On Constants for Cuttings in the Plane[*]

J. Matoušek

Department of Applied Mathematics, Charles University,
Malostranské nám. 25, 118 00 Praha 1, Czech Republic
matousek@kam.mff.cuni.cz

Abstract.   A theorem of Chazelle and Friedman with numerous applications in combinatorial and computational geometry asserts that for any set L of n lines in the plane and for any parameter $r > 1$ there exists a subdivision of the plane into at most Cr^2 (possibly unbounded) triangles, C a constant, such that the interior of each triangle is intersected by at most n/r lines of L. (Such a subdivision is called a $(1/r)$-cutting for L.) We give upper and lower bounds on the constant C. We also consider the canonical triangulation of the arrangement of a random sample of r lines from L. Although this typically is not a $(1/r)$-cutting, the expectation of the kth degree average of the number of lines intersecting a triangle is $O(n/r)$ for any fixed k. We estimate the constant of proportionality in this result.

1.   Introduction

Motivation.   Numerous algorithms developed in computational geometry use a geometric version of the divide-and-conquer paradigm. If the input is a collection of, say, n hyperplanes or simple algebraic surfaces in the d-dimensional Euclidean space, the space is partitioned into several pieces (of a possibly simple shape, such as simplices), and the situation within each of these pieces is dealt with separately, often by a recursive call of the same algorithm. Two survey papers considering such algorithms are [15] and [1].

This approach yields the asymptotically most efficient known algorithms for a variety of problems; many of these problems are of considerable interest. However, the constants of proportionality hidden in the asymptotic notation for the complexity of such geometric divide-and-conquer algorithms are seldom made explicit, and trying to estimate them by following the proofs seems to lead to very large, often astronomically large, values

(this phenomenon is not restricted to this particular type of algorithms, of course, but it appears especially pronounced for them). A consequence of this may be a widespread belief that such algorithms, although theoretically nice, are practically quite useless. This is probably completely true for complicated higher-dimensional algorithms which only bring a small theoretical saving in the complexity (improving a naive $O(n^2)$ algorithm to $O(n^{1.8})$, say). On the other hand, in some simple situations, such as divide-and-conquer in the plane, a closer inspection may show that the algorithms are not so bad or perhaps they can be made practical by some additional tricks (which are uninteresting for the asymptotic analysis). The final verdict in this matter should probably come from implementation experience, but we believe this should be preceded by further theoretical work focused on analyzing and improving the constants in the algorithms. In such an effort, encouragement may be gained by the fact that the idea of *binary space partitions*, which are data structures used frequently and successfully in practice, is similar in spirit to geometric divide-and-conquer.

In this paper we investigate quantitative bounds for geometric divide-and-conquer in the simplest setting: Given a set L of n lines in the plane, we want to subdivide the plane into a possibly small number of regions in such a way that each region is intersected by at most n/r of the lines, where r is a prescribed parameter (the factor by which the size of the resulting subproblems is reduced compared with the original problem). In the literature the regions are usually postulated to be triangles (possibly unbounded ones, i.e., intersections of three half-planes). In this case the subdivision with the above property is called a $(1/r)$-cutting for L. It is known that, for any L, a $(1/r)$-cutting consisting of $O(r^2)$ triangles can always be produced. This (asymptotically optimal) result was first proved by Chazelle and Friedman [8] by a probabilistic method, improving a previous slightly weaker bound of $O(r^2 \log^2 r)$ proved by Clarkson [9] and implicitly contained also in the paper Haussler and Welzl [16]. Two other, considerably different proofs of the $O(r^2)$ bound were given by Matoušek [17] and by Chazelle [6].

Size of Cuttings.   In this paper we investigate the constant of the leading quadratic term in the worst-case bound for the number of triangles in an optimal $(1/r)$-cutting, that is, we seek asymptotic upper and lower bounds of the form $Cr^2 + o(r^2)$ with numerical values of C.

For the purpose of obtaining a $(1/r)$-cutting with few regions, triangles do not seem to be the best shape. For algorithmic applications, trapezoids are equally suitable, and we can obtain a better bound for them. We also consider $(1/r)$-cuttings consisting of completely arbitrary sets. These are problematic to handle algorithmically, of course, but they may be interesting for some combinatorial geometry applications, and lower bounds obtained for them are strongest.

First we give a definition of a $(1/r)$-cutting which seems appropriate for arbitrary regions.

Let $\sigma \subseteq \mathbb{R}^2$ be a set, and let $\ell \subset \mathbb{R}^2$ be a line. We say that ℓ *crosses* σ if σ intersects both the open half-planes defined by ℓ. A $(1/r)$-*cutting* for a set L of n lines in the plane is a collection $\sigma_1, \sigma_2, \ldots, \sigma_m$ of subsets of the plane, such that $\bigcup_{i=1}^m \sigma_i = \mathbb{R}^2$, and each σ_i is crossed by at most n/r lines of L. (The reader may wonder, similarly as the referees did, if the author did not intend to assume that the regions σ be connected or something like that. Well, we really admit completely arbitrary regions, and the explanation is in

our definition of crossing: note that a line may cross a disconnected region without intersecting it, and that any region in a line arrangement crossed by few lines is bound to lie within a small portion of the arrangement.)

The following theorem summarizes our upper bounds:

Theorem 1. *Let $r > 1$ be a parameter, let L be a set of n lines in the plane, and suppose that n/r is an even integer* (*this technical assumption is made to avoid tedious calculations with integer parts*). *Then there exist*

  (i) *a $(1/r)$-cutting for L consisting of at most $4r^2 + 2r + 2$ connected sets, and*
  (ii) *a $(1/r)$-cutting for L consisting of at most $8r^2 + 6r + 4$ trapezoids.*

For a $(1/r)$-cutting consisting of triangles, we currently have nothing substantially better (in the worst case) than taking a cutting consisting of trapezoids as in (ii) and subdividing each trapezoid into two triangles, which produces $16r^2 + O(r)$ triangles.

The $(1/r)$-cuttings in this theorem are all constructed by a modification of the method of [17]. The other two known proofs of existence of asymptotically optimal cuttings currently seem to yield much worse bounds.

A simple argument showing that a $(1/r)$-cutting has to consist of at least $\Omega(r^2)$ sets is as follows (the idea is probably due to Haussler and Welzl (unpublished)): The arrangement of n lines in general position has $\Omega(n^2)$ two-dimensional cells. A set crossed by n/r lines can only meet $O((n/r)^2)$ of these cells, and hence $\Omega(r^2)$ sets are needed to cover all cells. (The argument is usually presented with counting vertices instead of cells but then there is a technical complication concerning the vertices on the boundaries of the sets in the $(1/r)$-cutting.) Our lower bound proofs are all based on this idea; for three types of shapes of sets in the cutting we refine it in different directions.

Theorem 2. *For any sufficiently large r there exist arbitrarily large sets L of lines in the plane such that*

  (i) *Any $(1/r)$-cutting for L consisting of arbitrary sets has at least $2(1 - o(1))r^2$ sets.*
  (ii) *Any $(1/r)$-cutting for L consisting of quadrilaterals has at least $2.54(1-o(1))r^2$ sets.*
  (iii) *Any $(1/r)$-cutting for L consisting of triangles has at least $4(1 - o(1))r^2$ sets.*

Cuttings with Good Averages. Let $\sigma_1, \sigma_2, \ldots, \sigma_m$ be sets covering the plane, and let $w(\sigma_i)$ be the number of lines crossing σ_i. In the definition of a $(1/r)$-cutting, we insist that $\max_i w(\sigma_i) \leq n/r$, but sometimes this may be an unnecessarily strong condition— for some applications, it might be sufficient that "most" of the regions are intersected by few lines. One way to formalize this would be to require that the kth degree average of the $w(\sigma_i)$ is at most n/r, for some suitable fixed k. The following definition adapts the concept of *semicuttings* of Chazelle [7] to our circumstances.[1]

---

[1] This concept has a longer history in computational geometry. The paper [12] uses the idea of replacing maxima by suitable averages in geometric random sampling in an essential way. The first existence proof for asymptotically optimal $(1/r)$-cuttings [8] essentially constructs a suitable semicutting first. Newer applications in computational geometry papers are abundant.

A collection $\sigma_1, \sigma_2, \ldots, \sigma_m$ of sets covering the plane is called a kth *degree* $(1/r)$-*semicutting* for a set L of n lines if

$$\left(\frac{1}{m} \sum_{i=1}^{m} w(\sigma_i)^k \right)^{1/k} \leq \frac{n}{r}.$$

Similarly as for $(1/r)$-cuttings, we seek estimates on the constant C_k in a bound of the form $C_k r^2 (1 + o(1))$ for the size of a kth degree $(1/r)$-semicutting for a finite set of lines. This question only makes a good sense for $k > 2$, since for $k \leq 2$, taking $m = r^2$, $\sigma_1 = \mathbb{R}^2$, and $\sigma_2 = \sigma_3 = \cdots = \sigma_m = \emptyset$ makes an excellent $(1/r)$-semicutting of size r^2 according to the definition, although it is certainly not something close to the intuitive notion of a "nearly $(1/r)$-cutting."

It can be checked that the lower bounds in Theorem 2 apply to the size of kth degree $(1/r)$-semicuttings for any fixed $k > 2$ (the observation needed to augment the proofs in Section 3 is that such a kth degree $(1/r)$-semicutting consisting of $O(r^2)$ sets must also be a $(1/r')$-cutting for $r' = \Omega(r^{1-2/k})$; we omit the details). We investigate upper bounds attained by the following construction, very popular in computational geometry: Fix a suitable probability p, and choose a random sample $R \subseteq L$ of the given lines by picking each line of L independently with probability p. Let $\sigma_1, \sigma_2, \ldots, \sigma_m$ be the triangles in a suitable triangulation of the arrangement of R. As follows from the work of Chazelle and Friedman [8], if we choose $p = Kr/n$ for an appropriate constant K, with a positive probability this yields a kth degree $(1/r)$-semicutting for L of size $O(r^2)$. Estimates for the constant of proportionality achievable by this method are studied in Section 4. For instance, we show that this gives a third degree $(1/r)$-semicutting of size $11.8r^2$ for a large enough r.

Open Problems.   Since we have not obtained matching upper and lower bounds, an obvious question is how to improve the bounds. It seems that for improving the upper bound, a considerably different construction is needed.

An interesting question, which we have left aside, is to give good numeric bounds for $(1/r)$-cuttings for small specific values of r, e.g., for $r = 2$. We remark that a construction of Yamamoto et al. [20] can be easily modified to yield a $(7/8)$-cutting consisting of four quadrants (defined by two lines).

Other obvious problems concern generalizations to other curves than lines and higher-dimensional generalization (for higher dimension, the cutting construction technique of [17] does not seem to be applicable anymore).

2.   Cutting Constructions

For simplicity, we suppose that the given set L of n lines is in general position, i.e., every two lines intersect in exactly one point, no three have a common point, and the x-coordinates of all intersections are pairwise distinct. (If not, we perturb L slightly to get general position and construct the $(1/r)$-cutting. A limit of these $(1/r)$-cuttings as the amount of perturbation goes to zero yields the desired $(1/r)$-cutting for L—we omit the details.)
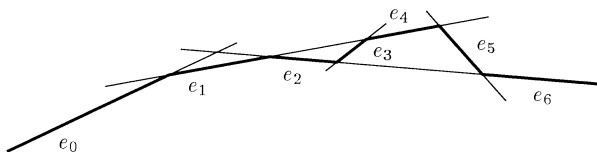
Fig. 1. A level in a line arrangement.

The constructions use the notion of a level in a line arrangement. The *level* of a point in the plane is the number of lines of L lying strictly below it. Consider the set E_k of all edges of the arrangement of L having level k (where $0 \leq k < n$). These edges form an x-monotone polygonal line, which looks as in Fig. 1. This polygonal line is called the *level k* of the arrangement of L. Let $e_0, e_1, \ldots, e_t$ be the edges of E_k (numbered from left to right; e_0 and e_t are the unbounded rays).

Cutting into Connected Sets.   First we describe a construction of a $(1/r)$-cutting consisting of connected sets, as in Theorem 1(i). Let r be the given parameter, and let $n = |L|$. We set $q = n/2r$ (by our assumption, this is an integer). We divide the levels 0 through $n - 1$ into q groups: the ith group contains all levels k with k congruent to i modulo q ($i = 0, 1, \ldots, q - 1$). Since the total number of edges in the arrangement is n^2, there is an i such that the ith group contains at most n^2/q edges. We fix one such i. Let s be the number of the levels in this group, and let the levels in this group be denoted by $\Lambda_0, \Lambda_1, \ldots, \Lambda_{s-1}$, i.e., Λ_j denotes the level $jq + i$. We have $s \leq \lceil n/q \rceil \leq 2r + 1$ (note that we do not assume r to be an integer). The Λ_j's divide the plane into $s + 1$ regions $R_0, R_1, \ldots, R_s$, where R_0 and R_s are the unbounded regions.

We further subdivide each region R_j by vertical segments. Namely, we let $x_1, x_2, \ldots$ be the x-coordinates of the vertices of the boundary of R_j (from both the top and bottom portions) listed in the increasing order, and we slice R_j by vertical segments with x-coordinates $x_{2q}, x_{4q}, x_{6q}, \ldots$; see Fig. 2.

Let t_j be the number of edges of the level Λ_j; by the choice of the Λ_j's we have $\sum_{j=0}^{s-1} t_j \leq n^2/q$. The levels Λ_{j-1} and Λ_j bounding the region R_j together have $t_{j-1} + t_j - 2$
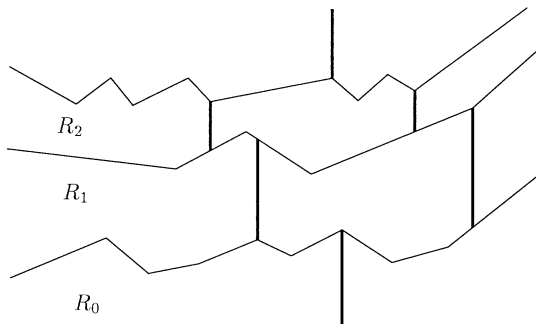

Fig. 2. Subdividing the regions $R_0, R_1, \ldots, R_s$ (with $q = 3$).
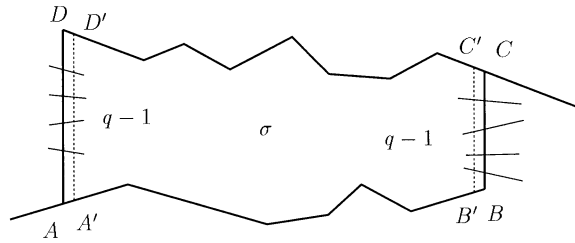
Fig. 3.   Bounding the number of lines crossing σ.

vertices, hence R_j is sliced into no more than $(t_{j-1} + t_j)/2q + 1$ sets (for consistence we define $t_{-1} = t_s = 0$). Therefore the total number of sets formed is at most $\sum_{j=0}^{s}[(t_{j-1} + t_j)/2q + 1] \le n^2/q^2 + s + 1 \le 4r^2 + 2r + 2$.

We claim that these sets form a $(1/r)$-cutting for L, that is, each of them is crossed by at most $n/r = 2q$ lines of L. To check this, we consider one such set, σ, as in Fig. 3 (we assume that it is bounded; modifications for the unbounded cases are straightforward). Each line crossing σ shares a segment with the upper or lower portion of the boundary of σ, or crosses one of the segments $A'D'$, $B'C'$ (where $A'D'$ is a vertical segment connecting the upper and lower boundaries of σ and lying very close to the segment AD, and similarly for $B'C'$). Moreover, each line contributes at least two such situations, one when entering the interior of σ and one when leaving it. The upper and lower portions of the boundary of σ together consist of $2q + 1$ segments, and each of the segments $A'D'$ and $B'C'$ is crossed by $q - 1$ lines. Together this gives the desired bound of at most $2q$ lines crossing σ. Part (i) of Theorem 1 is proved.

Cutting into Trapezoids.   Here we need the notion of *q-simplification* of a level, which has been introduced by Edelsbrunner and Welzl [14]. In order to avoid some technicalities and case analysis in the subsequent proof, we modify their original notion very slightly. Consider the level k with edges numbered $e_0, e_1, \ldots, e_t$ (from left to right), and choose a point p_i in the interior of the edge e_i. The q-simplification of level k is defined as the monotone polygonal line containing the part e_0 up to the point p_0, the segments $p_0 p_q$, $p_q p_{2q}, \ldots, p_{\lfloor(t-1)/q\rfloor q} p_t$, and the part of e_t from p_t on. Thus, the q-simplification has at most $t/q + 1$ vertices. Figure 4 illustrates this for $t = 9$, $q = 4$. The new feature of our definition is that we choose the vertices of the q-simplification inside edges, while in the literature they are chosen among vertices of the original level. We need the following property of q-simplification:

Lemma 3.   *The q-simplification of level k is contained in the (closed) strip between levels $k - \lceil q/2 \rceil$ and $k + \lceil q/2 \rceil$.*

Proof.   For the original notion of q-simplification, this was proved by Edelsbrunner and Welzl [14], and the proof remains almost the same for our modified notion (it becomes simpler if anything). We sketch it for completeness.
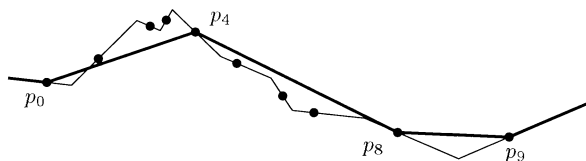
Fig. 4.   The 4-simplification of a level.

Imagine walking along some segment $p_j p_{j+q}$ of the q-simplification. We start at an endpoint, which has level k. Our current level may only change as we cross a line of L. Moreover, having traversed the whole segment we must be back to level k. Thus, to get from level k to $k + i$ and back to k we need to cross at least $2i$ lines on the way. However, each of these lines shares a segment (edge) with the portion of the level k between the points p_j and p_{j+q}, and this portion has $q + 1$ edges. From this we get $i \leq \lfloor (q + 1)/2 \rfloor = \lceil q/2 \rceil$.   □

For the construction of a $(1/r)$-cutting consisting of trapezoids, we again set $q = n/2r$ and define the levels $\Lambda_0, \Lambda_1, \ldots, \Lambda_{s-1}$ exactly as in the previous proof. We let P_j be the q-simplification of level Λ_j. We note that the polygonal chains P_j never intersect properly (if they did, a vertex of some P_j, which has level $qj + i$, would be strictly above P_{j+1}, and this is ruled out by Lemma 3).

We form the vertical decomposition for the P_j's, that is, we extend vertical segments from each vertex of each P_j upward and downward, until it hits P_{j-1} and P_{j+1}. The P_j's initially divide the plane into $s + 1$ regions. For each vertex of each P_j, the vertical subdividing segments erected from it increase the number of regions by two. The total number of vertices of the P_j's is no more than $\sum_{j=0}^{s-1}(t_j/q + 1) \leq n^2/q^2 + s$, so altogether the number of the resulting trapezoids is at most $2n^2/q^2 + 3s + 1 \leq 8r^2 + 6r + 4$.

It remains to prove that each trapezoid is crossed by at most $2q$ lines of L. So we look at some trapezoid $\sigma = ABCD$ in the strip between P_j and P_{j+1}; AB is the bottom side and CD the top side. Suppose AB is contained in an edge EF of P_j, and CD is an edge of P_{j+1} (see Fig. 5; few other possible cases are discussed similarly and are omitted). Let A_1 be the intersection of level Λ_j with the vertical line AD, and similarly let B_1 be the intersection of Λ_j with the vertical line BC.

Each line of L intersecting the segment CD, including those passing through C and D, shares a segment with the portion of level Λ_{j+1} between the points C and D, so the number of such lines is at most $q + 1$. Similarly we find that the segment EF is intersected by at most $q + 1$ lines; two of them intersect it at the endpoints, so the interior of the segment EF has at most $q - 1$ intersections. The interior of each of the segments DA_1 and CB_1 is intersected by exactly $q - 1$ lines. Hence if A_1 and B_1 lie both on or below the segment EF, we find that the total number of intersections of lines with the boundary of the considered trapezoid $ABCD$ is at most $4q - 3$.

More interesting cases arise when A_1 or B_1 lie above EF. Suppose for instance that A_1 lies above EF and B_1 on it or below (as in Fig. 5), and there are a lines intersecting the interior of the segment AA_1 (by a general position assumption, we may suppose that
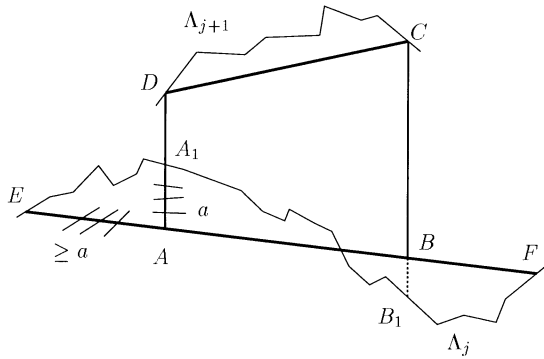
Fig. 5.   Bounding the number of lines intersecting a trapezoid.

no line passes through A). The point A_1 has level exactly $k = jq + i$, and so the point A has level $k - a$. At the same time, the point E has level k again, so the segment EA must be crossed by at least a lines, and thus the segment AB is intersected by no more than $q - 1 - a$ lines. Summarizing, the total number of intersections of the lines of L with the boundary of $ABCD$ is at most $q + 1$ (for the side CD) plus $2(q - 1)$ (for interiors of AA_1 and CB) plus $a + 1$ (for A_1A) plus $q - 1 - a$ (for AB). Together we get $4q$, and a similar counting for the case when both A_1 and B_1 lie above EF yields at most $4q + 1$ intersections. In all cases, the total number of lines crossing the considered trapezoid is at most $2q$. This proves Theorem 1(ii).   □

Remark.   One can also consider bounds for a $(1/r)$-cutting consisting of k-gons, i.e., sets whose boundary consists of at most k segments. By a construction similar to the two given above, an upper bound of $4((k - 2)/(k - 3))r^2 + O(r)$ can be obtained for the number of k-gons in such a $(1/r)$-cutting.

3.   Lower Bounds

General Sets and Triangles.   As was mentioned in the Introduction, our lower bound proofs are based on the idea that a single set in a $(1/r)$-cutting cannot cover too many cells (or vertices) of the arrangement. A set σ crossed by q lines cannot be incident to more cells than there are in an arrangement of q lines, i.e., $q^2/2 + O(q)$. However, if the set σ lies in a portion of the arrangement where the pattern of lines looks like a rectangular grid, then σ can only cover $q^2/4 + O(q)$ cells (if it is crossed by some a "vertical" and b "horizontal" lines, $a + b \leq q$, then it is incident to at most $(a + 1)(b + 1)$ cells). So we need to exhibit an arrangement where almost all cells lie within approximate affine copies of grid-like patterns.

We construct arrangements $L_0, L_1, \ldots$ inductively; L_0 consists of an arbitrary single line and L_k has 2^k lines $\ell_1, \ell_2, \ldots, \ell_{2^k}$. For the induction, we assume that, for L_k, there is a parallelogram P_k that is intersected by all lines of L_k in "regular stripes." More precisely, we require that the parallelogram P_k can be affinely mapped to the rectangle
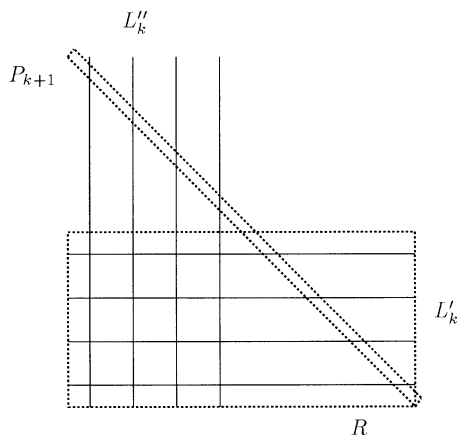
Fig. 6. Constructing L_{k+1} from L_k.

$R = [0, 2] \times [0, 1]$ so that each segment $P_k \cap \ell_i$ is mapped to a segment with Hausdorff distance at most ε_k from the horizontal segment $\{(x, y) \in R;\ y = (i - \frac{1}{2})/2^k\}$. Here $\varepsilon_k > 0$ is a real number which can be assumed as small as needed for the subsequent steps to work.

In the inductive step we take an affine copy L'_k of L_k under an affine mapping sending P_k to R as in the inductive assumption. Then we let L''_k be L_k reflected around the line $x = y$, and we set $L_{k+1} = L'_k \cup L''_k$ (see Fig. 6). It remains to specify the parallelogram P_{k+1}; clearly, a very thin rectangle whose axis is the line $x + y = 2$ has the required property.

In the inductive step we have created a $2^k \times 2^k$ grid pattern in the square $[0, 1]^2$ in L_{k+1}. Moreover, the arrangement of L_{k+1} contains affine copies of other, smaller grid patterns created in previous steps (the construction is such that the parallelograms enclosing these grid patterns are all disjoint).

Let $n = 2^k$, let $q = n/r$ be given, and suppose that $q = o(n)$. Call a cell in the arrangement of L_k *deep* if it is contained in one of the grid patterns in L_k and it is separated by at least $q + 1$ lines from each of the sides of the parallelogram bounding the considered grid pattern. An easy calculation shows that $(n^2/2)(1 - o(1))$ cells are deep ones.

Let σ be an arbitrary set crossed by at most q lines of L_k. We claim that it is incident to at most $q^2/4 + O(q)$ deep cells. Indeed, if σ is incident to at least one deep cell, then it is completely contained in the parallelogram P enclosing the grid pattern of that deep cell. We delete all lines that do not cross σ and then draw the boundary of P. The number of cells intersected by σ is unchanged. All these cells lie in P, and the number of cells there after the deletion is at most $q^2/4 + O(q)$. This shows that any $(1/r)$-cutting for L_k has $2r^2(1 - o(1))$ sets as claimed in Theorem 2(i).

If the considered σ is a triangle, it can only intersects $q^2/8 + O(q)$ cells of the grid: Map the parallelogram enclosing the considered grid pattern onto an axis-parallel square so that the grid lines have unit spacing. Then the length of the orthogonal projection of σ to the x-axis plus the length of the orthogonal projection to the y-axis are together

smaller than $q + 5$, and hence the area of σ is at most $q^2/8 + O(q)$. Thus σ encloses at most this many grid cells completely and its boundary intersects at most $O(q)$ more cells. This proves Theorem 2(ii).   □

Quadrilaterals.   For a cutting consisting of quadrilaterals, the above method based on grids only gives the same bound as for general sets. A random construction described next does slightly better.

Let S^2 be the unit sphere in $\mathbb{R}^3$ centered at the origin 0. For a point $x \in S^2$ we define x^* as the great circle whose plane is perpendicular to the vector $0x$. Let μ be the rotation-invariant probability measure on S^2 (the usual surface measure divided by the normalizing factor 4π). We define a probability measure τ on the set $\mathcal{G}$ of all great circles by setting $\tau(G) = \mu(\{x;\ x^* \in G\})$ for $G \subseteq \mathcal{G}$ (whenever the right-hand side is defined).

We identify $\mathbb{R}^2$ with the plane $z = 1$ in $\mathbb{R}^3$, and let U be the upper open hemisphere of S^2. Let $p\colon U \to \mathbb{R}^2$ be the central projection with center at 0. If g is a great circle distinct from the horizontal one, $p(g)$ is a line in $\mathbb{R}^2$. Given a sufficiently large parameter r, we choose n much larger than r, say $n \geq r^{10}$, and we let G be a random set of n great circles on S^2, obtained by n independent random draws from the distribution τ (the same circle may be drawn several times but this happens with probability 0). We let L be the set of the p-images of the great circles of G; this is the set of lines for which we want to show a lower bound on the size of a $(1/r)$-cutting.

The way to think about G is that these are "all the great circles"; only because cuttings are considered for finite sets of lines do we need to take a finite sample. We need a technical result showing that the random sample G approximates the set $\mathcal{G}$ of all great circles well enough in a suitable sense. Fortunately, the required approximation properties are special cases of known general results.

For a subset $X \subseteq S^2$ and a set $G \subseteq \mathcal{G}$ of great circles, we let G_X denote the set of all great circles of G intersecting X. Similarly we write G_X^2 for the set of all ordered pairs of distinct great circles of G such that X contains at least one of their two intersection points. We also recall that a set $X \subseteq S^2$ is called (*strongly*) *spherically convex* if it contains no two antipodal points and every great circle intersects it in a connected segment.

Lemma 4.   *Let $\varepsilon > 0$ be a real number with $C\varepsilon^{-2}\log(1/\varepsilon) < n$ for a sufficiently large absolute constant C (in particular, $\varepsilon = n^{-1/3}$ will do). Then the sample G has both of the following properties with a positive probability:*

(i) *For any spherically convex quadrilateral Q (bounded by at most four arcs of great circles) we have*

$$\left| \tau(\mathcal{G}_Q) - \frac{|G_Q|}{n} \right| \leq \varepsilon.$$

(ii) *Let $C \subseteq S^2$ be a spherically convex set. Then we have*

$$\left| \tau^2(\mathcal{G}_C^2) - \frac{|G_C^2|}{n^2} \right| \leq \varepsilon$$

(τ^2 *denotes the product measure on $\mathcal{G} \times \mathcal{G}$*).

Proof sketch. As was mentioned above, a proof follows from general results in the literature. Since this part has little to do with the main topics of this paper, we do not recall all the necessary definitions, and we only present brief remarks that should be sufficient for those familiar with the referred papers. Condition (i) says that G is an ε-*approximation* for the set system $\{\mathcal{G}_Q; Q$ is a spherically convex quadrilateral$\}$, and the required result follows from a theorem of Vapnik and Chervonenkis [19] plus a standard VC-dimension estimate (see, e.g., [16] for the definition of ε-approximation and related notions). Part (ii) is a special case of a result of Chazelle [6] on the so-called *product range spaces*. In that paper the result is formulated in a special setting of hyperplanes in $\mathbb{R}^d$, and a more general formulation directly implying the property we need is given in [4]. $\qquad\square$

In the rest of the proof suppose that a suitable set G with properties (i) and (ii) as in Lemma 4 has been fixed. Let $\sigma \subseteq \mathbb{R}^2$ be a quadrilateral crossed by at most $q = n/r$ lines of L. We want to show that σ cannot contain too many intersections of lines of L. Clearly, we may assume σ is convex. Let $Q = p^{-1}(\sigma) \subseteq U$ be the spherically convex pre-image of σ in S^2.

Since Q is intersected by at most $q + O(1)$ great circles of G, property (i) of G implies that $\tau(\mathcal{G}_Q) \leq (1/r)(1 + o(1))$ (the value of n was chosen so that the relative accuracy ε in (i) is much smaller than $1/r$, and even than $1/r^2$, which is needed in the next step). Further, by (ii), the number of ordered pairs of great circles in G with intersections in Q is at most $n^2(\tau^2(\mathcal{G}_Q^2) + \varepsilon)$, where $\varepsilon \ll 1/r^2$ is as in Lemma 4. Hence it suffices to show

Lemma 5. *Let $Q \subseteq S^2$ be a spherically convex quadrilateral (bounded by great circle arcs) with $\tau(Q) = \delta$, where $\delta < \delta_0$ for a suitable absolute constant $\delta_0 > 0$. Then we have $\tau^2(\mathcal{G}_Q^2) \leq \delta^2/2.54$.*

Proof. The set $(\mathcal{G}_Q)^2$ of all ordered pairs or great circles intersecting Q has τ^2-measure δ^2. Hence it suffices to show that if δ is small enough, then the probability that two great circles randomly chosen from $\mathcal{G}_Q$ intersect inside Q is at most $1/2.54$. We derive it from an analogous result in the plane (it is quite likely that a corresponding result for the sphere is known as well, but I could not find it in the literature).

Let ν denote a motion-invariant measure on the set of all lines in the plane. As is well known, this measure is unique up to a scaling factor. For a line ℓ not passing through the origin, let $\rho(\ell)$, $\varphi(\ell)$ be the polar coordinates of the point of ℓ closest to the origin. Then the ν-measure of a set of lines L is the Lebesgue measure of the image of L in the (ρ, φ)-plane, that is,

$$\nu(L) = \int_L \mathrm{d}\rho(\ell)\, \mathrm{d}\varphi(\ell).$$

For any convex set C, an appropriate scalar multiple of ν is a probability distribution on the set of all lines intersecting C, and it is known that if we pick two lines ℓ, ℓ' from this distribution randomly and independently, then the probability that ℓ and ℓ' intersect within C equals

$$\frac{2\pi \operatorname{area}(C)}{\operatorname{perimeter}(C)^2} \tag{1}$$

[18, p. 56]. Among convex quadrilaterals of given perimeter, the square has the largest area, and for the square the value of (1) is $\pi/8 < 1/2.54$.

It remains to transfer this result to the sphere. The considered quadrilateral Q with $\tau(\mathcal{G}_Q) = \delta$ has diameter $O(\delta)$ (since the τ-measure of the set of all great circles intersecting an arc of a great circle of angular length $2\pi x$ is $2x$ for $x \leq \frac{1}{2}$). By symmetry, we may suppose that Q lies in an $O(\delta)$-neighborhood of the "north pole" $(0, 0, 1)$ of S^2. Let $Q' = p(Q)$ be the projection of Q into the $z = 1$ plane. It suffices to show that for $\delta \searrow 0$, the probability distribution on the set of all lines intersecting Q' induced by the projection of τ converges uniformly to the distribution given by ν. This seems intuitively obvious, and it can be checked by calculation as follows.

We calculate the measure τ' induced on the set of all lines in $\mathbb{R}^2$ by the p-image of τ. We introduce the polar coordinates (ρ, φ) in the $z = 1$ plane (with origin at $(0, 0, 1)$), and spherical coordinates (ϑ, φ) on S^2 (where φ is the "latitude" and $\vartheta = \arccos z$ is the angle from the north pole). Then the surface measure element on S^2 is const. $\sin \vartheta \, \mathrm{d}\vartheta \, \mathrm{d}\varphi$ (the normalizing constant is unimportant so we leave it unspecified). If ℓ is a line in the $z = 1$ plane, let $v = v(\ell) \in U$ be the point of S^2 such that the projection of the great circle v^* is the line ℓ. It is easy to find that if the point of ℓ nearest to the origin has polar coordinates (ρ, φ), then v has spherical coordinates $(\operatorname{arccot} \rho, -\varphi)$. From this we get that the element of τ' is

$$\mathrm{d}\tau' = \text{const. } \sin(\operatorname{arccot} \rho) \, \mathrm{d}(\operatorname{arccot} \rho) \, \mathrm{d}\varphi = \frac{\text{const.}}{(1 + \rho^2)^{3/2}} \, \mathrm{d}\rho \, \mathrm{d}\varphi \, .$$

In a small enough neighborhood of the origin this converges to a suitable constant multiple of $\mathrm{d}\nu$. $\qquad\square$

Remark.   This lower bound proof can be applied also for cuttings consisting of convex k-gons for any fixed k. The resulting constant at r^2 is the reciprocal of the value of (1) for the regular convex k-gon, which is $(2k/\pi) \tan(\pi/k)$ (and with $k \to \infty$ it converges to 2).

4.   Upper Bounds for Semicuttings

As was remarked in the Introduction, we are going to analyze a simple probabilistic construction of a semicutting. This construction goes back to Clarkson [10], and our analysis is mainly based on the method of Chazelle and Friedman [8].

Since the alphabet does not seem to have sufficiently many convenient letters, the notation in this section reuses many letters with meaning quite different from the one in previous sections.

4.1.   *Canonical Triangulations*

First we define a *canonical triangulation* $CT(S)$ for a finite collection S of lines in the plane. We suppose that the lines of S are in general position (including the assumption that no line of S is horizontal). The triangulation is obtained by drawing, for each cell
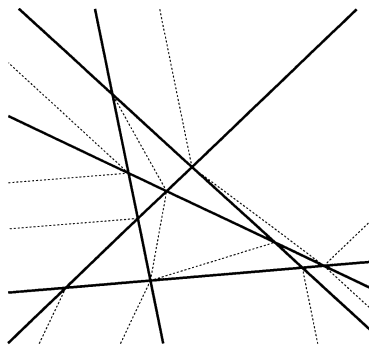
Fig. 7. An example of a canonical triangulation.

C of the arrangement of S, all the diagonals of C incident to the lowest vertex of C. In order to handle the unbounded cells properly (note that a cell unbounded from below has no lowest vertex), we may for instance formally add three lines "at infinity" to S, say the lines $x = -K$, $y = x - K$, and $y = -x + K$ for a very large $K > 0$ (we assume that these lines are also in general position relative to the lines of S). In reality this means that we specify how unbounded cells are subdivided into regions with at most three edges by drawing suitable semi-infinite rays; see Fig. 7 for an example of the resulting triangulation.

Note that $CT(S)$ may also contain unbounded convex polygons (with at most three sides). When we speak of triangles of $CT(S)$ we include also these unbounded polygons. Moreover, for the subsequent analysis, it is technically convenient to include the edges (sides) of the triangles in $CT(S)$ into $CT(S)$ as separate objects.

Let L be a given finite set of lines in general position. We let $\mathcal{T} = \bigcup_{S \subseteq L} CT(S)$ be the collection of all triangles and all edges that appear in the canonical triangulation for some subset of L. For $\sigma \in \mathcal{T}$, let $I(\sigma)$ be the set of lines of L intersecting the interior of σ (for an edge, we mean the relative interior), and we put $w(\sigma) = |I(\sigma)|$. Let $D(\sigma) \subseteq L$ be an inclusion-minimal subset of L with $\sigma \in CT(D(\sigma))$. As is not difficult to check, for lines in general position, $D(\sigma)$ is unique (these are the lines defining σ "geometrically") and it always consists of at most five lines. We put $b(\sigma) = |D(\sigma)|$. We write $\mathcal{T}_b$ for the set of all $\sigma \in \mathcal{T}$ with $b(\sigma) = b$. The triangles and edges in $\mathcal{T}_b$ are called b-*canonical*. Figure 8 shows examples of b-canonical triangles for various values of b. The reader may want to check that edges can be b-canonical for $b = 1, 2, 3, 4$.

A key property expressing the "canonicality" of $CT(S)$ is the following: For any set $S \subseteq L$ and any triangle or edge $\sigma \in \mathcal{T}$, we have

$$\sigma \in CT(S) \qquad \text{if and only if} \quad D(\sigma) \subseteq S \quad \text{and} \quad I(\sigma) \cap S = \emptyset \qquad (2)$$

(see [8]). A large part of our subsequent analysis will apply not only to the specific canonical triangulation in the plane, but to a general abstract framework. In this framework we have two finite sets L and $\mathcal{T}$ and two mappings $I, D \colon \mathcal{T} \to 2^L$, where $|D(\sigma)| \le b_0$ for all $\sigma \in \mathcal{T}$ (b_0 is some constant). In this abstract setting, the elements of $\mathcal{T}$ are often
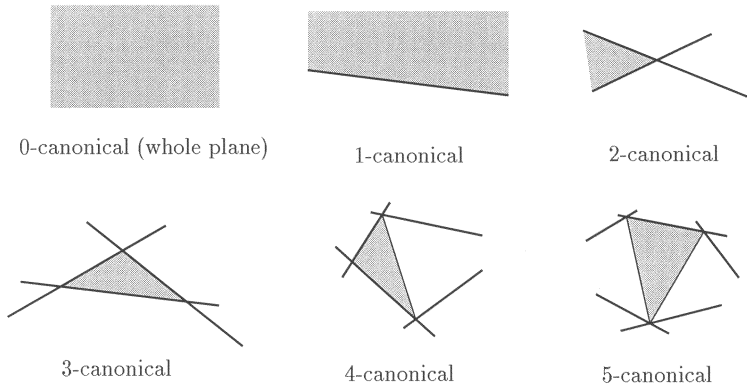
Fig. 8. Examples of b-canonical triangles.

called *regions*. Taking (2) as an axiom, we can then *define*

$$CT(S) = \{\sigma \in \mathcal{T}; \ D(\sigma) \subseteq S \text{ and } I(\sigma) \cap S = \emptyset\}$$

(see [2] for a more detailed discussion of abstract frameworks of this type in computational geometry).

Lemma 6. *Let S be a set of q lines in the plane, and let $CT(S)$ be the canonical triangulation of the arrangement of S.*

(i) *If F_b denotes the number of b-canonical edges in $CT(S)$, then $F_1 \leq 1$, $F_2 = O(q)$, $F_3 = q^2 + O(q)$, and $F_4 = q^2/2 + O(q)$.*

(ii) *If G_b denotes the number of b-canonical triangles in $CT(S)$, then $G_0 \leq 1$, $G_1 \leq 2$, $G_2 = O(q)$, $G_3 + G_4 + G_5 = q^2 + O(q)$, $|G_3 - G_5| = O(q)$, and $G_3 \leq q^2/3 + O(q)$.*

Proof. Concerning (i), a 1-canonical edge can only appear if S consists of a single line, and 2-canonical edges are always unbounded parts of the lines of S. The claimed bounds for b-canonical edges with $b = 3, 4$ have an $O(q)$ error term, so we can ignore unbounded cells safely. In a bounded cell, a 3-canonical edge is just a bounded edge of the arrangement of S, and there are $q^2 + O(q)$ of such edges. The 4-canonical edges are diagonals of the bounded cells. The drawing of all triangles in bounded cells determines a planar graph with $q^2/2 + O(q)$ vertices and with all faces triangular except for the outer one which has $O(q)$ edges, hence there are $\frac{3}{2}q^2 + O(q)$ edges in total, and $q^2/2 + O(q)$ of them are the diagonals.

In part (ii), we can ignore unbounded cells again. Then $G_3 + G_4 + G_5$ is the total number of triangles, which is asymptotically twice the number of vertices of the arrangement of S. Now 3-canonical triangles are just triangular facets of the arrangement of S, and it is known that this number is at most $q^2/3 + O(q)$ (and this can be attained) [13]. The 4-canonical triangles are those lying in a cell with at least four sides and sharing one side with that cell, and the remaining ones are 5-canonical. Hence a cell with i edges has

$i - 4$ 5-canonical triangles except when $i = 3$. If f_i denotes the number of cells with i edges, we thus have $G_5 = \sum_i (i-4) f_i + G_3$. On the other hand, $\sum_i f_i = q^2/2 + O(q)$ (the total number of cells) and $\sum_i i f_i = 2q^2 + O(q)$ (twice the total number of edges), hence $\sum_i (i-4) f_i = O(q)$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

4.2. Higher Moments for the Canonical Triangulation

Given a set L of n lines in the plane in general position, we fix a suitable probability $p \in (0, 1)$. We choose a random sample $R \subseteq L$ by picking each line of L independently with probability p, and the semicutting we want to investigate is the canonical triangulation $CT(R)$.

We assume that $p < n^{-\alpha}$ and $np > n^{\alpha}$ for some constant $\alpha > 0$. Then $|R| = np(1 + o(1))$ holds with high probability (the probability of this failing is exponentially small in n). This follows from a suitable tail estimate for the binomial distribution; see, e.g., the Appendix in [3]. Hence, when estimating the expectation of quantities lying between 1 and $n^{O(1)}$, we make a negligibly small error by restricting ourselves to the samples R with $|R| = np(1 + o(1))$.

We are interested in estimating the expectation of the kth moment

$$M_k^{\mathcal{T}}(p) = \mathbf{E}\left[\sum_{\sigma \in CT(R)} w(\sigma)^k\right].$$

In particular, $M_0^{\mathcal{T}}(p)$ is the expected size of $CT(R)$.

This definition refers to the abstract setting mentioned above. If we want to refer to the kth moment for the triangles in the canonical triangulation of a line arrangement, we write M_k^{triang}, and similarly M_k^{edge} refers to the kth moment for the edges of the canonical triangulation.

It has been known that for any $k > 0$ there exists a constant A_k such that $M_k^{\text{triang}}(p) \leq A_k(np)^2 p^{-k}$ for $p \ll 1$, $pn \gg 1$ [8]. This means that the kth degree average of $w(\sigma)$ over the canonical triangulation of a random sample is $O(1/p)$ with a positive probability. Here we investigate numerical bounds on A_k.

From the canonicality condition (2), we get that if $R \subseteq L$ is a random sample selected as above with probability p, then, for any $\sigma \in \mathcal{T}$,

$$\Pr[\sigma \in CT(R)] = (1 - p)^{w(\sigma)} p^{b(\sigma)}.$$

Hence if $N_w^{(b)}$ denotes the number of $\sigma \in \mathcal{T}_b$ with $w(\sigma) = w$, we can write $M_k^{\mathcal{T}}(p) = \sum_b M_k^{\mathcal{T},(b)}(p)$, where

$$M_k^{\mathcal{T},(b)}(p) = \sum_{w=0}^{\infty} N_w^{(b)} w^k p^b (1-p)^w. \qquad\qquad (3)$$

We could try to estimate the numbers $N_w^{(b)}$ directly, but this looks quite challenging already for the simplest planar cases (the number of 4-canonical segments in a line arrangement, say). Instead, following the Chazelle–Friedman method, we use information about $M_0^{\mathcal{T}}(p')$ for various p' to bound $M_k^{\mathcal{T}}(p)$.

4.3.   *Bounding Moments for b-Canonical Objects*

Here we consider bounding $M_k^{\mathcal{T},(b)}(p)$ for a single fixed value of b. We omit the super-scripts $^{\mathcal{T},(b)}$ in the notation in this part.

We consider the situation when we do not know anything about the N_w's directly, and we only assume certain estimates for $M_0(sp)$ for suitable values of a real parameter $s > 0$ (we only need to consider values of s not much larger than 1). The question is how well can we bound $M_k(p)$. In our application for semicuttings and in similar geometric situations, we usually have estimates of the form

$$M_0(sp) = s^d M_0(p) + \text{ smaller order terms}.$$

For instance, for the 4-canonical edges of the canonical triangulation of a line arrangement in the plane, Lemma 6(i) tells us that $M_0^{\text{edge},(4)}(sp) = \frac{1}{2}s^2(np)^2 + O(snp)$. Two different cases to consider are when we can estimate $M_0(sp)$ both from above and below, and when we only have an upper bound; the latter case gives a slightly weaker upper bound for $M_k(p)$ even if the upper bounds for $M_0(sp)$ are the same. The following proposition deals with these two cases.

Proposition 7.   *Let $k > 0$ and $b > d$ be fixed real numbers. Let $p \in (0, 1)$ be a parameter; the asymptotic notation "$o(1)$" in the subsequent formulas is with respect to $p \to 0$.*

(i)   *Set $s_0 = (b - d)/(b - d + k)$. Suppose that $M_0(s_0 p) \leq s_0^d M_0(p)(1 + o(1))$, and moreover that $M_0(s_0' p) = O(M_0(p))$ holds for some fixed $s_0' < s_0$. Then*

$$M_k(p) \leq B_{k,b-d}(1 + o(1)) \frac{M_0(p)}{p^k},$$

where

$$B_{k,\delta} = \frac{(\delta + k)^{\delta+k}}{e^k \delta^\delta}.$$

(ii)   *Let $k \geq 1$ be an integer, and let $\gamma > 0$ be a fixed real number. Suppose that*

$$M_0(sp) = s^d M_0(p)(1 + o(1))$$

holds for all $s \in (1 - \gamma, 1 + \gamma)$. Then

$$M_k(p) = \tilde{B}_{k,b-d}(1 + o(1)) \frac{M_0(p)}{p^k},$$

where $\tilde{B}_{k,\delta} = \delta(\delta + 1)(\delta + 2) \cdots (\delta + k - 1)$.

Remarks.   Table 1 shows some numerical values of the constants $B_{k,\delta}^{1/k}$ and $\tilde{B}_{k,\delta}^{1/k}$ (the kth roots seem more instructive in this context).

The estimate in part (i) is best possible in a certain sense (a precise formulation is given in Lemma 9 below).

Table 1. Constants in the estimates for the moments $M_k(p)$.

k	$\delta = 1$		$\delta = 2$		$\delta = 3$		$\delta = 4$	
	$B_{k,1}^{1/k}$	$\tilde{B}_{k,1}^{1/k}$	$B_{k,2}^{1/k}$	$\tilde{B}_{k,2}^{1/k}$	$B_{k,3}^{1/k}$	$\tilde{B}_{k,3}^{1/k}$	$B_{k,4}^{1/k}$	$\tilde{B}_{k,4}^{1/k}$
1	1.47	1.00	2.48	2.00	3.49	3.00	4.49	4.00
2	1.91	1.41	2.94	2.45	3.96	3.46	4.97	4.47
3	2.34	1.82	3.39	2.88	4.41	3.91	5.43	4.93
4	2.75	2.21	3.82	3.31	4.86	4.36	5.89	5.38
5	3.16	2.65	4.25	3.73	5.30	4.79	6.33	5.83

Part (ii) of this Proposition is somewhat related to a work of Clarkson [11]. For an arrangement of n lines (or hyperplanes), Clarkson considers a random sample R of exactly r lines (this is a technical difference to our setting). For a certain combinatorial quantity depending on R (not related to the canonical triangulation), he determines an analogue of our $M_0(p)$ exactly for all sample sizes r, and he reconstructs the underlying distribution (an analogue of our N_w's) from it exactly by matrix inversion.

Proof of Proposition 7.   In this proof we assume that $s > 0$ and $s = \Theta(1)$ (in other words, s lies in a fixed interval (s_1, s_2) with $s_1 > 0$). From (3) we have

$$M_k(sp) = \sum_{w=0}^{\infty} N_w w^k (sp)^b (1 - sp)^w.$$

First we want to manipulate this expression into a more convenient form; by a suitable rescaling we will be able to get fewer parameters. We introduce the substitution $w = t/p$, where t is a new variable. We let ν be the measure on $[0, \infty)$ concentrated at the points $0, p, 2p, 3p, \ldots$ assigning the mass of $(p^b/M_0(p))N_w$ to the point $pw, w = 0, 1, 2, \ldots$. With this notation, we obtain

$$
\begin{aligned}
M_k(sp) &= \frac{M_0(p)}{p^b} \sum_{w=0}^{\infty} \nu(wp) w^k (sp)^b (1 - sp)^w \\
&= \frac{M_0(p)s^b}{p^k} \int_0^{\infty} t^k (1 - sp)^{t/p} \, d\nu(t)
\end{aligned}
\tag{4}
$$

(for future considerations it seems more natural to write the finite sum formally as an integral). Using the inequality $1 - sp \le e^{-sp}$ we get

$$M_k(sp) \le \frac{M_0(p)}{p^k} s^b \int_0^{\infty} t^k e^{-st} \, d\nu(t). \tag{5}$$

We write $\mathbf{m}_k(s) = \int_0^{\infty} t^k e^{-st} \, d\nu(t)$. We want to show that (5) actually is an equality up to a $1 + o(1)$ factor, i.e., that replacing $1 - sp$ by e^{-sp} does no harm. This is expressed in the following lemma:

Lemma 8. *Let $s > 0$, $s = \Theta(1)$, and suppose that $\tilde{s} \in (0, s)$ is a number such that both $\tilde{s}$ and $s - \tilde{s}$ are bounded away from zero by a positive constant, and such that $M_0(\tilde{s}p) \leq K M_0(p)$ holds for some constant K. Then we have*

$$M_k(sp) = (1 + o(1)) \frac{M_0(p)}{p^k} s^b \mathbf{m}_k(s). \tag{6}$$

Proof of Lemma 8. These are rather routine estimates. In view of (5), we only need to show the inequality "$\geq$", which means showing

$$\int_0^\infty t^k (1 - sp)^{t/p} \, \mathrm{d}\nu(t) \geq (1 - o(1)) \int_0^\infty t^k e^{-st} \, \mathrm{d}\nu(t).$$

Using the asymptotics $1 - p = e^{-p + o(p)}$, we get $(1 - sp)^{t/p} \geq e^{-(1-o(1))st}$. For an arbitrarily large but fixed t_0, we have $e^{-st(1-o(1))} = (1 - o(1))e^{-st}$ for all $t \leq t_0$. Hence

$$\int_0^\infty t^k (1 - sp)^{t/p} \, \mathrm{d}\nu(t) \geq \int_0^{t_0} t^k (1 - sp)^{t/p} \, \mathrm{d}\nu(t)$$

$$\geq (1 - o(1)) \int_0^{t_0} t^k e^{-st} \, \mathrm{d}\nu(t).$$

It remains to show that, for large enough t_0, the integral $\int_{t_0}^\infty t^k e^{-st} \, \mathrm{d}\nu(t)$ is negligibly small, smaller than a prescribed $\varepsilon > 0$, and here we use the assumption on $\tilde{s}$. Set $\beta = (s - \tilde{s})/2s$, and let t_0 be so large that $e^{\beta s t_0} > (K/\varepsilon \tilde{s}^b) t_0^k$ (K is the constant from the assumption on $\tilde{s}$). We then have

$$\int_{t_0}^\infty t^k e^{-st} \, \mathrm{d}\nu(t) \leq \frac{\varepsilon \tilde{s}^b}{K} \int_{t_0}^\infty e^{-(1-\beta)st} \, \mathrm{d}\nu(t).$$

Since $e^{-(1-\beta)st} = e^{-(1+\beta)\tilde{s}t} \leq (1 - \tilde{s}p)^{t/p}$ for p small enough, the last displayed expression is certainly at most

$$\frac{\varepsilon \tilde{s}^b}{K} \int_0^\infty (1 - \tilde{s}p)^{\tilde{s}t/p} \, \mathrm{d}\nu(t) \leq \frac{\varepsilon \tilde{s}^b}{K} \frac{M_0(\tilde{s}p)}{M_0(p)\tilde{s}^b} \leq \varepsilon$$

by the assumption on $M_0(\tilde{s}p)$.

We continue the proof of Proposition 7. In view of (6) and of the assumptions on $M_0(sp)$ in Proposition 7, it is now sufficient to bound $\mathbf{m}_k(1)$ in terms of $\mathbf{m}_0(s)$ for suitable s's. This is done in the next lemma.

Lemma 9. *Let $k, \delta > 0$ be fixed real numbers and let ν be a measure on $[0, \infty)$. Put $\mathbf{m}_k(s) = \int_0^\infty t^k e^{-st} \, \mathrm{d}\nu(t)$, and write $\mathbf{m}_k = \mathbf{m}_k(1)$.*

(i) *Let $s_0 = \delta/(\delta + k)$. Suppose that ν is such that*

$$\mathbf{m}_0(s) \leq \frac{1}{s^\delta} \tag{7}$$

holds for $s = s_0$. Then $\mathbf{m}_k \leq B_{k,\delta}$, where $B_{k,\delta}$ is as in Proposition 7. This estimate cannot be improved even if we assume that (7) holds for all $s > 0$ simultaneously.

(ii) *Let $k \geq 1$ be an integer, and let $\gamma > 0$ be a fixed real number. Then, for every $\varepsilon > 0$, an $\eta = \eta(\varepsilon, \gamma, \delta, k) > 0$ exists with the following property: if we assume that*

$$\frac{1}{s^\delta} - \eta \leq \mathbf{m}_0(s) \leq \frac{1}{s^\delta} + \eta$$

holds for all $s \in [1, 1 + \gamma)$, then $\tilde{B}_{k,\delta} - \varepsilon \leq \mathbf{m}_k \leq \tilde{B}_{k,\delta} + \varepsilon$, where $\tilde{B}_{k,\delta}$ is as in Proposition 7.

Remarks. Part (ii) can be formulated with a general function $g(s)$ instead of $s^{-\delta}$. We need to assume that g is $(k+1)$-times differentiable in $(1, 1+\gamma)$, and the result is that $\mathbf{m}_k$ is close to $(-1)^k g^{(k)}(1)g(1)$; we omit a detailed formulation as we have no immediate application (the details can be recovered from the proof below).

We should also note that the lemma is essentially about the Laplace transform. Namely, we have some information about the Laplace transform of $d\nu$, and we are asking what bounds can be inferred for the Laplace transform of $t^k d\nu$.

Proof of Lemma 9. (i) The value of s_0 given in the statement comes from a solution to a simple minimization problem. For any given $s < 1$, we can calculate the smallest number $a = a(s, k)$ such that

$$t^k e^{-t} \leq a e^{-st} \tag{8}$$

holds for all $t \in [0, \infty)$. This a is the minimum value attained by the (convex) function $t \mapsto t^k e^{-(1-s)t}$. Elementary calculus shows that the minimum is at $t = k/(1-s)$, and we get $a = k^k/e^k(1-s)^k$. By integrating both sides of (8) according to $\nu(t)$, we get $\mathbf{m}_k \leq a\mathbf{m}_0(s) \leq a/s^\delta$. By minimizing the right-hand side as a function of s, we arrive at the value $s_0 = \delta/(\delta + k)$, for which the upper bound a/s^δ becomes just $B_{k,\delta}$.

The lower bound is equally simple (but it is a little surprising that the above straightforward method gives the best possible result). Namely, we consider the measure ν with mass $N = (e(k + \delta)/\delta)^\delta$ concentrated at a single point $t_0 = \delta + k$. For this ν, we obtain $\mathbf{m}_k = Nt_0^k e^{-t_0} = B_{k,\delta}$, while $\mathbf{m}_0(s) = Ne^{-st_0}$. With some more calculus one shows that $Ne^{-st_0} \leq s^{-\delta}$ holds for all $s > 0$, hence (7) is satisfied.

(ii) It seems that a similar result (in greater generality) must be known in probability theory, say, but having found no reference so far I present a proof.

Here is the simple underlying idea. Up to smaller order terms, we assume $\int_0^\infty e^{-st} d\nu(t) = s^{-\delta}$. If we were allowed to differentiate both sides as a function of the parameter s and move the differentiation operator inside the integration sign, we would get $\int_0^\infty t^k e^{-st} d\nu(t) = \delta(\delta + 1) \cdots (\delta + k - 1)s^{-\delta-k}$, and substituting $s = 1$ would give what we want. Of course, this does not quite work since we only assume an approximate equality $\int_0^\infty e^{-st} d\nu(t) \approx s^{-\delta}$ and so the derivatives of both sides might be entirely different.

To make everything work, we approximate differentiation by a suitable finite differencing operator. For instance, we may use the *forward differencing operator* Δ, which acts on a real function f as follows:

$$(\Delta_h f)(x) = f(x + h) - f(x)$$

($h > 0$ is a real parameter). The k-fold iteration of Δ can be expressed as

$$\Delta_h^k f(x) = \sum_{i=0}^{k} (-1)^{k-i} \binom{k}{i} f(x + ih). \tag{9}$$

For a function f having a $(k + 1)$st derivative on $(0, kh)$, we have

$$\Delta_h^k f(x) = h^k f^{(k)}(x) + O(h^{k+1} \max\{f^{(k+1)}(\xi); \ x < \xi < x + t\}) \tag{10}$$

(with an absolute constant in the O-notation; see, e.g., [5]). Let $h > 0$ be sufficiently small (in terms of ε), and set $s_i = 1 + ih$, $a_i = (-1)^{k-i} \binom{k}{i} h^{-k}$, $i = 0, 1, \ldots, k$. Then for any function f we have $\sum_{i=0}^{k} a_i f(s_i) = h^{-k} \Delta_h^k f(1)$, which is a good approximation to $f^{(k)}(1)$ for a smooth f by (10).

Applying this with $f(x) = e^{-xt}$, we get $\sum_{i=0}^{k} a_i e^{-s_i t} = (-t)^k e^{-t} + O(ht^{k+1} e^{-t}) = (-t)^k e^{-t} + O(he^{-t})$ (for all t). Integrating both sides according to $\nu(t)$ gives

$$\int_0^\infty t^k e^{-t} \, d\nu(t) = \mathbf{m}_k = (-1)^k \sum_{i=0}^{k} \left(a_i \int_0^\infty e^{-s_i t} \, d\nu(t) \right) + O \left(h \int_0^\infty e^{-t} \, d\nu(t) \right)$$

$$= \sum_{i=0}^{k} a_i \mathbf{m}_0(s_i) + O(h\mathbf{m}_0).$$

Using our assumption $s_i^{-\delta} - \eta \le \mathbf{m}_0(s_i) \le s_i^{-\delta} + \eta$, we can further write

$$\mathbf{m}_k = (-1)^k \sum_{i=0}^{k} \frac{a_i}{s_i^\delta} + O \left(h + \eta \sum_{i=0}^{k} |a_i| \right) = (-1)^k \sum_{i=0}^{k} \frac{a_i}{s_i^\delta} + O(h + \eta h^{-k}).$$

Finally, using the definition of the a_i's and the finite differencing formula (10) again we get

$$\mathbf{m}_k = (-1)^k \left(\frac{d^k(s^{-\delta})}{ds^k} \bigg|_{s=1} \right) + O(h + \eta h^{-k}) = \delta(\delta + 1) \cdots (\delta + k - 1) + O(h + \eta h^{-k}).$$

By choosing $h \ll \varepsilon$ and $\eta \ll h^k$ we get the claim of the lemma. $\square$

4.4. Planar Semicuttings

Estimate via Edges. We return to the specific planar situation, and consider estimating the moments $M_k^{\text{triang}}(p)$ for the canonical triangulation of a random sample $R \subseteq L$. Our first approach is based on estimating the moment M_k^{edge} first. Given a triangle $\sigma \in CT(R)$ with edges e_1, e_2, e_3, we note that $w(\sigma) = \frac{1}{2} \sum_{i=1}^{3} w(e_i)$ (since each line intersecting σ intersects two of its edges). From this we get

$$w(\sigma)^k = \frac{1}{2^k} \left(\sum_{i=1}^{3} w(e_i) \right)^k \le \frac{3^{k-1}}{2^k} \sum_{i=1}^{3} w(e_i)^k$$

(by Hölder's inequality $\sum_i x_i y_i \le (\sum_i x_i^k)^{1/k} (\sum_i x_i^{k/(k-1)})^{(k-1)/k}$ with $x_i = w(e_i)$,

$y_i = 1$). Every edge contributes to two triangles, hence $M_k^{\text{triang}}(p) \leq (3/2)^{k-1} M_k^{\text{edge}}(p)$. For estimating $M_k^{\text{edge}}(p)$ we combine Lemma 6 and Proposition 7(ii), which yields

$$M_k^{\text{edge}}(p) = (\tilde{B}_{k,1} + \tfrac{1}{2}\tilde{B}_{k,2})(1 + o(1))(np)^2 p^{-k}.$$

Write $q = np$. Since $CT(R)$ has almost surely $q^2(1 + o(1))$ triangles, we get that the expectation of the kth degree average of $w(\sigma)$ over $\sigma \in CT(R)$ is

$$(1 + o(1)) \left[\frac{M_k^{\text{triang}}(p)}{q^2} \right]^{1/k} \leq (1 + o(1)) A_k^{1/k} \frac{n}{q},$$

where $A_k = (3/2)^{k-1}(\tilde{B}_{k,1} + \tfrac{1}{2}\tilde{B}_{k,2})$. If this kth degree average should be at most n/r as required by the definition of kth degree $(1/r)$-semicutting, we must choose $q = A_k^{1/k} r$, and hence the resulting $(1/r)$-semicutting has size $(1 + o(1)) A_k^{2/k} r^2$. This yields

Theorem 10. *For any fixed integer $k \geq 3$, sufficiently large parameter r, and a sufficiently large collection L of lines in the plane, the above randomized construction yields, with a positive probability, a kth degree $(1/r)$-semicutting for L consisting of at most $C'_k r^2 (1+o(1))$ triangles, where $C'_k = \frac{9}{4}(k! \, (k+3)/3)^{2/k}$. Approximate numerical values are $C'_3 \approx 11.8$, $C'_4 \approx 16.8$, $C'_5 \approx 22.6$.*

We should perhaps stress that these are results concerning the particular randomized construction of a semicutting via canonical triangulation. For instance, already for $k = 4$, the construction of a $(1/r)$-cutting consisting of $16r^2 + O(r)$ triangles provides a better result than the randomized construction as in Theorem 10.

Direct Estimate via Triangles. For triangles, we cannot give precise bounds for the number of b-canonical triangles for each b; for instance, the number of 3-canonical and 5-canonical triangles can be about $q^2/3$ each, or, as another extreme, essentially all the triangles can be 4-canonical. The proportion of the 3-canonical and 5-canonical triangles for a random sample taken with probability sp may moreover vary with s (although this intuitively does not look very likely). Thus, rather than using Proposition 7 directly, we employ the method from its proof. As in the proof of Lemma 9, we bound the function $t^k e^{-t}$ by a function of the form ae^{-st} for the smallest possible $a = a(k, s)$. By the considerations in the proof of Proposition 7, we get that, for any fixed $s \in (0, 1)$,

$$M_k^{\text{triang}}(p) \leq (1 + o(1)) a(k, s) \sum_{b=3}^{5} \frac{1}{s^b p^k} M_0^{\text{triang},(b)}(sp).$$

By Lemma 6 we know that $M_0^{\text{triang},(b)}(sp) = g_b s^2 p^2 (1 + o(1))$ for some real numbers g_3, g_4, g_5 satisfying $g_b \geq 0$, $g_3 + g_4 + g_5 \leq 1$, and $g_3 = g_5 \leq \tfrac{1}{3}$, hence we obtain

$$M_k^{\text{triang}}(p) \leq (1 + o(1)) \frac{(np)^2}{p^k} a(k, s) \sum_{b=3}^{5} \frac{g_b}{s^{b-2}}.$$

For a fixed s, the expression $(g_3/s + g_4/s^2 + g_5/s^3)$ is maximal, under the above conditions on the g_b's, for $g_3 = g_4 = g_5 = \tfrac{1}{3}$. Hence we need to minimize the function

$s \mapsto a(s, k)(s^{-1} + s^{-2} + s^{-3})/3$. This can be done analytically but the expressions are fairly messy. The resulting numerical bounds for the size of the kth degree $(1/r)$-semicutting are approximately $13.4r^2$ ($k = 3$), $17.2r^2$ ($k = 4$), $21.3r^2$ ($k = 5$). This is slightly worse than the bounds obtained via edges in Theorem 10.

Acknowledgment

References

1. P. K. Agarwal. Geometric partitioning and its applications. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Computational Geometry*: *Papers from the DIMACS Special Year*, pages 1–37. American Mathematical Society, Providence, RI, 1991.
2. P. K. Agarwal, J. Matoušek, and O. Schwarzkopf. Computing many faces in arrangements of lines and segments. In *Proc*. 10*th Ann*. *ACM Symp*. *Comput*. *Geom*., pages 76–84, 1994. Also to appear in *SIAM J. Comput*.
3. N. Alon and J. Spencer. *The Probabilistic Method*. Wiley, New York, 1993.
4. H. Brönnimann, B. Chazelle, and J. Matoušek. Product range spaces, sensitive sampling, and derandomization. In *Proc*. 34*th Ann*. *IEEE Symp*. *Found*. *Comput*. *Sci*., pages 400–409, 1993. Revised version is to appear in *SIAM J. Comput*.
5. R. L. Burden, J. D. Faires, and A. C. Reynolds. *Numerical Analysis*. PWS Publishing, Florence, KY, 1989.
6. B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.
7. B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.
8. B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.
9. K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.
10. K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17:830–847, 1988.
11. K. L. Clarkson. A bound on local minima of arrangements that implies the upper bound theorem. *Discrete Comput. Geom.*, 10:427–433, 1993.
12. K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
13. H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. EATCS Monographs on Theoretical Computer Science, volume 10. Springer-Verlag, Heidelberg, 1987.
14. H. Edelsbrunner and E. Welzl. Constructing belts in two-dimensional arrangements with applications. *SIAM J. Comput.*, 15:271–284, 1986.
15. L. Guibas and M. Sharir. Combinatorics and algorithms of arrangements. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, pages 9–36. Algorithms and Combinatorics, volume 10. Springer-Verlag, New York, 1993.
16. D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
17. J. Matoušek. Construction of ε-nets. *Discrete Comput. Geom.*, 5:427–448, 1990.
18. L. A. Santaló. *Integral Geometry and Geometric Probability*. Addison-Wesley, Reading, MA, 1976.
19. V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16:264–280, 1971.
20. P. Yamamoto, K. Kato, K. Imai, and H. Imai. Algorithms for vertical and orthogonal L_1 linear approximation of points. In *Proc*. 4*th Ann*. *ACM Symp*. *Comput*. *Geom*., pages 352–361, 1988.