Computing D-convex hulls in the plane

Vojtěch Franěk

Department of Applied Mathematics
Charles University
Malostranské nám. 25, 118 00 Praha 1
Czech Republic

Jiří Matoušek

Department of Applied Mathematics and
Institute of Theoretical Computer Science
Charles University
Malostranské nám. 25, 118 00 Praha 1
Czech Republic

Abstract

A function $f: \mathbb{R}^d \to \mathbb{R}$ is called *D-convex*, where D is a set of vectors in $\mathbb{R}^d$, if its restriction to each line parallel to a nonzero $v \in D$ is convex. The *D-convex hull* of a compact set $A \subset \mathbb{R}^d$ is the intersection of the zero sets of all nonnegative D-convex functions that are 0 on A. Matoušek and Plecháč provided an algorithm for computing the D-convex hull of a finite set in $\mathbb{R}^d$ for D consisting of d linearly independent vectors (in this case one speaks about *separately convex* hulls). Here we present a (polynomial-time) algorithm for the D-convex hull of a finite point set in the plane for arbitrary finite D.

1 Introduction

Let D be a set of vectors (directions) in $\mathbb{R}^d$. We investigate the algorithmic problem of computing the *D-convex hull* of a given n-point set $A \subset \mathbb{R}^d$. Matoušek and Plecháč [MP98] provided an algorithm for the case where $D = \{e_1, e_2, \ldots, e_d\}$ is the standard basis in $\mathbb{R}^d$; in this case we speak of the *separately convex hull*. Here we deal with the case of an arbitrary finite D in the plane, i.e., $d = 2$.

***D*-convexity.** The notion of D-convex hull considered here is defined using functions, as follows. A *D-line* is a line in $\mathbb{R}^d$ parallel to a nonzero vector from D. A function $f: \mathbb{R}^d \to \mathbb{R}$ is called *D-convex* if its restriction to each D-line is convex. For $d = 2$ and $D = \{e_1, e_2\}$, i.e., separate convexity in the plane, a basic example of a D-convex function that is not convex in the usual sense is $f(x, y) = \max(x, 0) \max(y, 0)$.

A compact set $B \subset \mathbb{R}^d$ is D-convex if it is the zero set of some D-convex function $f: \mathbb{R}^d \to [0, \infty)$, and the D-convex hull of a compact set $A \subseteq \mathbb{R}^d$, denoted by $\mathrm{conv}^D(A)$, is the intersection of all D-convex sets containing A (we note that D-convex sets are closed under intersections, since the pointwise maximum of D-convex functions is D-convex).

There is another notion of convexity of a set with respect to a set D of directions, which looks more natural at first sight than the definition of D-convexity just given. Namely, we call a set $B \subseteq \mathbb{R}^d$ *D-lamination convex* if every D-line intersects B in a (possibly empty) interval. (Lamination convexity has also been studied in computational geometry under other names, such as *restricted-orientation convexity*; see, e.g., [FW96, FW98]).

Every D-convex set is obviously also D-lamination convex, but as was discovered independently by a number of authors in several contexts ([Sch74], [AH86], [NM91], [Tar93], [CT93]), the converse need not hold.[1] A basic example of this phenomenon is a 4-point planar configuration, commonly denoted by T_4, depicted in Fig. 1 on the left (black dots). The 4-point set itself is separately lamination-convex, but its separately convex hull, shown in the picture, consists of a square and four segments. Fig. 1 right shows an 18-point configuration in $\mathbb{R}^3$, which is again separately lamination-convex, and its separately convex hull [Mat01].

The definition of the D-convex hull of A talks about *all* nonnegative D-convex functions vanishing on A, but it is known [Mat01] that it suffices to consider one particular D-convex function. Namely, $\mathrm{conv}^D(A)$ equals the zero set of the D-convexification of the function δ_A, where $\delta_A(x)$ is the (Euclidean) distance of x from A, and the *D-convexification* of a function $f: \mathbb{R}^d \to \mathbb{R}$ is the pointwise supremum of all D-convex functions $g: \mathbb{R}^d \to \mathbb{R}$ that satisfy $g(x) \leq f(x)$ for all $x \in \mathbb{R}^d$. This suggests a possible approach to approximating $\mathrm{conv}^D(A)$: compute the D-convexification of δ_A approximately by an iterative algorithm, and look where it is very close

[1] In order to distinguish these two notions, what we call D-convex was called *functionally D-convex* in [MP98], while *set-theoretically D-convex* was used instead of our D-lamination convex. For $D = \{e_1, \ldots, e_d\}$, i.e. separate convexity, D-lamination convexity is also called *rectilinear convexity* or *ortho-convexity* in the literature.

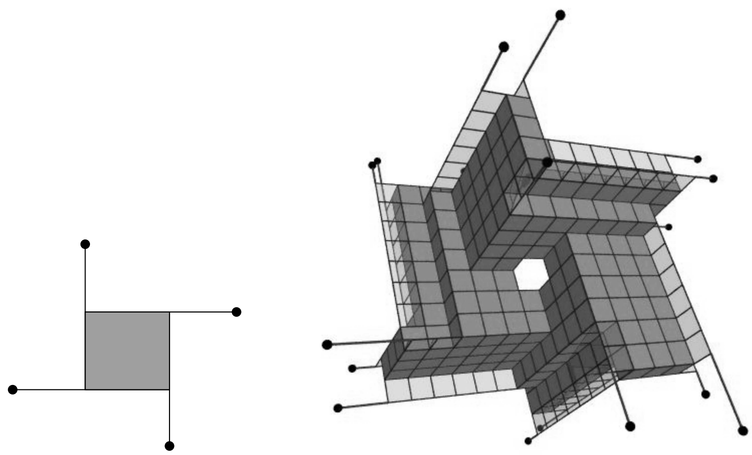

Figure 1: The separately convex hull of a 4-point configuration in $\mathbb{R}^2$ (left) and of an 18-point configuration in $\mathbb{R}^3$ (right). The hull on the right is drawn as a part of the integer grid, made of unit cubes, unit squares, and unit segments.

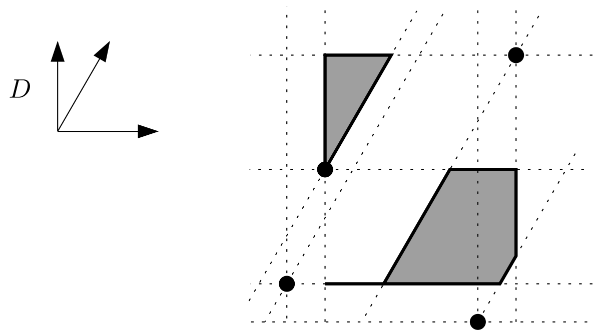

Figure 2: An example of a (D, A)-polygon (disconnected and not D-convex). Points of A are marked by black dots.

to 0. Unfortunately, an example in [Mat01] shows that in order to approximate $\mathrm{conv}^D(A)$ in this way, in some cases we need to compute the D-convexification of δ_A with unrealistically high precision (exponential in $|A|$). Hence a different kind of algorithm is needed, based on better understanding of the properties of the D-convex hull.

The result. To state the main result, we need some notions. Let $A \subset \mathbb{R}^2$ be a finite set and $D \subset \mathbb{R}^2$ a finite set of nonzero vectors. From now on, we will always assume that D spans $\mathbb{R}^2$. A (D, A)-*line* is a D-line passing through a point of A. A (D, A)-face is a face in the arrangement of all (D, A)-lines,[2] and a (D, A)-*polygon* is a set that can be expressed as the union of closures of (D, A)-faces; see Fig. 2.

Theorem 1.1 *Let* $A \subset \mathbb{R}^2$ *be a finite set and* $D \subset \mathbb{R}^2$ *a set of nonzero vectors. Then* $\mathrm{conv}^D(A)$ *is a* (D, A)*-polygon and it can be computed in time polynomial in* $|D|$ *and* $|A|$.

The complexity of the algorithm in the theorem is considered in the *Real RAM* model common in computational geometry, which permits unit-time arithmetic operations with real numbers. A brute-force estimate, with no sophistication in the implementation, leads to an $O(|D|^6|A|^6)$ complexity

[2]We recall that a face in the arrangement of a finite set L of lines in $\mathbb{R}^2$ is a *vertex*, i.e., an intersection of two or more lines of L, or an *edge*, i.e., a connected component of $\ell \setminus \bigcup_{\ell' \in L \setminus \{\ell\}} \ell'$ for some line $\ell \in L$, or a *cell*, i.e., a connected component of $\mathbb{R}^2 \setminus \bigcup_{\ell \in L} \ell$.

bound. Using standard data structures and tricks this bound could easily be reduced, but since our algorithm is unlikely to be optimal, we prefer not to elaborate on this issue.

Motivation and background. The main motivation for studying D-convexity comes from the particular case of *rank-one convexity.* This is D-convexity in $\mathbb{R}^d$ where $d = mn$, $\mathbb{R}^d$ is identified with the space of all real $m \times n$ matrices, and D corresponds to the set of all $m \times n$ matrices of rank 1. Rank-one convexity appears in several areas (calculus of variations, theory of partial differential equations, the existence of Lipschitz mappings with a prescribed set of gradients, and mathematical models of crystalline microstructure and phase transitions). Müler's lecture notes [Mül99] can serve as an introduction to the wider context and applications, Dolzmann [Dol03] focuses on models in physics, and [Szé06], [Szé05] are examples of recent works investigating the geometry of rank-one convexity per se.

Kirchheim, Müller and Šverák [KMŠ03] discuss rank-one convexity together with several other important notions of convexity, and relations among these notions. They also give interesting examples and results for D-convexity in $\mathbb{R}^3$ with $D = \{(x, y, 0) : x, y \in \mathbb{R}\} \cup \{(0, 0, z) : z \in \mathbb{R}\}$.[3] As one of the main open problems, they ask for an algorithm for deciding whether a given finite set of matrices is rank-one convex. We hope that investigating algorithmic aspects of D-convexity for "simpler" sets of directions D might eventually help in tackling this problem and other questions related to rank-one convexity.

2 The algorithm: Local biting

Before presenting the algorithm for D-convex hull in the plane, we first outline the algorithm from [MP98] for computing the separately convex hull, whose basic approach is somewhat similar.

We explain it on a concrete planar example, mainly by pictures. Fig. 3 left shows a set A, whose points are drawn as black dots. First we form the rectangular grid defined by A (dotted lines) and we initially mark all grid vertices not lying in A white.

Then we find a white point that, for each of the two coordinate directions, lacks at least one of the two neighbors in that direction, and we

[3] They refer to this case as "separate convexity" as well. This corresponds to considering $\mathbb{R}^3$ as the direct sum $\mathbb{R}^2 \oplus \mathbb{R}$. For the separate convexity with $D = \{e_1, e_2, e_3\}$ we regard $\mathbb{R}^3$ as $\mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R}$.

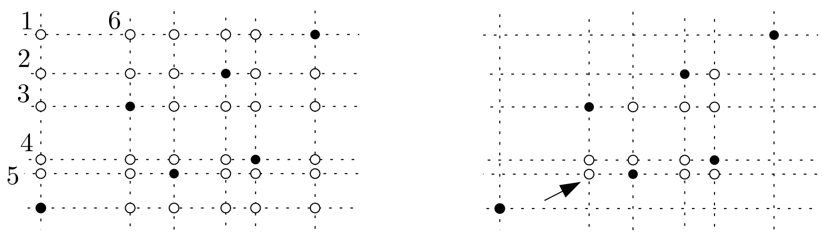

Figure 3: Computing the separately convex hull: setting up the grid (left); an intermediate stage (right).

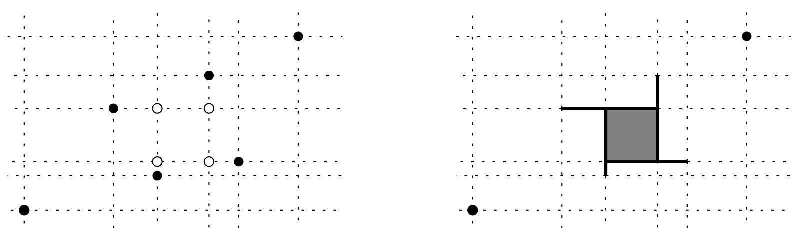

Figure 4: Final stage of the algorithm (left); the resulting separately convex hull (right).

remove it. In the first step we can take, for example, the upper left corner, labeled 1 (the other candidate would be the lower right corner). After removing 1, the white point 2 can be removed, etc. (the labels indicate the first 6 steps). Fig. 3 right shows an intermediate stage, where some white points can still be removed, say the one indicated by arrow, while Fig. 4 left depicts the final stage, where none of the four remaining white points can be removed (black points are never removed). Then the separately convex hull of A is reconstructed as the "box complex" induced by the remaining set of black and white points (the box complex consists of all boxes in the grid, of dimension 0, 1, or 2, for which all vertices are black or white—see Fig. 4 on the right).

The basic intuition of the algorithm described above, which we will retain, is that in each iteration we can "bite off" a "local corner". However, for more than two directions we have no reasonable analogy of the grid used in the algorithm just described, and we will have to work directly with polygons.

A "local corner" to be removed is just a convex vertex of the current polygon, but we also need to specify what portion of the polygon can be cut off. Conceptually, we divide this into two steps: First, depending on the local shape in a small neighborhood of the convex vertex, we define a suitable wedge-shaped "cutter" and we place its tip to the vertex. Then we imagine that the cutter starts moving and removing matter from the polygon, and we specify how far it can proceed before stopping—essentially, until its position is fixed by two points of A on the boundary.

Cutters. We proceed with some formal definitions. Let B be a (D, A)-polygon. A *convex vertex* of B is a point $p \in B$ such that there exists an open neighborhood U of p and a closed halfplane H whose boundary line passes through p, such that $(B \cap U) \cap H = \{p\}$. We may assume that the boundary of H is not a D-line (by slightly tilting it if needed); let us call such an H a *good supporting halfplane* of p.

For a convex vertex p of B we now define a *cutter for p*. We set $D^{\mathrm{sym}} := D \cup (-D)$, we let H be a good supporting halfplane of p, and we let the cutter for p be

$$T_0 := p + \operatorname{cone}\{v \in D^{\mathrm{sym}} : p + v \in \operatorname{int} H\},$$

where int denotes the interior and $\operatorname{cone}(W)$ denotes the convex cone spanned by W, that is, the set of all linear combinations of vectors of W with non-negative coefficients; see Fig. 5. Since D spans $\mathbb{R}^2$, T_0 is always a convex wedge with nonempty interior (and with tip at p). Let v_1 and v_2 denote the two vectors of D^{sym} corresponding to the two bounding rays of T_0, and let n_1 and n_2 denote the two inward normal vectors of these rays.

It is not difficult to check that the set $\mathbb{R}^2 \setminus \operatorname{int} T_0$ is the zero set of a nonnegative D-convex function,[4] namely, of the function

$$x \mapsto \max(\langle x - p, n_1 \rangle, 0) \max(\langle x - p, n_2 \rangle, 0).$$

For checking the D-convexity of this function, the key fact is that no D-line "cuts across" T_0; we omit the routine but slightly laborious details.

[4]This set is unbounded, and so we don't want to speak about its D-convexity, since we have defined only *compact* D-convex sets. Indeed, the usual definition of D-convexity for unbounded sets is generally not equivalent to being a zero set of a nonnegative D-convex function.

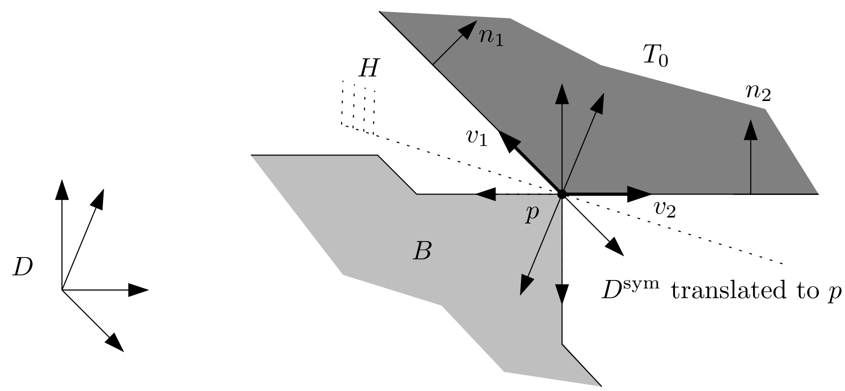

Figure 5: A cutter for a vertex v.

Fixing a cutter. In order to cut off a convex vertex p, we need to translate the cutter T_0 "into" B so that it gets stopped by a point of A on both of the rays. Formally, we set

$$\begin{aligned} \alpha_1 &:= \min\{\alpha > 0 : (T_0 - \alpha v_1) \cap A \neq T_0 \cap A\}, \\ \alpha_2 &:= \min\{\beta > 0 : (T_0 - \alpha_1 v_1 - \beta v_2) \cap A \neq (T_0 - \alpha v_1) \cap A\}, \\ T &:= T_0 - \alpha_1 v_1 - \alpha_2 v_2. \end{aligned}$$

Less formally, we first translate T_0 in the direction $-v_1$ until it swallows a new point of A, and then we translate the resulting set in the direction $-v_2$ until it swallows another point of A. It might happen that the first translation fails to reach a new point of A; i.e., $\alpha_1 = \infty$. In that case we consider $T_0 - \alpha_1 v_1$ as a halfplane. The case $\alpha_2 = \infty$ is handled similarly. So T can be a convex wedge, a halfplane, or the whole plane (the last case can't actually occur, but we prefer not to prove this), and it always contains p in the interior. We call T a *cutter for p in a fixed position.* Now we are ready to describe the algorithm.

The algorithm for planar *D*-convex hull. Given A and D, we let B_1 be a bounded convex (D, A)-polygon containing A; for example, we can take a suitable parallelogram. Then we repeat the following procedure for $s = 1, 2, \ldots$:

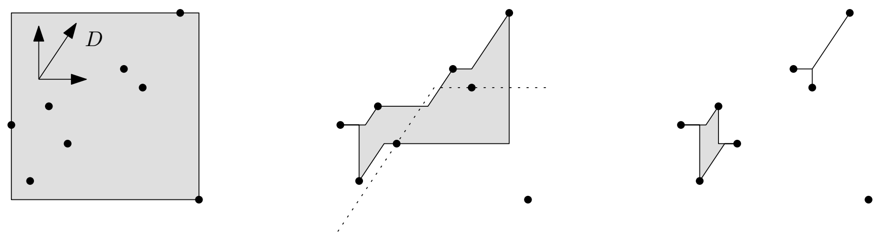

Figure 6: Illustration of the algorithm: The beginning; an intermediate step; the final result.

1. Suppose that a (D, A)-polygon B_s has already been computed. Find a convex vertex p of B_s not lying in A. If there is no such vertex, the algorithm ends with B_s as the output.

2. Let T be a cutter for p in a fixed position. Let R be the connected component of $B_s \cap T$ that contains p.

3. Set $B_{s+1} := B_s \setminus (R \cap \operatorname{int} T)$, $s := s + 1$, and continue with the next iteration.

Fig. 6 illustrates a possible execution of the algorithm with D containing three directions (indicated in the left picture). The left picture is the beginning, and the middle picture shows an intermediate step, where the lower right corner of the gray polygon is going to be cut by the cutter indicated by dotted line. We note that the polygon will be broken into two components. The rightmost picture is the final outcome.

The correctness of the algorithm relies on the following two lemmas:

Lemma 2.1 *Let $B_1, B_2, \ldots, B_{s_0}$ be the sequence of polygons constructed by the algorithm.*

(i) *Each B_s is a bounded (D, A)-polygon.*

(ii) *Each B_s is D-convex.*

(iii) *$A \subseteq B_s$ for all s.*

Lemma 2.2 *If B is a bounded D-convex (D, A)-polygon containing A such that all convex vertices of B lie in A, then $B = \operatorname{conv}^D(A)$.*

To see that these two lemmas imply Theorem 1.1, it suffices to check that the algorithm terminates in polynomially many iterations. This is because $B_1, B_2, \ldots$ is a strictly decreasing sequence of (D, A)-polygons (w.r.t. inclusion), and thus their number is bounded above by the total number of faces in the arrangement of all (D, A)-lines, which is $O(|D|^2|A|^2)$.

Part (i) of Lemma 2.1 is easy and we prove it next. Parts (ii) and (iii) take more work, and we deal with them in the subsequent two sections. Finally, Lemma 2.2 is proved in Section 5.

Proof of Lemma 2.1(i). We proceed by induction on s, assuming that B_s is a (D, A)-polygon. By a set-theoretic manipulation we can rewrite $B_{s+1} = B_s \setminus (R \cap \operatorname{int} T) = (B_s \setminus R) \cup (B_s \setminus \operatorname{int} T) = ((B_s \cap T) \setminus R) \cup (B_s \cap (\mathbb{R}^2 \setminus \operatorname{int} T))$ (for the last equality, we note that all points of B_s not contained in T are also in $B_s \setminus \operatorname{int} T$). Now T, a cutter in a fixed position, is a (D, A)-polygon, and so is $\mathbb{R}^2 \setminus \operatorname{int} T$. Since, clearly, (D, A)-polygons are closed under unions and intersections, we get that $B_s \cap (\mathbb{R}^d \setminus \operatorname{int} T)$ and $B_s \cap T$ are (D, A)-polygons. Then $(B_s \cap T) \setminus R$ is the union of all connected components of $B_s \cap T$ except for R, and thus a (D, A)-polygon too, and we can conclude that B_{s+1} is a (D, A)-polygon as claimed. □

3 The Cutter Lemma

Here we establish Lemma 2.1(ii), a key step in the proof of correctness of our algorithm. We actually formulate and prove a "Cutter Lemma" for an arbitrary dimension d, since we expect it to be a useful tool for a higher dimensional D-convex hull algorithm, and since it is not much harder than a planar version.

Lemma 3.1 (Cutter Lemma) *Let D be an arbitrary set of directions in $\mathbb{R}^d$, let $B \subset \mathbb{R}^d$ be a compact D-convex set, and let $T \subseteq \mathbb{R}^d$ (the "cutter") be closed and such that $\mathbb{R}^d \setminus \operatorname{int} T$ is the zero set of a nonnegative D-convex function. Let R be a subset of $B \cap T$ such that the distance of R from its complement in $B \cap T$, i.e. from $(B \cap T) \setminus R$, is strictly positive (see Fig. 7). Then the set $\check{B} = B \setminus (R \cap \operatorname{int} T)$ is D-convex.*

Lemma 2.1(ii) immediately follows from the Cutter Lemma by induction on s.

Proof. Let $f\colon \mathbb{R}^d \to [0, \infty)$ be a D-convex function witnessing the D-convexity of B (i.e., B is the zero set of f) and let $g\colon \mathbb{R}^d \to [0, \infty)$ be a

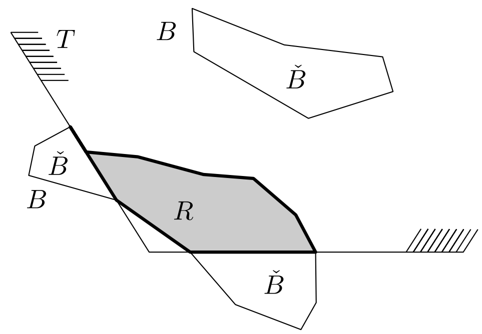

Figure 7: The Cutter Lemma.

D-convex function whose zero set is $\mathbb{R}^d \setminus \operatorname{int} T$. We may also assume that f and g are continuous (see, e.g., [MP98]).

For a set $R \subset \mathbb{R}^d$ and $\delta > 0$, let R^δ denote the closed δ-neighborhood of R.

Let $\delta > 0$ be half of the distance of R to $(B \cap T) \setminus R$, and let $\eta > 0$ be a sufficiently small real number, to be determined later. We define a function $h\colon \mathbb{R}^d \to [0, \infty)$ by

$$h(x) = \begin{cases} \max(f(x), \eta \cdot g(x)) & \text{for } x \in R^\delta \cap T \\ f(x) & \text{otherwise.} \end{cases}$$

Claim I: $\check{B}$ is the zero set of h.
Proof of Claim I. The function h is nonzero on the complement of B since $f > 0$ there, and it is nonzero on $R \cap \operatorname{int} T$ since $g > 0$ there. Thus $h(x) > 0$ for all $x \notin \check{B}$. Now let $x \in \check{B}$. Then $x \in B \setminus \operatorname{int} T$ or $x \in (B \cap T) \setminus R$. In the former case both $f(x) = 0$ and $g(x) = 0$, and so $h(x) = 0$ as well. For the latter case we get, by the condition on R in the lemma, that $x \notin R^\delta \cap T$, and hence $h(x) = f(x)$ and $f(x) = 0$ since $x \in B$. This proves Claim I.

Claim II: Let S be the closure of the set $(R^\delta \setminus R^{\delta/2}) \cap T$. Then the minimum of f on S is strictly positive.
Proof of Claim II (cf. Fig 8). Since S is compact, it suffices to show $f(x) > 0$ for all $x \in S$, and this is the same as $S \cap B = \emptyset$. Any point of S has distance at least $\delta/2$ to R, and hence an $x \in S \cap B$ would have to lie in $(B \cap T) \setminus R$, but this set is disjoint from R^δ. This proves Claim II.

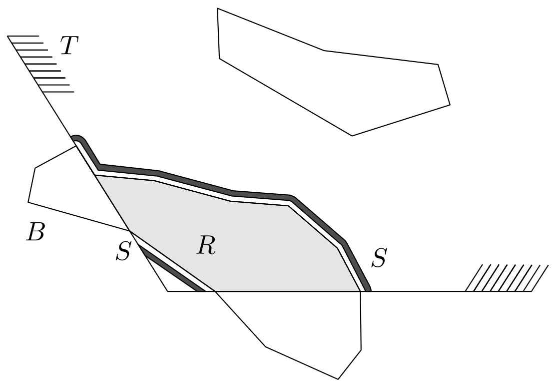

Figure 8: An illustration to Claim II.

Let us set $\mu = \min_{x \in S} f(x)$ and $M = \max_{x \in S} g(x)$. We have $\mu > 0$ by Claim II and $M < \infty$. We put $\eta = \mu/M$. Then $f \geq \eta \cdot g$ on S and hence $h(x) = f(x)$ for $x \in S$.

Let ℓ be a D-line. For a more convenient notation let us identify ℓ isometrically with the real axis, so that, for example, $a + \varepsilon$ is a well-defined point of ℓ for $a \in \ell$ and $\varepsilon \in \mathbb{R}$. Let us write h_ℓ for h restricted to ℓ, and similarly for f_ℓ and g_ℓ.

Claim III: Suppose that $\eta > 0$ is as above and ℓ is an arbitrary D-line. Then at least one of the following holds for every point $a \in \ell$:

(i) There exists $\varepsilon = \varepsilon(a) > 0$ such that $h_\ell(x) = f_\ell(x)$ for all $x \in (a - \varepsilon, a + \varepsilon)$.

(ii) There exists $\varepsilon = \varepsilon(a) > 0$ such that $h_\ell(x) = \max(f_\ell(x), \eta \cdot g_\ell(x))$ for all $x \in (a - \varepsilon, a + \varepsilon)$.

(iii) $f(a) = g(a) = h(a) = 0$.

Proof of Claim III. We distinguish several cases, depending on the position of a.

- For $a \notin R^\delta \cap T$ there is a neighborhood of a avoiding $R^\delta \cap T$, and so (i) applies.

- For a lying on the boundary of T we have $g_\ell(a) = 0$. If (i) doesn't apply, then arbitrarily small neighborhoods of a on ℓ contain points x with $h_\ell(x) \neq f_\ell(x)$. It follows from the definition of h that $f_\ell(x) \neq h_\ell(x) = \max(f_\ell(x), \eta \cdot g_\ell(x))$ and hence $f_\ell(x) < \eta \cdot g_\ell(x)$. Then we get $f_\ell(a) = 0$ by the continuity of f and g, and (iii) applies.
- For $a \in \operatorname{int}(T \cap R^\delta)$ there is a neighborhood of a contained in $\operatorname{int}(T \cap R^\delta)$, and hence (ii) applies.
- The remaining case is a lying in $\operatorname{int} T$ and on the boundary of R^δ. If we choose $\varepsilon = \delta/2$ then each $x \in (a-\varepsilon, a+\varepsilon)$ lies either in S or outside R^δ, and in both cases we have $h_\ell(x) = f_\ell(x)$. Hence (i) applies.

Claim III is proved.

Claim III shows that the function h_ℓ has the following property: For every point $a \in \ell$ there exists $\varepsilon = \varepsilon(a) > 0$ and an affine function $\sigma_a \colon \ell \to \mathbb{R}$ that is supporting for h_ℓ on $(a - \varepsilon, a + \varepsilon)$; that is, $\sigma_a(a) = h_\ell(a)$ and $\sigma_a(x) \leq h_\ell(x)$ for all $x \in (a-\varepsilon, a+\varepsilon)$. Indeed, in cases (i) and (ii) h_ℓ equals to a convex function on $(a-\varepsilon, a+\varepsilon)$ and hence it has a supporting tangent at a, and in case (iii) $\sigma_a(x) = 0$ is supporting since h_ℓ is nonnegative. The existence of a supporting affine function on some neighborhood of every point implies convexity of h_ℓ (see, e.g., Lemma 3.1 in [MR03]), and hence h is D-convex. Lemma 3.1 is proved. □

4 No point of A is cut

Here we prove Lemma 2.1(iii), namely, that $A \subseteq B_s$ throughout the algorithm.

Naturally we proceed by induction on s, where the case $s = 1$ is clear by the choice of B_1. For the induction step it suffices to show that $R \cap \operatorname{int} T$ never contains a point of A. This does need an argument since $\operatorname{int} T$ may contain points of A.

According to the procedure of translating the cutter T_0 to a fixed position T, we have $((\operatorname{int} T) \setminus T_0) \cap A = \emptyset$, so if a point $a \in R \cap \operatorname{int} T$ existed, it would have to lie in T_0; see Fig. 9. Since R is the connected component of $B_s \cap T$ containing p, such an a would have to be connected to p by a path γ within $B_s \cap T$.

Since there is a neighborhood U of p with $T_0 \cap (B_s \cap U) = \{p\}$, any path from p within B_s has to leave T_0 at the beginning.

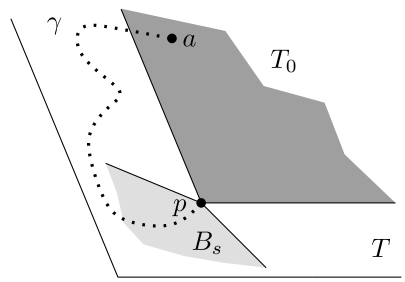

Figure 9: No points of A are ever cut off.

Moreover, the boundary rays of T_0 are parallel to directions in D, and since B_s is D-convex, they do not contain points of B_s distinct from p (for otherwise, one of them would contain a segment lying in B_s and ending at p, but this is not the case because $T_0 \cap (B_s \cap U) = \{p\}$). But then the path γ cannot cross the boundary of T_0 and reach a—a contradiction proving Lemma 2.1(iii).

5 Cutting all convex vertices suffices

Here we prove Lemma 2.2. For the purposes of this section we suppose that D consists of nonzero vectors and that none of them is a multiple of another.

First let B be a D-convex set. Following [MP98], we call a point $p \in B$ *D-extremal* if there exists no segment $s \subseteq B$ parallel to some nonzero $v \in D$ and containing p in its interior.

We are going to use the following result of [MP98]: If $A \subseteq B$ are compact sets in $\mathbb{R}^d$ such that B is D-convex and all D-extremal points of B belong to A, then $B = \mathrm{conv}^D(A)$.

In the setting of Lemma 2.2, $B \supseteq A$ is D-convex and has no *convex* vertices outside A. *If* we could prove that all D-extremal points of B are convex vertices, Lemma 2.2 would follow. Interestingly, it turns out that the statement just mentioned is not always true, but there is only a single type of an exception, which occurs only if D has exactly three directions.

In order to describe the exception, let us assume that D consists of three distinct directions v_1, v_2, v_3. We define a (D, A)-*tripod* as the union of three

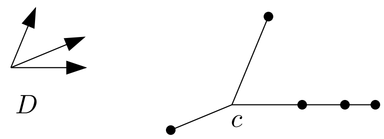

Figure 10: A (D, A)-tripod (dots mark points of A).

segments s_1, s_2, s_3, where

- s_i is parallel to v_i,
- s_1, s_2, s_3 have a common endpoint c, called the *center* of the tripod,
- the other endpoints of the s_i (the *tips* of the tripod) all belong to A, and
- c is in the convex hull of the tips;

see Fig. 10.

We now formulate two lemmas that together imply Lemma 2.2.

Lemma 5.1 *Let D consist of $t \geq 2$ distinct directions, let B be a bounded D-convex (D, A)-polygon with all convex vertices belonging to A, and let p be a D-extremal point of B. Then either p is a convex vertex of B, or $t = 3$ and p is the center of a (D, A)-tripod forming a connected component of B.*

Lemma 5.2 *Let a_1, a_2, a_3 be the tips of a (D, A)-tripod T. Then $T = \text{conv}^D(\{a_1, a_2, a_3\})$.*

Proof of Lemma 2.2. If $B \supseteq A$ is a D-convex bounded (D, A)-polygon with all convex vertices lying in A, then by Lemma 5.1, either all D-extremal points of B lie in A (and we are done by the result of [MP98] cited at the beginning of this section), or $|D| = 3$ and each D-extremal point of B not lying in A is the center of a (D, A)-tripod forming a connected component of B.

In this case we use the following result (a combination of Proposition 2.8 of [MP98] and Proposition 6.1 of [Mat01]):

> If $A \subseteq \mathbb{R}^d$ is a set contained in a D-convex set B that is the disjoint union of compact sets $B_1, B_2, \ldots, B_k$, then each B_i is D-convex and $\text{conv}^D(A) = \bigcup_{i=1}^k \text{conv}^D(A \cap B_i)$.

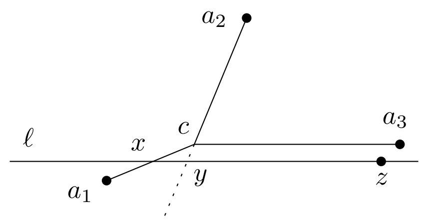

Figure 11: An argument showing that the center of a tripod lies in its D-convex hull.

So if we let $B_1, \ldots, B_k$ be the connected components of the (D, A)-polygon B considered in the lemma, each B_i either is a (D, A)-tripod, and in this case $\mathrm{conv}^D(A \cap B_i) = B_i$ by Lemma 5.2, or it is a compact D-convex set in which each D-extremal point belongs to A, and so $\mathrm{conv}^D(A \cap B_i) = B_i$ as well. Thus, $\mathrm{conv}^D(A) = \bigcup_{i=1}^k \mathrm{conv}^D(A \cap B_i) = \bigcup_{i=1}^k B_i = B$ as claimed. □

Proof of Lemma 5.2. The facts already proved about the correctness of our algorithm (Lemma 2.1) imply that $\mathrm{conv}^D(\{a_1, a_2, a_3\}) \subseteq T$ (apply the algorithm to $A = \{a_1, a_2, a_3\}$). For the reverse inclusion, it suffices to show $c \in \mathrm{conv}^D(\{a_1, a_2, a_3\})$, where c is the center of T.

Let f be a nonnegative D-convex function with $f(a_1) = f(a_2) = f(a_3) = 0$. For contradiction, we suppose that $f(c) > 0$. By continuity, we can choose a small open neighborhood U of a_3 on which f is bounded above by $f(c)$. Let us fix a line ℓ parallel to ca_3 passing through a point $z \in U$ and crossing the segment $s_1 = a_1c$ in an interior point x; see Fig. 11. Then ℓ intersects the line a_2c in a point y lying on the other side of c than a_2.

Since f is convex on the line a_2c and $f(a_2) = 0 < f(c)$, we have $f(y) > f(c)$. Similarly $f(x) < f(c)$, but then f cannot be convex on ℓ, since $f(x) < f(y) > f(c) \geq f(z)$. □

Proof of Lemma 5.1. If p is a D-extremal point of a (D, A)-polygon B, then it is necessarily a vertex of the arrangement of all (D, A)-lines (for otherwise, it would be contained in an open segment of a D-line).

Let us consider a neighborhood U of p containing no other vertices of the arrangement and look at the local shape of B inside U. The D-lines passing through p form $2t$ rays emanating from p, and we also have $2t$ open angular sectors between these rays. If B contains one of these sectors, then

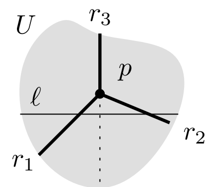

Figure 12: For more than three directions a D-extremal point is a convex vertex.

it also contains both of its boundary rays (within U, that is). Since p is D-extremal, we never have two opposite rays simultaneously present in B.

Assuming that p is not a convex vertex of B, there must be three of the rays, r_1, r_2, r_3, present in B such that p is contained in their convex hull.

Now if $t \geq 4$, there is a direction $v \in D$ distinct from the directions of r_1, r_2, r_3, and we can place a D-line so that it intersects two of the rays, say r_1 and r_2, near p; see Fig. 12. But this contradicts the D-convexity of B, since the ray opposite to r_3 is not in B. Hence for $t \geq 4$, every D-extremal point is a convex vertex.

Now let $t = 3$. Then $B \cap U$ consists exactly of the three rays r_1, r_2, r_3 (none of the sectors can be present, since otherwise, p would not be D-extremal). Let s_i be the intersection of the ray r_i (extended from p all the way to infinity) with B; by D-convexity, s_i is a segment. We let $T = s_1 \cup s_2 \cup s_3$.

We observe that there is an open neighborhood W of T with $W \cap B = T$. Indeed, if $x \in B$ were a point sufficiently near to s_1, say, but not in T, then the D-line through x parallel to s_1 would intersect s_2 or s_3, and we would get a contradiction to the D-convexity of B (here we use that $B \cap U \subseteq r_1 \cup r_2 \cup r_3$). Finally, the endpoints of the s_i distinct from p belong to A since they are convex vertices of B. Thus, T is a (D, A)-tripod forming a connected component of B as claimed. □

6 Concluding remarks

The planar algorithm has been implemented by the first author using the Computational Geometry Algorithms Library (CGAL) [CGA] and, in particular, the package for dealing with planar Nef polyhedra [See01].

An obvious challenge is extending the algorithm to dimensions 3 and higher. While some of the tools and ideas from the planar case may be useful for this, it seems that significant obstacles still have to be overcome.

Another problem is designing an efficient implementation of the presented planar algorithm and obtaining a good worst-case bound for its running time. A related purely combinatorial problem is the maximum possible complexity (number of vertices, say) of the D-convex hull of a planar set A, in terms of $|D|$ and $|A|$. At present, we have only the obvious $O(|D|^2|A|^2)$ upper bound that follows immediately from the fact established in this paper, namely, that the D-convex hull is a (D, A)-polygon.

Acknowledgment

We would like to thank two anonymous referees for careful reading and many corrections and useful comments.

References

[AH86] R. Aumann and S. Hart. Bi-convexity and bi-martinagles. *Israel. J. Math.*, 54(2):159–180, 1986.

[CGA] CGAL, Computational Geometry Algorithms Library. http://www.cgal.org.

[CT93] E. Casadio-Tarabusi. An algebraic characterization of quasiconvex functions. *Ricerche Mat.*, 42:11–24, 1993.

[Dol03] G. Dolzmann. *Variational methods for crystalline microstructure. Analysis and computation.* Lecture Notes in Mathematics 1803, Springer, Berlin, 2003.

[FW96] E. Fink and D. Wood. Fundamentals of restricted-orientation convexity. *Information Sciences*, 92:175–196, 1996.

[FW98] E. Fink and D. Wood. Generalized halfspaces in restricted-orientation convexity. *J. Geometry*, 62:99–120, 1998.

[KMŠ03] B. Kirchheim, S. Müller, and V. Šverák. Studying nonlinear PDE by geometry in matrix space. In *S. Hildebrandt (ed.) et al., Geometric analysis and nonlinear partial differential equations*, pages 347–395. Springer, Berlin, 2003.

[Mat01] J. Matoušek. On directional convexity. *Discr. Comput. Geom.*, 25(3):389–405, 2001.

[MP98] J. Matoušek and P. Plecháč. On functional separately convex hulls. *Discrete Comput. Geom.*, 19:105–130, 1998.

[MR03] E. Matoušková and S. Reich. The Hundal example revisited. *J. Nonlin. Convex Analysis*, 4(3):411–428, 2003.

[Mül99] S. Müller. Variational models for microstructure and phase transitions. In *Hildebrandt, S. (ed.) et al., Calculus of variations and geometric evolution problems, Lect. Notes Math. 1713*, pages 85–210. Springer, Berlin, 1999.

[NM91] V. Nesi and G. W. Milton. Polycristalline configurations that maximize electrical resistivity. *J. Mech. Phys. Solids*, 39:525–542, 1991.

[Sch74] V. Scheffer. Regularity and irregularity of solutions to nonlinear second order elliptic systems of partial differential equations and inequalities. Dissertation, Princeton University, 1974.

[See01] M. Seel. Implementation of planar Nef polyhedra. Research Report MPI-I-2001-1-003-1, Max-Planck-Institut für Informatik, Saarbrücken, 2001.

[Szé05] L. Székelyhidi jun. Rank-one convex hulls in $\mathbb{R}^{2\times 2}$. *Calc. Var. Partial Differ. Equ.*, 22(3):253–281, 2005. Erratum *ibid.* 28,4:545–546, 2007.

[Szé06] L. Székelyihidi jun. On the local structure of rank-one convex hulls. *Proc. Am. Math. Soc.*, 134(7):1963–1972, 2006.

[Tar93] L. Tartar. On separately convex functions. In *Microstructure and Phase Transition, The IMA volumes in mathematics and its applications, vol. 54 (D. Kinderlehrer et al., editors)*, pages 191–204. Springer, 1993.