Solution of Scott's Problem on the Number of Directions Determined by a Point Set in 3-Space*

János Pach† Rom Pinchasi‡ Micha Sharir§

December 31, 2004

Abstract

Let P be a set of n points in $\mathbb{R}^3$, not all in a common plane. We solve a problem of Scott (1970) by showing that the connecting lines of P assume at least $2n-7$ different directions if n is even and at least $2n-5$ if n is odd. The bound for odd n is sharp.

1 Introduction

According to the Gallai-Sylvester theorem [1, 6], for any set P of finitely many points in the plane, not all on a line, there exists a line that passes through precisely two elements of P. Erdős noticed that this result has the following corollary: Every n-element set P with the above properties, with $n \geq 3$, determines at least n connecting lines [6].

In 1970, Scott [16] raised the following questions: What is the minimum number of different directions assumed by the connecting lines of (1) n points in the plane, not all on a line, (2) n points in 3-space, not all on a plane?

*Work on this paper by János Pach and Micha Sharir has been supported by NSF Grants CCR-97-32101 and CCR-00-98246, and by a joint grant from the U.S.–Israel Binational Science Foundation. Work by János Pach has also been supported by PSC-CUNY Research Award 63382-0032 and by Hungarian Science Foundation grant OTKA T-032452. Work by Micha Sharir has also been supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. A preliminary version of this paper has appeared in *Proc. 20th Annu. ACM Sypos. Comput. Geom.*, 2004, pp. 76–85.

†City College, CUNY, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; pach@cims.nyu.edu

‡Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. Current address: Faculty of Mathematics, The Technion, Haifa 32000, Israel; room@math.technion.ac.il

§School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; michas@tau.ac.il

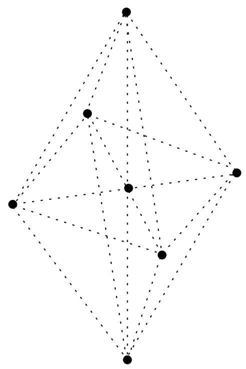

Figure 1: A set with $n = 7$ points that determines $2n - 5 = 9$ directions.

In 1982, after some initial results by Burton and Purdy [3], Ungar [18] solved the first problem, by verifying Scott's conjecture that in the plane the above minimum is equal to $2\lfloor n/2 \rfloor$, for any $n > 3$. For even n, this result is considerably stronger than the corollary of the Gallai-Sylvester theorem mentioned above. Ungar's proof is a real gem, a brilliant application of the method of *allowable sequences* invented by Goodman and Pollack [8], [9]. Moreover, it solves the problem in an elegant combinatorial setting, for "pseudolines", as was suggested independently by Goodman and Pollack and by Cordovil [4]. Interestingly, there is an overwhelming diversity of extremal configurations, for which equality is attained. Four infinite families and more than one hundred sporadic configurations were catalogued by Jamison and Hill [11]. See also [10] for an excellent survey by Jamison, and the monograph of Aigner and Ziegler [1], where Ungar's proof and some of its relatives are reproduced.

In lack of a natural ordering of all directions in 3-space, Ungar's method does not seem to generalize. This explains why until recently there had not been much progress concerning Scott's second question. Scott's construction of a double pyramid whose base is a regular polygon with an even number of edges, including the center of the base (see Figure 1), shows that the number of directions determined by n non-coplanar points can be as small as $2n - 5$ if n is odd. This bound was conjectured to be tight. Under the additional assumption that *no three points of the set are collinear*, Blokhuis and Seress [2] proved that the number of directions determined by $n \geq 6$ non-coplanar points in 3-space is at least $7n/4 - 2$. Using the same condition, we have recently succeeded in proving the tight bound $2n - 2$ if n is odd and $2n - 3$ if n is even [14].

In the present paper we solve Scott's second problem in full generality, by removing the assumption that no three points are collinear.

Theorem 1.1. *Every set of $n \geq 6$ points in $\mathbb{R}^3$, not all of which are on a plane, determine at least $2n - 5$ different directions if n is odd, and at least $2n - 7$ different directions if n is even. This bound is sharp for every odd $n \geq 7$.*

The case where n is even is handled by removing one point and applying the bound

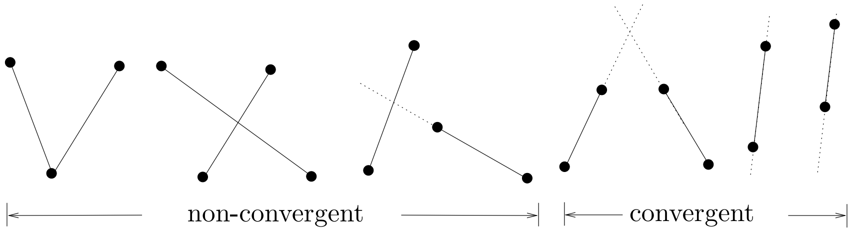

Figure 2: Convergent and non-convergent segments.

for odd n. Therefore, from this point on we assume that n is odd. Nevertheless, we believe that the bound for even n can be improved. We note that the double-pyramid construction in Figure 1, without the center, yields an upper bound of $2n-3$ for n even.

The idea of the proof is outlined in Section 2. A key new ingredient of our argument is Theorem 3.1, proved in Section 3, which is a far reaching "bipartite" variant of Ungar's aforementioned theorem.

Definition 1.2. *Two closed segments in $\mathbb{R}^d$ are called* convergent *if (i) they do not belong to the same line, and (ii) their supporting lines intersect, and the intersection point does not belong to any of the segments. See Figure 2.*

An alternative definition is that two segments are convergent if and only if they are disjoint and their convex hull is a planar quadrilateral. Two parallel segments that lie on distinct lines are also considered to be convergent (by regarding their lines to meet at infinity, or according to the alternative definition). Note also that condition (ii) rules out pairs of segments with a common endpoint.

Instead of Theorem 1.1, in Section 4 we establish the following significantly stronger result.

Theorem 1.3. *Every set of $n \geq 6$ points in $\mathbb{R}^3$, not all of which are on a plane, determine at least $2n-5$ segments if n is odd, and at least $2n-7$ segments if n is even, no two of which are convergent. This bound is sharp for every odd $n \geq 7$.*

We apply Theorem 1.3 in Section 5 to partially settle in the affirmative a conjecture of Blokhuis and Seress [2], showing (in Theorem 5.1) that any set P of n points in $\mathbb{R}^4$, not contained in a hyperplane and not having three collinear points, determine at least $3n-8$ different directions, if n is even, and at least $3n-10$ different directions if n is odd. The bound is sharp for every even $n \geq 8$.

Rédei's monograph on *lacunary polynomials* [15] was the starting point of many investigations related to algebraic variants of the above problem. For instance, it was proved in [15] that if n is a prime, then any set of n points in the affine plane $\mathrm{AG}(2,n)$ determines at least $(n+3)/2$ different directions. Lovász and Schrijver [13] characterized all sets for which equality is attained. In the finite projective plane PG$(2, N)$, a set P of $n > 4$ points, no three of which are collinear, is known to determine at least n different directions if N is odd and at least $n-1$ if N is even. Equality is attained here if and only if P spans a (properly defined) *affinely regular n-gon* (see [7, 12]). The last theorem, due to Wettl [19] answers a question of Gus Simmons in *cryptography.* For many similar results and applications in finite geometry, algebraic number theory, and group theory, consult the survey of Szőnyi [17].

2 Preliminaries

Let P be a set of n points in $\mathbb{R}^3$ such that not all of them lie in a common plane. Let p_0 be an *extreme* point of P, i.e., a vertex of the convex hull of P. Consider a supporting plane to P at p_0, and translate it to the side that contains P. Let π denote the resulting plane. Project from p_0 all points of $P \setminus \{p_0\}$ onto π. We obtain a *multiset* R of points in π, not all on a line, so that each point is the image of some points of P. We regard R as a set of weighted points, where the weight $w(r)$ of a point $r \in R$ is the number of points of $P \setminus \{p_0\}$ that project onto it.[1] The sum of the weights is $n-1$. For a subset $A \subseteq R$, we define $w(A) := \sum_{q \in A} w(q)$.

We assume that n is odd, thus $w(R) = n-1$ is even. We attempt to partition R into two subsets R^+, R^-, so that $w(R^+) = w(R^-) = (n-1)/2$ and all points of R^+ lie to the left of every point of R^- with respect to some generic coordinate frame in π, in which no two elements of R have the same x-coordinate.

For the choice of the coordinate frame and the partition, we begin with the following elementary geometric fact.

Lemma 2.1. *Let R be a set of non-collinear weighted points in the plane, with a total even weight m. Let r be any vertex of the convex hull of R whose weight is smaller than $m/2$. Then one of the following properties holds:*

(i) There exists a partition of R into two subsets, R^- and R^+, each of overall weight $m/2$, whose convex hulls are disjoint and which have a common inner tangent m_0 passing through r.

(ii) There exists a point $q \in R$ and a partition of $R \setminus \{q\}$ into two subsets, R_0^- and R_0^+, each of overall weight $< m/2$, so that the convex hulls of $R_0^- \cup \{q\}$ and $R_0^+ \cup \{q\}$ meet only at q, which is a common vertex of both hulls, and the line m_0 passing through r and q is an inner common tangent to the two hulls (supporting one of them in the edge qr).

Proof: See Figure 3. Rotate a directed line ℓ counterclockwise about r, starting with all the points of $R \setminus \{r\}$ lying to the left of ℓ, until the closed halfplane to the right of ℓ contains for the first time points with overall weight larger than $m/2$. Let R_0^-

[1] In the preceding paper [14], where it was assumed that no three points of P are collinear, R was a *set*, or, rather, the weight of each point was 1.

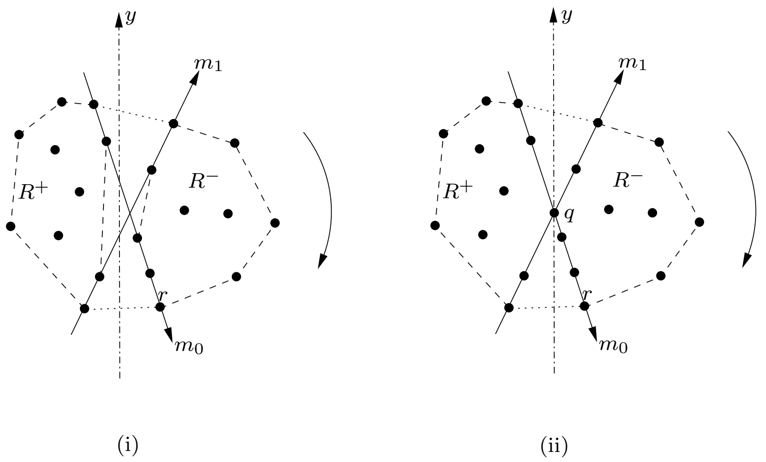

Figure 3: The primal construction of R^- and R^+: Case (i) (left) and Case (ii) (right).

denote the set R_0 of points in the open halfplane to the right of ℓ, plus the largest set of initial points of $\ell \cap R$ along ℓ (in their order along ℓ starting at r) whose overall weight does not exceed $m/2 - w(R_0)$.

If the overall weight of R_0^- is exactly $m/2$, we are in case (i). We define $R^- := R_0^-$, and $R^+ := R \setminus R_0^-$. See Figure 3(i). It is clear that the convex hulls of R^+ and R^- are disjoint, and that the final position of ℓ is the desired common inner tangent m_0.

If the overall weight of R_0^- is less than $m/2$, we are in case (ii). Let q be the next point of $\ell \cap R$ along ℓ, and define $R_0^+ := R \setminus (R_0^- \cup \{q\})$. See Figure 3(ii). It is easily seen that the properties asserted in (ii) hold, with m_0 being the final position of ℓ. $\Box$

We apply Lemma 2.1 to our set $R \subset \pi$, with $m = n-1$. In case (ii), we split q into two co-located points q^-, q^+, and distribute the weight $w(q)$ between them, so that $w(q^-) = (n-1)/2 - w(R_0^-)$ and $w(q^+) = (n-1)/2 - w(R_0^+)$. We set $R^- := R_0^- \cup \{q^-\}$ and $R^+ := R_0^+ \cup \{q^+\}$. We refer to q as the *central bichromatic point* of R.

Let m_1 denote the other inner tangent of the convex hulls of R^- and R^+. In case (ii), m_1 also passes through q and through at least one other point of one of the two sets. Now choose in π an orthogonal (x, y)-coordinate system whose y-axis is either a line that strictly separates R^- and R^+ in case (i), or a line through q that strictly separates R_0^- and R_0^+ in case (ii). We can carry out the construction so that (a) R^+ and R^- are to the left and to the right of the y-axis, respectively, (b) $r \in R^-$, and (c) m_0 is oriented from r away from the other contact point(s), and the positive y-direction lies counterclockwise to it. See Figure 3. This still leaves us with some freedom in fixing the coordinate frame. We will later impose further constraints on it to facilitate certain steps in our analysis.

The presence of q adds an extra level of complication to the proof. We note that in the configuration shown in Figure 1, choosing p_0 to be any vertex of the hull, say, the lowest point, the weighted set R has a central bichromatic point q, as shown in Figure 31(ii). As will follow from our analysis, the bounds in both Themorems 1.1

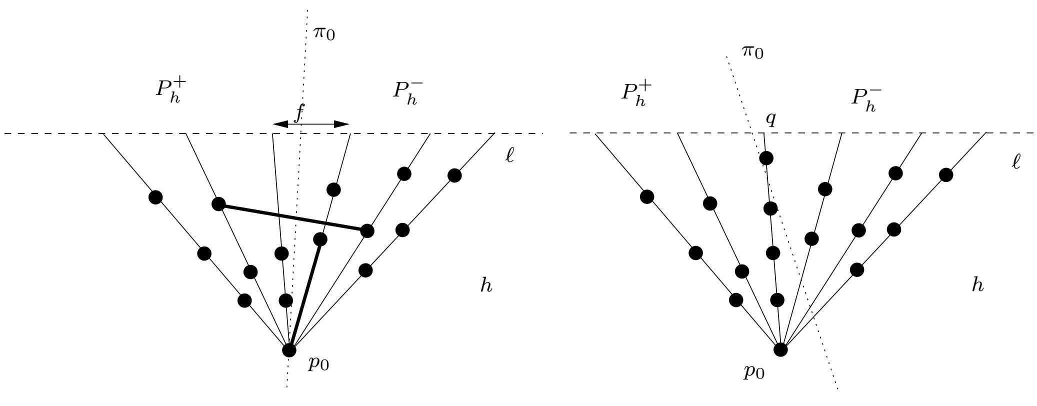

Figure 4: The sets P_h^+ and P_h^-. (a) The case where the central bichromatic point q (if it exists at all) does not lie on ℓ. (b) The case where q exists and lies on ℓ.

and 1.3 improve to $2n-2$, for n odd, when q does not exist.

Let P^+ (resp., P^-) denote the set of points of $P \setminus \{p_0\}$ that project from p_0 to points of R^+ (resp., R^-). Points projecting to q are split between P^+ and P^-. This can be best visualized by a plane π_0 that separates P^+ and P^-. If q does not exist, π_0 is the plane spanned by p_0 and the y-axis in π. If q exists, π_0 cuts the line containing the preimages of q into two pieces, one of which contains $w(q^+)$ preimages and the other contains $w(q^-)$ preimages. Without loss of generality, we may assume that the $w(q^+)$ preimages closest to p_0 belong to P^+. See Figure 4.

Let ℓ be a line in π, and let h be the plane spanned by ℓ and p_0. In case (i), the sets $P_h^+ = P^+ \cap h$ and $P_h^- = P^- \cap h$ are separated in h in such a way, that the cones with apex p_0 spanned by their convex hulls are disjoint (except for p_0); see Figure 4(a). In case (ii), if $q \in \ell$, then both cones are bounded by the common line λ connecting p_0 to q. See Figure 4(b). In this case, by construction, the first $w(q^+)$ points of $(P \setminus \{p_0\}) \cap \lambda$ (those nearest to p_0) belong to P^+, and the rest of the points on λ belong to P^-.

Here is a brief overview of the proof of Theorem 1.3. In Section 4, we collect a set F of mutually non-convergent segments in π whose endpoints belong to R. With few exceptions, the segments in F connect points of R^+ to points of R^-. Each segment $f \in F$ gives rise to a set $E(f)$ of pairwise non-convergent segments in 3-space, determined by P, and lying in the plane h spanned by p_0 and f. Each segment $e \in E(f)$ either connects a point of P_h^+ to a point of P_h^-, or connects p_0 to some point in P_h, such that the projection e^* of e from p_0 on the line supporting f, either fully contains f or is a point, outside the interior of f; see Figure 4(a). Let E denote the union of all the sets $E(f)$. Using a fairly intricate analysis based on Theorem 3.1 (stated and proved in Section 3), we show that (a) E contains at least $2n-5$ segments (Section 4.4), and (b) every pair of distinct segments in E are non-convergent and therefore non-parallel (Section 4.5). Once (a) and (b) are established, Theorems 1.3

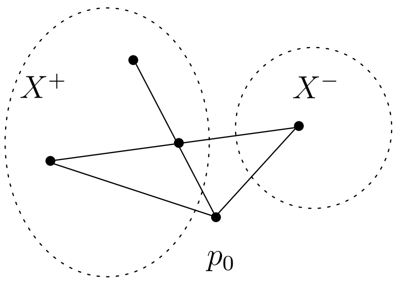

Figure 5: An example where Theorem 3.1 does not hold.

and 1.1 follow, because the directions of the segments in E are all different.

We emphasize once again that Theorem 1.3 is considerably stronger than Theorem 1.1. Besides being of independent interest, we expect this strengthening to be useful for extending our results to higher dimensions, using induction on the dimension; see the concluding section for more details.

3 A Bipartite Ungar-type Theorem

A crucial ingredient of our analysis is the following variant of Ungar's theorem, which we believe to be of independent interest.

Theorem 3.1. *Let X^+ and X^- be two finite sets of points in the plane, and let p_0 be a point in the plane, such that p_0 is an extreme point of $X^+ \cup X^- \cup \{p_0\}$, and there is a line through p_0 that strictly separates X^+ and X^-. We also assume that $0 < |X^-| \leq |X^+|$ and that the innermost ray from p_0 to a point of X^+ (forming the smallest angle with the separating line) contains* more than *$|X^+| - |X^-|$ points. Then one can select at least $|X^+| + |X^-| + 1$ pairwise non-convergent segments connecting points of $X^+ \cup \{p_0\}$ to points of $X^- \cup \{p_0\}$.*

We remark that the "+1" term in the above bound is crucial for our analysis, and that we may lose this term if the assumption on the points in the innermost ray does not hold, as is illustrated in Figure 5, where $|X^+| + |X^-| + 1 = 5$ but at most four pairwise non-convergent segments can be selected.

Corollary 3.2. *Assume the conditions of Theorem 3.1, with the difference that the innermost ray from p_0 to a point of X^+ contains exactly $|X^+| - |X^-|$ points. Then one can select at least $|X^+| + |X^-|$ pairwise non-convergent segments connecting points of $X^+ \cup \{p_0\}$ to points of $X^- \cup \{p_0\}$.*

We remark that, in the applications of Theorem 3.1 and Corollary 3.2 in Section 4, the sets X^-, X^+ will always be well balanced, so that the innermost ray will always contain at least $|X^+| - |X^-|$ points.

We also note that Ungar's theorem "almost" follows from Theorem 3.1 and its corollary. That is, let P be a set of n non-collinear points in the plane, where n is even. Pick an extreme point p_0 of P, and find a line that passes through p_0 and splits $P \setminus \{p_0\}$ into two subsets X^+, X^- whose sizes are as equal as possible. Suppose that $|X^+| \geq |X^-|$. Then the innermost ray from p_0 to points of X^+ must contain at least $|X^+| - |X^-|$ points, for otherwise we could have transferred these points to X^- and get a split with a smaller size difference. If the number of points on the innermost ray is *larger* than $|X^+| - |X^-|$, then Theorem 3.1 applies, and yields at least $|X^+| - |X^-| + 1 = |P|$ pairwise non-convergent segments connecting the points of P, which implies (and is much stronger than) Ungar's theorem. However, if the number of points on the innermost ray is *equal* to $|X^+| - |X^-|$, then only Corollary 3.2 can be applied, and it only yields $|X^+| - |X^-| = |P| - 1$ pairwise non-convergent segments connecting the points of P, one shorter of what Ungar's theorem gives. We leave it as an open problem to determine whether Ungar's theorem can always be deduced from Theorem 3.1 and Corollary 3.2.

Proof of Corollary 3.2: Remove one point from X^+ which is not on the innermost ray from p_0 to a point of X^+ (note that X^+ is not fully contained in that ray, since $|X^-| > 0$), and apply Theorem 3.1 to the resulting set of points. $\square$

Proof of Theorem 3.1: Fix an (x, y)-coordinate system in the plane. We apply a standard duality transform that maps a point $p = (p_1, p_2)$ to the line p^* with equation $y + p_1 x + p_2 = 0$. Vice versa, a non-vertical line l with equation $y + l_1 x + l_2 = 0$ is mapped to the point $l^* = (l_1, l_2)$. Consequently, any two parallel lines are mapped into points having the same x-coordinate. It is often convenient to imagine that the dual picture lies in another, so-called *dual*, plane, different from the original one, which is referred to as the *primal* plane.

The above mapping is incidence and order preserving, in the sense that p lies above, on, or below ℓ if and only ℓ^* lies above, on, or below p^*, respectively. The points of a segment $e = ab$ in the primal plane are mapped to the set of all lines in the closed *double wedge* e^*, which is bounded by a^* and b^* and does not contain the vertical direction. All of these lines pass through the point $q = a^* \cap b^*$, which is called the *apex* of the double wedge e^*. All double wedges used in this paper are assumed to be closed, and they never contain the vertical direction.

We call two double wedges *convergent* if their apices are distinct and the apex of neither of them is contained in the other. See Figure 6.

It is easy to see that, according to this definition, two non-collinear segments in the primal plane are convergent if and only if they are mapped to convergent double wedges.

Without loss of generality, we assume that p_0 is the origin, that X^+ lies to the left of the y-axis, that X^- lies to its right, and that both sets lie below the x-axis; see Figure 7(a). The duality maps p_0 into the x-axis, which we denote as ℓ_0, the lines connecting p_0 to points in X^+ (resp., X^-) to points on the negative (resp., positive) x-axis, and the points of X^+ (resp., X^-) to positively (resp., negatively)

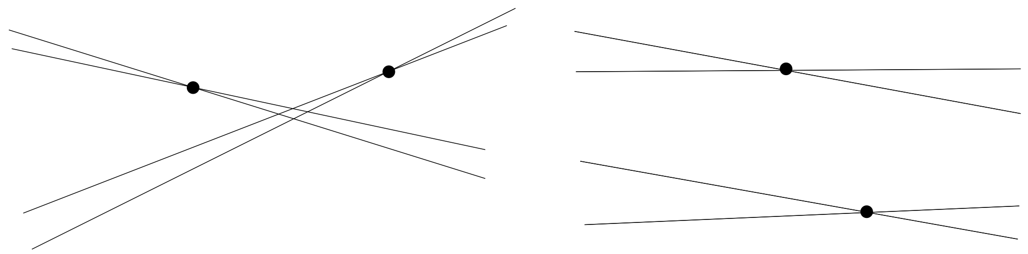

Figure 6: Two possible kinds of convergent double wedges.

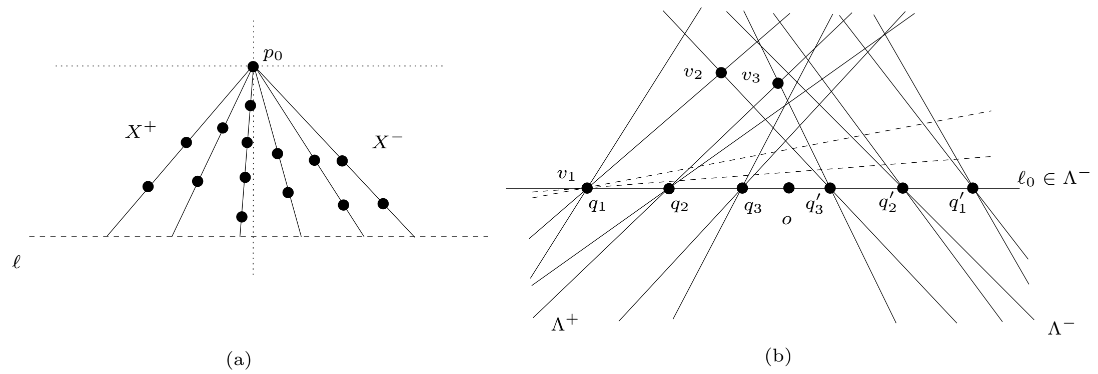

Figure 7: The setup in Theorem 3.1. (a) The primal configuration. (b) The dual configuration. Since $|X^+| = 8$ and $|X^-| = 6$, we have two excess lines, shown as dashed.

slopped lines; see Figure 7(b). Let Λ^+, Λ^- denote the set of lines dual to the points of X^+, X^-, respectively, and put $\Lambda = \Lambda^+ \cup \Lambda^- \cup \{\ell_0\}$. Enumerate the points dual to the lines connecting p_0 to the points of X^+ as $q_1, \ldots, q_k$ in this left-to-right order, and the points dual to the lines connecting p_0 to the points of X^- as $q'_1, \ldots, q'_\ell$ in this right-to-left order; thus q_1 is the leftmost point and q'_1 is the rightmost. Put $n^+ = |\Lambda^+| = |X^+|$, $n^- = |\Lambda^-| = |X^-|$.

We *redefine* the sets Λ^+, Λ^- as follows. First we set $\Lambda^- := \Lambda^- \cup \{\ell_0\}$, and then remove from Λ^+ the $n^+ - n^-$ lines that pass through q_1 and have the shallowest slopes; we refer to these lines as *excess lines*, and denote the set that they comprise by Λ^e. Redefine $\Lambda^+ := \Lambda^+ \setminus \Lambda^e$. We have $|\Lambda^-| = |\Lambda^+| + 1 = n^- + 1$. We note that by an appropriate choice of the coordinate frame in the duality transform, we may assume that the slopes of the excess lines are the smallest among all lines in Λ^+.

Constructing junctions. We apply an iterative pruning process that constructs a sequence of vertices ("junctions") $v_1, \ldots, v_m$ which are intersection points of lines from Λ^- and lines from Λ^+, and sets of intermediate vertices ("stations") between successive junctions, as well as a set of "termini" to the right of the rightmost junction. The sequence J of *junctions* $\langle v_1, v_2, \ldots, v_k \rangle$ is constructed as follows.

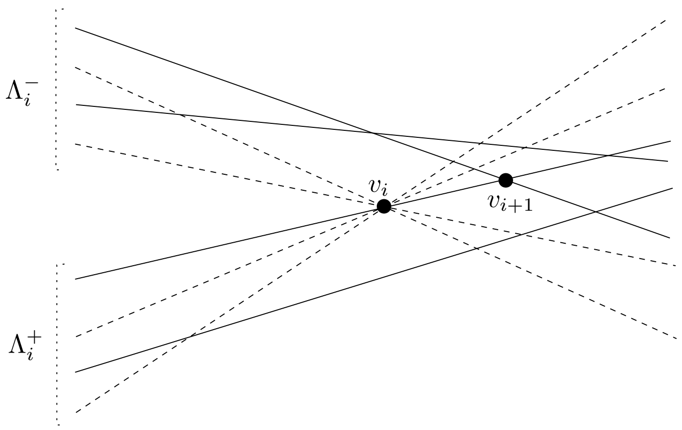

Figure 8: Constructing the junction v_i in J. The dashed lines, two from Λ^+ and two from Λ^-, are removed at v_i. The next junction v_{i+1} is also shown.

Step 1: Set $i := 1$ and $\Lambda_1^+ := \Lambda^+$, $\Lambda_1^- := \Lambda^-$.

Step 2: If $\Lambda_i^+ = \emptyset$, the construction of J terminates. Otherwise, as guaranteed by the construction, neither set is empty. Let v_i be the *leftmost* intersection point between a line in Λ_i^+ and a line in Λ_i^-. Let d_i^+ (and d_i^-) denote the number of lines of Λ_i^+ (and Λ_i^-, respectively) incident to v_i, and put $d_i := \min\{d_i^+, d_i^-\}$. Define Λ_{i+1}^+ (resp., Λ_{i+1}^-) as the set of lines obtained from Λ_i^+ (resp., Λ_i^-) by deleting from it the d_i lines that are incident to v_i and have the largest (resp., smallest) slopes among those incident lines. (That is, if $d_i^+ = d_i^-$, then all lines incident to v_i are deleted; otherwise, if, say, $d_i^+ > d_i^-$, we are left with $d_i^+ - d_i^-$ lines through v_i that belong to Λ_i^+ and separate the deleted elements of Λ_i^+ from the deleted elements of Λ_i^-. See Figure 8, where $d_i = d_i^- = 2$, $d_i^+ = 3$, and the dashed lines, two from Λ_i^+ and two from Λ_i^-, are removed at v_i.) Set $i := i + 1$, and repeat Step 2.

Note that, due to the special structure of the arrangement, we have $v_1 = q_1$ and $d_1 = 1$. See Figure 7(b). Recall also that the excess lines do not participate in the junction construction process.

It is easy to verify the following properties of this construction.

Claim 3.3. *(i)* $|\Lambda_i^-| = |\Lambda_i^+| + 1$, *for each* $i = 1, \ldots, k$; $\Lambda_{k+1}^+ = \emptyset$ *and* $|\Lambda_{k+1}^-| = 1$.

(ii) For every $1 \le i < j \le k$, *the junction* v_i *lies in the left unbounded face* f_j *of* $\mathcal{A}(\Lambda_j^+ \cup \Lambda_j^-)$ *which separates* Λ_j^+ *and* Λ_j^- *at* $x = -\infty$ *(and whose rightmost vertex is* v_j*).* v_i *lies in the interior of* f_j *if* $d_i^+ = d_i^-$; *otherwise it may lie on the boundary of* f_j.

(iii) $\sum_{i=1}^k d_i = |\Lambda^+| = n^-$. $\square$

Collecting stations. Next, between any two consecutive junctions v_i and v_{i+1}, for $1 \le i < k$, we specify $d_i + d_{i+1} - 1$ further vertices of $\mathcal{A}(\Lambda^+ \cup \Lambda^-)$, called *stations* (thus, the excess lines are still kept out of the construction).

Fix an index $1 \le i < k$, and consider the vertical slab between v_i and v_{i+1}. By Claim 3.3 (ii), v_i lies inside or on the boundary of the face f_{i+1} of $\mathcal{A}(\Lambda^+_{i+1} \cup \Lambda^-_{i+1})$, whose rightmost vertex is v_{i+1}. See Figure 9. Hence, the segment $e = v_i v_{i+1}$ is contained in the closure of f_{i+1}. Now at least one of the following two conditions is satisfied: (a) all the d_i lines removed from Λ^+_i and all the d_{i+1} lines removed from Λ^-_{i+1} pass strictly above e, or (b) all the d_i lines removed from Λ^-_i and all the d_{i+1} lines removed from Λ^+_{i+1} pass strictly below e.

Indeed, if v_i lies in the interior of f_{i+1} then the d_{i+1} lines of Λ^+_{i+1} (resp., of Λ^-_{i+1}) that are removed at v_{i+1} pass strictly below (resp., above) v_i. In this case, the validity of either (a) or (b) follows by considering the position of v_i among the lines of $\Lambda^+_i \cup \Lambda^-_i$ that are removed at v_i. If v_i lies on the boundary of f_{i+1} (as shown in Figure 9), then it has to lie on a line of $\Lambda^+_{i+1} \cup \Lambda^-_{i+1}$, say it lies on a line ℓ of Λ^+_{i+1} (as shown in the figure). Then all the d_{i+1} lines removed from Λ^-_{i+1} pass strictly above v_i and e. Now the line ℓ belongs to Λ^+_i and passes through v_i. Since it was not removed at v_i, all the d_i lines of Λ^+_i that were removed pass strictly above e, by construction, so (a) holds. If v_i lies on a line of Λ^-_{i+1}, a symmetric argument shows that (b) holds.

Assume, by symmetry, that (a) holds. Denote the lines removed from Λ^+_i by $\ell^+_1, \ldots, \ell^+_{d_i}$, listed according to increasing slopes, and those removed from Λ^-_{i+1} by $\ell^-_1, \ldots, \ell^-_{d_{i+1}}$, listed according to decreasing slopes. See Figure 9. Define the set of *stations* S_i in the vertical slab between v_i and v_{i+1} as the collection of all intersection points of $\ell^+_{d_i}$ with the lines $\ell^-_1, \ldots, \ell^-_{d_{i+1}}$, and all intersection points of $\ell^-_{d_{i+1}}$ with the lines $\ell^+_1, \ldots, \ell^+_{d_i}$. Clearly, we have $|S_i| = d_i + d_{i+1} - 1$ such points; see Figure 9. We refer to the grid-like crossing pattern between the lines $\ell^+_1, \ldots, \ell^+_{d_i}$ and the lines $\ell^-_1, \ldots, \ell^-_{d_{i+1}}$, as the *upper grid* between v_i and v_{i+1}. The collected stations lie on the "upper rim" of that grid. In complete analogy, when case (b) applies, we collect stations along the "lower rim" of the *lower grid* between v_i and v_{i+1}.

The description so far matches the one given in [14]. We now describe the new features of the present collection process. They involve (a) collecting "excess stations" for the excess lines, and (b) collecting vertices (that we refer to as "termini") to the right of v_k.

Collecting excess stations. The collection of excess stations proceeds as follows. As we collect the junctions v_i, we maintain a subset Λ^e_i of 'surviving' excess lines. For each i, the lines in Λ^e_i satisfy the property that they have passed below or through each of the junctions $v_1, \ldots, v_i$. Initially, $\Lambda^e_1 = \Lambda^e$, all of whose lines clearly satisfy this property (they pass through v_1). When we reach a new junction v_j, we remove certain lines from Λ^e_{j-1}. When an excess line is removed, we associate with it a new *excess station* that lies somewhere to the left of v_j. Typically, but not always, it will be a grid vertex between v_{j-1} and v_j. To disambiguate between the two kinds of stations, we will sometimes refer to the previously constructed stations as *standard stations*. The removal of excess lines and the construction of excess stations proceed

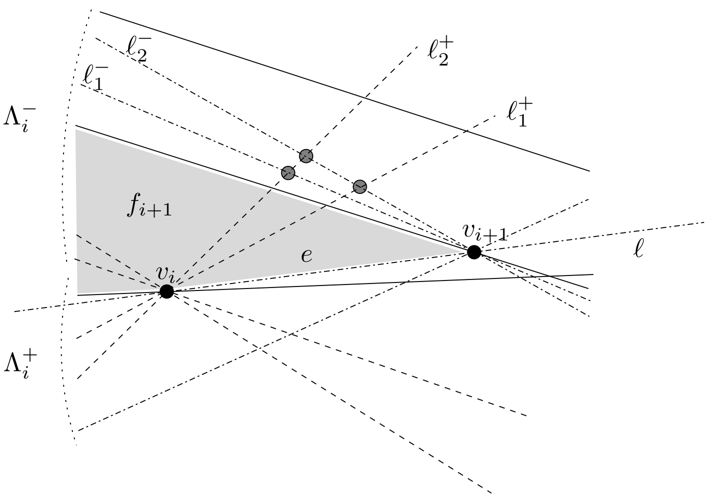

Figure 9: Collecting stations between v_i and v_{i+1}. We have $d_i = d_{i+1} = 2$. The lines removed at v_i are drawn as dashed, and those removed at v_{i+1} are drawn as dashed-dotted.

according to the following rule:

(i) Our default option is to use the upper grid for collecting intermediate (standard) stations between v_{j-1} and v_j. Recall that, for this to be possible, all lines of Λ^- incident to v_j and removed there have to pass strictly above v_{j-1}, and all lines of Λ^+ incident to v_{j-1} and removed there have to pass strictly above v_j. If the first condition is violated then the shallowest line of Λ^- incident to v_j and removed there also passes through v_{j-1} (by Claim 3.3(ii), it cannot pass below v_{j-1}), and if the second condition is violated then the shallowest line of Λ^+ incident to v_{j-1} and removed there passes through or below v_j. Thus, if none of the two latter conditions arise, we use the upper grid.

Assuming this to be the case, we remove each surviving excess line that passes *above* v_j. The removed excess lines meet the steepest line of Λ^- incident to v_j at points that lie along the upper grid and are further to the right of all the other grid points (and thus to the right of all the standard stations in S_{j-1}). This latter property is a consequence of the fact that all these excess lines pass below or through v_{j-1} and have slopes smaller than those of the lines of Λ^+ that are incident to v_{j-1}; see Figure 10.

In conclusion, each removed excess line is associated with a new upper grid vertex of the arrangement, and these are the excess stations that we have promised to collect. We set Λ_j^e to be the set of surviving excess lines, which still pass through or below v_j (so the invariant continues to hold), and continue the process with $j := j + 1$.

(ii) Suppose that we have to use the lower grid for collecting intermediate stations between v_{j-1} and v_j. As just mentioned, the lower grid has to be used either when (a) v_j lies on or above at least one of the d_{j-1} lines of Λ^+ incident to v_{j-1} and removed there, or when (b) v_{j-1} lies on the shallowest of the d_j lines of Λ^- incident to v_j and

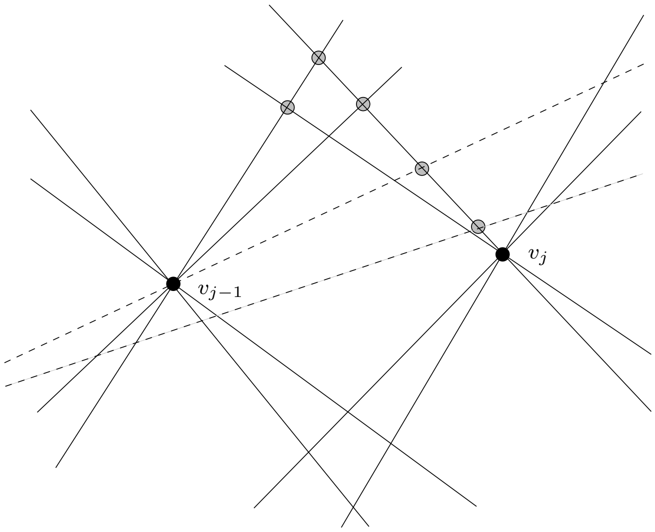

Figure 10: Charging excess lines of Λ^e_{j-1} that pass above v_j to excess stations in the upper grid between v_{j-1} and v_j.

removed there.

In case (a), let ℓ denote the shallowest line in Λ^+ through v_{j-1} that is removed at v_{j-1}. Refer to Figure 11(a), and note that v_j lies on or above ℓ. In this case, each excess line in Λ^e_{j-1} must pass below v_j, because it passes below or through v_{j-1} and its slope is smaller than that of ℓ. Hence, in this case we do not remove any excess line, and thus set $\Lambda^e_j := \Lambda^e_{j-1}$. In particular, the invariant property holds for Λ^e_j in this case, and we continue the collection process with $j := j+1$.

In case (b), which is depicted in Figure 11(b), let ℓ denote the shallowest line in Λ^- through v_j that is removed at v_j. ℓ passes also through v_{j-1}. We use the lower grid to construct excess stations for the excess lines of Λ^e_{j-1} that pass above v_j. These will be the intersection points of these lines with the steepest line of Λ^- incident to v_{j-1}. Because of the slope conditions, these points lie to the left of all the standard stations between v_{j-1} and v_j. However, if there exists an excess line λ through v_{j-1}, this procedure will fail to produce an additional excess station for λ. To gain such a station elsewhere, we observe that $j-1 \neq 1$ (since ℓ_0 is the only line of Λ^- through v_1, so that it is deleted there and does not belong to Λ^-_2), and that we must have used the upper grid between v_{j-2} and v_{j-1}. This holds because v_{j-2} must lie on or below ℓ and on or above λ. Hence, all the lines of Λ^+ incident to v_{j-2} must pass strictly above v_{j-1} (since they have slopes larger than that of λ), and all the lines of Λ^- incident to v_{j-1} and removed there must pass strictly above v_{j-2} (since they are all steeper than ℓ). Note that the number of lines of Λ^- through v_{j-1} is greater than d_{j-1}, because this set also contains ℓ, which has not been removed at v_{j-1}. Using that extra line, we can therefore gain one additional intersection point as the required excess station in the upper grid between v_{j-2} and v_{j-1}.

However, one such extra grid station may fail to exist if v_{j-2} also lies on ℓ. Refer

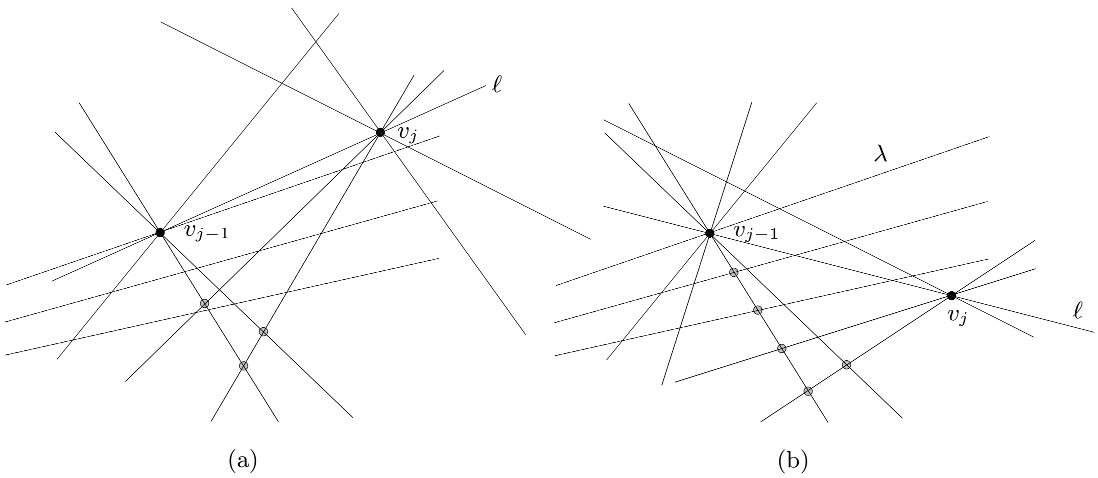

Figure 11: Using the lower grid between v_{j-1} and v_j.

to Figure 12.

In this case, it is easily seen that, as far as the collection of standard stations goes, we can use the lower grid between v_{j-2} and v_{j-1} instead of the upper grid. Indeed, all the lines of Λ^- incident to v_{j-2} and removed there pass below v_{j-1} (because ℓ passes through v_{j-1} and is not removed there), and all the lines of Λ^+ incident to v_{j-1} and removed there pass below v_{j-2} (because λ passes through v_{j-1} and below v_{j-2}). If v_{j-2} is not incident to an excess line, then all excess lines in Λ^e_{j-2} that pass through or above v_{j-1} (including λ) determine excess stations on the lower grid between v_{j-2} and v_{j-1}. Hence in this case we obtain on the lower grid one additional excess station, formed by λ, and can therefore quit this process. If v_{j-2} is incident to an excess line, we attempt to collect an extra excess station in the upper grid between v_{j-3} and v_{j-2}, expoiting, as above, the excess of lines of Λ^- at v_{j-2}. Again, this may fail if v_{j-3} also lies on ℓ, so we move to the lower grid between v_{j-3} and v_{j-2}, and we keep applying this backtracking process until we reach a junction v_s that lies strictly below ℓ. This will happen, if not earlier, when we reach v_1, since the only line of Λ^- incident to v_1 is ℓ_0, which is different from ℓ.

To recap, this process creates an excess station for each excess line removed at v_j. Note that if the construction had to backtrack from v_j through several preceding junctions, then ℓ is the *shallowest* line of Λ^-_j that passes through v_j. Hence, if backtracking will also be required at some later junction $v_{j'}$, for $j' > j$, then the process will have to terminate at a junction to the right of v_j (because no surviving line of Λ^-_j passes through v_j). That is, the backtracking processes are *independent* of each other, and none of them affects any of the preceding ones.

Collecting termini. Finally, consider the last junction v_k and the final set Λ^e_k of surviving excess lines. There are d_k lines of Λ^- as well as d_k lines of Λ^+ that pass through v_k, and there is another surviving line ℓ of Λ^-, which passes through or above

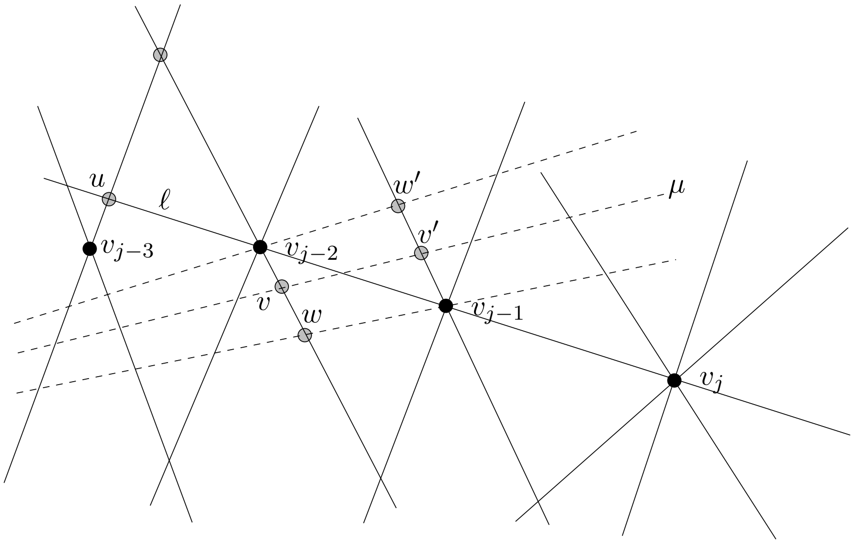

Figure 12: Handling the modified case (b), where ℓ passes through additional vertices preceding v_{j-1}.

v_k. Our goal is to collect $d_k + |\Lambda_k^e| + 1$ additional vertices of $\mathcal{A}(\Lambda)$ to the right of v_k, to which we refer as *termini*.

If ℓ passes above v_k, then we obtain on it the distinct intersection points with ℓ_0, with the excess lines in Λ_k^e, and with the d_k lines of Λ^+ through v_k (it is easy to verify that all these intersection points are indeed distinct); see Figure 13(a). Altogether, we collect $d_k + |\Lambda_k^e| + 1$ termini.

If ℓ passes through v_k (see Figure 13(b)), let ℓ' denote the steepest line of Λ^- incident to v_k. We charge each of the d_k+1 lines of Λ^- incident to v_k to its intersection with ℓ_0. In addition, each excess line in Λ_k^e, with the exception of the excess line λ_0 that passes through v_k (if there is such a line), meets ℓ' at a vertex, and we add these vertices to the set of collected termini; their x-coordinates are all distinct, and lie to

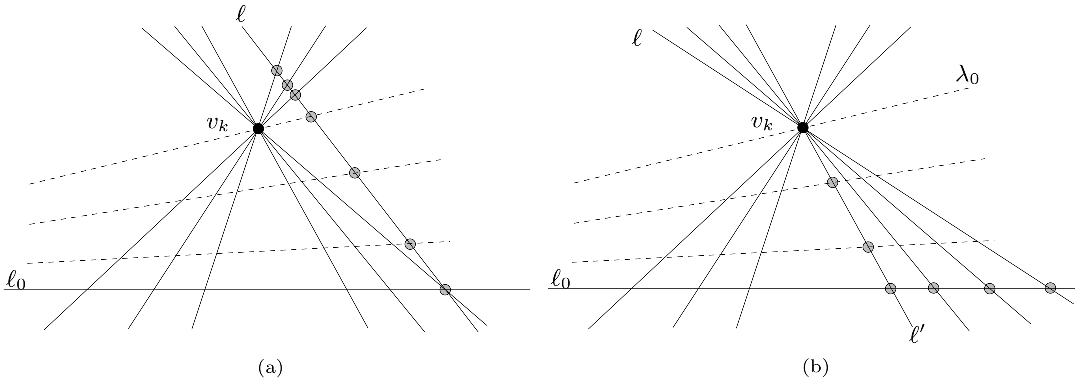

Figure 13: Collecting termini.

the right of v_k and to the left of any point q'_j charged by the lines of Λ^- incident to v_k. Altogether we collect at least $d_k + |\Lambda_k^e|$ termini. The only case in which we do not obtain $d_k + |\Lambda_k^e| + 1$ termini is when there is an excess line λ_0 through v_k. In this case we must have used the upper grid between v_{k-1} and v_k, which is argued as in case (ii) of the preceding analysis. As above, we can gain an extra excess station in this upper grid, because the number of lines of Λ^- through v_k is in fact at least $d_k + 1$. Again, the same technical difficulty that we faced earlier may arise here as well, when v_{k-1} also lies on ℓ. We resolve this exactly as before, backtracking to the left through junctions v_j that lie on ℓ, switch to lower grids between them without decreasing the number of collected stations, and gaining the desired extra station when we reach a junction v_j that lies strictly below ℓ or that is not incident to an excess line.

In both cases, we have managed to charge an extra terminus for every excess line left in Λ_k^e, and an additional terminus for the extra surviving line ℓ of Λ^-. Note that all termini, or all but one, lie to the right of v_k.

Adding these termini to the junctions and stations, we obtain, excluding the excess stations and termini, and recalling that $d_1 = 1$, a total of

$$(d_1 + d_2 - 1) + (d_2 + d_3 - 1) + \cdots + (d_{k-1} + d_k - 1) + k + (d_k + 1) =$$

$$1 + 2\sum_{i=1}^{k} d_i = 2n^- + 1$$

vertices. Hence, since we manage to collect one additional vertex for each excess line, we obtain a total of $2n^- + 1 + (n^+ - n^-) = n^+ + n^- + 1$ vertices. Observe that all the collected vertices are either on ℓ_0 or are intersection points of lines of (the original) Λ^+ with lines of (the original) Λ^-. In other words, each of the collected vertices represents a segment in the primal plane, connecting a point of $X^+ \cup \{p_0\}$ to a point of $X^- \cup \{p_0\}$.

Let Q denote the set of all collected junctions, stations, and termini. Associate with each element $q \in Q$ the *maximal* double wedge $W(q)$ (not containing the vertical line through q), which is bounded by a pair of lines passing through q.

To complete the proof of the theorem, we show that the collected segments are pairwise non-convergent.

Claim 3.4. *The set $\{W(q) \mid q \in Q\}$ of n double wedges has no two convergent elements.*

Proof: Let $u, v \in Q$ with u lying to the left of v. Recalling the definition of convergent double wedges, we need to show that either $u \in W(v)$ or $v \in W(u)$. We distinguish between several cases:

Case A: *Both u and v are junctions.*

Put $u = v_i$ and $v = v_j$, with $i < j$. Then $W(v)$ is bounded by a line $\ell \in \Lambda_j^+$ and by a line $\ell' \in \Lambda_j^-$. By Claim 3.3(ii), v_i lies between these two lines, and thus belongs to $W(v)$.

Case B: *u is a junction and v is a (standard or excess) station to the left of v_k.*

Put $u = v_i$ and let S_j be the set of stations that contains v, where $i \le j < k$. Then $W(v)$ is bounded by two lines ℓ, ℓ', where either $\ell \in \Lambda_j^+ \cup \Lambda_j^e$ and $\ell' \in \Lambda_{j+1}^-$ (if v lies on the upper grid), or $\ell \in \Lambda_{j+1}^+ \cup \Lambda_j^e$ and $\ell' \in \Lambda_j^-$ (if v lies on the lower grid). By construction, we have, in both cases, $\ell \in \Lambda_j^+ \cup \Lambda_j^e$ and $\ell' \in \Lambda_j^-$. If $\ell \in \Lambda_j^+$, the analysis is completed as in Case A. If $\ell \in \Lambda_j^e$, it passes through or below v_i, so the same analysis applies here as well.

Case C: *u is a (standard or excess) station to the left of v_k and v is a junction or a station to the left of v_k.*

Let S_i be the set of stations containing u. The arguments in Case A and Case B imply that $v_i \in W(v)$. If v is also a station in S_i or $v = v_{i+1}$ then it is easy to verify, by construction, that $W(u)$ and $W(v)$ are non-convergent (see Figure 9); this also holds if u and/or v are excess stations. Suppose then that v lies to the right of v_{i+1}. Consider first the case where u is a standard station. Then both v_i and v_{i+1} lie in the left wedge of $W(v)$, and u is incident to a line λ that passes through v_i and to a line λ' that passes through v_{i+1}. If $u \notin W(v)$ then a boundary line of $W(v)$ must separate u from v_i and v_{i+1}, in which case $v \in W(u)$; see Figure 14(a).

Suppose next that u is an excess station on the upper grid between v_i and v_{i+1}. If $u \notin W(v)$ then u must lie *above* $W(v)$. In this case u is incident to a line λ (an excess line) that passes through or below v_i and to a line λ' that passes through v_{i+1}. As above, it is easily seen that the line through v that bounds the left wedge of $W(v)$ from above must cross λ to the left of u and λ' to the right of u and to the left of v, again implying that $v \in W(u)$; see Figure 14(b).

A fully symmetric argument applies when u is an excess station on the corresponding lower grid.

Note that cases B and C also apply to excess stations constructed in the backtracking processes, starting either from some junction that precedes v_k or from v_k itself.

Case D: *u is a junction and v is a terminus to the right of v_k.*

Refer to Figure 13 to recall the types of termini that we construct. Consider first the case where v is the intersection point of an excess line λ that passes through or below v_k, with either the line ℓ (in the case depicted in Figure 13(a)), or the line ℓ' (in the case depicted in Figure 13(b)). By construction, λ passes through or below u and ℓ or ℓ' passes through or above u, so $u \in W(v)$.

Consider next the case where v is the intersection of ℓ with some line μ in Λ_k^+. Here too it is easily verified that u lies between the two lines, so $u \in W(v)$. The same argument applies to the last possible case, where v is the intersection of ℓ_0 with some line of Λ_k^-.

Case E: *u is a station and v is a terminus to the right of v_k.*

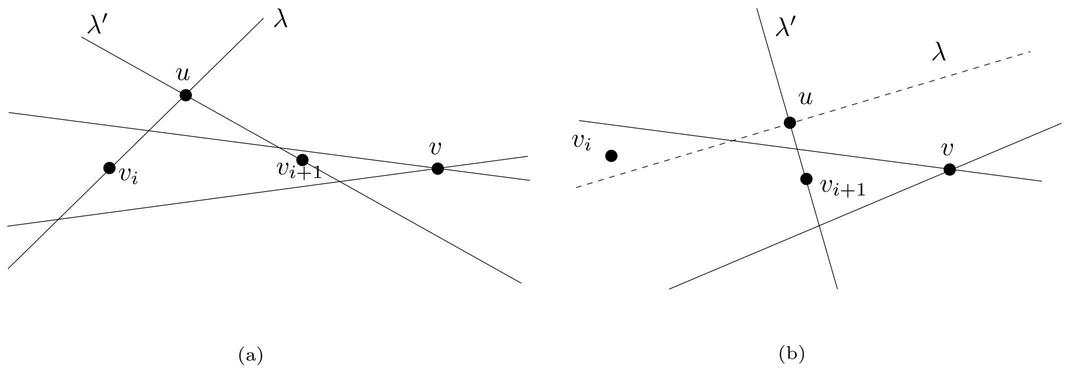

Figure 14: Illustrating Case C of the proof that $W(u)$ and $W(v)$ cannot be convergent. (a) u is a standard station. (b) u is an excess station on the upper grid.

Let S_i be the set of stations containing u. The arguments in Case D imply that $v_i, v_{i+1} \in W(v)$. If $u \notin W(v)$ then, arguing as in case C, we must have $v \in W(u)$.

Case F: Both u and v are termini to the right of v_k.

This case follows from a direct inspection of all the possible types of pairs of termini; see Figure 13.

This completes the proof of the claim, and thus of Theorem 3.1. $\square$

4 Constructing the Sets of Segments F and E in the Plane π and in 3-Space

Consider the projected set R of non-collinear points in the plane π, as defined in Section 2, and recall that we assume that its total weight $w(R) = n-1$ is *even*. Recall also that we have partitioned R into two sets, R^+ and R^-, by some vertical line which we choose to be the y-axis. Instead of selecting the suitable set of segments F in π, it will be more convenient to work in the dual plane, using the same duality transform as in the proof of Theorem 3.1, where segments correspond to double wedges. First, we will define the apices v of these double wedges $W(v)$, that are vertices in the arrangement of lines dual to the elements of R, and then we specify the boundary lines of each $W(v)$, which are the duals of the endpoints of the corresponding segment $f(v)$ in the 'primal' plane π.

The main part of the selection algorithm is an iterative pruning process that collects two types of different crossing points v, so-called *junctions* and *stations*, between the lines dual to the points of R. This process has many aspects similar to the one discussed in [14], and to the one given in the preceding section, but here the analysis is considerably more involved, because we have to handle *weighted lines*, and because the potential presence of the central bichromatic point q further complicates certain steps of the analysis.

After associating each collected vertex v with a certain double wedge $W(v)$ that has v as an apex, we consider the set F of segments $f(v)$ in the primal plane π that correspond to these double wedges, by duality. Each segment $f(v)$ connects two elements of R in π, and we show that these segments are pairwise non-convergent. Each segment $f(v) \in F$ spans with p_0 a plane $h(v)$ in $\mathbb{R}^3$, and we apply Theorem 3.1 to collect segments that connect pairs of points of P within $h(v)$. We denote by $E(f(v)) = E(v)$ (and sometimes also by $E(f)$) the set of segments in $\mathbb{R}^3$ that are spanned by P and are determined (in this manner) by $f(v)$, and we set $E := \bigcup_{f \in F} E(f)$.

4.1 Collecting junctions in the dual plane

Denote by L the set of lines dual to the elements of R. By choosing the directions of the coordinate frame sufficiently generic, we may assume that no two lines in L are parallel. (In the primal plane π, this would correspond to the requirement that no two points of R have the same x-coordinate.) Each line $\ell \in L$ has a weight $w(\ell)$ equal to the weight of its dual point, so $\sum_{\ell \in L} w(\ell) = n-1$. Let L^+, L^- denote respectively the sets of lines dual to R^+ and R^-. Since we have assumed in Section 2 that R^+ lies to the left of the y-axis and R^- lies to its right, it follows that all lines in L^- have negative slopes and all lines in L^+ have positive slopes. The central bichromatic point q, when it exists, is mapped to a *horizontal* line q^*, which we assume to be the x-axis. This line appears as two coincident copies, $(q^+)^* \in L^+$ and $(q^-)^* \in L^-$, with corresponding weights $w(q^+)$, $w(q^-)$.

We begin by constructing a sequence $J = \langle v_1, v_2, \ldots, v_k \rangle$ of vertices of $\mathcal{A}(L)$, called *junctions*.

Step 1: Set $i := 1$ and $L_i^+ := L^+$, $L_i^- := L^-$.

Step 2: If $L_i^+ = L_i^- = \emptyset$, the construction of J terminates. Otherwise, as we will see, neither set is empty. Let v_i be the *leftmost* intersection point between a line in L_i^+ and a line in L_i^-. Let d_i^+, d_i^- denote the overall weight of those lines of L_i^+, L_i^-, respectively, that are incident to v_i, and put $d_i := \min\{d_i^+, d_i^-\}$. Suppose, without loss of generality, that $d_i = d_i^+$. Remove from L_i^+ all its lines incident to v_i, and prune L_i^- as follows. Remove as many of the *steepest* lines of L_i^- (those with the smallest slopes) incident to v_i as possible, so that their overall weight c_i does not exceed d_i. If this weight is equal to d_i, we are done. Otherwise, we take the next steepest line ℓ and *reduce its weight* by $d_i - c_i$. The line ℓ *is not removed* from L_i^-. Note that each of the remaining lines of L_i^- incident to v_i separates the removed lines of L_i^+ from the removed lines of L_i^-. See Figure 15. Set L_{i+1}^+ and L_{i+1}^- to be the sets of surviving weighted lines of L_i^+ and L_i^-, respectively, where the line ℓ, if exists, has its new reduced weight. Set $i := i+1$ and repeat Step 2.

Since m_1 is the line with the *largest* slope connecting a point of R^+ and a point of R^-, our duality implies that m_1^*, the dual of m_1, is the *leftmost* intersection point

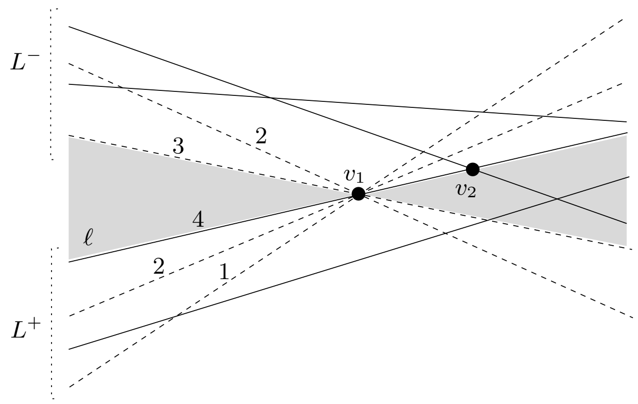

Figure 15: Choosing the first junction v_1. Lines are labeled with their weight. We have $d_1 = 5$. The dashed lines, two from L^+ and two from L^-, are removed, and the remaining line ℓ has its weight reduced by 2 at v_1. The double wedge $W(v_1)$ is shaded. The next junction v_2 is also shown.

between a line of L^+ and a line of L^-. Hence, we have $v_1 = m_1^*$.

If q exists, then $v_1 = m_1^*$ is the leftmost vertex along the line q^* (see Figure 3(ii)). At least one of the coincident copies $(q^+)^*$, $(q^-)^*$ of q^* contributes its full weight to d_1. Consequently, at least one of these copies is removed at v_1, which implies that q^* belongs from this point on to only one of the sets L_i^+, L_i^-. In other words, the presence of q will only affect the construction "at the vicinity" of v_1; see below for details.

As our construction sweeps the dual plane from left to right, we collect junctions (and stations) whose dual lines rotate clockwise from m_1 onwards (see Figure 3).

As in the proof of Theorem 3.1, it is easy to verify the following properties of the above construction (consult Figure 15):

Claim 4.1. *(i)* $w(L_i^+) = w(L_i^-)$, *for each* $i = 1, \ldots, k$.

(ii) For every $1 \le i < j \le k$, *the junction* v_i *lies in the left unbounded face* f_j *of* $\mathcal{A}(L_j^+ \cup L_j^-)$ *that separates* L_j^+ *and* L_j^- *at* $x = -\infty$, *and whose rightmost vertex is* v_j. *The point* v_i *lies in the interior of* f_j *if* $d_i^+ = d_i^-$; *otherwise it may lie on the boundary of* f_j.

(iii) $\sum_{i=1}^{k} d_i = (n-1)/2$. $\square$

(iv) At the time when v_i *is constructed, the weights of all lines that are removed or weight-reduced at* v_i, *are equal to their original weights (i.e., before being reduced at any preceding junction), with the only possible exception of the two* shallowest *lines in their respective sets, whose weights could have been earlier reduced.*

To see (iv), let ℓ^+ be a line of L_i^+ that is removed at v_i and is different from the shallowest such line ℓ_a^+. Then, by property (ii), ℓ^+ must pass *strictly below* each of the previously constructed junctions, so it did not participate in any preceding pruning step. The argument for L_i^- is fully symmetric. $\square$

We define, for each $1 \le i \le k$, the set of lines of L_i^+ (resp., L_i^-) that are incident to v_i and are either removed at v_i or have their weight reduced there, by D_i^+ (resp., D_i^-). We also put $D_i := D_i^+ \cup D_i^-$.

We associate with each junction v_i the double wedge $W(v_i)$ bounded by the *shallowest* lines in D_i^+ and D_i^-, respectively. See Figure 15.

4.2 Constructing segments in E from the junctions

In the primal plane π, each junction v_j, for $j = 1, \ldots, k$, corresponds to some line v_j^* in π, which contains projections (from p_0) of some points of P. Let h denote the plane spanned by v_j^* and p_0. We apply Theorem 3.1 to a certain subset of $P \cap h$, thereby obtaining a set of pairwise non-convergent segments determined by the points in that subset.

The presence of the central bichromatic point q may force us to modify the analysis at v_1. We first describe the analysis under the assumption that q does not exist (we sometimes refer to this situation as the *standard case*), and then discuss the modifications needed at v_1 when q exists.

The case where q does not exist. Fix an index $1 \le i \le k$. Let $\ell_1^+, \ldots, \ell_a^+$ denote the lines of D_i^+, and let $\ell_1^-, \ldots, \ell_b^-$ denote the lines of D_i^-. We enumerate the lines in the order of their slopes, so that ℓ_a^+ and ℓ_b^- are the innermost (shallowest) lines of D_i^+ and D_i^-, respectively. Consider the line v_i^* dual to v_i, and let h be the plane spanned by v_i^* and p_0. Let X^+ (resp., X^-) denote the set of points of $P \setminus \{p_0\}$ whose projection from p_0 is one of the duals of $\ell_1^+, \ldots, \ell_a^+$ (resp., $\ell_1^-, \ldots, \ell_b^-$) on π. By construction, $|X^+|, |X^-| \ge d_i$; either of $|X^+|$ and $|X^-|$ may consist of more than d_i points, in the case when either L_i^+ or L_i^- has a line whose weight is reduced (at v_i or in some earlier junction). By Claim 4.1(iv), only ℓ_a^+ and ℓ_b^- (which are the innermost lines of D_i^+ and D_i^-) may have reduced weight. By construction, the sets X^+ and X^- are separated by a line in h (consult with Figure 7(a)).

Clearly, $|X^+| = \sum_{j=1}^{a} w(\ell_j^+)$ and $|X^-| = \sum_{j=1}^{b} w(\ell_j^-)$, where the $w(\ell)$'s denote the *original weights* of the corresponding lines ℓ. We claim that Theorem 3.1 can be applied to the set $X = X^+ \cup X^-$ within the plane h. Indeed, assume without loss of generality that $|X^+| \ge |X^-|$. Since ℓ_a^+ is either deleted at v_i or has its weight reduced there, it follows that $\sum_{j=1}^{a-1} w(\ell_j^+) < \sum_{j=1}^{b} w(\ell_j^-)$. Therefore

$$|X^+| - |X^-| = \sum_{j=1}^{a} w(\ell_j^+) - \sum_{j=1}^{b} w(\ell_j^-) < w(\ell_a^+).$$

Theorem 3.1 can thus be applied to the set $X = X^+ \cup X^-$ within the plane h, and it yields a total of at least $|X^+| + |X^-| + 1$ pairwise non-convergent segments, each of which connects a point of $X^+ \cup \{p_0\}$ to a point of $X^- \cup \{p_0\}$. However, there may exist one segment that has to be excluded because of potential collinearity with segments generated at other junctions: This is a segment e along the ray from p_0 to the dual of the unique line ℓ among ℓ_a^+, ℓ_b^- whose weight is reduced at v_i but which is not removed there, if such a line exists. (Note that this ray is the innermost among those rays connecting p_0 to points of the corresponding set X^+ or X^-.) In that case, ℓ will also contribute weight to another subsequent junction $v_{i'}$, where a segment collinear with e may be generated in the primal plane, and these two segments cannot both be included in the output set E (whose elements have to be pairwise non-convergent). Reducing the count due to this potential double counting, we are therefore left with at least

$$w(D_i^+) + w(D_i^-) + 1 - \zeta_i = \sum_{j=1}^{a} w(\ell_j^+) + \sum_{j=1}^{b} w(\ell_j^-) + 1 - \zeta_i$$

pairwise non-convergent segments, where $\zeta_i = 1$ if there is a line whose weight has been reduced at v_i but which was not removed there, and $\zeta_i = 0$ otherwise. Here $w(D_i^+)$, $w(D_i^-)$ denote the total *original* weight of these sets.

Handling the central bichromatic point q. As noted, the presence of q may force us to modify the analysis at the first junction v_1, because the dual line q^* appears there as two coincident lines $(q^-)^* \in L_1^-$ and $(q^+)^* \in L_1^+$. Let d_0^+ (resp., d_0^-) denote the total weight of all the lines of $D_1^+ \setminus \{(q^+)^*\}$ (resp., of $D_1^- \setminus \{(q^-)^*\}$); recall that at least one of the sets D_1^+, D_1^- includes the respective copy of q^* with its *full* weight. We have $d_1 = \min\{d_0^+ + w(q^+),\ d_0^- + w(q^-)\}$; assume, without loss of generality, that $d_1 = d_0^+ + w(q^+) \le d_0^- + w(q^-)$.

Suppose first that $d_0^- \ge d_0^+ + w(q^+)$. Refer to Figure 16. Then $(q^-)^* \notin D_1^-$. Let X^- denote the set of all points $p \in P \setminus \{p_0\}$ that project to the points dual to the lines of D_1^-. If D_1^- contains a line ℓ whose weight is only reduced at v_1, let $b_1 > 0$ denote the surviving weight of ℓ. We have $|X^-| = d_1 + b_1$. Let X^+ denote the set of all points $p \in P \setminus \{p_0\}$ that project to the points dual to the lines of D_1^+, *including* q (with its full weight $w(q) = w(q^-) + w(q^+)$). We have $|X^+| = d_1 + w(q^-)$. If $b_1 \le w(q^-)$, then $|X^+| - |X^-| = w(q^-) - b_1 < w(q)$; the right-hand side is the number of points on the innermost ray from p_0 to the points of X^+ (see Figure 16). If $b_1 > w(q^-)$, then $|X^-| - |X^+| = b_1 - w(q^-) < w(\ell)$; the right-hand side is the number of points on the innermost ray from p_0 to the points of X^-. Hence, in either case, Theorem 3.1 is applicable to $X^+ \cup X^- \cup \{p_0\}$, and it yields a set $E(v_1)$ of at least

$$w(D_1^-) + w(D_1^+ \setminus \{q\}) + w(q) + 1 = w(D_1^-) + w(D_1^+) + w(q^-) + 1$$

pairwise non-convergent segments, where, as above, each line in $D_1^- \cup D_1^+$ is taken with its full original weight. Compared with the count in the standard case, we collect $w(q^-)$ additional segments in this case.

Suppose next that $d_0^- < d_0^+ + w(q^+) \le d_0^- + w(q^-)$. See Figure 17. In this case, D_1^- contains $(q^-)^*$ and D_1^+ contains $(q^+)^*$. We let X_0^+ (resp., X_0^-) denote the set

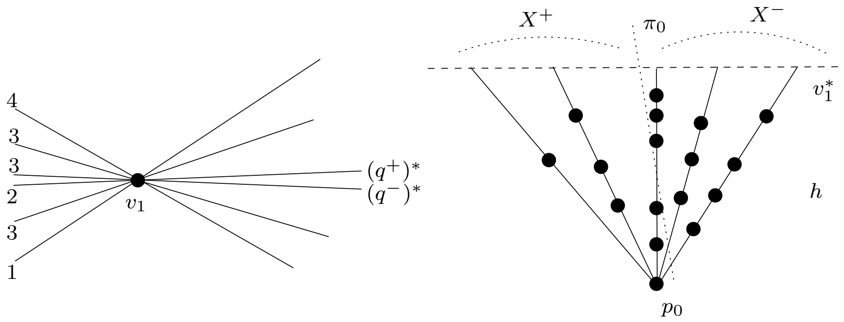

Figure 16: Collecting segments in $E(v_1)$ when q exists. Here $(q^-)^* \notin D_1^-$. We have $d_1 = 6$, $|X^-| = 7$ and $|X^+| = 9$. The lines $(q^-)^*$ and $(q^+)^*$ are coincident, but are drawn as separate lines for the purpose of illustration.

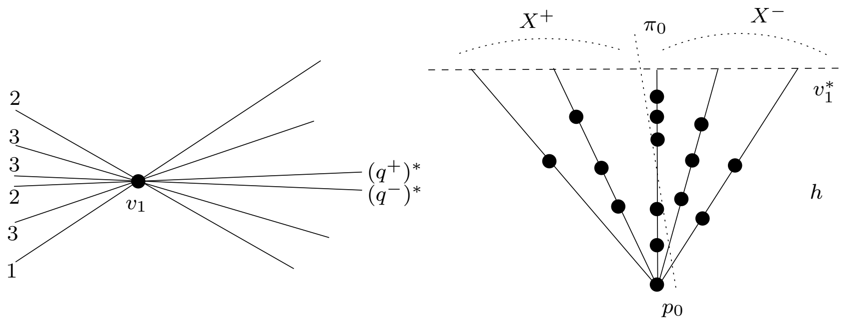

Figure 17: Collecting segments in $E(v_1)$ when q exists. Here $(q^-)^* \in D_1^-$, $d_1 = 6$, $|X^-| = 5$, $|X^+| = 9$, and Theorem 3.1 can be applied.

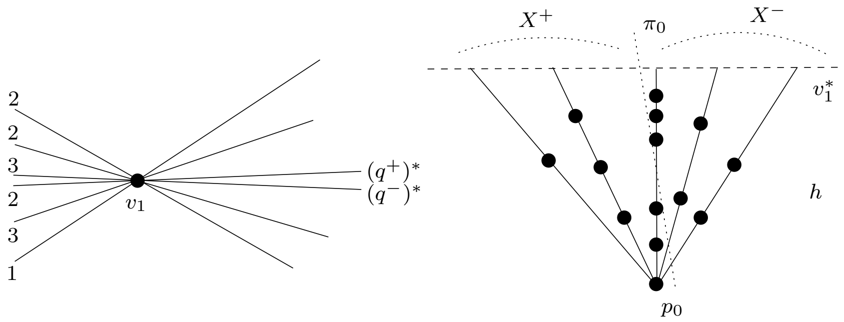

Figure 18: Collecting segments in $E(v_1)$ when q exists. Here $(q^-)^* \in D_1^-$, $d_1 = 6$, $|X^-| = 4$, $|X^+| = 9$, and $|X^+| - |X^-| = w(q)$, so only Corollary 3.2 can be applied.

of all points of $P \setminus \{p_0\}$ that project to the points dual to the lines of $D_1^+ \setminus \{(q^+)^*\}$ (resp., the lines of $D_1^- \setminus \{(q^-)^*\}$). We have $|X_0^+| = d_0^+$ and $|X_0^-| = d_0^-$.

First, assume further that $d_0^+ \neq d_0^-$, say $d_0^+ > d_0^-$. In this case we set $X^+ := X_0^+$ and take X^- to be the union of X_0^- with the set of all points of $P \setminus \{p_0\}$ that project to q. We have $|X^-| = d_0^- + w(q) > |X^+|$, and $|X^-| - |X^+| = d_0^- - d_0^+ + w(q) < w(q)$. A symmetric argument holds when $d_0^+ < d_0^-$. Hence, Theorem 3.1 is again applicable, and yields a set $E(v_1)$ of at least

$$d_0^- + d_0^+ + w(q) + 1 \geq w(D_1^+) + w(D_1^-) + 1$$

pairwise non-convergent segments, where, as above, each line in $D_1^- \cup D_1^+$ is taken with its full original weight. Here the lower bound is the same as the one yielded in the standard case.

The final, problematic case arises when $d_0^+ = d_0^-$. See Figure 18. In this case, the points of $P \setminus \{p_0\}$ that project to q can be added to either set X_0^-, X_0^+, say we add them to X_0^-. Then the resulting sets X^-, X^+ satisfy $|X^+| = d_0^+$, $|X^-| = d_0^- + w(q)$, and $|X^-| - |X^+| = w(q)$. In this case Theorem 3.1 is *not* applicable, and we can only apply Corollary 3.2, to obtain a set $E(v_1)$ with at least

$$d_0^- + d_0^+ + w(q) \geq w(D_1^+) + w(D_1^-)$$

pairwise non-convergent segments. That is, *we lose one segment* in $E(v_1)$, as compared with the standard situation. (Note that in this case *all* lines through v_1 are removed, except perhaps for $(q^-)^*$.)

In addition, as in the standard case, we need to subtract 1 from any of the bounds obtained above, in case D_1^+ or D_1^- has a line whose weight is only reduced at v_1, to accommodate the potential double counting due to collinear segments generated at subsequent junctions.

The double wedge $W(v_1)$ associated with v_1 is defined as in the standard case, except that in some of the above cases it may degenerate to the single line q^* (it always has $(q^-)^*$ or $(q^+)^*$ as one of its bounding lines). In this case, we still consider $W(v_1)$ to have its apex at v_1. In the primal plane, the corresponding segment $f(v_1)$ degenerates to the singleton point q, but it is still considered to lie along the line v_i^*.

Wrapping up. We repeat this collection process to each of the junctions v_i, and sum up the resulting bounds. This sum can be rearranged as follows. Let $\ell_1, \ldots, \ell_t$ denote an enumeration of all the lines in L, and put $w_j = w(\ell_j)$ (the original weight), for $j = 1, \ldots, t$. (In case q exists, the lines $(q^+)^*$ and $(q^-)^*$ appear as two separate lines in this enumeration, with their respective weights.) For each j, let κ_j denote the number of junctions v_s that are incident to ℓ_j, such that $\ell_j \in D_s$. Observe that if $\kappa_j > 1$, then in the first $\kappa_j - 1$ of these junctions v_s, the weight of ℓ_j is reduced at v_s but ℓ_j is not removed there. Hence $\zeta_s = 1$ at each of these junctions v_s, and we "blame" this reduction in the count on ℓ_j, making its effective weight contribution at v_s equal

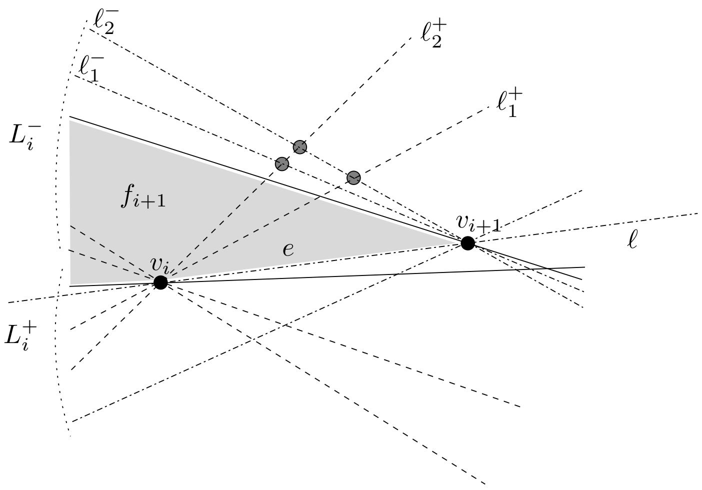

Figure 19: Collecting stations (shown highlighted) between v_i and v_{i+1}. The dashed lines are those removed at v_i, and the dashed-dotted ones are those removed at v_{i+1}.

to $w_j - 1$. ℓ_j is removed only at the last (i.e., the κ_j-th) of these junctions. Therefore, the overall number of segments in E generated at all the junctions $v_1, \ldots, v_k$ is at least

$$k + \sum_{j=1}^{t} (\kappa_j - 1)(w_j - 1) + \sum_{j=1}^{t} w_j - \varepsilon_0 = k + t + \sum_{j=1}^{t} \kappa_j (w_j - 1) - \varepsilon_0, \qquad (1)$$

where $\varepsilon_0 = 1$ if q exists and the problematic case $d_0^- = d_0^+$ arises at v_1, and $\varepsilon_0 = 0$ in all other cases.

4.3 Collecting stations in the dual plane and corresponding segments in E

In the next step we collect additional vertices, called *stations*, between pairs of successive junctions v_j, v_{j+1}. We first handle the standard case, in which either q does not exist, or q exists and $j \geq 2$, and then present a modified analysis for the case where q exists and $j = 1$.

The standard case. Fix an index $1 \leq i < k$, and consider the vertical slab between v_i and v_{i+1}. By Claim 4.1(ii), v_i lies inside or on the boundary of the face $f = f_{i+1}$ of $\mathcal{A}(L_{i+1}^+ \cup L_{i+1}^-)$ whose rightmost vertex is v_{i+1}; see Figure 19. Hence, the segment $e = v_i v_{i+1}$ is contained in (the closure of) f. We distinguish two cases:

Case 1: *e is contained in the interior of f (except for its right endpoint).*

This implies that the lines of D_{i+1}^- (resp., of D_{i+1}^+) pass strictly above (resp., below) e. Moreover, either all the lines of D_i^- pass below e, or all the lines of D_i^+ pass

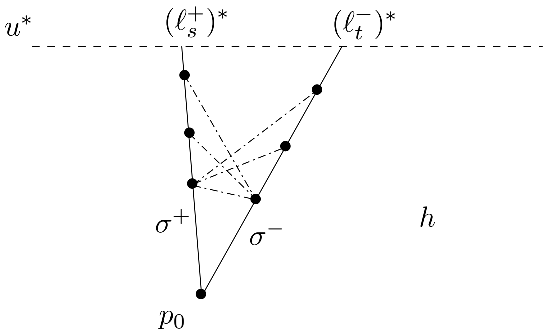

Figure 20: The set $E(u)$ of segments (drawn dashed-dotted) spanned by P that are determined by a station u.

above e. Suppose, without loss of generality, that the second case arises. Denote the lines of D_i^+ by $\ell_1^+, \ldots, \ell_{\nu_i}^+$, ordered according to *increasing* slope, and those of D_{i+1}^- by $\ell_1^-, \ldots, \ell_{\nu_{i+1}}^-$, ordered according to *decreasing* slope. See Figure 19 (which depicts this configuration, even though it illustrates the following Case 2).

Each of the lines ℓ_s^+ intersects every line ℓ_t^- in the slab between v_i and v_{i+1}, because ℓ_s^+ passes through the left endpoint of e, ℓ_t^- passes through the right endpoint of e, and they both lie above e. We refer to the points of intersection between these two sets of lines as the *upper grid* between v_i and v_{i+1}; the *lower grid* is defined analogously. Consider the vertices of $\mathcal{A}(L)$ where $\ell_{\nu_i}^+$ intersects the lines $\ell_1^-, \ldots, \ell_{\nu_{i+1}}^-$, and the vertices where $\ell_{\nu_{i+1}}^-$ intersects the lines $\ell_1^+, \ldots, \ell_{\nu_i}^+$. There are $\nu_i + \nu_{i+1} - 1$ distinct vertices of this kind (see Figure 19), and we let the set of stations S_i to consist of all these vertices. We associate with each station u the double wedge $W(u)$ between the two lines from D_i^+ and D_{i+1}^- that meet at u.

Each station u generates a set $E(u)$ of segments spanned by P in $\mathbb{R}^3$, as follows. Suppose that u is incident to some line ℓ_s^+ through v_i and to some line ℓ_t^- through v_{i+1} (where either $s = \nu_i$ or $t = \nu_{i+1}$). Consider the primal line u^* dual to u, and let h denote the plane in 3-space spanned by p_0 and u^*. The plane h contains two segments that connect p_0 to the two respective dual points $(\ell_s^+)^*$, $(\ell_t^-)^*$, both lying on u^*. The first segment σ^+ contains $w(\ell_s^+)$ points of R^+, and the second segment σ^- contains $w(\ell_t^-)$ points of R^-. We can easily collect here as many as $w(\ell_p^+) + w(\ell_q^-) - 1$ segments into $E(u)$, no two of which are convergent; for example, one can get that many distinct segments by taking all segments one of whose endpoints is either the nearest point to p_0 on σ^+ or the nearest point to p_0 on σ^-; See Figure 20.

These segments constitute the set $E(u)$. Hence, the total number of segments that are collected in this manner for all the new stations u is

$$\left|\bigcup_{u\in S_i} E(u)\right| = \sum_{t=1}^{\nu_{i+1}} \left[w(\ell_{\nu_i}^+) + w(\ell_t^-) - 1\right] + \sum_{s=1}^{\nu_i - 1} \left[w(\ell_s^+) + w(\ell_{\nu_{i+1}}^-) - 1\right].$$

Note that the sum $\sum_{s=1}^{\nu_i} w(\ell_s^+)$ is at least d_i; it may exceed d_i if it involves a non-deleted line with reduced weight, because in the sum we use the full weight of that line. Similarly, $\sum_{t=1}^{\nu_{i+1}} w(\ell_t^-) \geq d_{i+1}$. Therefore, the total number of segments that we collect this way is at least

$$d_i + d_{i+1} - 1. \tag{2}$$

We note that this estimate is rather conservative. In general, if the weights of the lines are greater than 1 and $\nu_i, \nu_{i+1} > 1$, we get a larger lower bound.

Case 2: *e is an edge of f.*

In this case, e is contained in a line ℓ which is incident to v_i but which was not removed when v_i was constructed (it could have been the one whose weight has been reduced there). Assume first that ℓ is not the line whose weight has been reduced at v_i. By construction, it then follows that the lines of D_i^- pass strictly below e, and the lines of D_i^+ pass strictly above e. Now either all the lines of D_{i+1}^- pass above e, or all the lines of D_{i+1}^+ pass below e. We can now repeat the preceding arguments, and obtain, as above, a set S_i of stations of $\mathcal{A}(L)$ along either the upper or the lower grid, which generate a total of at least $d_i + d_{i+1} - 1$ segments spanned by P, which are added to E. Figure 19 depicts this case of the analysis.

Suppose next that the line ℓ containing e is the (unique) weight-reduced line at v_i. If ℓ does not belong to D_{i+1}, then the first case of the analysis applies, and yields the same lower bound of $d_i + d_{i+1} - 1$ on the number of collected segments that are added to E. We thus assume that ℓ does belong to D_{i+1}.

Let a_i and a_{i+1} denote the contribution of ℓ to d_i and d_{i+1}, respectively. That is, the overall weight of the lines from the same family of ℓ (i.e., L^+ or L^-) that are *removed* at v_i (resp., at v_{i+1}) is $c_i = d_i - a_i$ (resp., $c_{i+1} = d_{i+1} - a_{i+1}$).

Claim 4.2. *In this case one can construct stations along either the upper or the lower grid between v_i and v_{i+1}, from which at least*

$$c_i + c_{i+1} = d_i + d_{i+1} - (a_i + a_{i+1}) \tag{3}$$

new segments can be collected in E (in the same manner as before).

Indeed, suppose, without loss of generality, that $\ell \in L^+$. Then the total weight of the lines of L^- that are incident to v_i (resp., to v_{i+1}) is d_i (resp., d_{i+1}), and the total weight of the lines of L^+ that are incident to v_i (resp., to v_{i+1}) and are removed there is c_i (resp., c_{i+1}). See Figure 21.

If both c_i and c_{i+1} are 0, the claim is trivial, so assume that, say, $c_i > 0$ (see Figure 21). In this case, the upper grid between v_i and v_{i+1} exists, and generates, arguing as above, at least

$$c_i + d_{i+1} - 1 = d_i + d_{i+1} - a_i - 1 \geq d_i + d_{i+1} - (a_i + a_{i+1})$$

new segments in E, as claimed. The case where $c_{i+1} > 0$ (and $c_i = 0$) is fully symmetric, except that in this case we use the lower grid (see Figure 21). This establishes our claim. $\square$

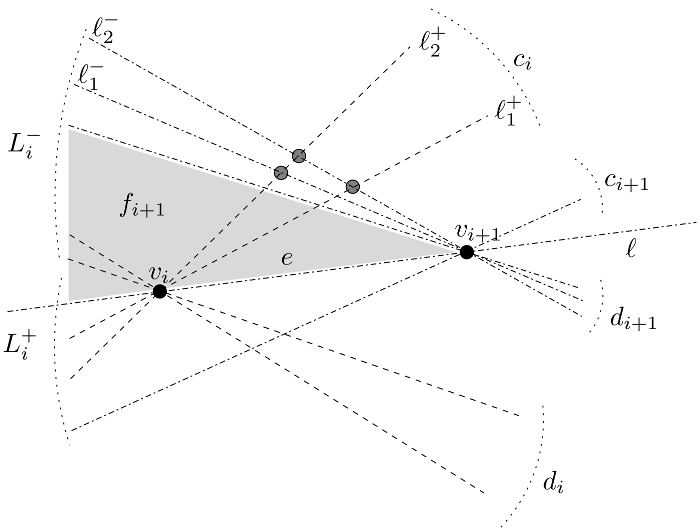

Figure 21: Illustrating the proof of Claim 4.2.

We denote by $E(S_i)$ the set of segments in E that are constructed from the stations collected between the two consecutive junctions v_i and v_{i+1}.

We have thus showed that $|E(S_i)| \geq d_i + d_{i+1} - 1$, if there is no line that contributes weight to both junctions. If on the other hand there is a line ℓ that contributes a weight of $a_i \geq 1$ to v_i and a weight of $a_{i+1} \geq 1$ to v_{i+1}, then $|E(S_i)| \geq d_i + d_{i+1} - a_i - a_{i+1}$.

The case where q exists and $j = 1$. Suppose next that the central bichromatic point q exists, and consider the construction of S_1. If q^* does not pass through v_2, then the analysis proceeds in essentially the same manner as in the standard case. For the sake of completeness, we repeat the details. Suppose, without loss of generality, that q^* passes below v_2. The case where $v_1 v_2$ is interior to the face f_2 is handled in exactly the same way as above: One can use the lower grid between v_1 and v_2, and construct stations that contribute a total of at least $d_1 + d_2 - 1$ segments to E; see Figure 22. The same holds for the case where there is a line (different from q^*) that passes through both v_1 and v_2 but it belongs to at most one of the sets D_1, D_2 (in this case, one of the upper or lower grids will generate stations with $|E(S_1)| \geq d_1 + d_2 - 1$). Consider then the case where there exists such a line ℓ whose weight has been reduced at both junctions (it may have been removed at v_2). Since the entire L^--weight of q must have been removed at v_1 (because no surviving line of L^- can pass below v_2), it follows that $\ell \in L_1^+$. Again, this case can be handled as in Claim 4.2, and yields at least $d_1 + d_2 - a_1 - a_2$ in E, according to the preceding notation. In summary, we can always collect from the stations in S_1 at least either $d_1 + d_2 - 1$ or $d_1 + d_2 - a_1 - a_2$ segments into E, depending on the cases considered above. (Note that in some of these cases we may actually gain $w(q^+)$ additional segments in $E(S_1)$.)

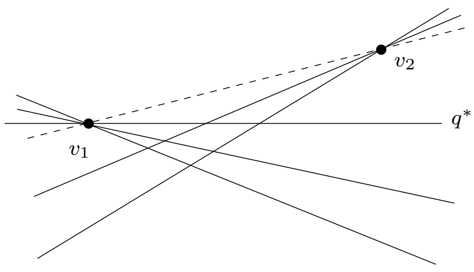

Figure 22: Collecting stations in S_1 when q^* passes below v_2.

Assume then that q^* passes through v_2. Without loss of generality assume that $(q^+)^*$ is fully removed at v_1. The following cases can arise:

Case 1. $(q^-)^* \in D_2^-$.

Let a_1 (resp., a_2) denote the weight removed from $(q^-)^*$ at v_1 (resp., at v_2). (It is possible that $a_1 = 0$.) We claim that one can collect at least $d_1 + d_2 - (a_1 + a_2)$ segments into $E(S_1)$ in either the upper or lower grid between v_1 and v_2. This is argued in much the same way as in the case where q does not exist. Specifically, let $c_1 = d_1 - a_1$ (resp., $c_2 = d_2 - a_2$) denote the L^--weight removed at v_1 (resp., v_2) from the lines in L^- other than $(q^-)^*$.

If both c_1 and c_2 are 0, then there is nothing to prove. If $c_1 > 0$, then, except for $(q^-)^*$, all other lines of D_1^- pass strictly below v_2, and have total weight at least $c_1 = d_1 - a_1$. By assumption, all lines of D_2^+ pass strictly below v_1, and have total weight at least d_2. Therefore the lower grid between v_1 and v_2 can be used to produce at least $c_1 + d_2 - 1 = d_1 + d_2 - 1 - a_1 \geq d_1 + d_2 - (a_1 + a_2)$ segments in $E(S_1)$.

If $c_1 = 0$ and $c_2 > 0$, then, except for $(q^-)^*$, all other lines of D_2^- pass strictly above v_1, and have total weight at least $c_2 = d_2 - a_2$. As $c_1 = 0$, there are no lines from L^- through v_1 other then $(q^-)^*$, and therefore at least one additional line (other than $(q^+)^*$) of L^+ must pass through v_1 (or else v_1 would not be a vertex of the arrangement). Therefore, the upper grid between v_1 and v_2 exists, and we may use it to collect at least $1 + c_2 - 1 = c_2 = d_2 - a_2 = (d_1 - a_1) + d_2 - a_2 = d_1 + d_2 - (a_1 + a_2)$ segments into $E(S_1)$.

Case 2. $(q^-)^* \notin D_2^-$ *(but it still passes through v_2).*

In this case, depicted in Figure 23, the sets D_2^+, D_2^- are both nonempty, and all lines of D_2^+ (resp., D_2^-) pass strictly below (resp., above) v_1; the total weight of either set is d_2. Set $w^+ = w(q^+)$, $w^- = w(q^-)$. There must exist either lines of L^+ or lines of L^- (other than q^*) that pass through v_1 and are removed there. In the former case, the total removed L^+-weight of these lines is $d_1 - w^+$, and we may use the upper grid between v_1 and v_2 (which necessarily exists), to collect at least $d_1 + d_2 - w^+ - 1$ segments into $E(S_1)$. In the latter case, arguing in an essentially symmetric manner, we may use the lower grid between v_1 and v_2 (which necessarily exists), to collect at least $d_1 + d_2 - w_1^- - 1$ segments into $E(S_1)$, where w_1^- is the weight that $(q^-)^*$

contributes at v_1.

As will follow from the later counting phase, given in Section 4.4, we need to compensate for the loss of either w^+ or w_1^- segments in $E(S_1)$, which we do by *including* the points on the ray p_0q in the set X^+ when we construct $E(v_2)$, even though neither $(q^+)^*$ nor $(q^-)^*$ belongs to D_2. Before doing so, the size of X^+ is exactly d_2, and the size of X^- is at least d_2; it can be larger if there is a negative line whose weight is only reduced at v_2. We add the points on the ray p_0q to X^+. In general, we can apply Theorem 3.1 to the modified sets X^-, X^+, except when $|X^-| = |X^+| = d_2$, in which case we can only apply Corollary 3.2. The modified $E(v_2)$ thus consists of at least $w(D_2^+) + w(D_2^-) + w(q) + 1 = 2d_2 + w(q) + 1$ pairwise non-convergent segments, if $|X^-| \neq |X^+|$, or of at least $w(D_2^+)+w(D_2^-)+w(q) = 2d_2+w(q)$ pairwise non-convergent segments, if $|X^-| = |X^+|$. In the most pessimistic scenario, we can only apply Corollary 3.2, whereas, when q was not included, we could have applied Theorem 3.1 to collect $2d_2 + 1$ segments in $E(v_2)$. We thus gain at least $w(q) - 1$ additional segments. However, we may have to subtract 1 extra segment from the count, because $(q^-)^*$ may contribute weight to a further junction. Thus, in the worst case, we can only guarantee $w(q) - 2 = w^+ + w^- - 2$ additional segments. In general, these suffice to compensate for the loss of $\max\{w^+, w_1^-\}$ a $E(S_1)$, unless $\min\{w^+, w_1^-\} = 1$. In this special case, we lose one segment in our count.

The price that we pay for including q is that the double wedge $W(v_2)$ has to shrink, and be bounded by q^* and by the shallowest line in D_2^-. However, as we will later show, in Section 4.5, the collected double wedges will remain pairwise non-convergent.

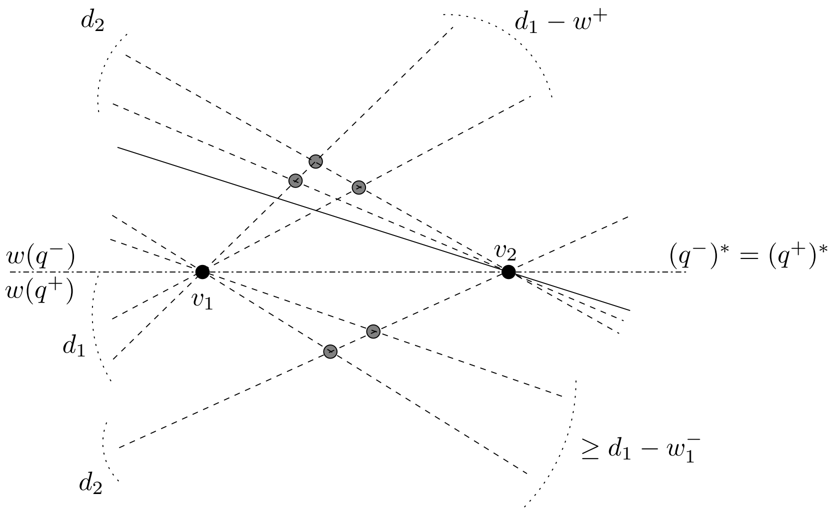

Figure 23: The problematic case in the construction of S_1 in the presence of a central bichromatic point q: Here $d_1 = d_1^+$, so $(q^+)^*$ is removed at v_1. We also assume that $(q^-)^* \notin D_2^-$. In this case we can only guarantee the generation of at least $d_1 + d_2 - \max\{w^+, w_1^-\} - 1$ segments in E, and we lose $\max\{w^+, w_1^-\}$ segments in the bound.

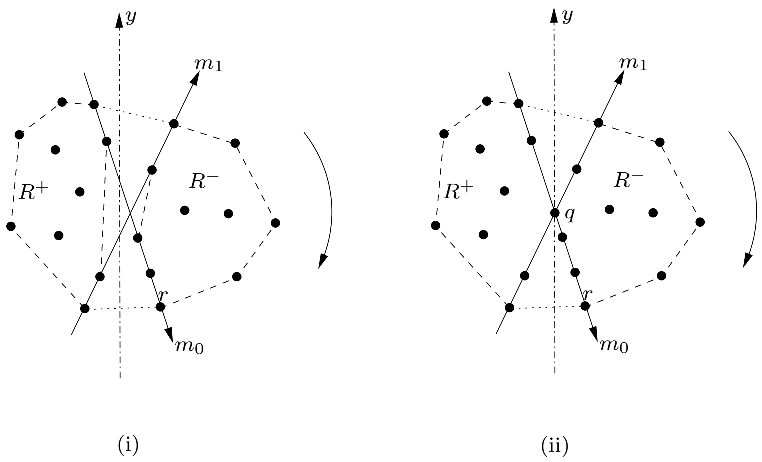

Figure 24: Reproducing the primal construction of R^- and R^+: Case (i) (left) and Case (ii) (right).

Collecting stations to the left of v_1 and to the right of v_k. We next define the last set of stations S_k, which are stations that lie to the right of v_k or to the left of v_1. Recall the specific partition of R into R^+ and R^-, as presented in Section 2. We will exploit certain features of this partition in the construction of S_k, and will find it convenient to "flip" between the primal and dual settings as we go. For the convenience of the reader, we reproduce here Figure 3 as Figure 24.

Claim 4.3. *At least one of the following two conditions will be satisfied:*

(i) The last junction, v_k, is identical to m_0^*, the dual of m_0.

(ii) r^*, the dual of $r \in R^-$, passes through v_k and is the unique element of L^- deleted during the procedure at v_k.

Proof: Suppose that during the procedure r^* is deleted at a junction v_j, for some $j \leq k$. Clearly, v_j^* passes through r and at least one point $t \in R^+$, whose dual line is also deleted, or has its weight reduced, at v_j^*.

If, in the primal plane π, v_j^* passes through another point $r' \neq r$ of R^-, then $v_j^* = m_0$ (otherwise it has to lie clockwise to m_0 and then it cannot meet any point of R^+). In this case, in the dual plane there cannot be any intersection point between a line of L^- and a line of L^+ to the right of v_j, so that $j = k$. That is, we have $v_k^* = m_0$, and (i) holds.

Suppose then that, in the primal plane π, v_j^* does not pass through any element $r' \in R^-$ other than r. If $j = k$, then condition (ii) is satisfied. So we can assume that $j < k$ and $v_k^* \neq m_0$. Refer to Figure 25). Take any two lines $\ell_- \in L^-$ and $\ell_+ \in L^+$ in the dual plane that are deleted during the procedure at the last junction v_k. By assumption and construction, we have $\ell_-^* \neq r$, and the slope of the segment $\ell_+^*\ell_-^* \subset v_k^*$ connecting their duals in the primal plane (i.e., the slope of v_k^*) is smaller than that of the segment tr. We claim that the two segments $\ell_+^*\ell_-^* \subset v_k^*$ and $tr \subset v_j^*$

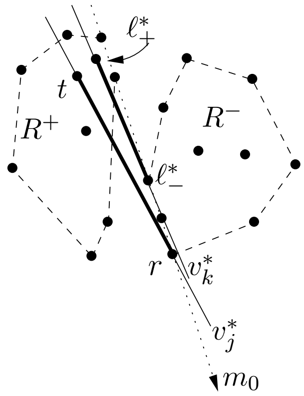

Figure 25: The segments tr and $\ell_+^*\ell_-^*$ must be convergent.

are convergent. Indeed, since m_0 (weakly) separates R^+ and R^-, the closed segment $\ell_+^*\ell_-^*$ must meet m_0, and this must happen at a point to the left of (and above) r, or else r would not be an extreme point of R (see Figure 25). For a similar reason, ℓ_-^* must lie above v_j^*. These facts, together with the slope relationship between v_j^* and v_k^*, imply that the two segments are convergent. This, in turn, implies that the double wedges dual to tr and to ℓ_-^*, ℓ_+^* are convergent. Since $W(v_k)$ and $W(v_j)$ are contained in these respective double wedges, they are also convergent. However, $W(v_k)$ is bounded by a line of L_k^+ and by a line of L_k^-, and Claim 4.1(ii) implies that v_j lies between these lines, and hence in $W(v_k)$, showing that $W(v_j)$ and $W(v_k)$ are non-convergent.[2] This contradiction completes the proof of the claim. $\square$

The above argument is valid for any coordinate system whose y-axis strictly separates the sets R^- and R^+, or, in case q exists, passes through q and strictly separates R_0^- and R_0^+. We specify a coordinate system with this property as follows.

Choose the y-axis to be very close to m_0, so that, in the dual plane, the slope of every line of L passing through m_0^* has smaller absolute value than the slope of any other line of L; that is, the x-coordinates of the points of $m_0 \cap R$ have smaller absolute values than those of any other point of R. See Figures 26(a), 27(a), and 28(a). In addition, if q exists, we make the y-axis pass through q, as already stated.

Now we are in a position to define the set of stations S_k. The reason for choosing the specific way of partitioning R, and the coordinate frame, is to force the stations in S_k to lie to the left of v_1, which will be useful when establishing the non-convergence of the segments in F and in E. With one possible exception, all stations in S_k do indeed lie to the left of v_1.

Pass to the dual plane. The first junction, v_1, lies inside or on the boundary of the face f_k of $\mathcal{A}(L_k^- \cup L_k^+)$, whose rightmost vertex is v_k, so that the segment $e = v_1v_k$ is contained in the closure of f_k. We distinguish the following cases:

[2]This is a special case of a more general argument, given in Lemma 4.4 below.

Case A: *Suppose first that $v_k = m_0^*$ and that the point $c := m_0 \cap m_1$ does not belong to R^-.*

Let $\ell_1^-, \ldots, \ell_{\nu_1}^-$ and $\lambda_1^-, \ldots, \lambda_{\nu_k}^-$ denote the lines of D_1^- and all the lines of $L_k^- = D_k^-$, respectively, listed in the decreasing order of their slopes. By the special choice of our coordinate system, each line ℓ_i^- intersects every line λ_j^- to the left of v_1. Indeed, all the lines of L_k^- pass above or through v_1, by Claim 4.1(ii), but no line passes through both v_k and v_1, because such a line is dual to the point $m_0 \cap m_1$, which we have assumed not to belong to R^-. The slope of the primal segment $(\lambda_j^-)^*(\ell_i^-)^*$ is larger than that of m_1, because, by what has just been argued, $(\lambda_j^-)^* \in m_0$ lies below m_1 and to the left of $(\ell_i^-)^* \in m_1$; see Figure 26(a). Hence the dual intersection point lies to the left of v_1. Define the last set of stations, S_k, as the collection of all intersection points of $\ell_{\nu_1}^-$ with the lines $\lambda_1^-, \ldots, \lambda_{\nu_k}^-$, and all intersection points of $\lambda_{\nu_k}^-$ with the lines $\ell_1^-, \ldots, \ell_{\nu_1}^-$. See Figure 26(b). Clearly, we have $|S_k| = \nu_k + \nu_1 - 1$ such stations, all lying to the left of v_1. Since the total (original) weight of the lines $\ell_1^-, \ldots, \ell_{\nu_1}^-$ is at least d_1, and the total (original) weight of the lines $\lambda_1^-, \ldots, \lambda_{\nu_k}^-$ is at least d_k, it follows, as in the construction of the other sets of stations, that the stations in S_k generate in this case at least $d_1 + d_k - 1$ segments in $E(S_k)$.

Case B: *Suppose next that $v_k = m_0^*$ and that the point $c := m_0 \cap m_1$ does belong to R^-.*

Note that if q exists it must coincide with c. We first consider the case where q does not exist, and then discuss the modifications that are needed when q exists. The dual line c^* passes through both v_1 and v_k. Since we assume for now that q does not exist, c does not belong to R^+. We thus switch to R^+, and collect stations using the dual lines in L^+, in a manner similar to that in case A. All lines in $D_k^+ = L_k^+$ pass strictly below v_1, and the lines of D_1^+ pass strictly above v_k. Arguing exactly as in case A, let $\ell_1^+, \ldots, \ell_{\nu_1}^+$ and $\lambda_1^+, \ldots, \lambda_{\nu_k}^+$ denote the lines of D_1^+ and the lines of $D_k^+ = L_k^+$, respectively, listed in the increasing order of their slopes. The special choice of our coordinate system implies that each line ℓ_i^+ intersects every line λ_j^+ to the left of v_1. Indeed, the slope of the primal segment $(\lambda_j^+)^*(\ell_i^+)^*$ is larger than that of m_1, because $(\lambda_j^+)^* \in m_0$ lies above m_1 and to the right of $(\ell_i^+)^* \in m_1$. Hence the dual intersection point lies to the left of v_1. In this case we define S_k as the collection of all intersection points of $\ell_{\nu_1}^+$ with the lines $\lambda_1^+, \ldots, \lambda_{\nu_k}^+$, and all intersection points of $\lambda_{\nu_k}^+$ with the lines $\ell_1^+, \ldots, \ell_{\nu_1}^+$. Clearly, we have $|S_k| = \nu_k + \nu_1 - 1$ such stations, all lying to the left of v_1, and they generate, as above, at least $d_1 + d_k - 1$ segments in $E(S_k)$.

Case C: *Suppose finally that $v_k \neq m_0^*$.*

In this case, according to Claim 4.3, v_k lies on r^* and $\nu_k = 1$. Refer to Figure 28. Again, let $\ell_1^-, \ldots, \ell_{\nu_1}^-$ denote the lines of D_1^-, listed in the decreasing order of their slopes. In the dual plane, the line r^* passes above v_1 and, by the choice of the coordinate system, it intersects every ℓ_i^- to the left of v_1, with the possible exception of ℓ_1^-. The intersection $r^* \cap \ell_1^-$ can lie to the right of v_1 (and of v_k) only if the point $c := m_0 \cap m_1$ belongs to R^- and is dual to a line in D_1^-, in which case that line must

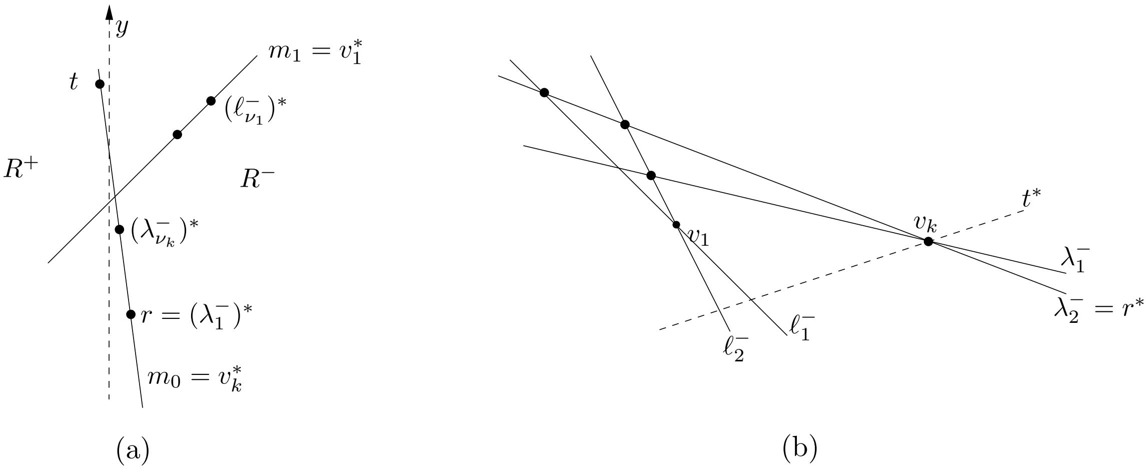

Figure 26: Case A of the construction of S_k, where $v_k = m_0^*$ and $m_0 \cap m_1 \notin R^-$. (a) The primal structure. (b) The stations in S_k (highlighted to the left of v_1).

be $\ell_1^- = c^*$. Note that in this case $r^* \cap \ell_1^- = r^* \cap c^*$ is identical to the point m_0^* dual to m_0, and the choice of the coordinate system implies that this is the rightmost vertex of $\mathcal{A}(L)$ on r^*. We define S_k to be the set of intersection points between the lines $\ell_1^-, \ldots, \ell_{\nu_1}^-$ and r^*. Thus, either all points of S_k, or all but one (namely, m_0^*) lie to the left of v_1. Clearly, we have $|S_k| = \nu_1 = \nu_k + \nu_1 - 1$, and, as above, these stations generate at least $d_1 + d_k - 1$ segments in E.

We associate with each station $u \in S_k$ the double wedge $W(u)$ formed by the two lines ℓ_i^-, λ_j^- (or ℓ_i^+, λ_j^+) that meet at u.

Constructing S_k when q exists. Suppose now that the central bichromatic point q exists. Examining the three cases in the construction of S_k, we note that Case A cannot arise, because in this case the point $m_0 \cap m_1$, which has to be equal to q, does not belong to R^-, contradicting the definition of q. In Case C we can proceed exactly as above, and collect at least $d_1 + d_k - 1$ segments in $E(S_k)$. (In fact, since we use the full weight of q, we get $w(q^+)$ additional segments in $E(S_k)$.)

It therefore remains to consider Case B, in which $v_k = m_0^*$ and $q = m_0 \cap m_1$. In the dual plane, q^* passes through both $v_1 = m_1^*$ and $v_k = m_0^*$. Suppose, without loss of generality, that $(q^+)^*$ was removed at v_1. See Figure 29.

We consider two subcases. In each of them the analysis becomes simpler if $(q^-)^* \in D_k^-$. Moreover, we assume in what follows that $k > 2$. The case $k = 2$ will be treated separately later.

Case 1. *Suppose first that $D_1^+ \setminus \{(q^+)^*\}$ is nonempty.*

In this case, we use the lower grid, to the left of v_1, formed by the lines of $D_1^+ \setminus \{(q^+)^*\}$ and the lines of L_k^+ (which, by assumption, do not include $(q^+)^*$). Clearly,

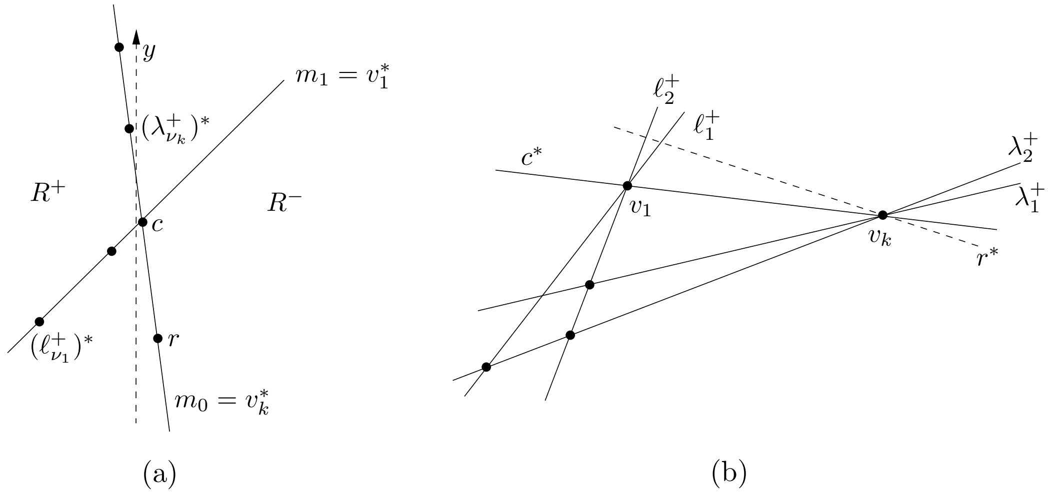

Figure 27: Case B of the construction of S_k, where $v_k = m_0^*$ and $c = m_0 \cap m_1 \in R^-$. (a) The primal structure. (b) The stations in S_k (highlighted to the left of v_1).

L_k^+ is not empty. Since D_1^+ contains lines other than $(q^+)^*$, the lower grid can indeed be used. The lines of L_k^+ contribute the full weight d_k, but the lines of $D_1^+ \setminus \{(q^+)^*\}$ contribute only $d_1 - w^+$ overall weight, so the grid generates (at least) $d_1 + d_k - w^+ - 1$ segments in $E(S_k)$.

If $(q^-)^* \in D_k^-$, the loss of w^+ segments in $E(S_k)$, as compared with the analysis in the standard case, will be automatically compensated in the construction of $E(v_k)$, since this analysis gains "for free" the weight w^+ of $(q^+)^*$ when it handles the line $(q^-)^*$. The case $(q^-)^* \notin D_k^-$ will be handled shortly.

Case 2. *$(q^+)^*$ is the only line in D_1^+.*

In this case, the lower grid does not exist, and we have $d_1 = w^+$. Clearly, there must exist lines of L_1^- through v_1 other than $(q^-)^*$. Moreover, the line r^* passes by construction through $v_k = m_0^*$ (see Figure 27). If $(q^-)^* \in L_k^-$, then r^* must also belong to L_k^-, for otherwise it would have to pass through some preceding junction v_j, so $(q^-)^*$ would pass below v_j, which is impossible for lines of L_k^-. If $(q^-)^* \notin L_k^-$, then all the lines of L_k^- (which is a nonempty set) pass above v_1. Hence we may use in this case the upper grid, which generates at least $d_1 - w_1^- + d_k - w_k^- - 1$ segments in $E(S_k)$, where w_1^-, w_k^- are the weights that $(q^-)^*$ contributes at the respective junctions v_1, v_k.

If $w_1^- = w_k^- = 0$, we obtain the standard bound $d_1 + d_k - 1$.

If $w_1^- > 0$ and $w_k^- = 0$, we are in a symmetric version of the situation in Case 1 that still needs to be treated. Both versions will be treated together below.

If $w_1^- = 0$ and $w_k^- > 0$, we automatically compensate for the loss of w_k^- segments in the count, in the construction of $E(v_1)$, which, similar to the argument in Case 1, gives us $w(q^-) \geq w_k^-$ extra segments "for free".

If $w_1^- > 0$ and $w_k^- > 0$, we interpret the bound in the context of Claim 4.2, except

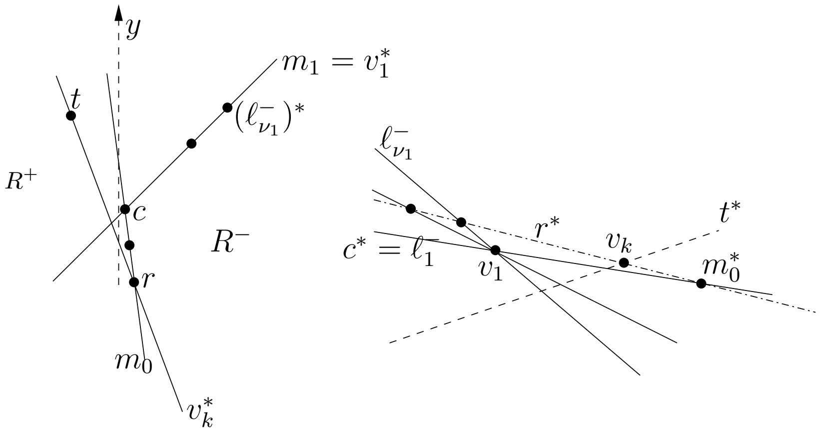

(a) (b)

Figure 28: Case C of the construction of S_k, where $v_k \neq m_0^*$. (a) The primal structure. (b) The dual picture.

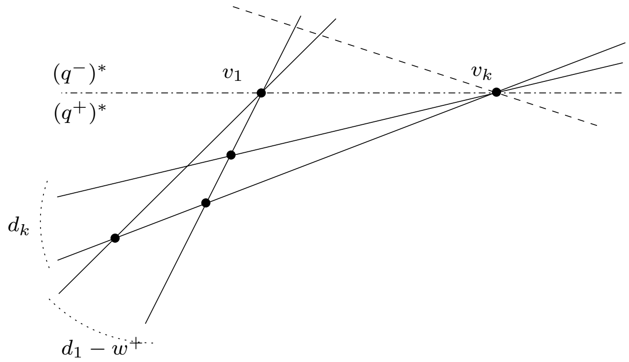

Figure 29: Constructing S_k in the presence of q^*.

that our bound is 1 *smaller* than what the Claim guarantees.

It remains to analyze the subcases where $(q^-)^* \notin D_k^-$, and where we still need to compensate for the loss of $\max\{w^+, w_1^-\}$ segments in $E(S_k)$.

Note that this loss is identical to the potential loss at $E(S_1)$, discussed above. We compensate for it in the same way—by including q in the construction of $E(v_k)$. The same analysis shows that we can always compensate for the loss, unless $\min\{w^+, w^-\} = 1$, in which case we lose one segment in the count.

(Note that, for this analysis to work, it is crucial that $k > 2$. Otherwise we need to compensate *twice* for the loss of $\max\{w^+, w_1^-\}$ segments, once in $E(S_1)$ and once in $E(S_k)$, but if $v_2 = v_k$ we can compensate for it only *once*.)

As in the case of S_1, here too we pay the price of replacing $W(v_k)$ by the narrower double wedge bounded by q^* and by the shallowest lineof the set among D_k^-, D_k^+ to which q^* was not adjoined. Nevertheless, we will show in Section 4.5 that this does not affect the pairwise non-convergence of the segments in F.

4.4 Counting the Number of Segments in E

The standard case. Let us first consider the standard case, where q does not exist. Combining the contributions in (2) and (3) with the contribution in (1), we obtain that E consists of a total of at least

$$k + t + \sum_{j=1}^{t} \kappa_j (w_j - 1) + 2\sum_{i=1}^{k} d_i - \sum_{i \in I_1} (a_i + a_{i+1}) - \sum_{i \in I_2} 1 \tag{4}$$

segments, where I_1 is the set of indices i for which there exists a (unique) line which contributes to both d_i and d_{i+1} (or to d_1 and d_k, for $i = k$), and I_2 is the complementary set.

Assume first that there is no line that contributes to all the k weights $d_1, \ldots, d_k$. Then each line ℓ_j can contribute to at most $\kappa_j - 1$ pairs of successive weights d_i, d_{i+1}, and each of the corresponding terms $(a_i + a_{i+1})$ is at most w_j. Even if there exists a line that contributes to all k weights d_i, it does not affect the construction of the segments from the stations of S_k, which always produces at least $d_1 + d_k - 1$ segments (when q does not exist). That is, we can always pretend that $k \in I_2$, so the analysis proceeds in the same way in this case, too. The remaining pairs of successive weights contribute -1 to the expression above (in the summation over $i \in I_2$). Therefore, an *overestimate* of the (absolute value of the) negative terms in (4) is $\sum_{j=1}^{t} (\kappa_j - 1) w_j + (k - \sum_{j=1}^{t} (\kappa_j - 1))$.

Using the fact that $2\sum_{i=1}^{k} d_i = \sum_{j=1}^{t} w_j = n - 1$, the bound in (4) is greater than or equal to

$$k + t + \sum_{j=1}^{t} \kappa_j (w_j - 1) + \sum_{j=1}^{t} w_j - k - \sum_{j=1}^{t} (\kappa_j - 1)(w_j - 1) = 2\sum_{j=1}^{t} w_j = 2n - 2.$$

The case where q exists. The differences between this case and the standard case are:

(i) We may lose one segment in $E(v_1)$.

(ii) We may lose one segment in $E(S_1)$. Even if we do not lose the segment, we may collect there only $d_1 + d_2 - (a_1 + a_2)$ segments, where one of a_1, a_2 is 0.

(iii) A similar situation may occur for S_k.

(iv) It is possible that $(q^-)^*$ or $(q^+)^*$ contributes weight to *all junctions* $v_1, \ldots, v_k$, which may cause the set I_1 to consist of all indices $1, \ldots, k$, and I_2 to be empty.

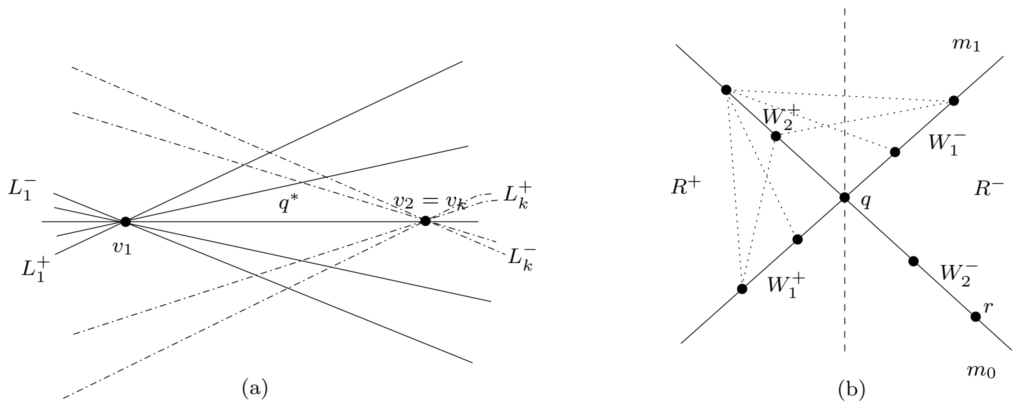

Figure 30: The case of only two junctions in the presence of q^*: (a) The dual configuration. (b) The primal configuration (in π).

Assume first that the situation in (iv) does not arise. Then the analysis proceeds as in the standard case, since, as is easily verified, it is not affected by having some of the a_i's vanish, except that we need to subtract 3 from the overall count, to accommodate the potential losses in (i)–(iii). Hence, in this case we have $|E| \geq 2n - 5$.

If the situation in (iv) arises, say, with $(q^-)^*$ being the line that contributes weight to all junctions, then $\sum_{i \in I_1}(a_i + a_{i+1}) = 2w^-$, and all lines $\ell_j \neq (q^-)^*$ have $\kappa_j = 1$. The total number of segments in E is therefore at least (without loss of generality, we assume that $(q^-)^*$ is the t-th line)

$$k(w^- - 1) + \sum_{j=1}^{t-1}(w_j - 1) + n - 1 - 2w^- + k + t - 3 =$$

$$(k-2)w^- + (n-1) + t - (t-1) + (n - 1 - w^-) - 3 = (k-3)w^- + 2n - 4.$$

Hence, if $k \geq 3$, we have $|E| \geq 2n - 4$. (Recall that the count so far actually *relied* on the assumption that $k \geq 3$.)

It thus remains to consider the case where only two junctions are generated. In this case, by construction, *all* the lines of L must pass either through v_1 or through v_2; see Figure 30(a). Hence, in the primal plane, all points of R must lie on the lines m_0 and m_1, with q lying on both lines; see Figure 30(b).

In this very degenerate case, we construct E explicitly, working in the primal plane π, as follows. Denote by $W_1^+, W_1^-, W_2^+, W_2^-$ the overall weight of all points of $R^+ \cap m_1$, $R^- \cap m_1$, $R^+ \cap m_0$, $R^- \cap m_0$, respectively, *excluding q in all four cases.*

(i) Apply Theorem 3.1 or Corollary 3.2 at m_0 and at m_1 (or, equivalently, at the dual junctions $m_1^* = v_1$ and $m_0^* = v_2$). To be on the safe side, we assume that only Corollary 3.2 can be used at either junction, and, as usual, we subtract 1 from the bound at v_1 to allow for potential double counting of a segment. This yields a total of at least

$$(W_1^+ + W_1^- + w(q) - 1) + (W_2^+ + W_2^- + w(q)) =$$

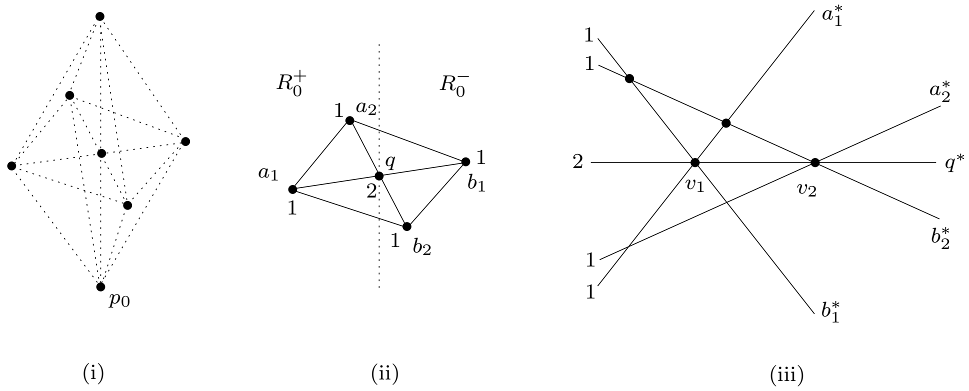

Figure 31: (i) The set of Figure 1 with $n = 7$ points that determines $2n - 5 = 9$ directions. (ii) The weighted set R in the primal plane π, obtained by projecting from p_0, with the central bichromatic point q of weight 2. (iii) The dual construction of junctions and stations. We collect 4 segments in $E(v_1)$ (using Corollary 3.2), 3 in $E(v_2)$ (using Corollary 3.2 and subtracting 1), 1 in $E(S_1)$, and 1 in $E(S_2)$, for a total of 9 pairwise non-convergent segments.

$$(W_1^+ + W_1^- + W_2^+ + W_2^- + w(q)) + w(q) - 1 = (n-1) + w(q) - 1$$

segments in E.

(ii) Suppose, without loss of generality, that $W_2^- \le W_2^+$. We then generate segments in F, in addition to those lying on m_0, m_1, as shown in Figure 30(b). That is, we connect the point of $R^+ \cap m_0$ farthest from q to all the points on m_1, excluding q, and connect the two points of $R^+ \cap m_1$, $R^- \cap m_1$ farthest from q to all the points of $R^+ \cap m_0$. Here it is easy to verify directly that all these segments, including the segments $f(v_1) \subset m_1$, $f(v_2) \subset m_0$ (where the first may degenerate to the singleton point q, but is still considered to lie along m_1), are pairwise non-convergent. The F-segments that we have constructed are dual to the stations in $S_1 \cup S_2$. The total number of segments in E that are generated from these stations, in the standard manner, is at least

$$(W_1^- + W_2^+ - 1) + (W_1^+ + W_2^+ - 1) \ge W_1^+ + W_1^- + W_2^+ + W_2^- - 2 = n - 1 - w(q) - 2.$$

Adding the bounds from (i) and (ii), we get $|E| \ge 2n - 5$. We note that the configuration in Figure 31 falls into this case.

4.5 Pairwise Non-convergence of the Collected Segments

To complete the proof, we have to show that no pair of segments in E are convergent. We first show:

Lemma 4.4. *Let Q denote the set of all junctions and stations that we have collected. For any $u, v \in Q$, the segments $f(u)$ and $f(v)$ associated with these vertices are non-convergent in the primal plane π.*

Proof: Let us first consider the standard case, where q does not exist. Let $u, v \in Q$ with u lying to the left of v. The property that $f(u)$ and $f(v)$ are non-convergent is dual to the property that $W(u)$ and $W(v)$ are non-convergent, that is, either u lies in (the closure of) $W(v)$ or v lies in (the closure of) $W(u)$. We distinguish between several cases:

Case A: *Both u and v are junctions.*

Put $u = v_i$ and $v = v_j$, with $i < j$. Then $W(v)$ is bounded by a line $\ell \in L_j^+$ and by a line $\ell' \in L_j^-$. By Claim 4.1(ii), $u = v_i$ lies between these two lines, and thus belongs to $W(v_j) = W(v)$.

Case B: *u is a junction and v is a station not in S_k.*

Put $u = v_i$ and let S_j be the set of stations that contains v, where $i \le j < k$. Then $W(v)$ is bounded by two lines ℓ, ℓ', where either $\ell \in L_j^+$ and $\ell' \in L_{j+1}^-$, or $\ell \in L_{j+1}^+$ and $\ell' \in L_j^-$. By construction, we have in both cases $\ell \in L_j^+$ and $\ell' \in L_j^-$, and the analysis is completed as in Case A.

Case C: *u is a station not in S_k and v is a junction or a station not in S_k.*

Let S_i be the set of stations containing u. The arguments in Case A and Case B imply that $v_i \in W(v)$. If v is also a station in S_i or $v = v_{i+1}$ then it is clear from the construction of S_i that $W(u)$ and $W(v)$ are non-convergent (see Figure 19). Suppose then that v lies to the right of v_{i+1}. Then both v_i and v_{i+1} lie in the left wedge of $W(v)$, and u is incident to a line ℓ that passes through v_i and to a line ℓ' that passes through v_{i+1}. If $u \notin W(v)$ then a boundary line of $W(v)$ must separate u from v_i and v_{i+1}, in which case $v \in W(u)$; compare with Figure 14(a).

Case D: *u is a station in S_k to the left of v_1 and v is a junction or station.*

If both u and v belong to S_k, then the claim follows easily from the construction of S_k. We thus suppose that $v \notin S_k$. Then we have $v \in \{v_i\} \cup S_i \cup \{v_{i+1}\}$, for some $1 \le i < k$.

We start with the case $v_k = m_0^*$. Refer to Figure 32. Suppose that $u \in S_k$ is the intersection point of two lines ℓ, λ, passing through v_1 and v_k, respectively, which, without loss of generality, we assume to belong to L^-. If v is contained in the double wedge bounded by ℓ and λ, then $v \in W(u)$, so $W(v)$ and $W(u)$ are non-convergent. Otherwise, since u lies to the left of v_1, v lies either above λ or below ℓ. If v is above λ, then it is not a junction, so it must be the crossing point of a line $\ell^+ \in D_i^+$ and a line $\ell^- \in D_{i+1}^-$. See Figure 32(a). Both v_i and v_{i+1} lie on or below λ, so the left portion of the double wedge bounded by ℓ^- and ℓ^+ contains u. Thus, we have $u \in W(v)$. If, on the other hand, v is below ℓ, as in Figure 32(b), then it is either a junction or a station, and it is the crossing point of a line $\ell^- \in L^-$ and a line $\ell^+ \in L^+$ which bound $W(v)$, such that either both ℓ^+ and ℓ^- are in D_i (if $v = v_i$ is a junction), or $\ell^- \in D_i^-$ and $\ell^+ \in D_{i+1}^+$ (if $v \in S_i$ is a station). Now ℓ^- must pass above (or through) v_1 and hence above u, while ℓ^+ must pass below u. Again we can conclude that the left portion of the double wedge bounded by ℓ^- and ℓ^+, and thus $W(v)$, contains u.

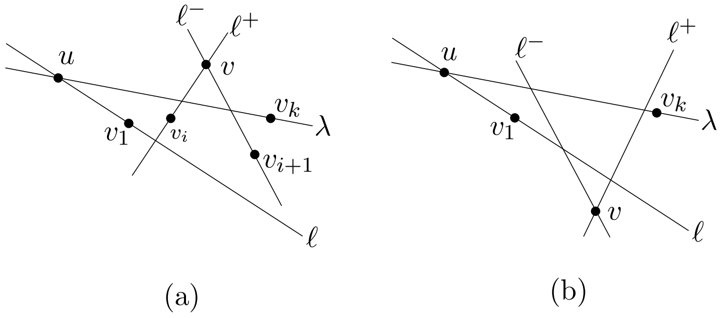

Figure 32: The proof that $W(u)$ and $W(v)$ are non-convergent when u is a station to the left of v_1. (a) v lies above $W(u)$. (b) v lies below $W(u)$.

The case where u lies on two lines of L^+ is handled in a fully symmetric manner.

If $v_k \neq m_0^*$, the above argument can be repeated *verbatim* for stations $u \in S_k$ *to the left* of v_1. If $v = m_0^*$, the sole station to the right of v_k, the claim is immediate from the construction of S_k. Hance, Case D of the Claim holds in either case.

Case E: *u is a junction or station not in S_k and v is a station in S_k (to the right of v_k).*

Case E can arise only when $v = m_0^* \in S_k$. Now it is simplest to establish the claim in the primal plane, by noting that the segment dual to $W(v)$ lies on the line m_0, and that, by construction (since $u \notin S_k$), the segment dual to $W(u)$ must connect a point of R^- to a point of R^+, and thus must intersect m_0, showing that these two segments are non-convergent.

Consider next the case where the central bichromatic point q exists, which requires a few modifications in the preceding analysis. First, if both $(q^-)^*$ and $(q^+)^*$ belong to D_1, the double wedge $W(v_1)$ degenerates to the single line q^*. (We still consider it to have v_1 as an apex. In the primal plane, the segment $f(v_1)$ degenerates to the singleton point q, but it is still considered to lie along the line v_1^*.) It is easily verified, though, by specializing Cases A,B,D,E to this configuration, that $W(v_1)$ and any other wedge $W(v)$ in our collection are still non-convergent.

The presence of q does not affect any other case in the preceding analysis, as long as we were not forced to include q in the construction of $E(v_2)$ or $E(v_k)$. Suppose then that we had to include q in the construction of $E(v_2)$ (even though neither $(q^+)^*$ nor $(q^-)^*$ belonged to D_2). In Case A, v_2 is contained in $W(v_j)$ for any $j > 2$ (the case where $W(v_k)$ was also shrunk will be treated below), so it only remains to consider the case $u = v_1$, $v = v_2$, which still works, since $v_1 \in q^*$, and thus v_1 still lies in the modified $W(v_2)$. Case B is not affected by the shrinking of $W(v_2)$. In Case C, we only need to consider the subcase when $u \in S_1$, and the property continues to hold since $v_2 \in W(u)$. In Case D, we have $v_2 \in W(u)$, which easily follows from the fact that $v_2 \in q^*$; see Figure 33. Case E is argued as in the standard case.

Suppose finally that we had to include q in the construction of $E(v_k)$ (even though

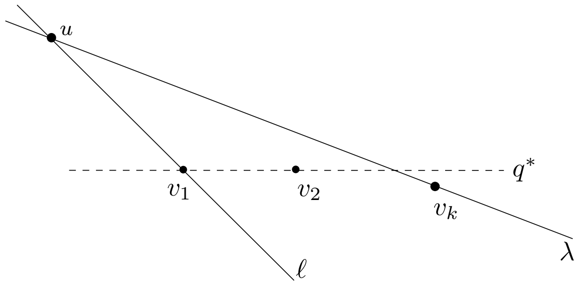

Figure 33: Case D of the analysis when $W(v_2)$ is shrunk: $v_2 \in q^*$ lies above ℓ and below λ and thus $v_2 \in W(u)$.

neither $(q^+)^*$ nor $(q^-)^*$ belonged to D_k). Recall that this can arise only in Case B of the construction of S_k, where $v_k = m_0^*$. Now, except for the stations in S_k, for any other vertex $u \in Q$, $f(u)$ connects a point of R^+ and a point of R^-, and thus the line containing $f(v_k)$, namely m_0, must cross $f(u)$, so $f(u)$ and $f(v)$ are non-convergent. If u is a station in S_k, then $v_k \in W(u)$, by construction.

Hence the lemma also holds when q exists. $\square$

Non-convergence of the elements of E. Recall that, for each $v \in Q$, the points of P that span the segments in $E(v)$ are those points that project to the line containing $f(v)$ in π, so that their projections are dual to lines in L that either were removed at v or had their weights reduced there (if v is a junction), or are the two lines incident to v (if v is a station).

Moreover, each segment e in $E(v)$ has the property that either its projection on π contains the segment $f(v)$ or it is a point not in the interior of $f(v)$; the latter case arises when e is contained in a ray emanating from p_0, a situation that can arise when we apply Theorem 3.1 or Corollary 3.2 at one of the junctions $v_1, \ldots, v_k$.

Let e_1 and e_2 be two segments in E. For $i = 1, 2$, let u_i denote the vertex in Q for which $e_i \in E(u_i)$, and set $f_i = f(u_i)$. It suffices to consider the case $u_1 \neq u_2$.

The segments f_1 and f_2 are non-convergent in π. If the projections e_1^*, e_2^* on π from p_0 of e_1 and e_2, respectively, are segments (so that they contain f_1 and f_2, respectively), then e_1^* and e_2^* are non-convergent in π, which is easily seen to imply that e_1 and e_2 are non-convergent in $\mathbb{R}^3$. If the projections of both e_1, e_2 are points on π, then e_1 and e_2 share p_0 as an endpoint and therefore are non-convergent.

We are left with the case in which, without loss of generality, e_1 projects from p_0 to a point $x \in \pi$ (which is on the line containing f_1 but not in the interior of f_1), whereas e_2 projects to a segment e_2' containing f_2. See Figure 34. The point x may be assumed to lie on the line containing f_2, for otherwise e_1 and e_2 are non-coplanar, and therefore non-convergent. If $x \in e_2'$ then clearly e_1 and e_2 are non-convergent, so we may assume that $x \notin e_2'$. It follows that $x \notin f_2$ and since f_1 and f_2 are non-convergent, x must be an endpoint of f_1 (otherwise f_1 and f_2 would be convergent,

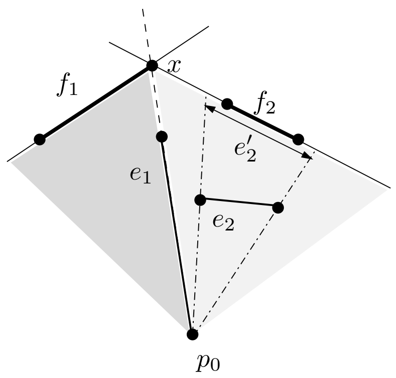

Figure 34: Non-convergence of e_1 and e_2: An impossible configuration.

because the lines on π that contain them meet at x which lies outside both segments).

We claim that this case is impossible. Indeed, we have already noted that if x is an endpoint of f_1, then the vertex u_1 dual to f_1 must be a junction v_i. The line ℓ dual to x passes through both vertices $u_1 = v_i$ and u_2, and it is the shallowest line in either D_i^+ or D_i^-. If the central bichromatic point q exists, and u_2 is not a station in S_k, then ℓ cannot be equal to q^*, i.e., x cannot be equal to q, because, by construction, q must lie inside f_2, which is delimited by a point of R^+ and a point of R^-. Hence, we may assume, without loss of generality, that $\ell \in L^+ \setminus \{(q^+)^*\}$. The case where q exists and u_2 is a station in S_k will be considered later.

Case 1: *$u_1 = v_i$ lies to the left of u_2.*

Since x lies outside f_2, the endpoint of f_2 nearer to x is dual to a line $\ell' \in L_i^+$ that passes through u_2 and has smaller slope than that of ℓ. (Since u_2 is constructed after u_1, the lines that define $W(u_2)$ must belong to L_i.) But then ℓ' must pass above v_i which is a contradiction since all the lines in L_i^+ must pass through or below v_i. See Figure 35(a).

Case 2: *$u_1 = v_i$ lies to the right of u_2.*

Assume first that u_2 is not a station in S_k. Let $1 \le j < i$ be the index such that either $u_2 = v_j$ or u_2 is a station in S_j. Since x lies outside f_2, the R^+-endpoint of f_2 (the one nearer to x) is dual to a line $\ell' \in D_j^+ \cup D_{j+1}^+$, which is shallower than ℓ (since x lies outside f_2). If $u_2 = v_j$ then $\ell' \in D_j^+$, and, by construction, ℓ must have also been removed at v_j or at an earlier junction, and thus it cannot be dual to an endpoint of f_1 (because such a point must be the dual of some line in L_i^+). See Figure 35(b). Hence this case is impossible. Suppose then that u_2 is a station in S_j. Regardless of whether $\ell' \in D_j^+$ or $\ell' \in D_{j+1}^+$, since ℓ is steeper than ℓ', v_{j+1} lies below ℓ. Hence, we must have $i > j + 1$, and, since $\ell \in L_i^+$, we obtain a contradiction to Claim 4.1(ii); see Figure 35(c).

Finally, assume that u_2 is a station in S_k to the left of v_1. Suppose first that u_2 lies on a line $\ell^- \in D_1^-$ and a line $\lambda^- \in D_k^- = L_k^-$. In this case, $\ell' = \lambda^-$ and ℓ passes

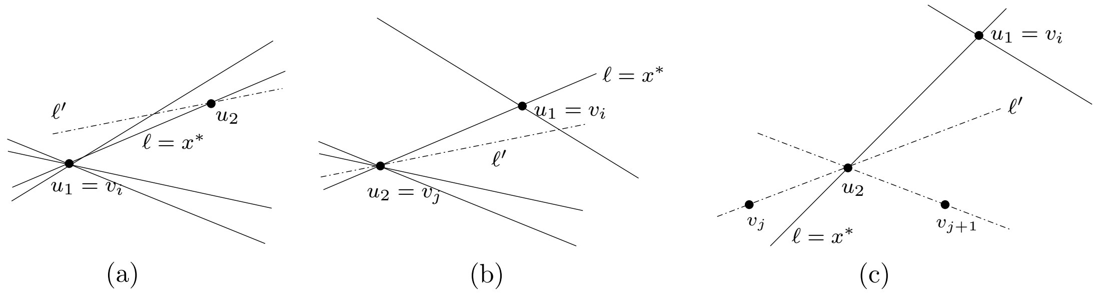

Figure 35: Showing the impossibility of the configuration in Figure 34. (a) u_1 is to the left of u_2. (b) u_2 is a junction to the left of u_1. (c) u_2 is a station to the left of u_1.

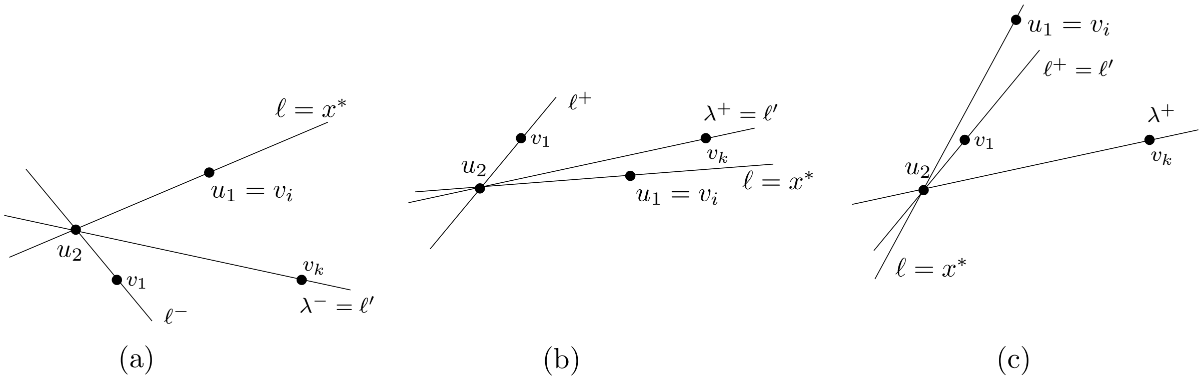

Figure 36: Showing the impossibility of the configuration in Figure 34 when u_2 is a station of S_k to the left of v_1, and $\ell \in L^+$. (a) u_2 is formed by two lines of L^-. (b) u_2 is formed by two lines of L^+, and ℓ is shallower than both of these lines. (c) u_2 is formed by two lines of L^+, and ℓ is steeper than both of these lines.

above v_1, which is impossible, since $\ell \in L_1^+$; see Figure 36(a). Suppose next that u_2 lies on a line $\ell^+ \in D_1^+$ and a line $\lambda^+ \in D_k^+ = L_k^+$. Since x lies outside f_2, ℓ is not contained in $W(u_2)$. If ℓ is shallower than λ^+ (see Figure 36(b)), then $\ell' = \lambda^+$, and v_i lies below $\lambda^+ \in L_k^+$, which is impossible. If ℓ is steeper than ℓ^+ (see Figure 36(c)), then $\ell' = \ell^+$, and v_1 lies below ℓ, which is impossible, since $\ell \in L_1^+$. (Note that in all three cases, ℓ cannot be equal to q^*, because q^* passes through v_1 and ℓ does not.)

All these contradictions show that Case 2 is also impossible.

This establishes the non-convergence of the segments in E, and thus, at long last, completes the proof of Theorem 1.3. $\square$

As mentioned earlier, we can get a better bound when q does not exist:

Theorem 4.5. *Let P be a set of $n \geq 6$ non-coplanar points in $\mathbb{R}^3$, such that n is odd, and there exists an extreme point p_0 of P, such that the projection of P from p_0 produces a set R without a central bichromatic point. Then P determines at least $2n-2$ segments, no two of which are convergent.*

This strengthens the bound $2n-3$ in Conjecture 1 of Blokhuis and Seress [?] for n odd.

5 Extensions and Open Problems

In this section we consider several extensions of our results, prove some of them, and leave the others as open problems.

The most obvious open problem is to obtain the exact worst-case bound for n even. Currently there is a small gap between our lower bound $2n-7$ and the best known construction, which gives $2n-3$ pairwise non-convergent segments.

Theorem 1.3 yields the following extension to four dimensions. It settles Conjecture 9 of Blokhuis and Seress [2] in the affirmative for $d=4$ and for even n.

Theorem 5.1. *Let P be a set of n points in $\mathbb{R}^4$, not contained in a hyperplane and not having three collinear points. Then P determines at least $3n-8$ different directions, if n is even, and at least $3n-10$ different directions if n is odd. The bound is sharp for every even $n \geq 8$.*

Proof: Let p_0 be the lowest point of P (in the x_4-direction). Let H be a horizontal hyperplane (parallel to the $x_1x_2x_3$-space) far above all the points of P. By applying a small rotation to P, we may assume that H is not parallel to any segment determined by P.

Project the points of $P \setminus \{p_0\}$ centrally from p_0 onto H, and color the projected images red. For each direction γ determined by P, let L_γ denote the line parallel to γ and passing through p_0. If a direction γ, determined by P, is not obtained through p_0, let b_γ denote the intersection point of L_γ with H. Color all such points b_γ green. Clearly, every red or green point on H gives rise to a different direction determined by P, and all these points are distinct. The number of red points on H is $n-1$.

Since P is not contained in a hyperplane, the red points on H are not coplanar. Therefore, by Theorem 1.3, they determine at least $2(n-1)-5 = 2n-7$ pairwise non-convergent segments, if $n-1$ is odd, and at least $2(n-1)-7 = 2n-9$ pairwise non-convergent segments, if $n-1$ is even.

We claim that along each line L in H passing through two or more red points there is a green point that lies outside the convex hull of the red points on L. Indeed, consider the 2-plane through p_0 and L. The direction γ in 4-space, determined by the two points of P that project to the two extreme red points on L, is not obtained through p_0, and thus yields the desired green point outside the convex hull of the red points on L. See Figure 37. Therefore, every collection of m pairwise non-convergent segments determined by the red points on H, gives rise to m distinct green points on H, formed in the manner just described. No two such green points can coincide, for that would make the corresponding red segments convergent.

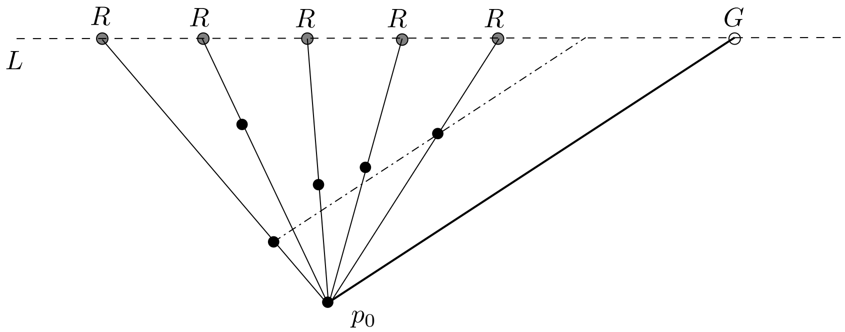

Figure 37: The green point determined by a red segment in H.

It follows that the number of points on H is at least $n-1+2n-7 = 3n-8$, if n is even, and at least $n-1+2n-9 = 3n-10$, if n is odd.

We next show that the bound is sharp for even $n \geq 8$. The construction extends the one depicted in Figure 1. Specifically, let P be the set of the vertices of a regular $(n-4)$-gon Q in the x_1x_2-plane, centered at the origin, and of the four points $\pm e_3 = (0,0,\pm 1,0)$, $\pm e_4 = (0,0,0,\pm 1)$. It is easy to see that P determines exactly $3n-8$ different directions: $n-4$ directions in the x_1x_2-plane, $n-4$ directions obtained by connecting the vertices of Q to e_3, $n-4$ directions obtained by connecting the vertices of Q to e_4, and 4 directions determined by $\pm e_3$, $\pm e_4$. $\square$

A grand generalization of Theorem 5.1, still in four dimensions, would be to establish the following conjecture:

Conjecture A: *Any* set P of n points in $\mathbb{R}^4$, not contained in a single hyperplane, determines at least $3n-c$ pairwise non-convergent segments, for some constant c (that might be larger than those in the theorem).

This conjecture would imply, by an appropriate extension of the preceding proof, that any set of n points in $\mathbb{R}^5$, not contained in a hyperplane, and not having three collinear points, determines at least $4n-(c+4)$ different directions.

The final grand challenge is to establish the following conjecture, which strengthen Conjecture 9 of Blokhuis and Seress [2]:

Conjecture B: *Any* set P of n points in $\mathbb{R}^d$, for $d \geq 4$, not contained in a single hyperplane, determines at least $(d-1)n - c_d$ pairwise non-convergent segments, for some constant c_d that depends (probably quadratically) on the dimension d.

References

[1] M. Aigner and G. Ziegler, *Proofs from The Book, 2nd ed.* Springer-Verlag, Berlin, 2001.

[2] A. Blokhuis and Á. Seress, The number of directions determined by points in the three-dimensional Euclidean space, *Discrete Comput. Geom.* **28** (2002), 491–494.

[3] G. R. Burton, G. B. Purdy, The directions determined by n points in the plane, *J. London Math. Soc. (2)* **20** (1979), 109–114.

[4] R. Cordovil, The directions determined by n points in the plane, a matroidal generalization, *Discrete Math.* **43** (1983), 131–137.

[5] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer Verlag, Heidelberg, 1987.

[6] P. Erdős, Solution to Problem Nr. 4065, *Amer. Math. Monthly* **51** (1944), 169.

[7] J. C. Fisher and R. E. Jamison, Properties of affinely regular polygons, *Geom. Dedicata* **69** (1998), 241–259.

[8] J. E. Goodman and R. Pollack, A combinatorial perspective on some problems in geometry, *Congr. Numer.* **32** (1981), 383–394.

[9] J. E. Goodman and R. Pollack, Allowable sequences and order types in discrete and computational geometry, in: *New Trends in Discrete and Computational Geometry (J. Pach, ed.), Algorithms Combin.* **10**, Springer, Berlin, 1993, 103–134.

[10] R. E. Jamison, Survey of the slope problem, in: *Discrete Geometry and Convexity, Ann. New York Acad. Sci.* **440**, New York Acad. Sci., New York, 1985, 34–51.

[11] R. E. Jamison and D. Hill, A catalogue of sporadic slope-critical configurations. in: *Proceedings of the Fourteenth Southeastern Conference on Combinatorics, Graph Theory and Computing (Boca Raton, Fla., 1983), Congr. Numer.* **40** (1983), 101–125.

[12] G. Korchmaros, Poligoni affin-regolari dei piani di Galois d'ordine dispari (English), *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Natur. (8)* **56** (1974), 690–697.

[13] L. Lovász and A. Schrijver, Remarks on a theorem of Rédei, *Studia Sci. Math. Hung.* **16** (1981), 449–454.

[14] J. Pach, R. Pinchasi, and M. Sharir, On the number of directions determined by a three-dimensional point set, *J. Combinat. Theory, Ser. A* **108** (2004), 1–16.

[15] L. Rédei, *Lacunary Polynomials Over Finite Fields*, North-Holland, Amsterdam, 1973.

[16] P. R. Scott, On the sets of directions determined by n points, *Amer. Math. Monthly* **77** (1970), 502–505.

[17] T. Szőnyi, Around Rédei's theorem, *Discrete Math.* **208/209** (1999), 557–575.

[18] P. Ungar, $2N$ noncollinear points determine at least $2N$ directions, *J. Combin. Theory, Ser. A* **33** (1982), 343–347.

[19] F. Wettl, On the nuclei of a pointset of a finite projective plane, *J. Geom.* **30** (1987), 157–163.