

Vertical Decomposition of Arrangements of Hyperplanes in Four Dimensions ^{*}

Leonidas J. Guibas[†] Dan Halperin[‡] Jiří Matoušek[§]
Micha Sharir[¶]

November 3, 1994

Abstract

We show that, for any collection \mathcal{H} of n hyperplanes in \mathbb{R}^4 , the combinatorial complexity of the *vertical decomposition* of the arrangement $\mathcal{A}(\mathcal{H})$ of \mathcal{H} is $O(n^4 \log n)$. The proof relies on properties of superimposed convex subdivisions of 3-space, and we also derive some other results concerning them.

1 Introduction

Let \mathcal{H} be a collection of n hyperplanes in \mathbb{R}^4 . The *vertical decomposition* $\mathcal{V}(\mathcal{H})$ of the arrangement $\mathcal{A}(\mathcal{H})$ of \mathcal{H} is defined in the following recursive manner. Denote the coordinates by x, y, z and w . For each cell C of $\mathcal{A}(\mathcal{H})$ and each 2-face g on ∂C , erect a 3-D vertical *wall* from g within C , defined as the union of maximal w -vertical line segments that have one endpoint on g and are contained in C . The collection of these walls decomposes C into vertical prisms, each bounded by two hyperplanes of

^{*}Work on this paper by Leonidas Guibas and Micha Sharir has been supported by a grant from the U.S.–Israeli Binational Science Foundation. Work by Leonidas Guibas was also supported by National Science Foundation Grant CCR-9215219. Work by Micha Sharir was also supported by National Science Foundation Grant CCR-91-22103, and by grants from the G.I.F. — the German Israeli Foundation for Scientific Research and Development, and the Fund for Basic Research administered by the Israeli Academy of Sciences. Work by Jiří Matoušek was done while he was visiting Tel Aviv University, and it was partially supported by a Humboldt Research Fellowship.

[†]Department of Computer Science, Stanford University, Stanford CA 94305

[‡]Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305 (Work on this paper was carried out while this author was at Tel-Aviv University.)

[§]Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czechoslovakia, and Institut für Informatik, Freie Universität Berlin, Arnimallee 2–6, W1000 Berlin 33, Germany

[¶]School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, U.S.A.

\mathcal{H} , one on its top and one on its bottom (sometimes, when C is unbounded, by just one hyperplane), and by some of the vertical walls. In the next stage we project each such prism onto the xyz -hyperplane, obtaining a 3-D convex polyhedron P , which we vertically decompose in an analogous manner. That is, we erect vertical walls (in the z -direction) within P from each edge h of P , where the wall of an edge h is the union of maximal z -vertical line segments that have one endpoint on h and are contained in P . These walls decompose P into z -vertical prisms, each bounded by two facets of P on the top and the bottom sides (or possibly just one if P is unbounded) and by some z -vertical walls. We next project each such prism onto the xy -plane, obtaining a convex polygon Q , which we now proceed to vertically decompose in a similar manner, erecting vertical segments (in the y -direction) within Q from each vertex of Q until they meet the boundary of Q again. We now complete the decomposition of P by erecting z -vertical walls within P from each of the y -vertical segments in the decomposition of each of the resulting polygons Q . Finally we complete the decomposition of each cell C of $\mathcal{A}(\mathcal{H})$ by erecting w -vertical walls from each newly created feature on each prism of C . Repeating this procedure over all cells C of $\mathcal{A}(\mathcal{H})$, we obtain the desired vertical decomposition $\mathcal{V}(\mathcal{H})$ of the arrangement.

We prove the following:

Theorem 1.1 *The number of cells in the vertical decomposition of an arrangement of n hyperplanes in four dimensions is $O(n^4 \log n)$.*

The notion of vertical decomposition can be extended to higher dimensions in an obvious manner. In fact, it can be extended to arrangements of algebraic surfaces in \mathbb{R}^d of bounded degree, as described in detail in [3]. The output of the decomposition are cells with “constant description complexity”—in the case of hyperplanes, each is a convex polyhedron with at most $2d$ facets, with two facets obtained at each recursive step. In the case of general algebraic surfaces the structure of cells is somewhat more involved, but each cell is still bounded by at most $2d$ surfaces of bounded degree, and thus has also constant description complexity.

The problem at hand is to obtain sharp bounds on the number of cells in the vertical decomposition. It is shown in [3] that, in the general algebraic case, the number of cells is $O(n^{2d-3}\beta(n))$, where $\beta(n)$ is a slowly growing function of n , depending also on d and on the degree of the given surfaces. In three dimensions this yields a nearly cubic bound on the size of the vertical decomposition (in the general case); for planes, a simpler argument gives a tight bound of $\Theta(n^3)$. Thus the first interesting case is $d = 4$, where the above bound is roughly $O(n^5)$ (even for hyperplanes), whereas the complexity of the arrangement, without vertical decomposition, is only $O(n^4)$.

Theorem 1.1 shows that the size of the vertical decomposition, for hyperplanes in 4-space, is only $O(n^4 \log n)$. This constitutes the first step towards obtaining a similar bound for general surfaces, and extending these bounds to higher dimensions.

The main motivation for studying vertical decompositions in arrangements of surfaces is in their applications to range searching, point location and many related problems; see [3] for some of these applications. We note that the exact shape of

the cells in the decomposition is irrelevant for these applications, as long as each cell has constant description complexity. Thus for arrangements of hyperplanes one can instead triangulate the arrangement into simplices in a standard manner, so that the number of simplices is only $O(n^d)$ [4]. Thus these applications, in the case of hyperplanes, have no real need for vertical decomposition. On the other hand, in the general case of algebraic surfaces, vertical decomposition seems to be the only known general decomposition scheme, so deriving sharper bounds on its complexity is an important problem that merits careful study; our analysis for hyperplanes can be seen as a preliminary step in this direction.

A main portion of the proof is based on an analysis of the overlay of convex subdivisions in 3-dimensional space. We establish properties of such superimposed subdivisions, which may be of independent interest. We also prove some other properties of such subdivisions, which are not needed for the vertical decomposition result.

2 A Reduction to a 3-Dimensional Problem

Let \mathcal{H} be a collection of n hyperplanes in 4-space, which we assume to be in general position. This involves no real loss of generality, because one can always perturb slightly the given hyperplanes so as to put them in general position, without decreasing the number of cells in the decomposition.

The heart of the proof of Theorem 1.1 is the following lemma:

Lemma 2.1 *Let C be a cell of the arrangement $\mathcal{A}(\mathcal{H})$ with a total of N_C faces (of all dimensions). Then the complexity of the vertical decomposition of C is $O(N_C^2)$.*

Assuming this lemma, Theorem 1.1 follows from a result of [1] on the sum of squares of cell complexities in arrangements of hyperplanes, which states that, in four dimensions, one has

$$\sum_C N_C^2 = O(n^4 \log n),$$

where the sum extends over all cells of $\mathcal{A}(\mathcal{H})$.

To prove the lemma, let C be a cell of the arrangement $\mathcal{A}(\mathcal{H})$, and assume, for the sake of clarity of exposition, that C is bounded (for unbounded cells the analysis is quite similar). Let us divide the boundary of C into the upper and lower portions, and let \mathcal{R} (resp. \mathcal{B}) denote the projection of the upper (resp. lower) portion into the xyz -hyperplane. We can regard \mathcal{R} and \mathcal{B} as convex subdivisions in 3-space which we refer to as the red and blue subdivisions, respectively¹.

¹Actually these are not convex subdivisions of the whole 3-space, but rather of the projection of C . However, the complement of the projection of C can be partitioned into convex cells whose total complexity does not exceed that of C , and thus \mathcal{R} and \mathcal{B} can be completed to convex subdivisions of the whole space. Or, alternatively, one can check that the restriction to the projection of C does not make a difference in the following analysis.

Since \mathcal{H} is in general position, \mathcal{R} and \mathcal{B} are *simple* decompositions. Let us remark that one cannot obtain all subdivisions of 3-space in this manner; as shown in [2], \mathcal{R} and \mathcal{B} must be *power diagrams* in 3-space.

We note that each feature (cell, face, edge, or vertex) of \mathcal{R} stands in a 1-1 correspondence with some feature (facet, 2-face, edge or vertex) of the top part of ∂C , and similarly for \mathcal{B} and the bottom part of ∂C . Let $N_{\mathcal{R}}, N_{\mathcal{B}}$ denote the total number of features of \mathcal{R}, \mathcal{B} , respectively. Thus $N_{\mathcal{R}} + N_{\mathcal{B}} \leq 2N_C$, where N_C is the total number of faces bounding C (the factor 2 comes from the fact that features on the silhouette of C appear both in the top part and in the bottom part of ∂C).

The first step of the vertical decomposition of C is equivalent to overlaying \mathcal{R} and \mathcal{B} to form one convex subdivision \mathcal{T} of 3-space. Each new feature of the decomposition \mathcal{T} corresponds to some intersection between a vertical wall erected upwards from the bottom part of ∂C and another wall erected downwards from the top part of ∂C . The remaining steps in the vertical decomposition of the (4-dimensional) cell C correspond to vertically decomposing each cell in the 3-dimensional subdivision \mathcal{T} .

Hence it suffices to establish the following lemma, whose proof is postponed to the following section:

Lemma 2.2 *Let \mathcal{R}, \mathcal{B} be simple convex subdivisions of 3-space with $N_{\mathcal{R}}, N_{\mathcal{B}}$ faces, respectively. Then the complexity of the vertical decomposition of the subdivision \mathcal{T} obtained by overlaying \mathcal{R} and \mathcal{B} is $O((N_{\mathcal{R}} + N_{\mathcal{B}})^2)$.*

Lemma 2.2 clearly implies Lemma 2.1, and thus completes the proof of Theorem 1.1.

3 Properties of Convex Subdivisions in 3-Space

We begin with the proof of Lemma 2.2:

Let P be one of the cells in the overlaid decomposition \mathcal{T} . The vertical decomposition of P can be obtained by projecting the top part and the bottom part of ∂P (relative to the z -direction) onto the xy -plane, and by overlaying these two convex subdivisions—every new vertex in the superimposed map, \mathcal{M} , corresponds to an intersection between two vertical walls, one coming upwards from an edge on the bottom part of ∂P , and one coming downwards from an edge on the top part of ∂P . This observation, together with Euler’s formula for planar maps, imply that the complexity of the vertical decomposition of P is proportional to the number of faces of \mathcal{M} (here and in the remainder of this proof, “face” means “2-dimensional face”). Note that we can ignore the last vertical decomposition step, namely that of planar vertical decomposition of each face of \mathcal{M} , because this step increases the overall complexity of the decomposition only by a constant factor.

For clarity of exposition, assume that P is a bounded polytope (which follows from our assumption that the cell C is bounded). Note that each face f of \mathcal{M} is the

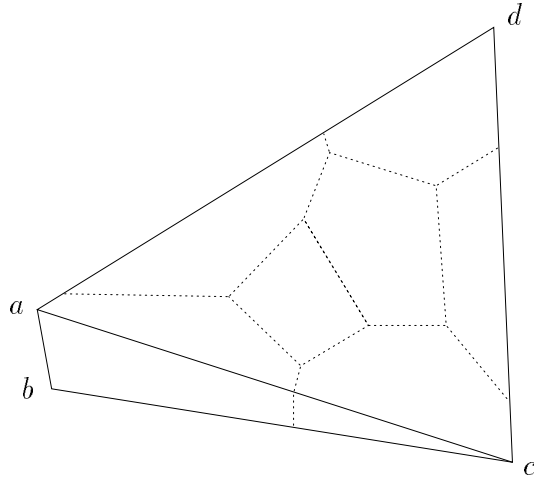


Figure 1: The red cell $abcd$ is a thin tetrahedron so that the red faces acd and bcd have roughly the same intersection pattern with the blue subdivision, denoted by dotted lines; thus the red faces acd and bcd are visible from one another inside several blue cells

intersection of the xy -projections of a face f^+ on the top part of ∂P and of a face f^- on the bottom part of ∂P . Since P is a cell in \mathcal{T} , each of f^+ , f^- is either a portion of a red face of \mathcal{R} or a portion of a blue face of \mathcal{B} . Our goal is to charge each face f of \mathcal{M} (or each *vertically visible* pair (f^+, f^-) of faces of P , which is equivalent) to a pair of features, each being a feature of either \mathcal{R} or \mathcal{B} , so that each such pair will be charged only a constant number of times over the entire collection of cells P . This will clearly imply Lemma 2.2.

Suppose first that f^+ is a portion of a red face r and f^- is a portion of a blue face b (the blue-red case is symmetric). There is a unique red cell R such that r lies on the top part of its boundary, and a unique blue cell B such that b lies on the bottom part of its boundary. Then both f^+ and f^- lie in the intersection $R \cap B$, which is thus the cell P . In other words, we can charge the pair (f^+, f^-) to the pair (r, b) of faces, and the above argument shows that this charge is unique. Thus the number of pairs of this kind is $O(N_{\mathcal{R}}N_{\mathcal{B}})$.

Next consider the case where both f^+ and f^- are portions of two respective red faces r^+ , r^- , necessarily appearing along the top and bottom parts of the boundary of some red cell R (the case of blue faces is fully symmetric). In this case we cannot charge (f^+, f^-) to (r^+, r^-) as above, because R may be split into several subcells by blue cells, and many of them might contain vertically visible pairs of appropriate portions of r^+ and r^- , so the charge need not necessarily be unique; see Figure 1 for an illustration.

Let r be the intersection of the xy -projections of r^+ and of r^- . Let B be the blue cell whose intersection with R is P . If $P = R$ then we can charge (f^+, f^-) to the pair

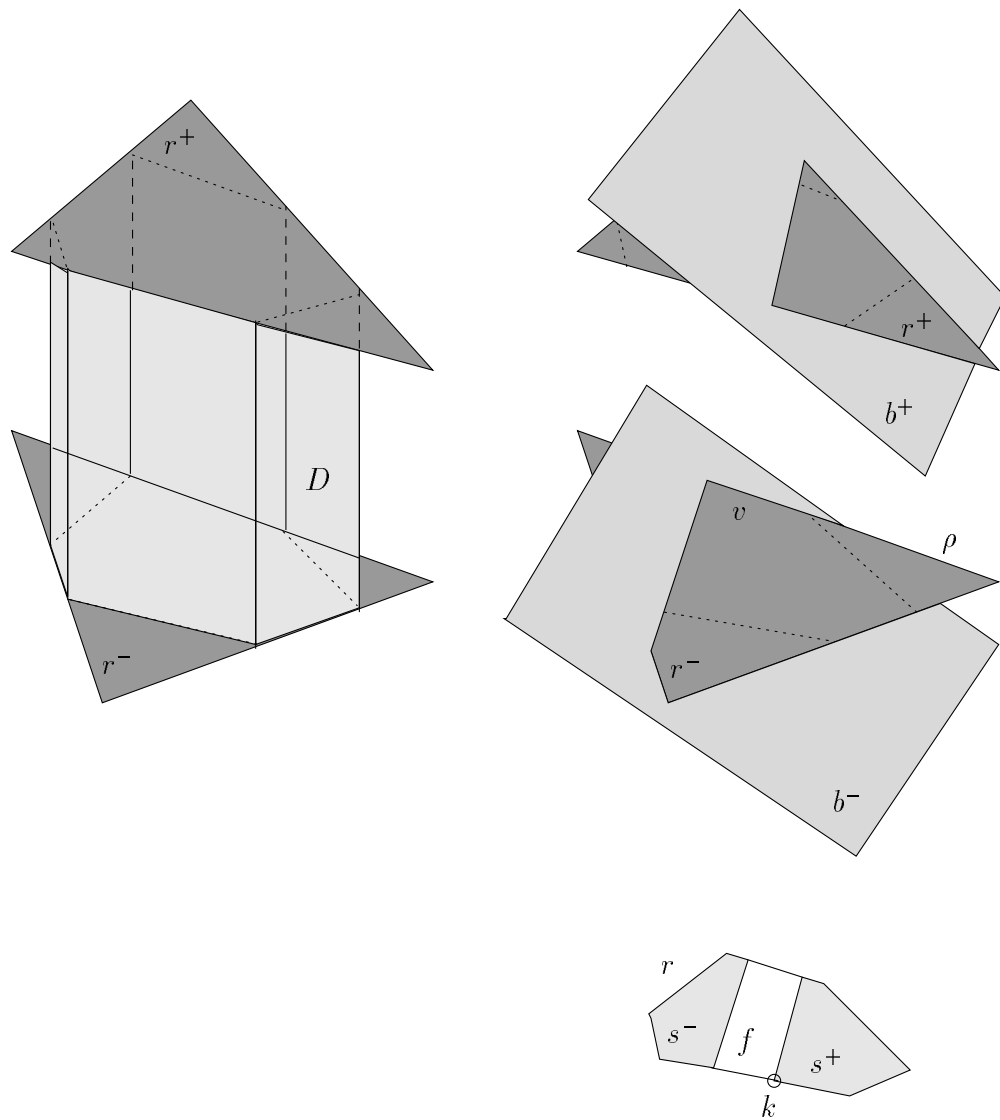


Figure 2: The final case of charging for a pair of vertically visible red faces

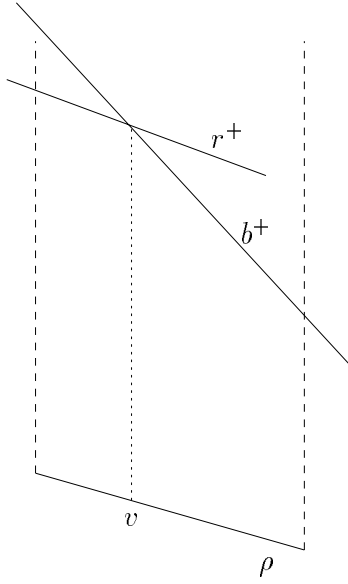


Figure 3: The face above v changes from r^+ to b^+

(r^+, r^-) as above in a unique manner (there will be $O(N_{\mathcal{R}}^2)$ such charges overall), so suppose that P is a proper subcell of R . Note that $f^+ = r^+ \cap B$, $f^- = r^- \cap B$. Let q be the intersection of the xy -projection of B with r . If q contains the projection of some vertex v of B , then we charge the pair (f^+, f^-) to the pair (v, R) , say, and observe that this pair is charged only a constant number of times, because, given v and R , there is a unique pair of red faces of R that the vertical line through v intersects, and (v, R) will be charged only by this pair interacting with the few blue cells incident to v . Similarly, if q is intersected by the projection of an edge β of the cell B , then either this edge has an endpoint that also projects into q , which is the case we have just discussed (namely, q contains the projection of this endpoint), or else β must cross some edge of r , which is the projection of either an edge of r^+ or of an edge of r^- . Suppose, with no loss of generality, that β crosses the projection of an edge ρ of r^+ . Then we charge (f^+, f^-) to the pair (β, ρ) , and again observe that such a pair will be charged only a constant number of times, because, given β and ρ , there is a unique vertical line passing through both β and ρ , and this line uniquely determines the other red face r^- . Thus charges to (β, ρ) can be made only by pairs (r^+, r^-) and cells B such that r^+ is incident to ρ and β is incident to B , and there is clearly only a constant number of such possible charges. Allowing for symmetric cases as well, we conclude that the total number of pairs (f^+, f^-) accounted for so far is $O((N_{\mathcal{R}} + N_{\mathcal{B}})^2)$.

The remaining case is thus when q does not contain the projection of any vertex or edge of B . In this case, assuming B is bounded, r is fully contained in the projection of one top face and of one bottom face of B ; let us denote these faces by b^+ , b^- , respectively. Let D denote the vertical cylinder whose xy -projection is r and which is bounded by r^+ on its top side and by r^- on its bottom side. The face b^+ intersects

D in a convex polygon whose xy -projection, s^+ , is bounded by some portion of ∂r and by at most two straight segments cutting across r (these are the projections of the segments $b^+ \cap r^+$ and $b^+ \cap r^-$; we assume that either at least one such segment exists, or that s^+ is empty, for otherwise $s^+ = r$, in which case b^+ makes f^+ and f^- vertically invisible within P , contrary to assumption). Similarly, b^- intersects D in another convex polygon whose projection, s^- , is also bounded by some portion of ∂r and by at most two other straight segments cutting across r (namely the projections of the segments $b^- \cap r^+$ and $b^- \cap r^-$ truncated to within r ; again at least one such segment must exist unless s^- is empty). The face f must be disjoint from both s^+ and s^- , and adjacent to at least one of them if they are not both empty. It follows from the discussion so far that, in all possible cases, f must extend all the way to the boundary of r . See Figure 2 for an illustration of (one possible case of) this configuration.

Let k be a point on $\partial f \cap \partial r$ which lies on the projection of one of these crossing segments ($b^+ \cap r^+$, $b^+ \cap r^-$, $b^- \cap r^+$, or $b^- \cap r^-$) that appears along the boundary of f , say the projection of $b^+ \cap r^+$. Then k is also the projection of some point lying on an edge ρ of either r^+ or r^- . If ρ is an edge of r^+ then ρ and b^+ intersect (at a point projecting to k), and we can charge (f^+, f^-) to the pair (ρ, b^+) , arguing as above that such a pair can be charged only a constant number of times (this is the situation illustrated in Figure 2). If ρ is an edge of r^- , and v is the point on ρ projecting to k , then, as we walk along ρ in the superimposed subdivision \mathcal{T} , the face directly above us (in the z -direction) changes at v from r^+ to b^+ ; see Figure 3. (If there are no crossing segments, that is if both s^+ and s^- are empty, take k to be a vertex of r which is either the projection of a vertex v^+ of r^+ , or the projection of a vertex v^- of r^- , or the intersection of the projections of an edge ρ^+ of r^+ and an edge ρ^- of r^- . We leave it to the reader to verify that we can charge (f^+, f^-) in each of these three cases respectively to v^+ , to v^- , or to (ρ^+, ρ^-) .)

Let σ be the z -vertical 2-D semi-infinite slab having ρ as its bottom edge, namely σ is the union of rays parallel to the z -axis whose endpoints lie on ρ and which extend in the positive z -direction. Given a collection S of line segments on σ , the *lower envelope* of S (with respect to ρ) is defined in the following manner: for every point p on ρ we extend a ray in the positive z -direction until it hits a segment of S at the point $\psi(p)$, or otherwise it extends to infinity. The lower envelope is defined whenever the ray hits S and it is the union of the points $\psi(p)$; see Figure 4. Note that if a segment in S has an endpoint on ρ then this point belongs to the lower envelope.

Let F denote the union of all the faces of \mathcal{T} . Let ψ denote the lower envelope of the restricted 2-D arrangement of segments $F \cap \sigma$. The analysis in the preceding paragraph suggests that we charge (f^+, f^-) to the ‘breakpoint’ of ψ directly above v . Indeed, this breakpoint, defined by the three features ρ , r^+ , b^+ , identifies the blue cell B and the two red faces r^+ , r^- , up to a constant number of possibilities. Note that ψ is the pointwise minimum of the two subenvelopes $\psi_{\mathcal{R}}$, $\psi_{\mathcal{B}}$, defined as the lower envelopes of the two respective arrangements $\mathcal{R} \cap \sigma$, $\mathcal{B} \cap \sigma$. It easily follows that the number of

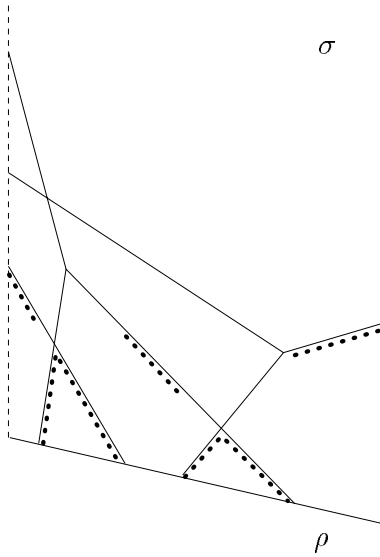


Figure 4: The lower envelope (dotted) of a set of segments on σ

breakpoints along ψ is proportional to the sum of the number of breakpoints along $\psi_{\mathcal{R}}$ and along $\psi_{\mathcal{B}}$ —if one merges the lists of breakpoints of these subenvelopes, sorted in their order along ρ , then there can be at most one new breakpoint of ψ between each pair of adjacent breakpoints in the merged list. The numbers of breakpoints of $\psi_{\mathcal{R}}$, $\psi_{\mathcal{B}}$ are clearly bounded by $N_{\mathcal{R}}$, $N_{\mathcal{B}}$, respectively. Applying this argument to all symmetric cases (obtained by interchanging top and bottom sides, red and blue, etc.), we conclude that the total number of vertically visible face pairs (f^+ , f^-) of the last kind, and thus also the overall number of vertically visible face pairs in \mathcal{T} , is $O((N_{\mathcal{R}} + N_{\mathcal{B}})^2)$. This finishes the proof of Lemma 2.2 and thus also of Lemma 2.1 and Theorem 1.1. \square

Recall that the first step of the vertical decomposition of a 4-dimensional arrangement cell corresponds to overlaying two convex subdivisions in 3-space. The complexity of the overlaid subdivision (before the second decomposition step) can be trivially estimated by $N_{\mathcal{R}}N_{\mathcal{B}}$, where $N_{\mathcal{R}}$, $N_{\mathcal{B}}$ are the total numbers of faces in the subdivisions. It turns out that one can derive a somewhat refined bound:

Theorem 3.1 *Let \mathcal{R} and \mathcal{B} be two simple convex subdivisions of 3-space, so that \mathcal{R} has $n_{\mathcal{R}}$ cells and a total of $N_{\mathcal{R}}$ faces, and \mathcal{B} has $n_{\mathcal{B}}$ cells and a total of $N_{\mathcal{B}}$ faces. Then the total number of faces in the superimposed decomposition is $N_{\mathcal{R}} + N_{\mathcal{B}} + O(n_{\mathcal{R}}n_{\mathcal{B}})$.*

Proof: If a face of \mathcal{R} lies fully within a cell of \mathcal{B} , then it contributes just one to the final face count, and similarly for faces of \mathcal{B} . Suppose r is a red face that intersects a blue cell B but does not lie completely inside it. Then r and ∂B intersect, so either an edge of r crosses a face of ∂B , or an edge of ∂B crosses r . In either case, we charge the pair (r, B) to the resulting vertex. Clearly, no vertex is charged more than

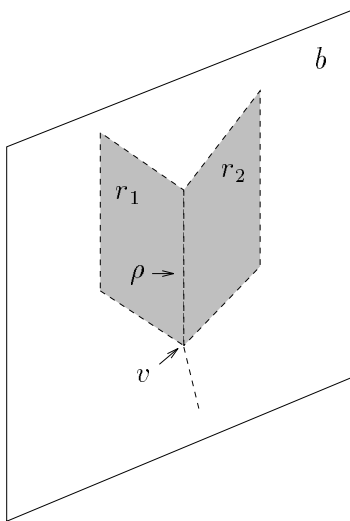


Figure 5: A cone with v as apex at the intersection of b , r_1 and r_2

a constant number of times, so it suffices to bound the number of vertices of these types. Consider, for example, the case of a vertex v formed by intersecting an edge ρ of r with a face b of ∂B . Since \mathcal{R} is simple, ρ is incident to just three faces of \mathcal{R} (one of which is r). These faces intersect b in a triple of segments incident to v . By slightly rotating the coordinate axes, as necessary, we can assume that neither ρ nor any of these segments is horizontal, and we thus may assume, with no loss of generality, that two of these segments increase in z as we traverse them away from v . Let r_1 and r_2 be the two red faces that form these ‘ascending’ segments. See Figure 5 for an illustration. Then r_1 and r_2 form a *convex* wedge and b slices across it producing two trihedral angles, or cones, in the neighborhood of v , with v as apex. It is easily verified that v is the lowest point in one of these cones. In other words, in the superimposed decomposition, which is of course convex, v is the lowest vertex of some cell. Hence the number of vertices under consideration is proportional to the number of cells in the superimposed decomposition, which is at most $n_{\mathcal{R}}n_B$. This argument implies the assertion of the theorem. \square

Corollary 3.2 *If \mathcal{R} and \mathcal{B} have (at most) n cells each, then the complexity of their superposition is $O(n^2)$.*

Proof: This follows immediately from Theorem 3.1 and from the observation that the complexity of each subdivision is $O(n^2)$. \square

Remark. The preceding corollary can be applied to solve problems that involve two distinct Voronoi diagrams in 3-space. For example, given two sets of point sites S_1, S_2 , each of size n , one might want to find a point that satisfies some relationship involving its nearest neighbor in S_1 and its nearest neighbor in S_2 . To find such a point, one may have to traverse all the cells of the subdivision obtained by superimposing the

two Voronoi diagrams of S_1 and of S_2 , and the corollary implies that such a traversal can be done in quadratic time.

Acknowledgement

The authors wish to thank Boris Aronov and Emo Welzl for several helpful discussions concerning the problems studied in this paper, and an anonymous referee for constructive comments.

References

- [1] B. Aronov, J. Matoušek and M. Sharir, On the sum of squares of cell complexities in hyperplane arrangements, *J. Comb. Theory Ser. A* 65 (1994), 311–321.
- [2] F. Aurenhammer, A criterion for the affine equivalence of cell complexes in R^d and convex polyhedra in R^{d+1} , *Discrete Comput. Geom.* 2 (1987), 49–64.
- [3] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Theoretical Computer Science.* 84 (1991), 77–105. Also in *Proc. 16th Int. Colloq. on Automata, Languages and Programming* (1989) pp. 179–193.
- [4] K. Clarkson, A randomized algorithm for closest point queries, *SIAM J. Computing* 17 (1988), 830–847.