ . Hence we can order the lines in K by traversing their boundary points on δ from left to right such that we first encounter corner lines of orientation down-left with increasing vertical length, then straight lines, and finally corner lines of orientation down-right with decreasing vertical length. Obviously this is also possible if some of the lines in K share the same boundary point on δ . Let the indices of the lines in $K = \{\lambda_1, \dots, \lambda_{|K|}\}$ denote their position in this order.

We will consider the lines 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 lines that substitute λ_2 such that the resulting set of lines is a k -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 lines $\Lambda_1(x, y)$ as those straight and corner lines for which $\mu \in \Lambda_1(x, y)$ if and only if μ is a segment line and

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

Consider the case when λ_1 and λ_2 are corner lines with orientation down-left. Let $(x, y) \notin \delta$ be the boundary point of λ_1 that does not lie on δ . We know that no line from L intersects $\mathcal{D}(\lambda_1)$ by the observations made above. This means that removing λ_1 and adding the lines in $\Lambda_1(x, y)$ yields a non-crossing k' -cut for some k' . The difference between k and k' equals the size of $\mathcal{D}(\lambda_1)$ and is hence less than the size of $\mathcal{D}(\lambda_2)$. Also the only line in L that intersects $\mathcal{D}(\lambda_2)$ is λ_1 . Hence, after replacing λ_1 , we can find a corner line μ of orientation down-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 k -cut. Notice that the total length of the lines 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 yield a k -cut of smaller cut-size than l . Also the number of boundary points on δ is reduced by one.

In case the lines $\lambda_{|K|}$ and $\lambda_{|K|-1}$ are corner lines of orientation down-right we can use an analogous argumentation as the one given above to derive a k -cut of cut-size smaller than l in which the number of boundary points on δ is reduced by one. By repeating this procedure, we can transform L into a k -cut of smaller cut-size in which there is at most one corner line of orientation down-left

and down-right each, that have boundary points on δ . We thus assume in the remainder of the proof that K contains at most one such corner line each.

Consider the case when λ_1 is a corner line of orientation down-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 lines overlap. If q denotes the corner of λ_1 , clearly we can introduce the corner line of orientation up-left and corner q , remove λ_1 and λ_2 , and thereby yield a k -cut of smaller cut-size than l . The number of boundary points on δ would then be reduced by two. In case λ_1 and λ_2 do not share the same boundary point on δ , we can again replace λ_1 with the lines in $\Lambda_1(x, y)$, exactly as above, yielding a non-crossing k' -cut. Let $\mathcal{L} \subseteq \mathcal{D}(\lambda_2)$ be the connected area of maximal size in the k' -cut, for which $\mathcal{D}(\lambda_1) \subseteq \mathcal{L}$. 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 lies on δ and the other boundary point can either lie on the boundary of \mathcal{P}_G or on a line $\lambda' \in L \setminus \{\lambda_1, \lambda_2\}$. In the former case we can exchange λ_2 with σ and again yield a k -cut which has a smaller cut-size and one boundary point less on δ . Otherwise note that λ' lies in $\mathcal{D}(\lambda_2)$ and hence must be a corner line of orientation up-left, since any other straight or corner line would either cross λ_1 or λ_2 , or would have a boundary point on δ which 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 $\mu = \sigma \cup \sigma'$ of orientation down-left that has a corner on the horizontal bar line of λ' . Removing λ_2 and λ' and introducing μ instead will yield a k -cut with a smaller cut-size and one boundary point less on δ .

Since we assumed that there is at most one corner line of orientation down-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. 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 line from L that intersects $\mathcal{D}(\lambda_1)$ must be a corner line of orientation up-left. Let $L' \subseteq L$ be the lines 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 lines 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 exchange λ_1 with $\Lambda_1(x_\delta, y_{\lambda_1})$. In both cases we yield a k' -cut for some k' . Analogous to the case when λ_1 is a corner line of orientation down-left and λ_2 is a straight line, we can find a line with which to exchange λ_2 , where we possibly also remove some other line in L' , to yield a k -cut of smaller cut-size 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 has orientation down-left, the latter down-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 lines and introduce a straight line that consists of the horizontal bars of λ_1 and λ_2 instead, and thereby yield a k -cut with smaller cut-size than l and with two boundary points less on δ . Consider the case when the two lines share the same boundary point on δ , they have different vertical lengths, and there is a corner line $\lambda' \in L$ of orientation up-left having the same corner as λ_1 . Then we can remove λ_1 , λ_2 , and λ' and introduce the corner line of

orientation up-right that has the same corner as λ_2 , and thereby yield a k -cut of smaller cut-size in which two boundary points on δ are removed. In all other cases let $(x, y) \notin \delta$ be the boundary point of λ_1 that does not lie on δ . Exchanging λ_1 with $\Lambda_1(x, y)$ results in a k' -cut for some k' as before. Let $\mathcal{L} \subseteq \mathcal{D}(\lambda_2)$ be the connected area of maximal size in the k' -cut, for which $\mathcal{D}(\lambda_1) \subseteq \mathcal{L}$. 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 we excluded all cases where $\mathcal{D}(\lambda_1) = \mathcal{L}$. In case there is a corner line μ of orientation down-right that has σ as its vertical bar line and overlaps with λ_2 , we can exchange λ_2 with μ and yield a k -cut that has a cut-size 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 that does not lie on δ is either part of the boundary of \mathcal{P}_G , or it lies on a line $\lambda' \in L$ that must be a corner line of orientation up-left. In the former case we exchange λ_2 with σ and yield a k -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 has orientation down-left, it has σ as its vertical bar line, and its horizontal bar line overlaps with λ' . The second line has orientation up-right, 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 a k -cut with cut-size at most l , again 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 lines in D is never increased. Hence by repeating the above procedure for the lines in K and then in the same manner for all lines in D that include more than one boundary point of lines in the resulting k -cut, we yield a k -cut that has a smaller cut-size than l in which any line in D includes at most one boundary point of a line in the k -cut. \square

Similarly, 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 accordingly, to compensate for the area difference. This ends in a situation with at most one non-grid line in the cut. During the whole process, the cut length does not increase (it might even decrease, since the cut was not necessarily optimal).

Lemma 28. *For any grid G and any non-crossing corner k -cut L of cut-size l in \mathcal{P}_G , there is a non-crossing corner k -cut M of cut-size at most l in \mathcal{P}_G such that all lines in M except at most one are grid lines.*

Proof. According to Lemma 27 we can assume that L contains no two lines that have boundary points that lie on the same line from D . Let $K \subseteq L$ be the set of lines that are not grid lines and assume that $|K| \geq 2$. For a straight line $\lambda \in K$, any line from K that intersects the corridor \mathcal{K}_λ must have a boundary point on the same line from D as λ . Hence no set of lines 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 line in K that intersects $\mathcal{X}(\lambda)$ must have a boundary point on the same line in D as λ . Therefore no line in L intersects with $\mathcal{X}(\lambda)$.

We can thus conclude that if $\lambda_1, \lambda_2 \in K$ is a pair of lines with intersecting corridors, then λ_1 and λ_2 must be corner lines and λ_i , for $i \in \{1, 2\}$, must intersect $\overline{\mathcal{X}}(\lambda_j)$, where $j \in \{1, 2\} \setminus \{i\}$. Assume w.l.o.g. that λ_1 has orientation down-left. Observe that this means that λ_2 has orientation up-right and no other line in K can intersect the corridors of λ_1 or λ_2 , since otherwise there would be lines in L that have boundary points on the same line from D . 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 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 27 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 exchange λ_1 with $\Lambda_1(x, y)$, as defined in Lemma 27, where (x, y) is the boundary point of the horizontal bar of λ_1 , and yield a non-crossing k' -cut for some k' . If $l_1 < l_2$ there exists a corner line λ'_2 of orientation up-right that contains the horizontal bar line of λ_2 and intersects $\mathcal{D}(\lambda_2)$, such that exchanging λ_2 with λ'_2 will yield a non-crossing k -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 k -cut has a cut-size of at most l since $l_1 \leq l_2$ and hence σ_1 was “moved” farther (or equally far) to the left than σ_2 . Also note that the new cut has at least one line less in D containing a boundary point since the vertical bar lines in $\Lambda_1(x, y)$ are grid lines.

For any two lines in K with non-intersecting corridors we can use an analogous transformation as above. Since each transformation yields a k -cut in which there is at least one line less in D with a boundary point, by repeating the above procedure we can transform L into the k -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 18), 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 suitable short extra grid lines exist. In the end, this will preserve the property that the cut out area lies in the interval defined by k and ε , but will increase the cut-size by a constant 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 $\lceil \log_2(\lceil a \rceil) \rceil$ 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 19). 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 lines to cut out area a is limited to $3a$.

Definition 29. 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 $\tilde{x} > x$ and $(\tilde{x}, y) \notin \mathcal{P}_G$, such that $\lambda \in \Lambda$ if and only if

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

We call the part $\{(\tilde{x}, y') \in \mathcal{P}_G \mid y' \in [y - 1, y]\}$ the *bend*, 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 k on the upper right side of its corner, we say that it is a virtual pseudo-corner line *for* k .

Hence 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.

Lemma 30. 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 lines in $\Lambda \cup L$ are non-crossing and the cut-size 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 31 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 lines 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—the transformation is described below. 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} . Hence if this

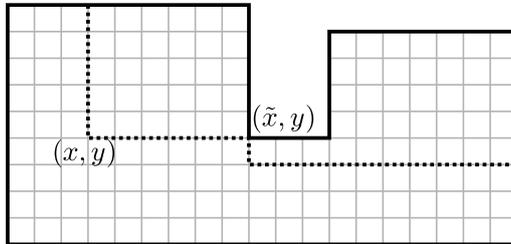


Figure 19: A virtual pseudo corner line with its corner at (x, y) and its unit sized step at \tilde{x} .

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 $\lceil \log_2(a) \rceil$ many 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 many 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 line 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 conclude that $a'_i \leq a/2^i$ which means that $l_i \leq a/2^i + 2$. Therefore the cut-size l of L is

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

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

What remains to be shown is that we can convert the line 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. Hence we can 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 a 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 the boundary of $\mathcal{A}(\Lambda'_{i+1})$ then we include all segment lines $\lambda \subseteq \gamma \setminus \beta$ 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 lines from Λ_i and Λ'_{i+1} . However, since these sets contain only grid lines and the length of a bend is 1, $\gamma \setminus \beta$ can contain at most one bend 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

We still need to convince ourselves that such a 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. This is because such a set of lines can be seen as a virtual pseudo corner line, as the unit step of the underlying staircase line of the latter can entirely lie outside of the polygon. In the other case a suitable set of lines can be constructed using three virtual corner lines at some integer points (x^*, y^*) , $(x^* + 1, y^*)$, and $(x^*, y^* + 1)$ (Figure 20). 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 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 lines cuts out an area between a and $3a/2$, and the upper most point of the unit sized step is on the boundary or outside of the polygon.

Lemma 31. *For any grid G with n vertices and any $b \in \{0, \dots, n\}$, for some $a \in [b, \frac{3}{2}b]$ there exists a virtual pseudo-corner line Λ for a in \mathcal{P}_G .*

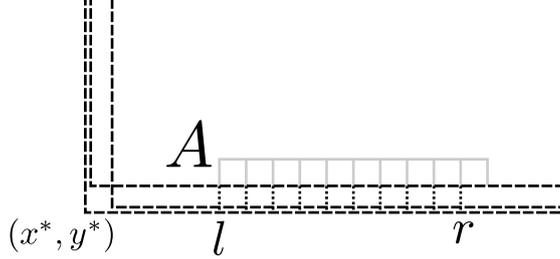


Figure 20: The virtual corner line at (x^*, y^*) cuts out an area of size more than $\frac{3}{2}b$ while those at $(x^* + 1, y^*)$ and $(x^*, y^* + 1)$ cut out at most $b - 1$. If all points $(x^* + i, y^* + 1)$, where $i \in \{l, \dots, r\}$, are inside the polygon, then all unit squares from A are also inside. Since there are $r - l + 1 > b/2$ of them, the area $\mathcal{P}(x^* + 1, y^* + 1)$ has a 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 , where $\mathbb{H} := \{x - \frac{1}{2} \mid x \in \mathbb{N}\}$. 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. Thus 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 $\frac{3}{2}b + 1$ 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 $\frac{3}{2}b + 1$ 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 $\lfloor \frac{3}{2}b \rfloor$ and the size of $\mathcal{P}_{01} \cup \mathcal{P}_{r0}$ equals b . If for a value $c \in \{l, \dots, r\}$ there is a pair of crossing lines in $\Lambda_{01} \cup \Lambda_{c0}$ their crossing point is $p_{c1} := (x^* + c, y^* + 1)$. If however there exists a corresponding value for c such that there are no lines in $\Lambda_{01} \cup \Lambda_{c0}$ that cross then let Λ include the lines to the left of p_{c1} from the first set together with the lines below p_{c1} from the second set, i.e.

$$\Lambda = \{\lambda \in \Lambda_{01} \mid \forall (x, y) \in \lambda : x < x^* + c\} \cup \{\lambda \in \Lambda_{c0} \mid \forall (x, y) \in \lambda : y < y^* + 1\}.$$

Clearly the set Λ fulfils the lemma. Hence it remains to show that we can always find a corresponding value c such that p_{c1} is not in \mathcal{P}_G .

Assume this is not the case, i.e. for any value $c \in \{l, \dots, r\}$ it holds that $p_{c1} \in \mathcal{P}_G$. This means that any unit square that has one of these points as a

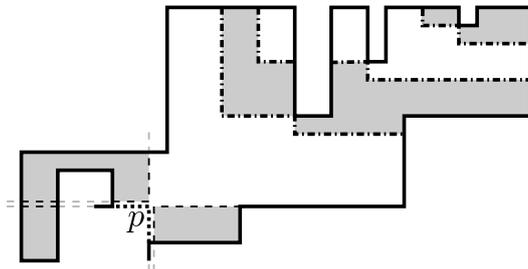


Figure 21: A tail with its corner line at p (black dotted). The excess area in grey is cut out using the four virtual corner lines at p (thin dashed) together with the recursive method that uses virtual pseudo-corner lines (dashed and dotted).

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_{c1} 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_{c1} . Hence 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 $\frac{1}{2}b + 2$. 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 $\frac{1}{2}b + 1$. 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 $\frac{1}{2}b - 2$. However this contradicts the lower bound derived above and hence the lemma holds. \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 21), as for the staircase line argument in the previous section. We have to make sure that there is a tail that is big enough to support the area of size a . For a non-crossing corner cut L containing only one line 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 $\mathcal{K}_\lambda \cap \mathcal{A}(L)$. Obviously the proof of the following lemma is analogous to the proof of Lemma 23.

Lemma 32. *For a grid G , let L be a non-crossing corner k -cut in \mathcal{P}_G with cut-size l containing exactly one line λ that is not a grid line. There exists a non-crossing corner k -cut M in \mathcal{P}_G which contains exactly one line that is not a grid line, has cut-size of at most l , and the line of any tiny tail cut out by M equals the line that is not a grid line.*

However, we need to make sure that no additional lines are produced while cutting out the area of size a which would increase the cut-size 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 line 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 difference still needed 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 33. *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 lines in $M \cup L$ are non-crossing. Furthermore, the cut-size of M is at most $3a$.*

Proof. Let $\lambda \in L$ be the line 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. Definition 12 only covers the case when the orientation of the virtual corner line is up-right. However it is straightforward to define virtual corner lines having any of the four possible orientations and we will make use of such lines in this proof. Thus 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. Hence assume that $0 < a < b$. Thus there exists a (possibly empty) subset $S \subset I$ such that the size b_S of the union area $\mathcal{T}_S := \bigcup_{i \in S} \mathcal{T}_i$ is at most a while for any $j \in I \setminus S$ the size of the area $\mathcal{T}_S \cup \mathcal{T}_j$ is greater than a . For each $i \in S$ the set M contains the lines 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}_S includes λ we also include λ in M . So far these lines cut out an area of size b_S from \mathcal{P}_G .

Since all involved lines 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 line 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-size of $\bigcup_{i \in S} \Lambda_i$ is at most $b_S + |S|$.

Let $j \in I \setminus S$. According to Lemma 30 we can find a set M' of non-crossing corner grid lines in \mathcal{P}_G that cut out an area of size $a - b_S$ from \mathcal{T}_j such that the cut-size of M' is at most $3(a - b_S)$. Hence if we also include M' in M we cut out an area of size a from \mathcal{T} without crossing a line in L . Furthermore the cut-size of M is at most

$$b_S + |S| + 3(a - b_S) = 3a + |S| - 2b_S \leq 3a,$$

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

Thus the main result of this article, as stated in Theorem 1, follows from the following theorem which summarises the results of this section.

Theorem 34. *Let l be the cut-size of an optimal k -cut L , for some $k \in \mathbb{N}$, in the polygon \mathcal{P}_G of a grid G . For any $\varepsilon \in]0, 1]$ there exists a non-crossing corner k' -cut L' for some $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, such that all lines in L' are grid lines and the cut-size is at most $(216\sqrt{7/\varepsilon} + 260) \cdot l$.*

Proof. We can apply Corollary 25 and Lemma 28 to L , i.e. we know that there exists a non-crossing corner k' -cut M , for some $k' \in [(1 - \varepsilon)k, (1 + \varepsilon)k]$, with cut-size at most $(54\sqrt{7/\varepsilon} + 63) \cdot l$ in \mathcal{P}_G such that M contains at most one line

that is not a grid line. If all lines in M are grid lines we are done. If not then let $\lambda \in M$ be the line that is not a grid line. In case there exists a tail cut out by M such that λ is not its line, let \mathcal{T} be this tail and assume w.l.o.g. that $\mathcal{T} \subseteq \mathcal{A}(M)$. By Lemma 32 we can assume that the size of \mathcal{T} is at least the size of $\mathcal{K}_\lambda \cap \mathcal{B}(M)$. Notice that, if a denotes the size of $\mathcal{K}_\lambda \cap \mathcal{B}(M)$, a is not necessarily an integer since k' might not be a natural number. However we can furthermore conclude that the size of \mathcal{T} is at least $\lceil a \rceil$ since the line 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 exchanging the line λ with the lines in Λ we yield a k'' -cut M' where $k'' = k' + a$.

We attempt to cut out the excess area of size a in \mathcal{T} using only grid lines. Notice that k'' must be an integer since M' contains only grid lines. If we assume w.l.o.g. that $k' \geq k$, since k is also an integer this means that $k'' - \lceil a \rceil$ is a natural number in the interval $[k, k']$ which is contained in $[(1 - \varepsilon)k, (1 + \varepsilon)k]$. Using Lemma 33 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 $[(1 - \varepsilon)k, (1 + \varepsilon)k]$. Hence it remains to show (for the case when λ is not the line of \mathcal{T}) that the cut-size 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-size l' of M we can conclude that $\lceil a \rceil \leq a + 1 \leq l' + 2$. By Lemma 33 this means that the cut-size 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-size of M' is at most the cut-size of M plus 2. Therefore the cut-size of L' is at most $l' + 2 + 3(l' + 2) \leq (216\sqrt{7/\varepsilon} + 252) \cdot l + 8$. Clearly 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-size l of at least 1. Hence in this case the latter bound on the cut-size 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 line. Clearly this can only mean that there are two tails which both have λ as their line and λ is the only line in M . Let \mathcal{T} be the tail that corresponds to the area $\mathcal{A}(M)$. Exchanging λ with Λ as before, we yield a k'' -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 and even though $\mathcal{A}(M')$ might not be a tail since Λ may contain more than one line, proving an analogous statement as Lemma 33 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. \square

References

- [1] P. Arbenz and R. Müller. Microstructural finite element analysis of human bone structures. *ERCIM News*, 74:3132, 2008.

- [2] J. Díaz, M. J. Serna, and J. Torán. Parallel approximation schemes for problems on planar graphs. *Acta Informatica*, 33(4):387–408, 1996.
- [3] A. E. Feldmann, S. Das, and P. Widmayer. Simple cuts are fast and good: Optimum right-angled cuts in solid grids. In *Proceedings of the 4th Annual International Conference on Combinatorial Optimization and Applications*, 2010.
- [4] A. E. Feldmann, S. Das, and 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*, 2011.
- [5] M. R. Garey and D. S. Johnson. *Computers and Intractability: A guide to the Theory of NP-completeness*. W.H. Freeman and Co., San Fransisco, California, 1979.
- [6] M. Goldberg and Z. Miller. A parallel algorithm for bisection width in trees. *Computers & Mathematics with Applications*, 15(4):259–266, 1988.
- [7] R. M. MacGregor. *On partitioning a graph: a theoretical and empirical study*. PhD thesis, University of California, Berkeley, 1978.
- [8] C. Papadimitriou and M. Sideri. The bisection width of grid graphs. *Theory of Computing Systems*, 29:97–110, 1996.
- [9] H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, 2008.
- [10] G. Strang. Maximum area with Minkowski measures of perimeter. *Proceedings of the Royal Society of Edinburgh: Section A Mathematics*, 138(01):189–199, 2008.