On Degrees in Random Triangulations of Point Sets*

Micha Sharir
School of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel
Courant Inst. of Math. Sci.
New York University
New York, NY 10012, USA
michas@tau.ac.il

Adam Sheffer
School of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel
sheffera@tau.ac.il

Emo Welzl
Institut für Theoretische
Informatik, ETH Zürich
CH-8092 Zürich, Switzerland
emo@inf.ethz.ch

ABSTRACT

We study the expected number of interior vertices of degree i in a triangulation of a point set S, drawn uniformly at random from the set of all triangulations of S, and derive various bounds and inequalities for these expected values. One of our main results is: For any set S of N points in general position, and for any fixed i, the expected number of vertices of degree i in a random triangulation is at least $\gamma_i N$, for some fixed positive constant γ_i (assuming that $N > i$ and that at least some fixed fraction of the points are interior).

We also present a new application for these expected values, using upper bounds on the expected number of interior vertices of degree 3 to get a new lower bound, $\Omega(2.4317^N)$, for the minimal number of triangulations any N-element planar point set in general position must have. This improves the previously best known lower bound of $\Omega(2.33^N)$.

Categories and Subject Descriptors

G.2.1 [**Discrete Mathematics**]: Combinatorics—*Counting Problems*

General Terms

Theory

Keywords

random triangulation, counting, degree sequences, number of triangulations, plane graphs, crossing-free geometric graphs, charging

*Work on this paper was partially supported by Grants 155/05 and 338/09 from the Israel Science Fund. Work by Micha Sharir was also supported by NSF Grants CCF-05-14079 and CCF-08-30272, by Grant 2006/194 from the U.S.-Israel Binational Science Foundation, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Research by Emo Welzl was supported by the Swiss National Science Foundation, SNF Project 200020-125027.

1. INTRODUCTION

A *planar graph* is a graph that can be drawn in the plane in such a way that its edges intersect only at their endpoints. A *plane straight-line graph* is an embedding of a planar graph in the plane such that its edges are mapped to non-crossing straight line segments. In this paper, we consider only plane *straight-line* graphs, but refer to them as plane graphs or non-crossing graphs, for simplicity.

A *triangulation* of a finite point set S in the plane is a maximal plane graph on S. Let $\mathcal{T}(S)$ denote the set of all triangulations of S and let $\mathcal{P}(S)$ denote the set of all plane graphs of S. Moreover, $\mathsf{tr}(S) := |\mathcal{T}(S)|$ and $\mathsf{pg}(S) := |\mathcal{P}(S)|$.

Improved bounds on the number of plane graphs of several special kinds on a set S of N points have been obtained by studying the properties of a graph chosen uniformly at random from the set of all such graphs on a fixed set S.

• $\mathsf{pg}(S) \leq \mathsf{tr}(S) \cdot 8^N$ holds for any planar point set S of N points, because each triangulation has at most $3N$ edges and each plane graph is contained in some triangulation. Razen et al. [9] proved that this relation is not tight with an exponentially better bound of $\mathsf{pg}(S) = O\left(\mathsf{tr}(S) \cdot 7.98^N\right)$. The new bound is established by showing that the expected number of edges in a plane graph, uniformly chosen from $\mathcal{P}(S)$, is at least $\frac{M}{2} + \frac{N-4}{16}$, where $M \leq 3N-6$ is the number of edges in a triangulation of S.

• Upper bounds for the maximal value of $\mathsf{tr}(S)$, for sets S of N points in the plane, have been studied during the past three decades ([1, 3, 10, 11, 14, 13]). The best known bound is $\mathsf{tr}(S) < 30^N$ from [12]. It is obtained by showing that the expected number of vertices of degree 3 in a triangulation, drawn uniformly at random from $\mathcal{T}(S)$, exceeds $N/30$.

(Abstract random planar graphs of a given size have been considered, for example, in [4, 6, 8].) Here we continue the study of random triangulations of planar point sets, initiated in [13], investigating the number of vertices of degree i in a random triangulation. For $T \in \mathcal{T}(S)$, let $v_i(T)$ denote the number of interior vertices (i.e., points which are not vertices of the convex hull of S) of degree i in T. Furthermore, let

$$\hat{v}_i = \hat{v}_i(S) := \mathbb{E}(v_i(T)) = \frac{\sum_{T\in\mathcal{T}(S)} v_i(T)}{\mathsf{tr}(S)},$$

i.e., the expected number of interior vertices of degree i in a random triangulation of S. Due to linearity of expectation, any linear identity or inequality in the $v_i(T)$'s will also be satisfied by the $\hat{v}_i$'s. However, as noted in [13], the $\hat{v}_i$'s are more constrained than the v_i's. Charging schemes seem to provide an efficient way to bound the $\hat{v}_i$'s; they are the main tool used throughout this paper. In [12, 13] considerations were restricted to point sets with triangular convex hull. Here we abandon this restriction. For this purpose, we define S^o as the set of interior points of S, and write $N = n + h$, with h the number of hull vertices and $n := |S^o|$.

Section 2 establishes the upper bounds $\hat{v}_3 \le \frac{2n+h/2}{5}$ and $\hat{v}_3 \le n/2$ (for $N \ge 7$), and it is shown how these bounds imply a lower bound of $\Omega\left(2.4317^N\right)$ on the minimal number of triangulations any set of N points in general position must have. This improves the previous bound of $\Omega\left(2.33^N\right)$ in Aichholzer et al. [2]. A construction with $O(3.455^N)$ triangulations is presented by Hurtado and Noy [5], the currently best known upper bound on the minimum number of triangulations. McCabe and Seidel [7] proved that when h is constant, there are $\Omega\left(2.63^N\right)$ triangulations (full proof not yet published). This is sharper than our bound. More specifically, our bound depends on the ratio between n and N, and attains the minimum when $n \approx 0.89901N$. When $N - n = O(1)$, as in [7], our bound is $\Omega\left(2.5^N\right)$, short of their bound.

In Section 3, we derive inequalities relating the $\hat{v}_i$'s and we show that the inequalities imply the bound

$$\hat{v}_4 \ge \max\left\{\tfrac{1}{340}\left(n + 15 - \tfrac{8h}{3}\right),\ \tfrac{1}{1360}\left(n + 18 - 2h\right)\right\};$$

(this will prove useful in the later sections).

Following the lower bounds on $\hat{v}_3$ in [12, 13] and on $\hat{v}_4$ (just stated), we complete the picture in Section 4, where we derive linear lower bounds on $\hat{v}_i$ for any $i \ge 5$. Specifically, we show that for each $i \ge 5$, $\varepsilon > 0$, and $3 \le h \le \left(\frac{1}{2} - \varepsilon\right) n$, there exists a positive constant $\delta_{i,h}$ such that, for any N-element point set S with h hull vertices, we have $\hat{v}_i \ge \delta_{i,h} n$. That is, a random triangulation of S is expected to contain many interior vertices (at least a constant fraction of the points) of degree i, for *every* $3 \le i < N$.

We end Section 4 with observations on improved lower bounds on the expected number of low-degree vertices in a random triangulation, which are sharper than the bounds that we can get for individual degrees: $\max\{\hat{v}_3, \hat{v}_4, \hat{v}_5, \hat{v}_6\} \ge \frac{n-h+9}{10}$ and $\max\{\hat{v}_4, \hat{v}_5, \ldots, \hat{v}_{11}\} \ge \frac{12n-9.5h+45}{180}$.

1.1 Notations

Vints and bints. As in [12, 13], we consider $S^o \times \mathcal{T}(S)$ and call each of its elements a *vint* (*v*ertex *in* *t*riangulation), i.e. a vint is an instance of a vertex in a specific triangulation. The degree of a vint (p, T) is the degree (number of neighbors) of p in T; a vint of degree i is called an *i-vint*. Note that hull vertices do not participate in this definition.

The *link* of a vint (p, T) is the face obtained by removing p with incident edges from T, a star-shaped polygon with respect to p; its number of edges equals the degree of p.

Let B denote the set of edges of the convex hull of S. Similarly to the set of vints, we consider $B \times \mathcal{T}(S)$ and call each of its elements a *bint* (*b*oundary edge *in* *t*riangulation).

Catalan numbers. $C_m := \frac{1}{m+1}\binom{2m}{m} = \Theta(m^{-3/2}4^m) = \Theta^*(4^m)$, $m \in \mathbb{N}_0$, denotes the mth *Catalan number*. (In $O^*()$, $\Theta^*()$, and $\Omega^*()$, we neglect polynomial factors.) The number of triangulations of h points in convex position is C_{h-2} and the number of subtrees of the complete binary tree, that contain exactly k nodes, is C_k (see [15, section 5.3]), properties useful later in the paper.

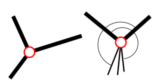

Figure 1: Separable edges.

Separable edges. Let $w = (p, T)$ be a vint. We call an edge e incident to p in T *separable at w* if it can be separated from the other edges incident to p by a line through p. Equivalently, the two angles between e and its clockwise and counterclockwise neighboring edges (around p) have to sum up to more than π. We observe the easy following properties.

(S0) No edge is separable at both vints induced by its endpoints.

(S1) If w has degree 3, every edge incident to its point is separable at w (recall that points of vints are interior).

(S2) If w has degree at least 4, at most two incident edges can be separable at w; if two edges are separable at w, they must be consecutive.

External chords. Let o be a link of a vint. An edge, which is openly disjoint from o but its endpoints are vertices of o, is called an *external chord* of o (see Figure 2). Observe that a link with i edges can have at most $i - 3$ external chords.

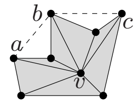

Figure 2: External chords ab and bc.

"Flips-down-to" relation. For vints $u = (p_u, T_u)$ and $v = (p_v, T_v)$, we define $u \to v$ if $p_v = p_u$ and T_v is obtained by flipping one edge incident to p_u in T_u. We see that u and v are associated with the same point but in different triangulations; u is an $(i+1)$-vint and v an i-vint, for some $i \ge 3$. We let $\to^*$ denote the transitive reflexive closure of $\to$. If $u \to^* v$, we say that u *can be flipped down to* v.

2. UPPER BOUNDS FOR $\hat{v}_3$

We derive upper bounds on $\hat{v}_3$. Upper bounds of this kind translate into lower bounds on the number of triangulations every N-element point set in general position must have.

The following lemma with proof is taken from [13] (with "3" there replaced by "h" here). The main raison d'être of this first lemma is to introduce its proof technique, which we will adapt in two lemmas to obtain the improvements relevant for our lower bounds on the number of triangulations.

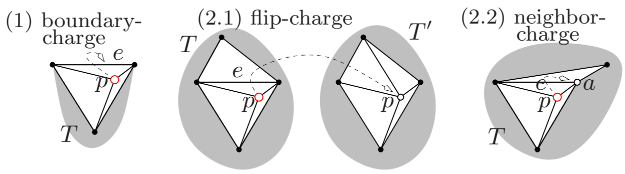

Figure 3: The various types of charges of a 3-vint in the proof of Lemma 2.1.

LEMMA 2.1. *$\hat{v}_3 \le \frac{2n+h}{5}$ holds for every planar point set S in general position with $|S^o| = n$ and $|S| = n + h$.*

Proof. We apply a scheme where every 3-vint charges 3 units to vints of larger degrees or to bints. No vint will be charged more than 2 units, no bint more than 1. Hence

$$3\hat{v}_3 \le h + 2\textstyle\sum_{j\ge 4}\hat{v}_j = h + 2(n - \hat{v}_3), \qquad (1)$$

which yields the asserted inequality.

Let $v = (p, T)$ be a 3-vint, and let o_v denote its link, which is a triangle. For each edge e of o_v we do the following, depending on the nature of e; see Figure 3.

(1) *e is an edge of the hull.* Then we let v charge 1 to bint (e, T); we call this a *boundary-charge*.

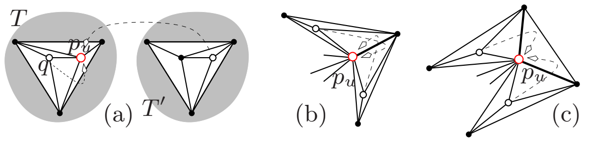

Figure 5: (a) 4-vint (p_u, T) with non-convex link is charged once by a flip-charge (right) and once by a neighbor-charge (left). (b) Two neighbor-charges to vint u with one separable edge. (c) Neighbor-charges to vint u with two separable edges.

(2) *There is a triangle t incident to e on its other side:*

(2.1) *t forms with p a convex quadrilateral.* We flip e to get a 4-vint (in a different triangulation!) to which v charges 1; we call this a *flip-charge.*

(2.2) *t forms with p a non-convex quadrilateral.* Let a be the endpoint of e that is reflex; a cannot be a hull vertex and it has to be of degree at least 4, since interior vertices of degree 3 are never adjacent. Here v charges 1 to vint (a, T); we call this a *neighbor-charge.* (e must be separable at (a, T).)

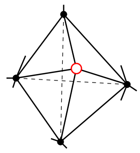

Figure 4: A 4-vint with two flippable incident edges.

The fact that no bint is charged more than once is obvious, so we turn to show that no vint u can be charged more than twice. Consider first the case of a 4-vint $u = (p_u, T_u)$. Let o_u denote the quadrangular link of p_u. We note that at most two edges incident to p_u are flippable: One out of each pair of opposite edges is separable at u (regardless of whether the link of v is convex or not), and thus unflippable; see Figure 4. We distinguish between the following two cases:

(a) *o_u is a convex quadrilateral*, as depicted in Figure 4. In this case, u receives exactly two flip-charges. Moreover, u cannot be charged as a neighbor, since no vertex of o_u can be interior and of degree 3.

(b) *o_u has a reflex vertex*, as depicted in Figure 5(a). Here u receives exactly one flip-charge. This obvious fact is a special case of a more general analysis, given in Lemma 3.2. Moreover, u can be charged at most once as a neighbor. Indeed, if q is a vertex of o_u of degree 3, then it must be a reflex vertex of o_u and there can be at most one such vertex. (q, T_u) cannot charge u twice through two edges of the link of (q, T_u), for then these two edges have to be separable at p, but they are not consecutive around p; cf. (S2).

Consider next the case where $u = (p_u, T_u)$ is a vint of degree at least 5. Each flip-charge is to a 4-vint and therefore u receives neighbor-charges only. Neighbor charges are made within the *same* triangulation and claim that in this case p_u can be a neighbor of at most two points of degree 3 that charge it as a neighbor (as just noted, no point can charge u twice in this manner). Recall the ingredients necessary for such a neighbor-charge to be made to u: (i) an edge e that is separable at u, and (ii) a neighbor a of p_u that has degree 3 so that the edges e and $p_u a$ are consecutive around p_u. Clearly, if there is only one edge separable at u then there are at most two such constellations; see Figure 5(b). If there are two separable edges at u, then they have to be consecutive around p_u, cf. (S2). This rules out the possibility that any of these two edges is involved in more than one neighbor-charge, since an edge cannot be both separable at p_u and connect to an interior point of degree 3 (Figure 5(c)). $\square$

With slight changes to the charging scheme, the two following lemmas improve the result of Lemma 2.1.

LEMMA 2.2. *$\hat{v}_3 \leq \frac{n}{2}$ holds for every point set S in general position with $|S^o| = n$ and $|S| \geq 7$.*

Proof. A refined version of Inequality (1) in the proof of Lemma 2.1 is $3\hat{v}_3 \leq h_c + 2(n - \hat{v}_3)$, where h_c is the expected number of charged bints in a uniformly chosen triangulation. We now change the charging scheme so that we can show $h_c \leq \hat{v}_3$; combining this bound with the inequality implies the assertion of the lemma.

We change the charging scheme as follows. For each 3-vint that charges two bints, we move the charge from one of these bints to a vint of degree at least 5, such that the overall number of charges made to a vint is at most two.

First, consider a point p such that any triangle (spanned by S) that contains only p in its interior, is incident to at most one boundary edge. Such a case is depicted in Figure 6(a), where p has to be above the dotted edges. Equivalently, the condition says that, for each triangle t, two of whose edges are consecutive edges of the hull, either p lies outside this triangle t or it lies in t with at least one additional interior point. In this case, each 3-vint with p as a vertex charges at most one bint; thus we leave the charges made by p as in the preceding scheme.

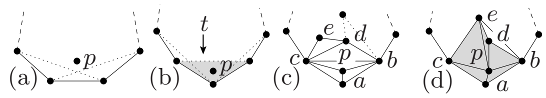

Figure 6: (a) No triangle incident to two boundary edges contains p in its interior. (b) Only the shaded triangle t is incident to two boundary edges and contains p in its interior. (c) bc and cd are flippable. (d) The link of the 5-vint of p is not convex.

Next, assume that there is a single triangle t that is incident to two boundary edges and contains only p in its interior, as depicted in Figure 6(b). We notice that the 3-vints of p which charge two bints are exactly those for which the link of p is t. Consider such a 3-vint $v = (p, T)$, and denote the vertices of the link of v as a, b, and c, such that ab and ac are boundary edges (as depicted in Figure 6(c)). It is easily checked that the edge bc must be flippable; let us denote the third vertex of the other triangle incident to bc as d. Moreover, we notice that, after flipping bc, at least one of the edges bd and cd is flippable. Without loss of generality, we assume that cd is flippable (after flipping bc), and denote the third vertex of the other triangle incident to cd as e (as depicted in Figure 6(c)).

Consider the pentagon $abdec$, which is the link of a 5-vint of p (in a different triangulation). First, assume that this is a convex pentagon, as depicted in Figure 6(c). In triangulations where the pentagon is the link of a 5-vint of p, it cannot contain any 3-vints on its boundary (recall that a vint can use only interior vertices, so we are claiming that neither d nor e can be a 3-vint in such a triangulation), and thus the 5-vint is not charged at all in the charging scheme of Lemma 2.1. We move the charge from one of the bints of v to this 5-vint. There are two 3-vints which can charge the 5-vint in this new manner—v and the symmetric 3-vint of p, with the same link and with the edge be replacing cd. Therefore, the 5-vint gets charged exactly twice overall.

Assume next that the pentagon is not convex, which implies that e cannot "see" b from within the pentagon, and perhaps also a (notice that e cannot "hide" from a behind c, since c is a hull vertex). In such a case, v is the only 3-vint that will charge the 5-vint in the new manner (since the quadrilateral $bced$—the complement of the link of v inside the pentagon—has a single triangulation). Moreover, there may be a 3-vint of d adjacent to the 5-vint, which may charge the 5-vint (only once) by a neighbor-charge, as depicted in Figure 6(d), where the link of the 5-vint is shaded. No other 3-vint can charge the 5-vint by a neighbor-charge, which implies that the 5-vint is charged, in the modified scheme, at most twice.

Finally, assume that there are two distinct triangles, each incident to two boundary edges and containing only p, as depicted in Figure 7(a) (notice that there cannot exist more than two such triangles). Let us denote the vertices of these triangles as a, b, c, and d, appearing in this order along the convex hull, so the two triangles are Δabc and Δbcd, and the non-boundary edges are ac and bd (as depicted in Figure 7(a)). Consider a 3-vint $v = (p, T)$ with Δabc as its link (the case where the link is Δbcd is handled symmetrically). v can be analyzed as before, flipping edges to turn it into a 5-vint and charging this 5-vint, except for the case where the resulting pentagon contains the quadrilateral $abcd$ (as depicted in Figure 7(b)). In such a case, the pentagon must be convex, and four 3-vints of p that charge two bints have their links contained in the pentagon—two with Δabc as the link of p and two with Δbcd as the link, and they all charge the 5-vint.

We thus need to find additional "victims" to distribute charges to. To this end, we denote the fifth vertex of this convex pentagon as e, as depicted in Figure 7(b), and note that, after flipping ac and ad, at least one of the edges ae and de is flippable (since $|S| \geq 7$). Without loss of generality, we assume that ae is flippable (after flipping ac and ad), and denote the third vertex of the other triangle incident to ae as f.

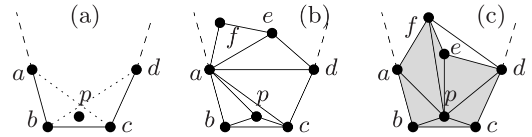

Figure 7: (a) Two triangles that are incident to two boundary edges and contain p only in their interiors. (b) ad and ae are flippable; the resulting hexagon is convex. (c) 6-vint with non-convex link.

Consider the hexagon $abcdef$, which is the link of a 6-vint of p (in an appropriate triangulation). First, assume that this is a convex hexagon, as depicted in Figure 7(b). Such a hexagon contains the links of ten 3-vints of p that charge two bints (five with Δabc as their link and five with Δbcd as their link; the number 5 is C_3, the number of triangulations of the pentagon completing the link of the 3-vint into the hexagon). Moreover, in triangulations where such a hexagon is the link of a 6-vint of p, it cannot contain any 3-vints on its boundary, and thus the 6-vint is not charged at all in the charging scheme of Lemma 2.1. This also applies to the four convex 5-vints obtained by removing either a, d, e, or f from the hexagon. (Note that the link of each of the "truncated" 5-vints contains at least one of the triangles Δabc, Δbcd, and that p lies in both of these triangles, so the link of each of these 5-vints does indeed contain p in its interior, as it should.) Therefore, we have five vints which are not charged at all. We can therefore shift ten boundary charges from the above ten 3-vints to those five vints, so each of those latter vints is charged exactly twice.

Assume next that the hexagon is not convex, which implies that f cannot "see" d within the hexagon and perhaps also c, as depicted in Figure 7(c), where the link of the 6-vint is shaded. In such a case, the link of the hexagon contains the links of at most six 3-vints of p that charge two bints (three with Δabc as their link and three with Δbcd as their link; here 3 is the maximum number of triangulations of the non-convex pentagon completing the link of the 3-vint into the hexagon). There may be a 3-vint of e adjacent to the 6-vint, as depicted in Figure 7(c). However, since both fp and dp are flippable, the 6-vint is not charged at all by the old scheme. A similar analysis also applies to the 5-vints obtained by removing either a or d from the 6-vint. (Again, the links of these 5-vints do indeed contain p.) The 5-vint obtained by removing f also does not get charged, since its link is a convex pentagon. Therefore, we can shift one boundary charge from each of the (at most) six 3-vints to one of the four larger vints, with enough room to conclude that each of these latter vints gets charged at most twice.

We have thus established the bound $h_c \leq \hat{v}_3$, and this completes the proof of the lemma. ◻

LEMMA 2.3. *$\hat{v}_3 \leq \frac{2n+h/2}{5}$ holds for every point set S in general position with $|S^o| = n$ and $|S| = n + h \geq 6$.*

Proof. We use the same inequality $3\hat{v}_3 \leq h_c + 2(n - \hat{v}_3)$ as in the preceding proof, h_c is the expected number of charged bints in a uniformly chosen triangulation. Again we modify the charging scheme, now in order to ensure $h_c \leq h/2$; combining this bound with the above inequality implies the assertion of the lemma. This bound is achieved by decomposing the collection of all bints into sets of at most nine bints each, such that the total charge in each set is at most half its size. The sets are not necessarily disjoint, but only charged bints can appear in more than one set, which is only to our advantage in establishing the bound $h_c \leq h/2$.

A bint β gets charged if and only if the third vertex of the triangle incident to β is a 3-vint. Consider such a bint $\beta = (b, T)$ that is charged by a 3-vint $v = (p, T)$, as depicted in Figure 8(a) (where b is the edge ac). Following the notations in the figure, we notice that at least one of the edges ad and cd must be flippable, and assume, without loss of generality, that cd is flippable (as depicted in the figure). Let T' denote the triangulation obtained by flipping cd in T. Notice that the bint $\beta' = (b, T')$ is not charged, since it is adjacent to the 4-vint $v' = (p, T')$ (depicted in Figure 8(b)). If v' cannot flip-down to another 3-vint with a link that contains b, we create the set $\{\beta, \beta'\}$, which has two elements, only one of which is charged.

Suppose that v' can flip down to another 3-vint with a link containing b, as in Figure 8(c) (note here that the link of v', the quadrilateral $aced$, has to be convex and p has to lie below its diagonals). Since the link of v' is convex and contains a boundary edge, one of its edges must be flippable. Let us denote one of the 5-vints obtained by flipping such an edge as $v'' = (p, T'')$ and let f be the vertex that is in the link

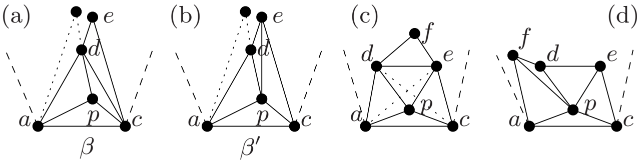

Figure 8: (a) The link of the 3-vint has at least one flippable edge. (b) A 4-vint that flips down to v. (c) A link of a 4-vint that flips down to two 3-vints that charge β. (d) A 5-vint with a non-convex link.

of v'' but not in the link of v'. First, assume that the link of v'' is a convex pentagon, as depicted in Figure 8(c), where we flip de in T' to get v''. The link of v'' contains the links of five 3-vints (there are five ways to triangulate a convex pentagon and to insert p as a 3-vint into it) and three 4-vints of p (by removing either d, e, or f from the pentagon—note that any of these removals keeps p inside the resulting quadrilateral) that are adjacent to the edge b. Since the link of v'' is convex, it is not charged at all in the original scheme, and thus we can charge it twice. Notice that ad and bc might be boundary edges, so they could induce bints that also get charged by some 3-vints of p. Nevertheless, the assumption on the location of p (depicted in Figure 8(c)) implies that, even if there is a 3-vint of p which charges bints of ad or ce, the respective 4-vints cannot flip down to another 3-vint of p that charges the same edge, which implies that such bints belong to the previous case. Therefore, there is a unique case where v'' gets charged. In conclusion, we have a set of nine bints, all involving edge b (five adjacent to 3-vints, three to 4-vints, and one to a 5-vint), with a total modified charge of $5 - 2 = 3 < 9/2$ (only 3-vints charge b in the original scheme).

Finally, assume that the link of v'' is not convex. This implies that, in the link of v'', f cannot see either one or two of the other vertices. Figure 8(d) depicts such a case, where f is adjacent to ad and only e is invisible from f (the following analysis does not refer specifically to this example, though). Let x denote the number of 3-vints of p that are adjacent to b and have their link contained in the link of v''. Notice that $x \le 3$, since there are at most three ways to triangulate the non-convex 5-gon and then insert p as the 3-vint. We can always obtain a 4-vint of p that is adjacent to b by removing f from the 5-vint, and an additional 4-vint by removing either d or e (or both, depending on the position of f). Therefore, we have a set of at least $x + 3$ bints, all involving the edge b (x bints adjacent to 3-vints, one adjacent to a 5-vint, and at least two adjacent to 4-vints), with $x \le \frac{x+3}{2}$ of them charged.

We have shown $h_c \le h/2$ as claimed ◻

Let $\mathsf{tr}^-(n,h)$ denote the minimal number of triangulations for point sets with n interior points and h boundary points and set $\mathsf{tr}^-(N) := \min_{n+h=N} \mathsf{tr}^-(n,h)$. We now employ the upper bounds for $\hat{v}_3$ for a lower bound for $\mathsf{tr}^-(N)$. The following is a generalization of [13, Lemma 2.1(ii)].

LEMMA 2.4. *For $n \ge 1$, let $\delta_{n,h} > 0$ be a real number, such that $\hat{v}_3 \le \delta_{n,h} n$ holds for any set of n interior points and h boundary points in general position. Then,*

$$\mathsf{tr}^-(n,h) \ge \tfrac{1}{\delta_{n,h}} \mathsf{tr}^-(n-1,h) \;.$$

Proof. Let S be a set that minimizes $\mathsf{tr}(S)$ among all sets with n interior points and h boundary points in general position. As easily seen and argued for [13, Lemma 2.1(ii)], we have

$$\hat{v}_3 \cdot \mathsf{tr}(S) = \textstyle\sum_{T \in \mathcal{T}(S)} v_3(T) = \sum_{q \in S^o} \mathcal{T}(S \setminus \{q\}).$$

The leftmost expression equals $\hat{v}_3 \cdot \mathsf{tr}^-(n,h)$, the rightmost one is at least $n \cdot \mathsf{tr}^-(n-1,h)$. Hence, with $\hat{v}_3 \le \delta_{n,h}\, n$, $\mathsf{tr}^-(n,h) \ge \frac{n}{\hat{v}_3} \cdot \mathsf{tr}^-((n-1,h) \ge \frac{1}{\delta_{n,h}} \cdot \mathsf{tr}^-(n-1,h)$. ◻

THEOREM 2.5. $\mathsf{tr}^-(N) = \Omega\left(2.4317^N\right)$.

Proof. We know $\mathsf{tr}^-(0,h) = C_{h-2} = \Theta^*(4^h)$. Moreover, combining the results of Lemmas 2.2, 2.3, and 2.4 yields

$$\mathsf{tr}^-(n,h) \ge 2 \cdot \mathsf{tr}^-(n-1,h) \quad \text{and} \tag{2}$$

$$\mathsf{tr}^-(n,h) \ge \tfrac{5n}{2n+h/2} \cdot \mathsf{tr}^-(n-1,h) = \tfrac{10n}{3n+N} \cdot \mathsf{tr}^-(n-1,h) \tag{3}$$

for $n \ge 1$ and $N \ge 7$. (3) is stronger iff $n > N/2$, and therefore, for any point set S with $n \le N/2$ and $N \ge 7$,

$$\mathsf{tr}(S) \ge 2^n \cdot \mathsf{tr}^-(0, N-n) = \Omega^*\left(2^{2N-n}\right) = \Omega^*\left(2^{3N/2}\right) \tag{4}$$

with $2^{3/2} = 2.828\ldots$. Next, assume that $n > N/2$, so (3) is the stronger inequality. Let x be the maximal number of interior points we can remove before (2) becomes the stronger; x is the maximal with $\frac{10(n-x)}{(N-x)+3(n-x)} \ge 2$, that is, $x = 2n - N$. Using the above, we derive the bound (assuming that $N+3n$ is divisible by 4, which, if true initially, remains true as we remove interior points)

$$\begin{aligned}\mathsf{tr}^-(N) &\ge \tfrac{10n}{N+3n} \cdot \tfrac{10(n-1)}{N+3n-4} \cdots \tfrac{10(n-x+1)}{N+3n-4(x+1)} \cdot 2^{n-x} \cdot \Theta^*(4^{N-n}) \\ &= \tfrac{10^x \cdot n!}{(n-x)!} \cdot \tfrac{((N+3n)/4-x)!}{4^x \cdot ((N+3n)/4)!} \cdot 2^{n-x} \cdot \Theta^*(4^{N-n})\,.\end{aligned}$$

In order to simplify this bound, we use Stirling's approximation. Since we are only interested in the exponential part of the bound, we can simply replace $m!$ by $(m/e)^m$. Therefore, $\mathsf{tr}^-(N)$ is lower bounded by

$$\Omega^*\Big(2^{2N-n-2x} \cdot 5^x \cdot \tfrac{n^n}{e^n} \cdot \tfrac{e^{n-x}}{(n-x)^{n-x}} \cdot \tfrac{((N+3n)/4-x)^{(N+3n)/4-x}}{e^{(N+3n)/4-x}} \cdot \tfrac{e^{(N+3n)/4}}{((N+3n)/4)^{(N+3n)/4}}\Big),$$

and after some cleanup, we get a lower bound of

$$\Omega^*\left(2^{2N-n} \cdot 5^x \cdot \frac{n^n}{(n-x)^{n-x}} \cdot \frac{(N+3n-4x)^{(N+3n)/4-x}}{(N+3n)^{(N+3n)/4}}\right).$$

After replacing x with $2n - N$, substituting $n = tN$, $0.5 < t < 1$, and performing some additional cleanup, we get

$$\mathsf{tr}^-(N) = \Omega^*\left(2^{2N-n} \cdot 5^{(3n+N)/4} \cdot n^n \cdot \frac{(N-n)^{(N-n)/4}}{(N+3n)^{(N+3n)/4}}\right)$$

$$= \Omega^*\left(\left(2^{2-t} \cdot 5^{(3t+1)/4} \cdot t^t \cdot \frac{(1-t)^{(1-t)/4}}{(1+3t)^{(1+3t)/4}}\right)^N\right). \tag{5}$$

Finding the t minimizing this expression (for given N) can be done either numerically or through differentiation. The latter approach produces the quartic equation $t^4 - 288t^2 - 128t - 16 = 0$, whose solution is $t \approx 0.89901$, which implies a minimum of $\Omega\left(2.4317^N\right)$. We have thus shown that for every $n + h = N$, we have $\mathsf{tr}^-(N) = \Omega\left(2.4317^N\right)$. ◻

Remark. The analysis provides lower bounds on $\mathsf{tr}^-(n,h)$, for any n and h. Collecting the bounds in (4) and (5), with $t = N/n$, $0 \le t < 1$, we have lower bounds on $\mathsf{tr}^-(n,h)$ of

$$\Omega^*\left(\left(2^{2-t}\right)^N\right) \qquad 0 \le t \le 0.5$$
$$\Omega^*\left(\left(2^{2-t}\cdot 5^{(3t+1)/4}\cdot t^t \cdot \frac{(1-t)^{(1-t)/4}}{(1+3t)^{(1+3t)/4}}\right)^N\right) \qquad 0.5 < t < 1.$$

The base in the bound starts at 4 for $t = 0$ and ends at the limit 2.5 for $n = N - 3$ (where t is almost 1). This latter value is still not as large as the base of 2.63 in [7].

3. RELATING THE $\hat{v}_i$'S

We derive inequalities among the $\hat{v}_i$'s which we then manipulate for a lower bound on $\hat{v}_4$. These facts are required for the proof of Theorem 4.1 in Section 4, which yields linear lower bounds for all $\hat{v}_i$'s, $i \ge 4$. We first recall the notion of a flip-tree studied in [12, 13] (and implicitly in [10]).

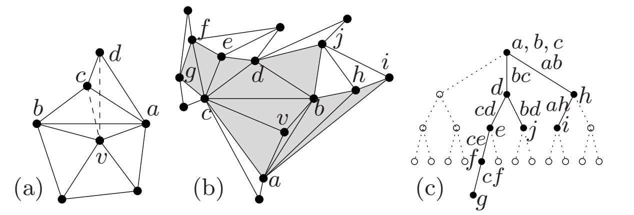

Figure 9: (a) A 5- and a 6-vint that can be flipped down to the same 4-vint. (b) The shaded area is dual to the flip tree of the 3-vint v. (c) The flip-tree itself (solid edges only).

The flip-tree of a vint. How do we find the vints that flip down to a given i-vint $v = (p_v, T_v)$? Clearly, there is v itself. Consider a flippable edge e (in T_v) that is not incident to p_v but is part of the boundary of its link. Flipping e yields an $(i+1)$-vint $u = (p_v, T_u)$ that can be flipped down to v (by reversing the preceding flip). Similarly, if in the triangulation T_u there is a flippable edge that is not incident to p_v but is part of its link, then we can flip this edge to get an $(i+2)$-vint that can be flipped down to v, and so on. Figure 9(a) depicts a 4-vint v, that, by flipping ab to cv, turns into a 5-vint that can be flipped down to v (and which, by flipping ac to dv, turns into a 6-vint that can also be flipped down to v).

In order to represent this structure, we associate with an i-vint $v = (p_v, T_v)$ a *flip-tree* $\tau(v)$, defined as follows. The root of the tree is labeled by the pair (o_v, N_v), where o_v is the link of v (an i-gon) and N_v is the set of vertices of o_v (the neighbors of p_v in T_v). Any other node of the tree is associated with a pair (t, q), where t is a face of T_v and q is a vertex of that face, which does not belong to the union of the faces labeling the ancestors of the present node (note, though, that o_v from the root is not a face of T_v—it is the union of the i faces incident to p_v). The associated faces represent a duality between the flip-tree and part of the triangulation T_v (which will be explained momentarily). While defining the structure of the flip-tree in the following paragraphs, we refer to an example depicted in Figures 9(b) and 9(c). These figures depict a 3-vint v and its flip-tree, and the nodes of this flip-tree are labeled only by their vertex (and not by their triangle).

(i) Every edge e of o_v gives rise to a child if it can be flipped in T_v. If so, this child is labeled by the triangle incident to e that is not incident to p_v, and by the vertex of this triangle which is not incident to e. Therefore, the root has at most i children. In our example, the root has two children—$(\Delta bcd, d)$ (since bc is flippable) and $(\Delta abh, h)$ (since ab is flippable). In what follows, as in Figure 9, we will often suppress the triangle t in the label (t, q) of a node of the flip-tree, and just use the vertex q. The triangle t is the unique triangle of T_v into which the segment qp_v enters as we trace it from q.

(ii) Consider now a non-root node of the tree labeled by (t, q) and an edge e of t incident to q. If e is a boundary edge, no child will be obtained via e. Otherwise, let t' be the other triangle incident to e. If t' together with the triangle formed by e and p_v is a convex quadrilateral (in which e can be flipped), then this gives rise to a child of (t, q) labeled by (t', q'), where q' is the vertex of t' that is not incident to e. Therefore, a non-root node has at most two children. In our example, the node corresponding to h has the single child i, since the quadrilateral $vhia$ is convex, but the other potential quadrilateral $vbjh$ is not.

Note that the union of all triangles of the nodes of any subtree of $\tau(v)$ (containing the root) form a polygon which is star-shaped with respect to p_v; this follows easily by the inductive definition of $\tau(v)$. The triangles of the original T_v form a triangulation of the polygon, and the subtree is actually the dual tree of this triangulation. The shaded area in Figure 9(b) is the portion of the triangulation dual to the entire flip-tree of v. Also, an edge in the flip-tree incident to two nodes that are dual to (i.e., labeled by) the triangles Δ_1, Δ_2 in T_v, can be regarded as dual to the edge in T_v incident to both Δ_1 and Δ_2. If we retriangulate this polygon in T_v by connecting p_v to all vertices of the polygon, we get a vint that flips down to v. Moreover, every vint u that flips down to v can be obtained in this way (by taking the subtree dual to the link of u). That is:

LEMMA 3.1. *The subtrees of $\tau(v)$ containing its root are in bijective correspondence to the vints that flip down to v.*

We recall a basic fact about flippable edges.

LEMMA 3.2. *Each i-vint, $i \ge 4$, is incident to a flippable edge.*

Proof. The link of a vint (p, T) has at least three vertices with a convex angle (less than π) and at most two edges are separable at p. Hence, there is an edge incident to p that is separable at none of its endpoints – thus flippable. ◻

The next lemma is from [13, Lemma 4.1] with its proof based on [10, Lemma 4]. (Note that h does not play a role.)

LEMMA 3.3. *For all integers $3 \le i < j$ there is a positive integer $\delta_{i,j}$ such that $\hat{v}_i \ge \frac{\hat{v}_j}{\delta_{i,j}}$. In particular, $\hat{v}_i \ge \frac{\hat{v}_{i+1}}{i}$, $\hat{v}_3 \ge \frac{\hat{v}_j}{C_{j-1}-C_{j-2}}$ for $j \ge 4$, and $\hat{v}_4 \ge \frac{\hat{v}_j}{C_{j-1}-2C_{j-2}}$ for $j \ge 5$.*

Proof. For a proof of $\hat{v}_i \ge \frac{\hat{v}_{i+1}}{i}$, we let every $(i+1)$-vint charge some i-vint it can be flipped down to, by Lemma 3.2 this is possible, since $i + 1 \ge 4$. In this way an i-vint can be charged at most i times, so the first inequality holds.

For the general inequality we let every j-vint charge some i-vint it can be flipped down to. By Lemma 3.1, every j-vint that flips down to an i-vint v corresponds to a subtree of the flip-tree of v. More precisely, since the root of the flip-tree of

v corresponds to an i-gon in the triangulation of v, every j-vint corresponds to a subtree with $j-i+1$ nodes. Therefore, an i-vint can be charged at most $t_{i,j-i+1}$ times, where $t_{i,k}$ denotes the number of (ordered) binary trees with k nodes and with an exceptional root of degree i; that is, the root has i potential children pointers, but not all of them need to be used (just like binary nodes distinguish between a left and a right child, the root discriminates its children via an index in $\{1, 2, \ldots, i\}$). For example, $t_{i,1} = 1$ and $t_{i,2} = i$. Hence, as in the case of $j = i + 1$, we can take $\delta_{i,j} = t_{i,j-i+1}$. The number of ordered binary trees is known to be $t_{2,k} = C_k$ (see Subsection 1.1), which also implies that $t_{1,k} = C_{k-1}$. Furthermore, a recurrence of $t_{i,k} = t_{i-1,k+1} - t_{i-2,k+1}$ can be derived, cf. [10]. This allows us to choose

$$\begin{aligned}\delta_{3,j} &\leq t_{3,j-2} = t_{2,j-1} - t_{1,j-1} = C_{j-1} - C_{j-2}, \text{ and}\\ \delta_{4,j} &\leq t_{4,j-3} = t_{3,j-2} - t_{2,j-2} = C_{j-1} - 2C_{j-2}.\end{aligned}$$ ◻

We will now improve the bound for $\delta_{4,j}$, which we will use to derive a reasonably large lower bound on $\hat{v}_4$.

LEMMA 3.4. *$\hat{v}_4 \geq \frac{3\hat{v}_j}{C_{j-1}-C_{j-2}}$ holds for all integers $j \geq 5$.*

Proof. In the previous proof we made a j-vint charge a single 4-vint, a scheme now modified. For u a vint, let $supp_4(u) := \left|\{v \mid v \text{ 4-vint with } u \to^* v\}\right|$, called *4-support of u*. We let every j-vint split a charge of 1 evenly among the 4-vints it can be flipped down to, i.e., $\frac{1}{supp_4(u)}$ each.

Given a 4-vint v, let w be a 3-vint such that $v \to w$, and let τ be the flip-tree of w. The subtree of τ which corresponds to v consists of a single level-1 edge e (i.e., an edge emanating from the root). Therefore, there is a bijective correspondence between the vints that flip down to v and subtrees of τ that contain e. Counting the number of such subtrees of size $j-2$ (i.e., the number of j-vints that flip down to v) implies our previous result of $t_{3,j-2} - t_{2,j-2} = C_{j-1} - 2C_{j-2}$.

τ might have two additional level-1 edges e_2 and e_3. Let v_2 (v_3) denote the 4-vint corresponding to the subtree of τ which consists of edge e_2 (e_3, resp.) only. A subtree containing e_2 can flip down to v_2, a subtree containing e_3 can flip down to v_3. Hence, a vint with a subtree that contains two level-1 edges has a 4-support of at least 2, and a vint with a subtree that contains all three level-1 edges has a support of at least 3. For example, out of the four possible 5-vints that can flip down to a certain 4-vint, if they all exist, two contain two level-1 edges, and thus have a support of at least 2. Figure 10(a) depicts a subtree of a 5-vint that might have a support of 1, and Figure 10(b) depicts a subtree of a 5-vint that has a support of at least 2.

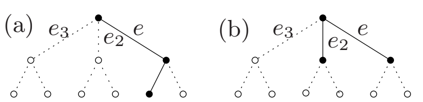

Figure 10: (a) One of two subtrees that correspond to a 5-vint with a support of at least 1. (b) One of two subtrees of a 5-vint with a support of at least 2.

Therefore, the number of j-vints that flip down to v and have a support of 1, is at most $t_{1,j-2}$. The number of j-vints that flip down to v with a subtree having exactly two level-1 edges, is at most $2\left(t_{2,j-2} - 2t_{1,j-2}\right)$ (choose one of e_2, e_3 and count subtrees containing both e and the chosen edge but not the third). The number of j-vints that flip down to v with their corresponding subtrees having three level-1 edges, is at most $t_{3,j-2} - 3t_{2,j-2} + 3t_{1,j-2}$ (inclusion-exclusion principle). Therefore, v receives a charge of at most

$$t_{1,j-2} + \tfrac{1}{2}2(t_{2,j-2} - 2t_{1,j-2}) + \tfrac{1}{3}(t_{3,j-2} - 3t_{2,j-2} + 3t_{1,j-2})$$

which equals $\frac{1}{3}t_{3,j-2} = \frac{1}{3}\left(C_{j-1} - C_{j-2}\right)$. ◻

LEMMA 3.5. *If $|S^o| \geq 3$, we have $\sum_i i\hat{v}_i \leq 6n + h - 9$.*

Proof. The number of edges in a triangulation of S is $3n + 2h - 3$, therefore $\sum_i iv_i = 2(3n + 2h - 3) - D$, where D is the sum of the degrees of the hull vertices. One can show $D \geq 3h + 3$ – details omitted –, provided there are at least 3 interior points. This yields the bound for the v_i's which carries over to the $\hat{v}_i$'s via linearity of expectation. ◻

LEMMA 3.6. *For $n \geq 3$, we have $\hat{v}_4 \geq \frac{1}{340}\left(n + 15 - \frac{8h}{3}\right)$. In particular, when S has a triangular convex hull, $\hat{v}_4 > \frac{n}{340}$.*

Proof. $\sum_i (9-i)\hat{v}_i \geq 9n - (6n + h - 9) = 3n - h + 9$ (by Lemma 3.5), $\hat{v}_3 \leq \frac{2n+h/2}{5}$ (Lemma 2.3), and by Lemma 3.4

$$\begin{array}{ll}\hat{v}_5 \leq \frac{C_4-C_3}{3}\hat{v}_4 = 3\hat{v}_4 & \hat{v}_7 \leq \frac{C_6-C_5}{3}\hat{v}_4 = 30\hat{v}_4\\ \hat{v}_6 \leq \frac{C_5-C_4}{3}\hat{v}_4 = \frac{28}{3}\hat{v}_4 & \hat{v}_8 \leq \frac{C_7-C_6}{3}\hat{v}_4 = 99\hat{v}_4.\end{array} \tag{6}$$

Hence,

$$\begin{aligned}3n + 9 - h &\leq 6\hat{v}_3 + 5\hat{v}_4 + 4\hat{v}_5 + 3\hat{v}_6 + 2\hat{v}_7 + \hat{v}_8\\ &\leq \tfrac{6(2n+h/2)}{5} + \hat{v}_4\left(5 + 4\cdot 3 + 3\cdot\tfrac{28}{3} + 2\cdot 30 + 99\right)\\ &= \tfrac{12n+3h}{5} + 204\hat{v}_4\ ,\end{aligned}$$

implying that $\hat{v}_4 \geq \frac{1}{340}\left(n + 15 - \frac{8h}{3}\right)$, as asserted. ◻

Here is an alternative lower bound whose dependence on h is better, while the dependence on n is worse.

LEMMA 3.7. *For $n \geq 4$, we have $\hat{v}_4 \geq \frac{1}{1360}(n + 18 - 2h)$.*

Proof. Lemma 3.5, again, delivers $\sum_i (10-i)\hat{v}_i \geq 4n - h + 9$. From Lemma 2.2, we employ $\hat{v}_3 \leq \frac{n}{2}$ and from Lemma 3.4, we use (6) and $\hat{v}_9 \leq \frac{C_8-C_7}{3}\hat{v}_4 = \frac{1001}{3}\hat{v}_4$. Hence, we get

$$\begin{aligned}4n - h + 9 &\leq 7\hat{v}_3 + 6\hat{v}_4 + 5\hat{v}_5 + 4\hat{v}_6 + 3\hat{v}_7 + 2\hat{v}_8 + \hat{v}_9\\ &\leq \tfrac{7n}{2} + \hat{v}_4\cdot\left(6 + 5\cdot 3 + 4\cdot\tfrac{28}{3} + 3\cdot 30 + 2\cdot 99 + \tfrac{1001}{3}\right)\\ &= \tfrac{7n}{2} + 680\hat{v}_4\ .\end{aligned}$$ ◻

4. LOWER BOUNDS FOR ALL $\hat{v}_i$'S

In this section, we establish lower bounds for each of the quantities $\hat{v}_i$. The bound $\hat{v}_3 > \frac{n}{30}$ was proved in [12] (for a triangular convex hull), and in Section 3 we derived the bound $\hat{v}_4 \geq \frac{1}{1360}(n + 18 - 2h)$. We now present a generalized bound which holds for each $i \geq 4$.

THEOREM 4.1. *For each $N > i \geq 4$, $\varepsilon > 0$, and $3 \leq h \leq \left(\frac{1}{2} - \varepsilon\right)n$, there exists a constant $\gamma_{i,h} = \gamma_{i,h}(\varepsilon)$, which depends on i, h, and ε, such that, for any set S with h hull vertices and n interior vertices, we have $\hat{v}_i \geq \gamma_{i,h}n$.*

Proof. The proof is by induction on i, where the base case $i = 4$ has already been established. More precisely, Lemma 3.7 implies that, when $h \leq \left(\frac{1}{2} - \varepsilon\right)n$, there is a constant $c_\varepsilon > 0$ such that $\hat{v}_4 \geq c_\varepsilon n$, and we use this inequality as our induction basis, with $\gamma_{4,h} = c_\varepsilon$.

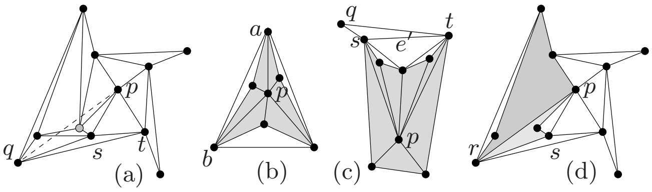

Figure 13: (a) The grey vertex is a vertex of the link that is contained in Δpsq. (b) No edge of the link is incident to a triangle with a third vertex not contained in o. (c) There is a path from the link of p to the triangle containing q. (d) The link is split into two corridors.

For the general induction step, we assume $\hat{v}_{i-1} \geq \gamma_{i-1,h} n$ and proceed to establish a similar inequality on $\hat{v}_i$. This will be done by charging each $(i-1)$-vint v to various vints of degree at least i. This will yield an inequality involving the quantities $\hat{v}_k$, for $k \geq i-1$, which we will then combine with the inequalities of Lemma 3.3, to replace all $\hat{v}_k$, for $k > i$, by $\hat{v}_i$, and thereby obtain the desired lower bound on $\hat{v}_i$.

Let $v = (p, T)$ be a fixed $(i-1)$-vint, and let o denote its link. The charging that v makes depends on the structure of o and of the triangles of T in its vicinity. The charging is performed depending on the following cases:

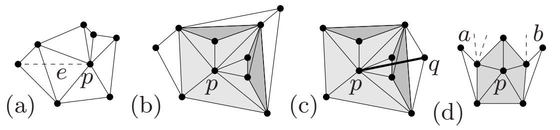

Figure 11: (a) A link with a flippable edge. (b) The lightly shaded portion is the link, and the union of the shaded portions is the extended link. (c) The edge pq crosses only edges of the extended link. (d) Two vertices "hiding" from p.

(a) *o has a flippable bounding edge e,* as depicted in Figure 11(a). In this case we flip e to turn v into an i-vint, and charge that i-vint. Clearly, any i-vint can be charged in this way at most i times.

(b) *No edge of o is flippable, and all the vertices of o that are interior to the hull have degree at most $i-1$.* We argue that the following property holds in such a case.

LEMMA 4.2. *In case (b), we can connect p to some point $r \in S$ outside o, such that pr crosses at most $(i-2)(i-4)+1$ edges of T.*

Proof. Recall $N > i$, so there is at least one point not connected to p. First, assume that there exists an edge e of the link, such that e is incident to a triangle with a third vertex that is not a vertex of o (see below for an illustration and analysis of the complementary situation). Let s and t be the endpoints of e, and let q be the third vertex of the triangle. Without loss of generality, assume that the edge sq is shorter than the edge tq (i.e., that q is "hiding" from p behind s). Such a case is depicted in Figure 12 (for now, ignore the non-solid lines and the shading).

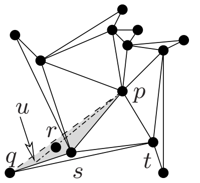

Figure 12: Case (b).

We consider the triangle Δpsq, and notice that it may contain additional vertices of o, as well as other points of S. Such a case is depicted in Figure 13(a), where the grey vertex is a vertex of o that is contained in Δpsq. Let $S_{\Delta psq}$ be the set of points of S contained in Δpsq, but not in o (including q, but not p). Let r be the vertex in $S_{\Delta psq}$ that minimizes the angle $\angle spr$ (r is well defined, because $S_{\Delta psq}$ is nonempty—it contains q). Let ℓ denote the line containing the edge pr, and let u denote the point where ℓ crosses sq (by definition, ℓ must indeed cross sq; the crossing could be at q if $r = q$). Our choice of r implies that the interior of the triangle Δspu does not contain any point of $S_{\Delta psq}$, so the only points of S it can contain are vertices of o. We again refer the reader to Figure 12, where Δpsq is shaded.

We now show that pr can cross at most $(i-2)(i-4)+1$ edges of T. For an edge of T to cross pr, exactly one of its vertices must lie in the triangle Δspu (this is only a necessary condition). However, we have just argued that this triangle can only contain vertices of o. Since t lies outside Δspu, there are at most $i-2$ vertices in this triangle (including s). Each of these vertices is of degree at most $i-1$, one of its incident edges is connected to p, and two other edges are part of the boundary of o. Therefore, excluding the single crossing between pr and o, each of the $i-2$ vertices can participate in at most $i-4$ edges that cross pr.

Next, assume that there is no edge e with an incident triangle on the other side which has a vertex not in o (such a case is depicted in Figure 13(b)). Every edge of o interior to the hull must be incident to a triangle that has an external chord as an edge.

We walk through T, starting at some edge e of o and walking away from o, crossing from each visited triangle to an adjacent one through a common external chord, until we get to a triangle incident to an external chord e' and to a vertex q which is not contained in o. It is easily seen that the rules for the walk are well-defined: we can either find an external chord to cross into the next triangle, or get stuck with a terminal triangle as above. Moreover, since $N > i$, the walk will always end in such a terminal triangle. Such a walk is depicted in Figure 13(c), where the link is shaded, and there is a path that leads to q. Now, denoting the endpoints of e' as s and t, we can apply the same analysis as above. That is, assume first that q hides behind s, as above, and denote by r the vertex in $S_{\Delta psq}$ that minimizes the angle $\angle spr$. In this case one can argue, as above, that pr can cross at most $(i-2)(i-4)+1$ edges. However, here (unlike in the preceding analysis) it is also possible for q not to hide at all, that is, pq can cross e'. In this case, we can choose q as the point r, and notice that pq can only cross external chords and a single edge of o, which implies that there are at most $(i-3)+1 = i-2$ crossings ($< (i-2)(i-4)+1$ for $i \geq 5$). ◻

We now explain how to deal with an $(i-1)$-vint v that falls under case (b) of the present analysis. We charge v to a vint $v' = (p, T')$ obtained as follows. Let r be the point provided in Lemma 4.2. Delete from T all the $\mu \leq i^2 - 6i + 9$ edges that cross pr, and add pr to the new graph. This leaves two links, referred to as "corridors", with pr as a common edge. We triangulate each of the corridors in an arbitrary manner, leaving the rest of T untouched, to obtain a triangulation T', and then charge v to $v' = (p, T')$. In Figure 13(d), the shaded areas are the two untriangulated corridors obtained from the vint depicted in Figure 13(a).

Note that the degree of v' is at least i (because we have added pr as an edge). It can be larger than i, if the triangulations of the corridors use additional edges incident to p.

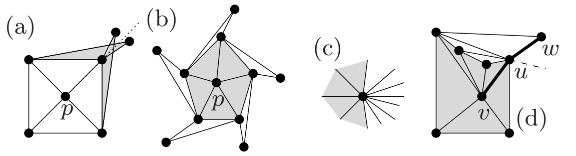

Figure 14: (a) Two vertices cannot hide behind the same vertex of o'. (b) All the vertices of S' hide in a clockwise manner. (c) The shaded area is a dominating sequence of size 4. (d) There is a dominating sequence of u from pu clockwise to uw.

Since we have removed μ edges from T, and then inserted the edge pr, the re-triangulation of the corridors requires $\mu - 1$ additional edges. The maximal degree of v' is obtained when all of these additional edges are incident to p. Hence, the degree of p is at most $i+(\mu-1) \leq i+(i^2-6i+9-1) = i^2-5i+8$.

We next show that the number of $(i-1)$-vints, that fall under case (b) and charge the fixed vint v', is at most some (large) quantity M_i that depends only on i. By the preceding discussion, we can assume that the degree d of v' is at most $i^2 - 5i + 8$. Given v', we can reconstruct v as follows. We first choose the vertex r from the d neighbors of p in T'. Next, we choose the two corridors bounded by pr. To do so, we recall that, together, these corridors consisted of at most $i^2 - 6i + 10$ triangles (in the original triangulation T), so T' uses the same number of triangles to fill them up. Thus, starting from the two triangles of T' that are adjacent to pr, we append to them up to $i^2 - 6i + 8$ additional triangles, in a breadth-first manner. Specifically, we maintain a queue of triangles to be appended to the corridors. When appending a triangle, we have already used one of its edges to reach it from its predecessor (the same holds for the two initial triangles, which cannot reach other triangles through their common edge pr). We therefore have up to two other edges that we can cross to reach other triangles. Hence, when appending a triangle, we have up to four choices: append its two neighboring triangles, append only the first of them, append only the second, or not append any of them. (Some of these neighbor triangles may have already been appended; this only limits our choices.) We continue this process until we collect the desired number of triangles. Hence, a crude upper bound for the number of such choices is $4^{i^2-6i+10}$.

Finally, having guessed these triangles, we remove from T' all the inner edges of their union (i.e., edges adjacent to two of these triangles), and re-triangulate the resulting link in an arbitrary manner (using only edges that cross pr). Notice that the number of ways to triangulate this polygon is maximized when it is convex (since every triangulation of a non-convex polygon is combinatorially equivalent to a triangulation of a convex polygon with the same number of edges). The polygon has at most $i^2 - 6i + 12$ vertices, and thus it has $M_i' = O^*\left(4^{i^2-6i+10}\right)$ triangulations.

In conclusion, the number of $(i-1)$-vints, that fall under case (b) and charge the fixed vint v', is at most

$$M_i = (i^2 - 5i + 8) \cdot 4^{i^2-6i+10} \cdot M_i'.$$

In the remaining cases, we assume that no edge of o is flippable and that o contains a vint u of degree at least i. Define the *extended link* o' of o by iteratively repeating the following process — find in T a triangle Δ incident to two (consecutive) edges of the current link o^* and lying in the exterior of o^*, and append Δ to o^*. When we can no longer find such a triangle, we have obtained the extended link $o' = o^*$; see Figure 11(b). Note that all the edges of the triangles Δ encountered in this process are either edges of o or extended chords of o.

(c) *There exists a vertex q which is not a vertex of o, and the line segment pq crosses only edges contained in o'.* Such a case is depicted in Figure 11(c). Notice that pq crosses a single edge of o, and possibly some external chords. Therefore, pq cannot cross more than $(i-4)+1 = i-3$ edges of T. Similarly to case (b), we can remove the edges that cross pq, insert pq, retriangulate the two resulting "corridors", and charge the resulting j-vint of p (where $j \geq i$). As in the previous case, such a j-vint cannot get charged more than L_i times in this manner, for some L_i depending only on i.

(d) *For every $q \in S$, either q is a vertex of o or the line segment pq crosses at least one edge not contained in o'.* By definition, every triangle of T not contained in o' is incident to at most one edge of o'. Let S' be the set of third vertices (not contained in o) of triangles incident to edges of o'.

For each $s \in S'$, its corresponding vint in T must "hide" from p, either in a *clockwise* manner (such as a in Figure 11(d)), or in a *counterclockwise* manner (such as b in Figure 11(d)). Notice that two such vints cannot hide behind the same vertex of o, one in a clockwise manner and the other in a counterclockwise manner, since this would imply that their corresponding triangles overlap (see Figure 14(a)). Therefore, either all of the vertices of S' hide in a clockwise manner, or they all hide in a counterclockwise manner; see Figure 14(b). This also implies that o' cannot contain an edge of the convex hull.

Consider a j-vint w. Its link o_w consists of j triangles, all incident to w, such that the sum of the angles at w over all j triangles is 2π. We say that a set D of consecutive triangles is a *dominating sequence* if the sum of the angles in those triangles is larger than π. Figure 14(c) depicts an 11-vint, and the four shaded triangles form a dominating sequence. Note that a separable edge defines a dominating sequence of size 2 (see Figure 1).

Lemma 4.3. *For every $i \in \mathbb{N}$, every vint has fewer than $3i^2$ dominating sequences of size at most i.*

Proof. Consider a j-vint w with at least one dominating sequence D of size at most i. A subset of the complementary set D' cannot be dominating. Thus every dominating sequence of w must include at least one triangle of D. There are less than i^2 contiguous subsequences of D. For sequences of size at most i not completely contained in D, we note that there are at most $2i$ such sequences with one element from D, at most $2i$ such sequences with two elements from D, ..., for a total of at most $2i^2$ sequences. □

Recall that we assume that there exists a vint $u = (p_u, T)$ in o of degree at least i. From the above, there is another vint $w = (p_w, T)$ hiding from v behind u (Figure 14(d)). Note that, because of the 'hiding', there is a dominating sequence of v starting from uv and going clockwise up to uw. The edges between the triangles of this sequence are external chords and a single edge of o. Thus the size of the set is at most $i-2$. In this case, we let v charge u. Since uv is on the boundary of a dominating sequence of u of size at most $i-2$, u is charged fewer than $6i^2$ times in this manner.

Recurrence with solution. Summing up the charges in all three cases and averaging over all $(i-1)$-vints, we obtain

$$\begin{aligned}\hat{v}_{i-1} &\le i\hat{v}_i + L_i \sum_{k\ge i}\hat{v}_k + 6i^2\sum_{k\ge i}\hat{v}_k + M_i\sum_{k\ge i}\hat{v}_k\\ &= (i+M_i+L_i+6i^2)\hat{v}_i + (M_i+L_i+6i^2)\sum_{k\ge i+1}\hat{v}_k\\ &= A_i\hat{v}_i + B_i\sum_{k=i+1}^{t}\hat{v}_k + B_i\sum_{k>t}\hat{v}_k,\end{aligned}$$

where $A_i = i + M_i + L_i + 6i^2$, $B_i = M_i + L_i + 6i^2$, and where t is chosen so that $t > 14B_i/\gamma_{i-1,h}$. Note that t too depends only on i (and on h). By Lemma 3.5, $\sum_{k\ge 3} k\hat{v}_k \le 6n + h - 9 < 7n$, and thus

$$B_i\sum_{k>t}\hat{v}_k \le \frac{B_i}{t}\sum_{k>t}k\hat{v}_k < \frac{B_i\cdot 7n}{t} < \frac{1}{2}\gamma_{i-1,h}n.$$

By the induction hypothesis, $\gamma_{i-1,h}n \le \hat{v}_{i-1} \le A_i\hat{v}_i + B_i\sum_{k=i+1}^{t}\hat{v}_k + \frac{1}{2}\gamma_{i-1,h}n$. By Lemma 3.3, for each $k \ge i+1$ there is a constant $\delta_{i,k}$ such that $\hat{v}_k \le \delta_{i,k}\hat{v}_i$. Putting $D_i = \sum_{k=i+1}^{t}\delta_{i,k}$, we get

$$\frac{1}{2}\gamma_{i-1,h}n \le (A_i + B_iD_i)\hat{v}_i \text{ , and } \hat{v}_i \ge \frac{\gamma_{i-1,h}}{2(A_i+B_iD_i)}n.$$

This establishes $\gamma_{i,h} = \frac{\gamma_{i-1,h}}{2(A_i+B_iD_i)}$ for induction on i and so completes the proof of Theorem 4.1. ◻

4.1 Large $\hat{v}_i$'s must always exist

The lower bounds for the v_i's, presented above come with small constants. We complement the analysis by showing, using a simple counting argument, that, for every point set with sufficiently many interior points, there are $\hat{v}_i$'s with much larger values, for small values of the index i. The "catch" is that we use an averaging argument, so we do not know which specific $\hat{v}_i$ has to be large.

LEMMA 4.4. *For every planar set of N points in general position, where n of them are interior and h are hull vertices, we have* $\max\{\hat{v}_3,\hat{v}_4,\hat{v}_5,\hat{v}_6\} \ge \frac{(n-h+9)}{10}$. *In particular, when the convex hull is triangular,* $\max\{\hat{v}_3,\hat{v}_4,\hat{v}_5,\hat{v}_6\} > \frac{N}{10}$ *holds.*

Proof. Consider a charging scheme for a triangulation T, where each i-vint of T is charged $7-i$. By Lemma 3.5, we get the following lower bound on the total charge in T:

$$\begin{aligned}\sum_i (7-i)v_i(T) &= \sum_i 7v_i(T) - \sum_i iv_i(T)\\ &\ge 7n - (6n+h-9) = n-h+9.\end{aligned}$$

Charges of i-vints with $i \ge 7$ are non-positive, thus ignoring them can only increase the total charge. Hence

$$4v_3(T)+3v_4(T)+2v_5(T)+v_6(T) \ge \sum_i (7-i)v_i(T) \ge n-h+9.$$

By linearity of expectation, $4\hat{v}_3 + 3\hat{v}_4 + 2\hat{v}_5 + \hat{v}_6 \ge n-h+9$. Letting $\hat{m} = \max\{\hat{v}_3,\hat{v}_4,\hat{v}_5,\hat{v}_6\}$, we have

$$10\hat{m} \ge 4\hat{v}_3 + 3\hat{v}_4 + 2\hat{v}_5 + \hat{v}_6 \ge n-h+9,$$

so $\hat{m} \ge \frac{n-h+9}{10}$, as asserted. ◻

Looking at the preceding lemma, one might suspect that the larger lower bound that it yields is due to $\hat{v}_3$ being large (as also suggested by the lower bound of [12], even though this latter bound is only $N/30$), and that the other $\hat{v}_i$'s are probably much smaller. As the following lemma shows, this is not the case, and some other large $\hat{v}_i$'s must also exist.

LEMMA 4.5. *For every point set in general position, with parameters N, n, and h, as above,* $\max\{\hat{v}_4,\hat{v}_5,\dots,\hat{v}_{11}\} \ge \frac{12n-9.5h+45}{180}$. *In particular, when the convex hull is triangular, we have* $\max\{\hat{v}_4,\hat{v}_5,\dots,\hat{v}_{11}\} \ge \frac{N-39/24}{15}$.

Proof. Similarly to Lemma 4.4, we consider a charging scheme for a triangulation T, where each i-vint of T is charged $12-i$. Once again, Lemma 3.5 provides the following lower bound on the total charge in T (with $v_i = v_i(T)$):

$$\sum_i (12-i)v_i \ge 12n - (6n+h-9) = 6n-h+9.$$

Ignoring the charge of i-vints with $i \ge 12$ can only increase the total charge. Let $m := \max\{v_4, v_5,\dots,v_{11}\}$ and obtain

$$9v_3+36m \ge 9v_3+8v_4+\cdots+v_{11} \ge \sum_i (12-i)v_i \ge 6n-h+9.$$

Using linearity of expectation and arguing as in the previous proof, we have $\hat{m} \ge \frac{6n-h+9-9\hat{v}_3}{36}$, where we set $\hat{m} := \max\{\hat{v}_4,\hat{v}_5,\dots,\hat{v}_{11}\}$. With the inequality $\hat{v}_3 \le \frac{2n+h/2}{5}$ from Lemma 2.3 this yields the asserted bound. ◻

5. REFERENCES

[1] M. Ajtai, V. Chvátal, M. M. Newborn, and E. Szemerédi, Crossing-free subgraphs, *Annals Discrete Math.* **12** (1982), 9–12.

[2] O. Aichholzer, F. Hurtado, and M. Noy, A lower bound on the number of triangulations of planar point sets, *Comput. Geom. Theory Appl.* **29**(2) (2004), 135–145.

[3] M.O. Denny and C.A. Sohler, Encoding a triangulation as a permutation of its point set, *Proc. 9th Canadian Conf. on Computational Geometry* (1997), 39–43.

[4] O. Giménez and M. Noy, The number of planar graphs and properties of random planar graphs, Proc. International Conf. on Analysis of Algorithms, *Discrete Math. & Theoretical Computer Sci.* **AD** (2005), 147–156.

[5] F. Hurtado, M. Noy, Counting triangulations of almost-convex polygons, *Ars Combin.* **45** (1997), 169–179.

[6] V. A. Liskovets, A pattern of asymptotic vertex valency distributions in planar maps, *J. Combinat. Theory*, Ser. B **75** (1999), 116–133.

[7] P. McCabe and R. Seidel, New lower bounds for the number of straight-edge triangulations of a planar point set, *Proc. 20th European Workshop Comput. Geom.* (2004).

[8] C. McDiarmid, A. Steger, and D.J.A. Welsh, Random planar graphs, *J. Combinat. Theory*, Ser. B **93** (2005), 187–205.

[9] A. Razen, J. Snoeyink, E. Welzl, Number of crossing-free geometric graphs vs. triangulations, *Electronic Notes in Discrete Math.* **31** (2008), 195–200.

[10] F. Santos and R. Seidel, A better upper bound on the number of triangulations of a planar point set, *J. Combinat. Theory*, Ser. A **102**(1) (2003), 186–193.

[11] R. Seidel, On the number of triangulations of planar point sets, *Combinatorica* **18**(2) (1998), 297–299.

[12] M. Sharir and A. Sheffer, Counting triangulations of planar point sets, 0911.3352v2 (improves [13]).

[13] M. Sharir and E. Welzl, Random triangulations of planar point sets, *Proc. 22nd ACM Symp. on Computational Geometry* (2006), 273–281.

[14] W. S. Smith, *Studies in Computational Geometry Motivated by Mesh Generation*, Ph. D. Thesis, Princeton University, 1989.

[15] R. P. Stanley, *Enumerative Combinatorics*, vol. **2**, Cambridge University Press, Cambridge, 1999.