

The Flip Diameter of Rectangulations and Convex Subdivisions*

Eyal Ackerman¹ Michelle M. Allen² Gill Barequet³ Maarten Löffler⁴
 Joshua Mermelstein² Diane L. Souvaine² Csaba D. Tóth^{2,5}

¹ Department of Mathematics, Physics, and Computer Science, University of Haifa at Oranim, Tivon, Israel

² Department of Computer Science, Tufts University, Medford, MA, USA

³ Department of Computer Science, Technion—Israel Inst. of Technology, Haifa 32000, Israel

⁴ Department of Computing and Information Sciences, Utrecht University, The Netherlands

⁵ Department of Mathematics, California State University Northridge, Los Angeles, CA, USA

received 19th May 2015, revised 30th Jan. 2016, accepted 7th Mar. 2016.

We study the configuration space of rectangulations and convex subdivisions of n points in the plane. It is shown that a sequence of $O(n \log n)$ elementary *flip* and *rotate* operations can transform any rectangulation to any other rectangulation on the same set of n points. This bound is the best possible for some point sets, while $\Theta(n)$ operations are sufficient and necessary for others. Some of our bounds generalize to convex subdivisions of n points in the plane.

Keywords: rectangulation, combinatorial geometry

1 Introduction

The study of rectangular subdivisions of rectangles is motivated by VLSI floorplan design [29] and cartographic visualization [14, 26, 33]. The rich combinatorial structure of rectangular floorplans has also attracted theoretical research [6, 15]. Combinatorial properties lead to efficient algorithms for the recognition and reconstruction of the *rectangular graphs* induced by the corners of the rectangles in a floorplan [22, 32], the *contact graphs* of the rectangles [25, 40], and the contact graphs of the horizontal and vertical *line segments* that separate the rectangles [11]. The number of combinatorially different floorplans with n rectangles is known to be $B(n) = \Theta(8^n/n^4)$, the n th Baxter number [43].

Rectangular subdivisions in the presence of points have also been studied in the literature. Given a finite set P of points in the interior of an axis-aligned rectangle R , a *rectangulation* of (R, P) is a subdivision of R into rectangles by pairwise noncrossing axis-parallel line segments such that every point in P lies in the relative interior of a segment (see Figure 1).

*A preliminary version of this work appeared in the Proceedings of the 11th Latin American Theoretical INformatics Symposium (LATIN 2014), LNCS 8392, Springer, 2014, pp. 478–489. Research on this paper was partially supported by the NSF grant CCF-0830734 and Netherlands Organization for Scientific Research (NWO) under grant 639.021.123.




Fig. 1: Two different rectangulations of a set of four points.

Finding a rectangulation of minimum total edge length has attracted attention [7, 8, 13, 19, 20, 21, 28, 30] due to its applications in VLSI design and stock cutting in the presence of material defects. This problem is known to be NP-hard [30], however, its complexity is unknown when the points in P are in general position in the sense that they have distinct x - and y -coordinates, that is, the points are *noncorectilinear*. It is not hard to see that in this case the minimum edge-length rectangulation must consist of exactly n line segments [7]. For the rest of this paper, we consider only noncorectilinear point sets P , and rectangulations determined by $|P|$ line segments, one containing each point in P .

The space of all the rectangulations of a point set P (within a rectangle R) can be explored using the following two elementary operators introduced in [2] (refer to Figure 2).




Fig. 2: A rectangulation r_1 of a set of 6 points, $r_2 = \text{FLIP}(r_1, p)$, and $r_3 = \text{ROTATE}(r_2, c)$.

Definition 1 (Flip) Let r be a rectangulation of P and let $p \in P$ be a point such that the segment s that contains p does not contain any endpoints of other segments. The operation $\text{FLIP}(r, p)$ changes the orientation of s from vertical to horizontal or vice-versa.

Definition 2 (Rotate) Let r be a rectangulation of P . Let $s_1 = ab$ be a segment that contains $p \in P$. Let $s_2 = cd$ be a segment such that c lies in $ap \subset s_1$ and ac does not contain any endpoints of other segments. The operation $\text{ROTATE}(r, c)$ shortens s_1 to cb , and extends s_2 beyond c until it reaches another segment or the boundary of R .

For a finite set of noncorectilinear points $P \subset R$, we denote by $G(P) = (V, E)$ the *graph of rectangulations of P* , where the vertex set is $V = \{r : r \text{ is a rectangulation of } P\}$ and the edge set is $E = \{(r_1, r_2) : \text{a single flip or rotate operation on } r_1 \text{ produces } r_2\}$. Since both operations are reversible, $G(P)$ is an undirected graph. It is not hard to show that $G(P)$ is connected [2], and there is a sequence of $O(n^2)$ flip and rotate operations between any two rectangulations in $G(P)$ when P is a set of n points in R . It is natural to ask for the diameter of $G(P)$, which we call the *flip diameter* of P for short. For

every set of n points in R , the flip diameter is at least $\Omega(n)$, since every point set admits a rectangulation with all horizontal segments and one with all vertical segments, and each operation modifies at most two segments.

Results

In this paper, we show that the flip diameter of P is $O(n \log n)$ for every n -element noncorectilinear point set P (Section 2), and it is $\Omega(n \log n)$ for some n -element noncorectilinear point sets (Section 3). However, there are n -elements noncorectilinear points sets for which the flip diameter is $\Theta(n)$ (Section 4). That is, the flip diameter is always between $O(n \log n)$ and $\Omega(n)$, depending on the point configuration, and both bounds are the best possible.

We extend the flip and rotate operations and the notion of flip diameter to convex subdivisions (Section 5). A *convex subdivision* of a set $P \subset \mathbb{R}^2$ of points is a subdivision of the plane into convex faces by pairwise noncrossing line segments, halflines, and lines, each of which contains exactly one point of P . We show that the flip diameter for the convex subdivisions of n points is always $O(n \log n)$ and sometimes $\Theta(n)$.

Related work

Determining the exact number of rectangulations on n noncorectilinear points remains an elusive open problem in enumerative combinatorics [2, 3]. Recently, Felsner [16] proved that every combinatorial floorplan with $n + 1$ rooms can be embedded into every set of n noncorectilinear points, hence every set of n noncorectilinear points has at least $B(n) = \Theta(8^n/n^4)$ rectangulations.

The currently best known upper bound, $O(18^n/n^4)$ by Ackerman [1], uses the so-called “cross-graph” charging scheme [35, 36], originally developed for counting the number of (geometric) triangulations on n points in the plane. This method is based on elementary “flip” operations that transform one triangulation into another. Lawson [27] proved that every triangulation on n points in the plane in general position (i.e., no three on a line) can be transformed into the Delaunay triangulation with $O(n^2)$ flips, and this bound is the best possible by a construction due to Hurtado et al. [24]. However, for n points in convex position, $2n - 10$ flips are sufficient, due to a bijection with binary trees with $n - 2$ internal nodes [37]. Hence the flip diameter of every triangulation on n points in the plane is always between $\Theta(n)$ and $\Theta(n^2)$ depending on the point configuration. Eppstein et al. [14] and Buchin et al. [5] define two elementary flip operations on floorplans, in terms of the directed dual graph, and solve optimization problems on floorplans by traversing the flip graph.

Generalizing simultaneous edge flips in a triangulation, Meijer and Rappaport [31] considered *simultaneous edge flips* in convex decompositions of a point set P . A *convex decomposition* of P is a subdivision of the convex hull of P into convex polygons whose joint vertex set is P ; and two convex decompositions of P are related by a simultaneous edge flip if their union contains no crossing edges.

2 An Upper Bound on the Flip Diameter of Rectangulations

In this section, we show that for every set P of n noncorectilinear points in a rectangle R , the diameter of $G(P)$ is $O(n \log n)$.

Theorem 1 *For every noncorectilinear set P of n points in the plane, the diameter of the graph $G(P)$ is $O(n \log n)$.*

Given a rectangulation r of P , we construct a sequence of $O(n \log n)$ operations that transforms r into a rectangulation with all segments vertical (a *canonical* rectangulation). Our method relies on the concept of independent sets, defined in terms of the bar visibility graph. Let r be a rectangulation of P . The *bar visibility graph* [12, 38, 42] on the horizontal segments of r is defined as a graph $H(r)$, where the vertices correspond to the horizontal segments in r ; and two horizontal segments s_1 and s_2 are adjacent in $H(r)$ if and only if there are points $a \in s_1$ and $b \in s_2$ such that ab is a vertical segment (not necessarily in r) that does not intersect any other horizontal segment in r . It is clear that the bar visibility graph is planar.

Observe that we can always change the orientation of any line segment s with $O(n)$ operations: successively shorten s using rotate operations until s contains no other segment endpoints, and then flip s . This simple procedure is formulated in the following subroutine.

`Shorten&Flip(r, s)`. Let s be a segment in a rectangulation r . Assume $s = ab$ and $p \in P$ is in the relative interior of s . While s contains the endpoint of some other segment, let $c_1 \in s$ and $c_2 \in s$ be the endpoints of some other segments closest to a and b , respectively (possibly $c_1 = c_2$). If $p \notin ac_1$, then apply `ROTATE(r, c_1)` to shorten $s = ab$ to c_1b . Else, apply `ROTATE(r, c_2)` to shorten s to ac_2 . When s does not contain the endpoint of any other segment, apply `FLIP(r, p)`.

The proof of Theorem 1 follows from a repeated invocation of the following lemma.

Lemma 1 *Let r be a rectangulation of a set of n pairwise noncorectilinear points in a rectangle R . There is a sequence of $O(n)$ flip and rotate operations that turns at least one quarter of the horizontal segments into vertical segments, and keeps vertical segments vertical.*

Proof: By the four color theorem [34], $H(r)$ has an independent set I that contains at least one quarter of the horizontal segments in r . The total number of endpoints of vertical segments that lie on some horizontal segment in I is $O(n)$. Successively call the subroutine `Shorten&Flip(r, s)` for all horizontal segment $s \in I$.

The horizontal segments in I are shortened and flipped into vertical orientation. All operations maintain the invariants that (1) the segments in I are pairwise disjoint, and (2) the remaining horizontal segments in I form an independent set in the bar visibility graph (of all horizontal segments in the current rectangulation). It follows that each operation either decreases the number of horizontal segments in I (flip), or decreases the number of segment endpoints that lie in the relative interior of a segment in I (rotate). After $O(n)$ operations, all segments in I become vertical. Since only the segments in I are flipped (once each), all vertical segments in r remain vertical, as required. \square

Proof of Theorem 1.: Let P be a set of n pairwise noncorectilinear points in a rectangle R . Denote by r_0 the rectangulation that consists of n vertical line segments, one passing through each point in P .

We show that every rectangulation r_1 of P can be transformed into r_0 by a sequence of $O(n \log n)$ flip and rotate operations. By Lemma 1, a sequence of $O(n)$ operations can decrease the number of horizontal segments by a factor of at least $4/3$. After at most $\log n / \log(4/3)$ invocations of Lemma 1, the number of horizontal segments drops below 1, that is, all segments become vertical and we obtain r_0 . \square

Remark. The proof of Theorem 1 is constructive, and provides an algorithm for transforming a rectangulation on n noncorectilinear points to the rectangulation with n vertical segments. If the rectangulations are maintained in a doubly connected edge list (DCEL) data structure, then a flip or rotate operation can be implemented in $O(1)$ time, and the bar-visibility graph can be computed in $O(n)$ time. The bottleneck of the overall running time is the 4-coloring of the bar-visibility graph. The current best algorithm for 4-coloring an m -vertex planar graph runs in $O(m^2)$ time [34], and a repeated call to this algorithm to exponentially decaying bar-visibility graphs takes $O(n^2)$ time. If we use an $O(m)$ -time 5-coloring algorithm [10, 18, 41], up to $\log n / \log(5/4)$ times, the overall running time improves to $O(n)$.

3 A Lower Bound on the Flip Diameter of Rectangulations

We show that the diameter of the graph $G(P)$ is $\Omega(n \log n)$ when P is an n -element *bit-reversal point set* (alternatively, Halton-Hammersley point set) [9, Section 2.2]. For every integer $k \geq 0$, we define a point set P_k of size $n = 2^k$ with integer coordinates lying in the square $R = [-1, n]^2$. For an integer m , $0 \leq m < 2^k$, with binary representation $m = \sum_{i=1}^k b_i 2^{i-1}$, the bit-reversal gives $y(m) = \sum_{i=1}^k b_i 2^{k-i}$. The bit-reversal point set of size $n = 2^k$ is $P_k = \{(m, y(m)) : m = 0, 1, \dots, n - 1\}$. By construction, no two points in P_k are corectilinear.




Fig. 3: The sets P_2 , P_3 , and P_4 . The edges connect pairs of points whose binary representations differ in a single bit, showing that P_k is a projection of a k -dimensional hypercube. The grey rectangles are spanned by point pairs whose binary representations differ exactly in the last coordinate.

We establish a lower bound of $k2^{k-3}$ for the diameter of $G(P_k)$ using a charging scheme. We define $k2^{k-1}$ empty rectangles (called *boxes*) spanned by P_k , and charge one unit for “saturating” a box with vertical segments (as defined below). We show that when a rectangulation with all horizontal segments is transformed into one with all segments vertical, each box becomes saturated. We also show that each rotate (resp., flip) operation contributes a total of at most 2 (resp., 4) units to the saturation of various boxes in our set. It follows that at least $(k2^{k-1})/4 = k2^{k-3} = n \log n / 8$ operations are required to saturate all $k2^{k-1}$ boxes.

Consider the point set P_k for some $k \in \mathbb{N}$. We say that a rectangle $B \subset [-1, n]^2$ is *spanned* by P_k if two opposite corners of B are in P_k ; and B is *empty* if its interior is disjoint from P_k .

Let \mathcal{B} be the set of closed rectangular boxes spanned by point pairs in P_k in which the binary representation of the x -coordinates (b_1, \dots, b_k) differ in precisely one bit. See Figure 3 for examples. Each point in P_k is incident to k boxes in \mathcal{B} , since there are k bits. Every box in \mathcal{B} is spanned by two points of P_k , thus $|\mathcal{B}| = k \cdot |P_k|/2 = k2^{k-1}$.

Each point is incident to k boxes of sizes $2^{i-1} \times 2^{k-i}$ for $i = 0, \dots, k-1$, since changing the i th bit b_i incurs an 2^{i-1} change in the x -coordinate and an 2^{k-i} change in the y -coordinate. However, if we change several bits successively, then either the x -coordinate changes by more than 2^{i-1} or the y -coordinate changes by more than 2^{k-i} , and so every box in \mathcal{B} is empty. Note also that the boxes of the same size are pairwise disjoint.

We now define the ‘‘saturation’’ of each box $B \in \mathcal{B}$ with respect to a rectangulation of P_k . Let $B \in \mathcal{B}$ and let r be a rectangulation of P_k . The vertical *extent* of B is the orthogonal projection of B into the y -axis. Consider the vertical segments of r clipped in B (i.e., the segments $s \cap B$ for all vertical segments s in r). The *saturation* of B with respect to r is the percentage of the vertical extent of B covered by projections of vertical segments of r clipped in B . See Figure 4 for examples. By definition, the




Fig. 4: A rectangulation of P_4 . The saturation of box B_1, \dots, B_8 is $\frac{1}{2}, 0, 1, 0, 1, 1, 1,$ and 1 , respectively.

saturation of B is a real number in $[0, 1]$. For every $B \in \mathcal{B}$, we have that the saturation of B is 0 when r is a rectangulation with all horizontal segments, and it is 1 when r consists of all segments vertical. If we transform an all-horizontal rectangulation into an all-vertical one by a sequence of operations, the total saturation of all $k2^{k-1}$ boxes in \mathcal{B} increases from 0 to $k2^{k-1}$. The key observation is that a single operation increases the total saturation of all boxes in \mathcal{B} by at most a constant.

It remains to bound the impact of a single operation on the saturation of a box in \mathcal{B} . Consider first an operation $\text{ROTATE}(r, c)$ that increases the saturation of some box $B \in \mathcal{B}$. A rotate operation shortens a segment s_1 and extends an orthogonal segment s_2 . The saturation of a box B can increase only if a vertical segment grows, so we may assume that s_1 is horizontal and s_2 is vertical. Denote by s the newly inserted portion of s_2 . Note that s lies in a single face of the rectangulation r . Similarly, if an operation $\text{FLIP}(r, p)$ increases the saturation of a box in \mathcal{B} , then it replaces a horizontal segment by a vertical segment passing through p . The new vertical segment lies in *two* adjacent faces of r , separated by the original horizontal segment through p . We represent the new vertical segment as the union of two collinear vertical segments $s \cup s'$ that meet at point p . In summary, an operation $\text{ROTATE}(r, c)$ inserts one vertical segment s that lies in the interior of a face of r , and an operation $\text{FLIP}(r, p)$ inserts two such segments. We show now that if such a new vertical segment s increases the saturation of some box in $B \in \mathcal{B}$, then s must lie in B .

Lemma 2 *Suppose that an operation inserts a vertical segment s that lies in a face f of the rectangulation r . If the insertion of s increases the saturation of a box $B \in \mathcal{B}$, then $s \subset B$.*

Proof: Suppose, to the contrary, that $s \not\subset B$. Let $p, q \in P_k$ denote the two opposite corners of points that span B , such that p is the upper left or upper right corner of B , and q is the opposite corner of B . Since s increases the saturation of B , it must intersect B . Hence $f \cap B \neq \emptyset$. Since $s \not\subset B$, at least one of the endpoints of s lies in the exterior of B . Assume, without loss of generality, that the upper endpoint of s lies outside B , and p is the upper left corner of B . Then, the top side of f is strictly above the top side of B . Since point p cannot be in the interior of f , the left side of the face f intersects the top side of B . Note that s and the left side of f have the same orthogonal projection on the y -axis. Therefore, the insertion of s cannot increase the saturation of B , contradicting our assumption. We conclude that both endpoints of s lie in B , and $s \subset B$. \square

Lemma 3 *A rotate (resp., flip) operation increases the total saturation of all boxes in \mathcal{B} by at most 2 (resp., 4).*

Proof: Suppose that an operation $\text{ROTATE}(r, c)$ inserts a vertical segment s , or an operation $\text{FLIP}(r, p)$ inserts two collinear vertical segments $s \cup s'$ that meet at p . By Lemma 2, the insertion of s increases the saturation of a box $B \in \mathcal{B}$ of height h by $|s|/h$ if $h \geq |s|$, and does not affect the saturation of boxes of height $h < |s|$. Recall that the boxes in \mathcal{B} have only k different sizes, $2^{i-1} \times 2^{k-i}$ for $i = 1, \dots, k$, and the boxes of the same size are pairwise disjoint. Let $j \in \{1, 2, \dots, k\}$ be the largest index such that $|s| \leq 2^{k-j}$. For $i = 1, \dots, j$, segment s increases the saturation of at most one box of height $h = 2^{k-i}$, and the increase is at most $|s|/h = |s| \cdot 2^{i-k}$. So s increases the total saturation of all boxes in \mathcal{B} by at most $\sum_{i=1}^j |s| 2^{i-k} \leq \sum_{i=1}^j 2^{i-j} < 2$, as required. \square

Theorem 2 *For every $n \in \mathbb{N}$, there is an n -element point set $P \subset [-1, n]^2$ such that the diameter of $G(P)$ is $\Omega(n \log n)$.*

Proof: First assume that $n = 2^k$ for some $k \in \mathbb{N}_0$. We have defined a set P_k of $n = 2^k$ points and a set \mathcal{B} of $k2^{k-1} = n \log n / 2$ boxes spanned by P_k . The total saturation of all boxes in \mathcal{B} is 0 in the rectangulation with horizontal segments, and $|\mathcal{B}| = n \log n / 2$ in the one with all segments vertical. By Lemma 3, a single flip or rotate operation increases the total saturation by at most 4. Therefore, at least $n \log n / 8$ operations are required to transform the horizontal rectangulation to the vertical one, and the diameter of $G(P_k)$ is at least $n \log n / 8$.

If n is not a power of two, then put $k = \lfloor \log_2 n \rfloor$ and let $P \subset [-1, n]^2$ be the union of P_k and $n - 2^k$ arbitrary (noncorectilinear) points in $[2^k, n]^2$. All axis-parallel segments containing the points in $P \setminus P_k$ are in the exterior of $[-1, 2^k]^2$. Therefore $k2^{k-3} = \Omega(n \log n)$ operations are required when all segments containing the points in $P_k \subset P$ change from horizontal to vertical. \square

4 The Flip Diameter for Diagonal Point Sets

We say that a point set P is *diagonal* if all points in P lie on the graph of a strictly increasing function (e.g., $f(x) = x$). In this section we show that the flip diameter is $O(n)$ for any n -element diagonal set. We denote by $\overline{p_i}$ the segment that contains p_i .

Theorem 3 For every $n \in \mathbb{N}$, the diameter of $G(P)$ is at most $11n$ when P is a diagonal set of n points.

Proof of Theorem 3: Without loss of generality, we may assume that the diagonal set is $P = \{p_i : i = 1, \dots, n\}$, where $p_i = (i, i)$. Given a rectangulation r for P , we construct a sequence of at most $5.5n$ flip and rotate operations that transforms r into a rectangulation with all segments vertical. Our algorithm consists of the following four phases.

Phase 1: No three consecutive segments are parallel. We describe an algorithm that, given a rectangulation r of P , applies less than $3n$ operations and returns a rectangulation in which no three consecutive points are contained in parallel segments.

`NoThreeParallel`(r). While there are three consecutive points in P contained in parallel segments, let $p_{i-1}, p_i, p_{i+1} \in P$ be arbitrary points such that the segments \mathcal{p}_{i-1} , \mathcal{p}_i , and \mathcal{p}_{i+1} are parallel, and call `Shorten&Flip`(r, \mathcal{p}_i).

Note that `Shorten&Flip`(r, \mathcal{p}_i) changes the orientation of the middle segment \mathcal{p}_i only. Hence, after one invocation of `Shorten&Flip`(r, \mathcal{p}_i), segments \mathcal{p}_{i-1} , \mathcal{p}_i , and \mathcal{p}_{i+1} have alternating orientations, and will never be in the middle of a consecutive parallel triple. This implies that `NoThreeParallel`(r) terminates after changing the orientation of at most $n/2$ segments, and no three consecutive points are in parallel segments in the resulting rectangulation.

It remains to bound the number of flip and rotation operations: we have already seen that there are at most $n/2$ flip operations, and we shall charge each rotation operation to the segment extended by that operation. Note that after an invocation of `Shorten&Flip`(r, \mathcal{p}_i), the two endpoints of \mathcal{p}_i lie on \mathcal{p}_{i-1} and \mathcal{p}_{i+1} , respectively, and so segment \mathcal{p}_i is not extended after a flip in Phase 1. Thus, each segment can be extended in two possible directions (left/right or up/down). We claim that each segment is extended at most once in each direction. Suppose to the contrary that a segment \mathcal{p}_k is extended twice in the same direction (refer to Figure 5, right). By symmetry, we may assume that \mathcal{p}_k is vertical, and




Fig. 5: Left-Middle: If the segment through $p_{i-1}, p_i,$ and p_{i+1} are parallel, we perform `Shorten&Flip`(r_0, \mathcal{p}_i). Right: If a rotation extends a vertical segment \mathcal{p}_k up from \mathcal{p}_i to \mathcal{p}_j , then \mathcal{p}_{i+1} cannot be horizontal.

it is extended upward in two rotate operations involving some horizontal segments \mathcal{p}_i and \mathcal{p}_j (each of which is sandwiched between two horizontal segments). In the first rotation, \mathcal{p}_k is extended from its intersection with \mathcal{p}_i to an intersection with \mathcal{p}_j . Hence, the order of the corresponding three points is $k < i < j$. This step creates a rectangle bounded by \mathcal{p}_i on the top, \mathcal{p}_j on the bottom, and \mathcal{p}_k on the left. The right boundary of this rectangle cannot be \mathcal{p}_{i+1} , since it is horizontal; and it cannot be any segment \mathcal{p}_ℓ , $\ell > i + 1$, otherwise the rectangle would contain p_{i+1} . The contradiction confirms the claim. We conclude that this phase involves at most $2n$ rotations and at most $n/2$ flip operations.

Phase 2: Creating a staircase. We define a *staircase* in a rectangulation of P as a monotone increasing path along the segments that does not skip two or more consecutive points (refer to Figure 6). In the second




Fig. 6: The left figure shows a rectangulation with a valid staircase, the right figure shows a rectangulation which does not have a valid staircase.

phase of our algorithm, we transform a rectangulation r with no three consecutive parallel segments into a rectangulation r that contains a staircase from the lower left to the upper right corner of R .

We construct a staircase π incrementally, starting from the lower left corner. In each step we successively extend π with one horizontal and one vertical segment to a point of P or to the upper right corner of R . Let π be a current staircase from the lower left corner to a point $a \in P$. Denote by b and c the next two points of the diagonal set P . If there is a monotone path from a to b or from a to c , then append this path to π to obtain a longer staircase. Suppose there is no such path. Then the segments \overline{ab} and \overline{ac} must be parallel, otherwise there would be a monotone path from a to b . The segment \overline{ac} must be perpendicular to the segment \overline{ab} since there are no three consecutive parallel segments.

We would like to perform a rotation at the intersection of segments \overline{ab} and \overline{ac} , and then we can extend π from a to c (along \overline{ab} and \overline{ac}). However, other operations may be necessary before we can rotate at $\overline{ab} \cap \overline{ac}$. Successively rotate all intersection points of \overline{ab} with other segments that are to the right of b if \overline{ab} is horizontal, and that are above b if \overline{ab} is vertical (rotating at $\overline{ab} \cap \overline{ac}$ eventually).

We claim that this phase requires at most n rotations. Let a , b , and c be defined as before, such that there is no path from a to either b or c . Assume, by symmetry, that a and b are horizontal and c is vertical. Let d_1, \dots, d_k be segment endpoints on \overline{ab} to the right of b . The bottom endpoints of the vertical segments through $d_1 \dots d_k$ are now below b , and so these endpoints will never be involved in another rotation in our incremental procedure. Similarly, we perform a rotation at the left endpoint of each horizontal segment at most once. Therefore, the total number of rotations is bounded by the number of segments, which is n .

Phase 3: Extending all vertical segments of the staircase to full length. The third phase of our algorithm returns a rectangulation in which every vertical segment of the staircase extends to the top and bottom sides of the bounding box R . The staircase subdivides R into two regions, one above and one below (see Figure 7). We handle the two regions independently without modifying the staircase. To handle the region above, we sweep from top to bottom and successively shorten each horizontal segment we encounter until further shortening is impossible or would modify the staircase. If we encounter a horizontal segment s that is not part of the staircase, we use subroutine `Shorten&Flip(r, s)` to shorten it as much as possible and then flip it to vertical position. Similarly, we sweep the region below the staircase bottom-up.

Each flip or rotate operation in Phase 3 increases the number of segment endpoints lying on the top or bottom side of R . Therefore, the number of operations is bounded by the total number of segment endpoints, which is $2n$.




Fig. 7: Left: a rectangulation with a valid staircase. Middle: the result of sweeping the region above the staircase. Right: the result of sweeping the regions on both sides of the staircase.

Phase 4: Final steps. The vertical segments of the staircase extend to the top and bottom of the bounding box and partition it into vertical strips (Figure 7, right). Each strip contains a horizontal segment of the staircase, including a point in P . Because the staircase does not skip two or more consecutive points, there are at most three points in each strip. The total number of points in all strips combined is less than n .

If there is just one point in a strip, then we flip the horizontal segment that contains it. If there are two points in a strip, then one lies on a horizontal and one on a vertical segment. In this case, perform a rotation so the vertical segment reaches maximum length, and then flip the horizontal segment. If there are three points in a strip, then the middle point lies on a horizontal segment and the other two lie on vertical segments above and below the staircase, respectively. In this case, we perform two rotations so that both vertical segments reach full length, and then flip the horizontal segment. Similarly to Phase 3, each operation in this phase increases the number of segment endpoints lying on the top or bottom side of R . Therefore, $2n$ is an upper bound on the number of operations in Phases 3 and 4 combined. \square

Remark. Felsner et al. [17] showed that the rectangulations of a diagonal point set are in bijection with *twin pairs of binary trees*. Visually, the part of the boundaries of the rectangles lying above (resp., below) the points in P form a binary tree whose vertices are the points in P and the top-right (resp., bottom-left) corners of R . This observation allows determining the number of rectangulations for diagonal point sets in terms of Baxter permutations (cf. [17]). Theorem 3 implies the following.

Corollary 1 *For any two pairs of twin binary trees with n leaves, there is a sequence of twin binary trees $(t_1, t'_1), \dots, (t_k, t'_k)$ of size $k = O(n)$ that starts with the first pair, ends with the second pair, and for every i ($1 \leq i < k$), either $t_i = t_{i+1}$ or t_{i+1} is obtained from t_i by a single tree-rotate operation (the same for t'_i and t'_{i+1}).*

5 Generalization to Convex Subdivisions

Given a set P of n points in the plane, a *convex subdivision* for P is a subdivision of the plane into convex cells by n pairwise noncrossing line segments (possibly lines or half-lines), such that each segment contains exactly one point of P and no three segments have a point in common.

The flip and rotate operations can be interpreted for convex subdivisions of a point set P , as well (see Figure 8). The definition of the operation $\text{ROTATE}(r, c)$ is identical to the rectilinear version. The operation $\text{FLIP}(r, p)$ requires more attention, since a segment may have infinitely many possible orientations.




Fig. 8: A convex subdivision r_1 of 6 points, $r_2 = \text{FLIP}(r_1, p, \sigma)$, and $r_3 = \text{ROTATE}(r_2, c)$.

Definition 3 (Flip) Let r be a convex subdivision of P , let $p \in P$ be a point such that the segment s containing p does not contain any endpoints of other segments, and let $\sigma \in \mathbb{S}^1$ be a unit vector. The operation $\text{FLIP}(r, p, \sigma)$ replaces s by a segment of direction σ containing p .

Similarly to the graph of rectangulations $G(P)$, we define the *graph of convex subdivisions* of P , $\widehat{G}(P) = (V, E)$, where the vertex set is $V = \{r : r \text{ is a convex subdivision of } P\}$ and the edge set is $E = \{(r_1, r_2) : \text{a single flip or rotate operation on } r_1 \text{ produces } r_2\}$. Our main result in this section is that even though $\widehat{G}(P)$ is an infinite graph, its diameter is $O(n \log n)$, where $n = |V|$.

Theorem 4 For set P of n points, the graph $\widehat{G}(P)$ is connected and its diameter is $O(n \log n)$.

We show that any convex subdivision can be transformed into a subdivision with all segments vertical through a sequence of $O(n \log n)$ operations. Subroutine $\text{Shorten\&Flip}(r, s)$ from Section 2 can be adapted almost verbatim: for a unit vector $\sigma \in \mathbb{S}^1$, subroutine $\text{Shorten\&Flip}(r, s, \sigma)$ shortens segment s maximally by rotate operations, and then flips it to direction σ .

Lemma 4 Let r be a convex subdivision of a set of n points in the plane with distinct x -coordinates. There is a sequence of $O(n)$ flip and rotate operations that turns at least a $\frac{1}{36}$ fraction of the nonvertical segments vertical, and keeps all vertical segments vertical.

Before proving Lemma 4, we need to introduce a few technical terms. Consider a convex subdivision r of a set of n points with distinct x -coordinates. We say that a segment s_1 *hits* another segment s_2 if an endpoint of s_1 lies in the relative interior of s_2 . An *extension of s_1 beyond s_2 hits s_3* if s_1 hits s_2 and s_1 is contained in a segment s'_1 , such that s'_1 hits s_3 and s'_1 crosses at most one segment (namely, s_2). Note that the operation $\text{ROTATE}(r, s_1 \cap s_2)$, if applicable, would extend s_1 beyond s_2 to hit segment s_3 . We define the *extension visibility digraph* $\widehat{H}(r)$ on all segments in r , where the vertices correspond to the segments in r , and we have a directed edge (s_2, s_3) if s_2 hits s_3 or there is a segment s_1 such that an extension of s_1 beyond s_2 hits s_3 .

The graph $\widehat{H}(r)$ is not necessarily planar: it is not difficult to construct a convex subdivision r for a set of $O(t^2)$ points where $\widehat{H}(r)$ contains a subgraph isomorphic to the complete graph K_t (Figure 9). It is enough to describe the segments of such a construction (the points may lie anywhere in the interior of the segments). Start with the sides of a convex t -gon, and extend the sides in counterclockwise direction to obtain a convex subdivision. The extension visibility digraph of these t segments contains a cycle (s_1, \dots, s_t) . Add short segments in the exterior of the initial t -gon: If a segment hits s_i and its supporting line passes through s_j , it generates an edge (s_i, s_j) in the extension visibility digraph.




Fig. 9: A convex subdivision r of a set of 15 points. The five bold segments induce K_5 in the extension visibility graph $\widehat{H}(r)$.

Note that the number of edges in $\widehat{H}(r)$ is at most $4n$, since each segment hits at most two other segments, but some segments extend to infinity; and the extension of each segment beyond each of its endpoints hits at most one other segment. If P contains n points, the average degree in $\widehat{H}(r)$ is less than 8. Therefore, $\widehat{H}(r)$ has an independent set of size at least $n/9$ (obtained by successively choosing minimum-degree vertices [23, 39]).

Proof of Lemma 4.: Let r be a convex subdivision of a set of n points with distinct x -coordinates. Let I_0 be an independent set in the extension visibility graph $\widehat{H}(r)$ induced by all nonvertical segments. As noted above, I_0 contains at least $1/9$ of the nonvertical segments in r . Let $I_1 \subseteq I_0$ be an independent set in the bar visibility graph of the segments in I_0 (two nonvertical segments in I_0 are mutually visible if there is a vertical segment between them that does not cross any segment of the subdivision). Since the bar visibility graph is planar, we have $|I_1| \geq |I_0|/4$, and so I_1 contains at least a $1/36$ fraction of the nonvertical segments in r . The total number of segment endpoints that lie in the relative interior of segments in I_1 is $O(n)$. An invocation of subroutine `Shorten&Flip`($r, s, (0, 1)$) for each segment $s \in I_1$ changes their orientation to vertical.

The operations maintain the invariants that (1) the segments in I_1 are pairwise disjoint; and (2) the nonvertical segments in I_1 form an independent set in both \widehat{H} and the bar visibility graph of all segments in I_0 . It follows that each operation either decreases the number of nonvertical segments in I_1 (flip), or decreases the number of segment endpoints that lie in the relative interior of a nonvertical segment in I_1 (rotate). After performing $O(n)$ operations, all segments in I_1 become vertical. Since only the segments in I_1 change orientation (each of them is flipped to become vertical), all vertical segments in r remain vertical, as required. \square

Proof of Theorem 4.: Let P be a set of n points in a bounding box. We may assume, by rotating the point set if necessary, that the points in P have distinct x -coordinates. Denote by r_0 the convex subdivision given by n vertical line segments, one passing through each point in P .

Consider a convex subdivision r_1 of P . By Lemma 4, $O(n)$ operations can decrease the number of nonhorizontal segments by a factor of at least $36/35$. After at most $\log n / \log(36/35)$ invocations of Lemma 4, the number of nonvertical segments drops below 1, that is, all segments become vertical and we obtain r_0 , as claimed. \square

A linear upper bound for collinear points

We show that the bound $O(n \log n)$ on the diameter of the flip graph $\widehat{G}(P)$ from Theorem 4 can be improved to $O(n)$ for some simple point configurations.

Theorem 5 *For every $n \in \mathbb{N}$, the diameter of $\widehat{G}(P)$ is $O(n)$ when P is a set of n collinear points.*

Proof: We may assume that $P = \{p_i : i = 1, \dots, n\}$, $p_i = (i, 0)$. Let r be a convex subdivision of P . No segment in r is horizontal, since each segment contains a unique point. We show that there is a sequence of $O(n)$ operations that transforms r into a convex subdivision with all segments vertical. Suppose that not all segments in r are vertical, and let $m > 0$ be the minimum absolute value of the slope of a nonvertical segment. Our algorithm proceeds in two phases.

Phase 1: Building a staircase. In this phase, we transform r into a convex subdivision in which every \mathcal{p}_i is a nonvertical ray with a left endpoint at infinity, and has slope m/n when i is odd and $-m/n$ when i is even. We process the points p_1, \dots, p_n successively in this order. We maintain the following invariant (refer to Figure 10): When we start processing p_i , segments $\mathcal{p}_1, \dots, \mathcal{p}_{i-1}$ are already rays with




Fig. 10: Transforming a convex subdivision for n points lying on the x -axis into rays of alternating slopes m/n and $-m/n$.

the desired properties, and they do not contain the endpoints of any segment $\mathcal{p}_1, \dots, \mathcal{p}_n$ in their relative interiors.

Suppose that points p_j , for all $1 \leq j < i$, have already been processed. We use the subroutine $\text{Shorten\&Flip}(r, \mathcal{p}_i, \sigma)$ to shorten \mathcal{p}_i and flip it into the desired orientation (of slope m/n or $-m/n$ depending on the parity of i). For $i = 1$, this already guarantees that segment \mathcal{p}_1 is a ray with a left endpoint at infinity. If $i > 1$, however, the left endpoint of \mathcal{p}_i lies on the ray \mathcal{p}_{i-1} , whose slope is different from that of \mathcal{p}_i . Due to the invariant, \mathcal{p}_{i-1} contains no segment endpoints to the right of the intersection point $\mathcal{p}_{i-1} \cap \mathcal{p}_i$. We can now apply a rotation at $\mathcal{p}_{i-1} \cap \mathcal{p}_i$. Since the slopes of all segments are at least m/n in absolute value, \mathcal{p}_i extends to infinity, and our invariant is established for segments $\mathcal{p}_1, \dots, \mathcal{p}_i$.

We argue that Phase 1 uses $O(n)$ operations. For each point p_i , we invoke $\text{Shorten\&Flip}(r, \mathcal{p}_i, \sigma)$, which performs only one flip operation, followed by at most one rotation. It is enough to bound the number of rotations performed in the invocations of subroutines Shorten\&Flip . We claim that these invocations extend each segment at most once in each direction. Suppose, to the contrary, that a segment \mathcal{p}_k is extended in the same direction twice (when processing points p_i and p_j). Since the points are processed from left to right, we have $i < j < k$. Consider the step in which a rotation extends \mathcal{p}_k from its intersection with \mathcal{p}_i to an intersection with \mathcal{p}_j . Then, the x -axis and segment \mathcal{p}_k intersect \mathcal{p}_i and

p_j in different orders. Hence, p_i and p_j must cross each other, contradicting the fact that they are part of a convex subdivision. This proves our claim. It follows that Phase 1 uses at most $5n$ operations.

Phase 2: Orienting all lines to be vertical. Refer to Figure 11. We are given a convex subdivision in




Fig. 11: The first four passes over the rays yield vertical lines through all even points, separating the odd points.

which every p_i is a ray with a left endpoint at infinity, and has slope m/n when i is odd and $-m/n$ when i is even. Note that only consecutive segments intersect. (This is still not a “canonical” subdivision, as m depends on the initial convex subdivision.) In Phase 2, we transform all segments into vertical lines. We make five passes over the odd or even points: (1) For every odd point p_i from right to left, rotate ray p_i into a line of slope m/n . As a result, the lines through the odd points separate the even points. (2) We can now flip the segments through the even points independently into vertical segments. (3) For every even point p_i from left to right, rotate the top endpoints to infinity. (4) For every even point from right to left, rotate the bottom endpoint to infinity. We obtain vertical lines through the even points, that separate the odd points. (5) Finally, flip the segments through the odd points independently into vertical lines. We made three passes over even segments and two passes over odd segments, so the total number of operations in Phase 2 is $2n + \lceil n/2 \rceil \leq 3n$. The two phases combined use at most $8n$ operations. \square

6 Conclusions

We have shown that the diameter of the flip graph $G(P)$ is between $\Omega(n)$ and $O(n \log n)$ for every n -element point set P , and these bounds cannot be improved. The diameter is $\Theta(n)$ for diagonal point sets, and $\Theta(n \log n)$ for the bit-reversal point set. The flip graph $G(P)$ of a noncorectilinear set P is uniquely determined by the permutation of the x - and y -coordinates of the points [2] (e.g., diagonal point sets correspond to the identity permutation). It is an open problem to find the average diameter of $G(P)$ over all n -element permutations. It would already be of interest to find broader families of point sets with linear diameter: Is the diameter of $G(P)$ linear if P is in convex position, unimodal, or corresponds to a separable permutation (see [4])?

We have shown that the diameter of the flip graph is also $O(n \log n)$ for the convex subdivisions of n points in the plane. We do not know whether this bound is tight. It is possible that the flip diameter is $\Omega(n \log n)$ for the bit-reversal point set defined in Section 3, but our proof of Theorem 2 heavily relies on axis-aligned boxes and does not seem to extend to convex subdivisions.

Given a convex subdivision r of a point set P , the flip and rotate operations can be thought of as a continuous deformation; refer to Figure 12: $\text{FLIP}(r, p, \sigma)$ rotates the segment containing p continuously to position σ ; and $\text{ROTATE}(r, c)$ rotates continuously a portion of the segment containing c into the extension of the segment that currently ends at c . The *weight* of an operation can be defined as the number of




Fig. 12: Left: a convex subdivision r_1 of 6 points. Middle: operation $\text{FLIP}(r_1, p, \sigma)$ performed as a continuous deformation of the segment containing p . Right: operation $\text{ROTATE}(r_2, c)$ performed as a continuous deformation of part of the segment containing c .

vertices swept during this continuous deformation. By Theorem 4, a sequence of $O(n \log n)$ operations can transform any convex subdivision to any other convex subdivision on n points. A single operation, however, may have $\Omega(n)$ weight. We conjecture that the weighted diameter of the graph $\widehat{G}(P)$ is also $O(n \log n)$ for every n -elements point set P .

Acknowledgments. Research on this paper was initiated at the Workshop on *Counting and Enumerating of Plane Graphs* held in March 2013 at Schloss Dagstuhl. We thank Sonia Chauhan, Michael Hoffmann, and André Schulz for insightful comments and stimulating conversations about the problems discussed in this paper.

References

- [1] E. Ackerman, *Counting Problems for Geometric Structures: Rectangulations, Floorplans, and Quasi-Planar Graphs*, Ph.D. Thesis, Dept. of Computer Science, Technion—Israel Inst. of Technology, Haifa, Israel, 2006.
- [2] E. Ackerman, G. Barequet, and R.Y. Pinter, On the number of rectangulations of a planar point set, *J. Combinatorial Theory, Series A*, 113:6 (2006), 1072–1091.
- [3] A. Asinowski, G. Barequet, M. Bousquet-Mélou, T. Mansour, and R.Y. Pinter, Orders induced by segments in floorplans and (2-14-3,3-41-2)-avoiding permutations, *Electronic J. Combinatorics*, 20 (2), 35 pp., 2013.
- [4] P. Bose, J. Buss, and A. Lubiw, Pattern matching for permutations, *Information Processing Letters*, 65 (1998), 277–283.
- [5] K. Buchin, D. Eppstein, M. Löffler, M. Nöllenburg, and R.I. Silveira, Adjacency-preserving spatial treemaps, *Proc. 12th Workshop on Algorithms and Data Structures (WADS)*, vol. 6844 of *Lecture Notes in Computer Science*, 2011 Springer, pp. 159–170.
- [6] A.L. Buchsbaum, E.R. Gansner, C.M. Procopiuc, and S. Venkatasubramanian, Rectangular layouts and contact graphs, *ACM Trans. on Algorithms*, 4 (2008), #8, 28 pp.
- [7] F.C. Calheiros, A. Lucena, and C.C. de Souza, Optimal rectangular partitions, *Networks*, 41 (2003), 51–67.

- [8] M. Cardei, X. Cheng, X. Cheng, and D.Z. Du, A tale on guillotine cut, in: *Novel Approaches to Hard Discrete Optimization* (P.M. Pardalos and H. Wolkowicz, eds.), vol. 37 of *Field Institute Communications*, 2003, AMS, pp. 41–54.
- [9] B. Chazelle, *The Discrepancy Method*, Cambridge University Press, 2000.
- [10] N. Chiba, T. Nishizeki, and N. Saito, A linear 5-coloring algorithm of planar graphs, *J. Algorithms* 2 (1981), 317–327.
- [11] H. de Fraysseix, P.O. de Mendez, and J. Pach, A left-first search algorithm for planar graphs, *Discrete & Computational Geometry*, 13 (1995), 459–468.
- [12] P. Duchet, Y. Hamidoune, M. Las Vergnas, and H. Meyniel, Representing a planar graph by vertical lines joining different levels, *Discrete Mathematics*, 46 (1983), 319–321.
- [13] D.Z. Du, L.Q. Pan, and M.T. Shing, Minimum edge length guillotine rectangular partition, Technical Report MSRI 02418-86, University of California, Berkeley, CA, 1986.
- [14] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek, Area-universal and constrained rectangular layouts, *SIAM J. Computing*, 41 (2012), 537–564.
- [15] S. Felsner, Rectangle and square representations of planar graphs, in: *Thirty Essays in Geometric Graph Theory* (J. Pach, ed.), 2013, Springer, New York, pp. 213–248.
- [16] S. Felsner, Exploiting air-pressure to map floorplans on point sets, *J. Graph Algorithms and Applications*, 18 (2014), 233–252.
- [17] S. Felsner, É. Fusy, M. Noy, and D. Orden, Bijections for Baxter families and related objects, *J. Combinatorial Theory, Series A*, 118 (2011), 993–1020.
- [18] Greg N. Frederickson, On linear-time algorithms for five-coloring planar graphs, *Inf. Proc. Lett.*, 19 (1984), 219–224.
- [19] T.F. Gonzalez and S.-Q. Zheng, Bounds for partitioning rectilinear polygons, *Proc. 1st Sympos. on Computational Geometry*, 1985, ACM Press, pp. 281–287.
- [20] T.F. Gonzalez and S.-Q. Zheng, Improved bounds for rectangular and guillotine partitions, *J. Symbolic Computation*, 7 (1989), 591–610.
- [21] T.F. Gonzalez and S.-Q. Zheng, Approximation algorithms for partitioning a rectangle with interior points, *Algorithmica*, 5 (1990), 11–42.
- [22] M. Hasan, M. Rahman, and M. Karim, Box-rectangular drawings of planar graphs, *J. Graph Algorithms and Applications*, 17 (2013), 629–646.
- [23] D.S. Hochbaum, Efficient bounds for the stable set, vertex cover, and set packing problems, *Discrete Applied Mathematics*, 6 (1983), 243–254.
- [24] F. Hurtado, M. Noy, and J. Urrutia, Flipping edges in triangulations, *Discrete & Computational Geometry*, 22 (1999), 333–346.
- [25] K. Koźimiński and E. Kinnen, Rectangular duals of planar graphs, *Networks*, 15 (1985), 145–157.
- [26] M. van Kreveld and B. Speckmann, On rectangular cartograms, *Computational Geometry: Theory and Applications*, 37 (2007), 175–187.
- [27] C. Lawson, Software for c_1 surface interpolation, in: *Mathematical Software III* (J. Rice, ed.), 1977, Academic Press, New York, pp. 161–194.

- [28] C. Levcopoulos, Fast heuristics for minimum length rectangular partitions of polygons, *Proc. 2nd Sympos. on Computational Geometry*, 1986, ACM Press, pp. 100–108.
- [29] C.C. Liao, H.I. Lu, and H.C. Yen, Compact floor-planning via orderly spanning trees, *J. Algorithms*, 48 (2003), 441–451.
- [30] A. Lingas, R.Y. Pinter, R.L. Rivest, and A. Shamir, Minimum edge length rectilinear decompositions of rectilinear figures, *Proc. 20th Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, 1982, pp. 53–63.
- [31] H. Meijer and D. Rappaport, Simultaneous edge flips for convex subdivisions, *Proc. 16th Canadian Conference on Computational Geometry*, Montréal, Québec, Canada, 2004, pp. 57–59.
- [32] M. Rahman, T. Nishizeki, and S. Ghosh, Rectangular drawings of planar graphs, *J. Algorithms*, 50 (2004), 62–78.
- [33] E. Raisz, The rectangular statistical cartogram, *Geographical Review*, 24 (1934), 292–296.
- [34] N. Robertson, D. P. Sanders, P. Seymour, and R. Thomas, The four-colour theorem, *J. Combinatorial Theory, Series B*, 70 (1997), 2–44.
- [35] F. Santos and R. Seidel, A better upper bound on the number of triangulations of a planar point set, *J. Combinatorial Theory, Series A*, 102 (2003), 186–193.
- [36] M. Sharir and E. Welzl, Random triangulations of planar point sets, *Proc. 22nd Sympos. on Computational Geometry*, 2006, ACM Press, pp. 273–281.
- [37] D. Sleator, R. Tarjan, and W. Thurston, Rotations distance, triangulations and hyperbolic geometry, *J. AMS*, 1 (1988), 647–682.
- [38] R. Tamassia and I. G. Tollis, A unified approach to visibility representations of planar graphs, *Discrete & Computational Geometry*, 1 (1986), 321–341.
- [39] P. Turán, On an extremal problem in graph theory (in Hungarian), *Math. Fiz. Lapok*, 48 (1941), 436–452.
- [40] P. Ungar, On diagrams representing graphs, *J. London Math. Soc.*, 28 (1953), 336–342.
- [41] M. H. Williams, A linear algorithm for colouring planar graphs with five colours, *Computer Journal*, 28 (1985), 78–81.
- [42] S.K. Wismath, Characterizing bar line-of-sight graphs, in *Proc. 1st Sympos. on Computational Geometry*, 1985, ACM Press, pp. 147–152.
- [43] B. Yao, H. Chen, C.K. Cheng, and R. Graham, Floorplan representations: Complexity and connections, *ACM Trans. on Design Automation of Electronic Systems*, 8 (2003), 55–80.