

1 651 908



90C27
68R05

Geometric Lower Bounds for Parametric Matroid Optimization

D. Eppstein

1-CR3-C

Department of Information and Computer Science,
 University of California, Irvine, CA 92717, USA
<http://www.ics.uci.edu/~eppstein/>
eppstein@ics.uci.edu

Abstract. We relate the sequence of minimum bases of a matroid with linearly varying weights to three problems from combinatorial geometry: k -sets, lower envelopes of line segments, and convex polygons in line arrangements. Using these relations we show new lower bounds on the number of base changes in such sequences: $\Omega(nr^{1/3})$ for a general n -element matroid with rank r , and $\Omega(m\alpha(n))$ for the special case of parametric graph minimum spanning trees. The only previous lower bound was $\Omega(n \log r)$ for uniform matroids; upper bounds of $O(mn^{1/2})$ for arbitrary matroids and $O(mn^{1/2}/\log^2 n)$ for uniform matroids were also known.

1. Introduction

In this paper we study connections between combinatorial geometry and matroid optimization theory, as represented by the following problem.

Parametric Matroid Optimization. Given a matroid for which the elements have weights that vary as a linear function of a parameter t , what is the sequence of minimum weight bases over the range of values of t ? How many different bases can occur in such a sequence?

For example, an important special case of this problem would ask for the sequence of minimum spanning trees in a graph with linearly varying edge weights. This parametric spanning tree problem has applications including the stochastic spanning tree problem studied by Ishii *et al.* [25]; other applications are discussed by Hassin and Tamir [23].

* This work was supported in part by NSF Grant CCR-9258355 and by matching funds from Xerox Corp. A preliminary version of this paper appeared at the 27th ACM Symposium on Theory of Computing, 1995.

Although related parametric spanning tree problems have also been studied for geometric point set inputs [28], we are interested here in the general graph problem. Nevertheless as we shall see, computational geometry has an interesting role to play in this problem.

It has been known for some time that, in a matroid with n elements and rank k , the number of different bases occurring as the solutions to a parametric matroid optimization problem is $O(n \min(k, n - k)^{1/2})$ [18], [27]. Various authors have studied related questions of computing an optimal value for the time parameter in this sequence of bases [3], [5], [12], [15], [16], [31], [33], [35]. Katoh (personal communication) already noticed one connection between parametric matroid optimization and a classical problem of computational geometry:

***k*-Sets.** Given a set of n points in the plane, in how many combinatorially distinct ways can subsets of exactly k points be covered by half-spaces?

The k -set problem has been studied since 1971, when Lovász [29] proved an $O(nk^{1/2})$ bound. As Katoh noted (and as we describe below) the k -set problem is equivalent to the special case of parametric matroid optimization for *uniform matroids*, so this result follows from the more general matroid bound. Since Lovász' result, computational geometers have struggled with little success to improve this bound. The best result to date is $O(nk^{1/2}/\log^* k)$ [32].

Even less is known about lower bounds for these problems than is known about upper bounds, and it is this question we study here. The best lower bound for the k -set problem is $\Omega(n \log k)$ [13], quite far from the upper bound. No other lower bound was known for any other version of the parametric matroid problem. For graphic matroids (the minimum spanning tree problem discussed above) no nontrivial bound was known. In this paper we prove new lower bounds on the general parametric matroid optimization problem and on its special case of parametric minimum spanning trees by showing connections between these problems and two more seemingly unrelated problems from computational geometry: envelopes of segments and polygons in arrangements.

1.1. *New Results*

We prove the following bounds:

- There can be $\Omega(m\alpha(n))$ different minimum spanning trees in a graph with m edges, n vertices, and edge weights linearly varying with time, where $\alpha(n)$ denotes the inverse Ackermann function.
- There can be $\Omega(nr^{1/3})$ different minimum weight bases in a matroid with n elements, rank r , and element weights linearly varying with time.

The first bound comes from a reduction from the problem of *lower envelopes of line segments*, for which a $\Theta(n\alpha(n))$ bound is known. We use lower envelopes to construct a graph with $O(n)$ edges and the same $\Omega(n\alpha(n))$ bound on the number of minimum spanning trees; the overall bound is found by forming the union of $O(m/n)$ such graphs on the same vertex set.

We prove the second bound by proving that the problem is equivalent up to constant factors (for $r \leq n/2$) to the following new geometric problem:

Polygons in Arrangements. Given r convex polygons with sides formed by n lines, no two polygons overlapping in a set of nonzero length, what is the maximum total number of polygon vertices?

Although this problem has not been studied before, various special cases have. If the polygons form faces of the arrangement, have disjoint interiors, or do not cross each other, tight bounds of $\Theta(n + n^{2/3}r^{2/3})$ vertices are known [6], [9], [21], [24]; these are sufficient to prove a bound of $\Omega(n^{2/3}r^{2/3})$ on the matroid problem. Our new lower bound is somewhat stronger than this for $r = o(n)$, and shows that the general problem of polygons in arrangements has somewhat different behavior than these special cases. However our bounds and the lower bounds above are the same when $r = n$, and it may perhaps be possible for such r to extend the upper bounds for the special cases above to similar bounds for the general problem of polygons in arrangements. This could possibly lead to an $O(nr^{1/3})$ bound on the parametric matroid optimization problem and its special case the k -set problem.

1.2. Notation and Definitions

A matroid [36] consists of a set of *elements* and a family of finite sets of elements (the *independent sets* of the matroid) satisfying the following two axioms:

1. Any subset of an independent set is independent.
2. If I and J are independent, with $|I| < |J|$, then, for some $j \in J$, the set $I \cup \{j\}$ is independent.

The *rank* of a set in a matroid is the cardinality of its largest independent subset. We usually let n denote the number of elements in a matroid, and r denote its rank. A *base* of a matroid is a maximal independent set, or equivalently an independent set of cardinality equal to the matroid's rank. If the matroid elements have real-valued weights, we are particularly interested in the *minimum weight base*, that is, the base minimizing the sum of the element weights.

Standard examples of matroids include the following.

- The *uniform matroid* U_r^n . The set of elements is taken to be any set of cardinality n ; its independent subsets are those of cardinality at most r , and its bases are the subsets of cardinality exactly r . Finding a minimum weight base in such a matroid is equivalent to selecting the k th smallest element weight, so uniform matroids are related to median selection algorithms.
- The *graphic matroids* $M(G)$ for some graph G . $M(G)$ is defined to have the edges of G as elements, the forests of G as independent sets, and the spanning forests of G as bases. If G has m edges, n vertices, and c connected components, $M(G)$ has m elements and rank $n - c$. The minimum weight base in $M(G)$ is just the

minimum spanning tree of G , so graphic matroids are related to minimum spanning tree computation as well as to testing the connectivity of graphs.

- The *transversal matroids* $T(F)$ for some finite family of sets $F = \{S_i\}$. The set of elements of $T(F)$ is the union of the sets S_i . A set in a transversal matroid is independent if its members can be matched one-for-one with sets S_i containing them. Another way of viewing this is to consider a bipartite graph with edges connecting each set to its members; a collection of vertices on one side of the bipartition is independent if there is a matching in the graph connecting the collection to the same number of vertices on the other side. Therefore transversal matroids are related to bipartite graph matching.

In any matroid, a minimum weight base can be found by a *greedy algorithm*. The elements are sorted by weight and then considered one at a time in sorted order. We maintain an independent set S , and at each step add the element e under consideration to S if $S + e$ remains independent. For instance, this algorithm restricted to graphic matroids is just Kruskal's algorithm for minimum spanning trees.

2. Uniform Matroids and k -Sets

In this section we outline the connection between the k -set problem and the problem of parametric optimization for uniform matroids, as pointed out by Katoh. Some of the tools described here (in particular the correspondence between sets of element weights and line arrangements) are reused later.

A k -set of a set of n points is a subset of exactly k points such that some half-space covers exactly those points. Let $KS(n, k)$ denote the maximum number of k -sets among sets of n points, and let $UM(n, r)$ denote the maximum length of the sequence of minimum weight bases for a uniform matroid U_r^n with linearly varying edge weights.

Theorem 1. *Let $UM(n, r)$ and $KS(n, r)$ measure the complexities of parametric uniform matroid optimization and k -sets, as defined above. Then $UM(n, r) \leq 1 + KS(n, r)$.*

Proof. Let M be a linearly weighted uniform matroid achieving the maximum value $UM(n, r)$. For some sufficiently large value of x replace each weight function $w = at + b$ by $w = (a+x)t + b$; this affine transformation does not change the sequence of minimum weight bases, but causes all weight functions to have positive slope.

Plot the edge weight functions of the elements of M as an arrangement A of nonvertical lines in the (w, t) plane. We can perform a small perturbation of the weight functions of M , so that no three lines in the arrangement meet in a single point, without decreasing the number of base changes in the sequence of minimum weight bases. The minimum weight base at any time t_0 can be found by selecting the elements corresponding to the r lowest intersections of the lines in A with the vertical line $t = t_0$, as depicted in Fig. 1. As we sweep this vertical line from left to right, a change in the base occurs when it crosses a vertex v of A that has exactly $r - 1$ lines of A passing below it; in the figure, such a base change will happen at the next vertex swept over by the line.

Now if v is formed by two lines l_1 and l_2 , let C be the cell between them on the left




Fig. 1. Minimum weight base in a uniform matroid with positive-slope weight functions.

side of v . Since both lines have positive slope, v is the topmost vertex of C and C has exactly r lines below it.

Thus the number of base changes is at most the number of cells with r lines below them. We now show that this quantity is bounded by $KS(n, r)$. We use projective duality to transform the lines of A into points, and vice versa: for each line $w = at + b$ we place a point (b, a) , and for each point (w, t) we draw a line $b = -ta + w$. This transformation has the property that if a point is above a line in the (w, t) plane, the corresponding line is above the corresponding point in the (b, a) plane. Therefore the cell C corresponds to a set of lines, all above the same set of r points in the (b, a) plane; this is exactly an r -set and therefore the number of cells is bounded by the number of such sets. \square

Theorem 2. *Let $UM(n, r)$ and $KS(n, r)$ measure the complexities of parametric uniform matroid optimization and k -sets, as defined above. Then $KS(n, k) \leq 2(UM(n, k) - 1)$.*

Proof. Let S be a configuration of points realizing $KS(n, k)$; without loss of generality S is placed above the horizontal axis. We can assume by symmetry that at least $KS(n, k)/2$ of the k -sets are covered by the half-space below some line, rather than above a line. (Some k -sets may be covered by half-spaces of both types, but this only works to our advantage.) By applying the reverse of the duality transformation used above, we can form an arrangement of positive-slope lines such that at least $KS(n, k)/2$ cells have exactly k lines passing below them. Let these lines be formed as above from the weight functions of a uniform matroid U_k^n . Then the top vertices of each cell correspond to base changes of this matroid. \square

As a consequence, any bound of the form $O(n^a k^b)$ for k -sets can be transformed into a similar bound for parametric uniform matroid optimization, and vice versa. In particular the known $\Omega(n \log k)$ bounds for k -sets lead to a bound of the same form for uniform matroid optimization, and hence for parametric matroid optimization in general.




Fig. 2. The lower envelope of line segments.

3. Graphic Matroids and Lower Envelopes

We next consider the parametric graph minimum spanning tree problem, which is the restriction of parametric matroid optimization to graphic matroids. No lower bound was known for this case; not even the $\Omega(n \log k)$ bound above can be made to work here, as graphic matroids do not have nontrivial uniform submatroids.

We relate this problem to the following seemingly unrelated problem from computational geometry. Given n line segments s_i in the (x, y) plane, all above the x axis, the *lower envelope* of the segments is the function $f(x)$ giving the minimum value of y such that (x, y) is on some segment. (If no such y exists, let $f(x) = +\infty$.) Then $f(x)$ is a piecewise linear function of x . The *complexity* of $f(x)$ is the minimum number of intervals we need to partition the real line into, so that $f(x)$ is linear in each interval. In other words, it is the number of contiguous pieces of line segments that can be connected by vertical lines to the x axis (Fig. 2).

Lemma 1 [22], [37]. *The maximum complexity of a lower envelope of line segments is $\Theta(n\alpha(n))$, where α is the inverse Ackermann function.*

We now give the basic construction connecting this geometric concept with parametric minimum spanning trees. The construction produces a sparse graph from a line segment arrangement; we later show that a graph with any desired density can be formed by combining several copies of these sparse graphs.

Lemma 2. *Let S be a collection of n line segments in the plane, with lower envelope complexity c . Then there is a graph G with $3n$ edges and $2n + 2$ vertices, and a set of linearly varying edge weights on G , such that G has at least $c - n$ different trees in its sequence of minimum spanning trees.*

Proof. Let G be formed by taking two vertices s and t , and connecting them by n three-edge paths (Fig. 3(a)). Each path will correspond to a single segment s_i in S ; if a path has three edges e_1 , e_2 , and e_3 we let e_1 have an edge weight function plotted by the line through s_i , we let e_2 have an edge weight function of very large negative slope, such that the weights of e_1 and e_2 are equal at the left endpoint of s_i , and we let e_3 have a weight function of large positive slope crossing the weight of e_1 at the right endpoint




Fig. 3. Lower bound construction for graphic matroids: (a) the graph G ; (b) weight functions for the three edges corresponding to a line segment.

of s_i . These three weight functions and the segment s_i they correspond to are depicted in Fig. 3(b).

Any spanning tree of G is formed by choosing all three edges from one of the n paths, and any two edges from each of the remaining $n - 1$ paths. In the minimum spanning tree, these pairs of edges are always the minimum two edges in each path. The path with three edges is therefore chosen to minimize the weight of the maximum weight edge in the path.

Within the values of t covered by segment s_i , the maximum edge in the corresponding path is e_1 ; outside this range it is e_2 or e_3 . Therefore the function graphing the largest weight in the path is closely approximated by the lower envelope of the single segment s_i . The function graphing the minimum of these largest weights is closely approximated by the lower envelope of all segments. Any breakpoint between finite-valued segments of the lower envelope corresponds to a situation in G where the path with three chosen edges changes; each such situation involves a change to the minimum spanning tree of G . The infinite-valued segments of the lower envelope contribute two breakpoints, but perhaps only one change to the tree; however, there can only be $n - 1$ such segments. \square

As a corollary, for any n there is a graph with n vertices and $m = O(n)$ edges, having $\Omega(n\alpha(n))$ trees in its sequence of minimum spanning trees. We now extend this result to allow m and n to differ.

Theorem 3. *For any m and n there is a graph with m edges and n vertices, having $\Omega(m\alpha(n))$ trees in its sequence of minimum spanning trees.*

Proof. We choose a suitable $x = O(n)$ and $y = O(m/n)$, and form a graph with $2(x + y)$ vertices: $2y$ vertices a_i and b_i , $0 \leq i < y$, and $2x$ vertices c_i and d_i , $0 \leq i < x$. For any i we can form a copy of the graph G in the lemma above by connecting a_i to each c_j , b_i to each d_j , and connecting each c_j to $d_{(j+i) \bmod x}$. This gives a total of y edge-disjoint copies of G .

By applying an affine transformation to the weight functions of the graph G in the lemma above, we can cause all of its edges to have slope within some ϵ of each other without changing the sequence of minimum spanning trees. In this way we “flatten”




Fig. 4. Combining several flattened line arrangements into a single convex chain.

the weights in each of the copies of G to different ranges of slopes. We then use linear translations to arrange the flattened sets of weight functions to form a large convex chain (Fig. 4). Within each segment of the chain the sequence of minimum spanning trees undergoes $\Omega(x\alpha(x))$ changes. Therefore overall there are $\Omega(xy\alpha(x)) = \Omega(m\alpha(n))$ changes. We then add additional edges or vertices to produce the desired total numbers of edges and vertices, with sufficiently large edge weight functions so as not to change the sequence of minimum spanning trees found above. \square

4. General Matroids and Polygons in Line Arrangements

We next show a relation between the general parametric matroid optimization problem and the problem of polygons in arrangements defined in the Introduction. We use this relation to prove an $\Omega(nr^{1/3})$ bound on the matroid optimization problem.

Define $MO(n, r)$ to be the maximum length of the sequence of minimum weight bases among all matroids with n elements and rank r . Define $AP(n, m)$ to be the maximum number of vertices in a collection of m convex polygons with edges drawn from a collection of n lines, no two polygon edges overlapping in a set of nonzero length (Fig. 5(a)). Our first result gives an upper bound on $MO(n, r)$ in terms of $AP(n, m)$;




Fig. 5. (a) Three nonoverlapping polygons in a line arrangement. (b) Adding steep tangents to left and right sides of convex polygons.

together with the lower bound that is our main result of this section, this shows the equivalence for $r \leq n/2$ of these two problems.

Theorem 4. *Let $MO(n, r)$ and $AP(n, r)$ measure the complexity of matroid optimization and polygons in arrangements as defined above. Then $MO(n, r) \leq AP(n + r, r) - 2r + 1$.*

Proof. Given a matroid M with linearly varying element weights, form an arrangement A of the lines in the (w, t) plane traced out by the element weights, as in the proof of Theorem 1. As in that same proof, we assume that no three lines of A meet in a point, for otherwise we could perturb the weight functions slightly to achieve this property while not decreasing the number of base changes in M .

We now use a token-passing scheme to trace out polygons in this arrangement. Consider sweeping the t parameter from $-\infty$ to $+\infty$. At a value of t prior to any base change, assign a token to each member of the minimum weight base for t . Whenever the sweep reaches a value of t at which a base change occurs, reassign the token that had been assigned to the changed member of the old base, so it becomes assigned to that member's replacement in the new base.

This token-assignment process can also be interpreted geometrically. Consider sweeping a vertical line left-to-right through the arrangement, with each placement of the line corresponding to a different parameter t , as depicted in Fig. 1. As this sweep line progresses from left to right, we maintain the placement of r tokens, on points where the sweep line crosses arrangement lines representing the minimum weight base of M for that value of t . As the sweep line moves left to right, the tokens move with it along the arrangement lines they are placed on. At points in the sweep corresponding to values of t at which a base change occurs, a single element of the old base is removed and a new element is added to replace it; in this case we move the token, from the arrangement line corresponding to the old element, to the arrangement line corresponding to the new element.

At the time of the base change, the old and new elements must have equal weights; before the change the old element's weight is smaller than the new element's weight and after the change the new element's weight is smaller than the old element's weight. Therefore each base change corresponds to a convex vertex in the piecewise linear motion of a single token, and each token traces out a convex chain in the arrangement A . By adding r new lines to A , one per chain, we can complete these chains to convex polygons of the same complexity. Each chain has one vertex per base change, plus two added when we completed the chain to a polygon; therefore the total number of base changes is equal to the number of vertices minus $2r$. \square

We next show a bound in the other direction, a lower bound on MO in terms of AP .

Theorem 5. *Let $MO(n, r)$ and $AP(n, r)$ measure the complexity of matroid optimization and polygons in arrangements as defined above. Then $AP(n, m) \leq 2(MO(n + 2m, m) - m - 1)$.*

Proof. Let A be an arrangement of n lines, with m nonoverlapping polygons in A having a total of $AP(n, m)$ vertices. By symmetry we can assume that the number of vertices visible from above in each polygon is $m + AP(n, m)/2$.

We add to A a set of $2m$ more lines: one line of large positive slope tangent to each polygon at its left extremum, and one line of large negative slope tangent to each polygon at its right extremum (Fig. 5(b)). We use these lines to replace each polygon by a convex chain similar to the chains in the theorem above, having the same vertices as the polygon.

Define a matroid M having this set of $n + 2m$ lines as elements, with the lines giving the weight function of these elements. To define M , form a set S_i for each of the convex chains formed above, consisting of those lines having a nonzero-length segment in the chain. We then let M be the transversal matroid of these sets S_i .

We now show that the set of chains traced out by the token-passing procedure described in the previous proof is exactly the set of chains constructed above. We show that any time t_0 such that no vertex of the arrangement has t -coordinate t_0 , the minimum weight basis can be found by choosing the element in each convex chain crossed by the vertical line $t = t_0$. This follows by induction from the greedy algorithm for minimum weight bases, as follows.

Assume we have chosen so far a set X of elements, with $|X| < m$, that are the $|X|$ least weight elements crossed by the line $t = t_0$ among the collection of chains. Let C_i be a chain crossed by $t = t_0$ above all members of X ; then no member of X can be in S_i , and so any set $e + X$ where $e \in S_i$ is independent. Conversely suppose some set $e + X$ is independent. Then this set can be matched against sets S_i defining matroid M . By the pigeonhole principle, some matched S_i corresponds to a chain C_i above all crossings in X . No element of X is in S_i , so e must be in S_i . Thus the independent extensions $e + X$ are exactly those for which e is a member of some S_i corresponding to a chain C_i crossed above X . Since the extension chosen by the greedy algorithm is the one minimizing the weight of e , e must be the element corresponding to the next crossing of a chain with the line $t = t_0$. This completes the induction and shows that except at arrangement vertices, the minimum weight basis follows the chains. By continuity of the functions traced out by the token-passing process, the same fact must be true at the vertices of A . Therefore all vertices of the set of chains correspond to basis changes in the transversal matroid M . \square

These two results combine to give the following consequence, which may be useful in proving further bounds on parametric matroid optimization.

Corollary 1. *The maximum complexity of parametric matroid optimization can be realized to within constant factors by a transversal matroid.*

For completeness, we state an upper bound on polygons in arrangements. This is similar to known bounds on parametric matroid optimization [18], [27] but holds also for $r > n$.

Theorem 6. *Let $AP(n, r)$ measure the complexity of polygons in arrangements as defined above. Then $AP(n, r) = O(nr^{1/2})$.*

Proof. Given a collection of polygons in an arrangement A , we show that it has the above complexity. Perturb A if necessary so no three lines meet at a point. Sort the lines of A by slope. We say that a vertex of A is *sharp* if the lines forming the vertex have slopes separated by more than $r^{1/2}$ positions apart in the sorted sequence, and *dull* otherwise. Since the slopes of a polygon's edges form a monotone sequence, each polygon can have at most $n/r^{1/2}$ sharp vertices, so the total number of such vertices in all the polygons is $O(nr^{1/2})$. The remaining vertices are dull, but there are $O(nr^{1/2})$ dull vertices in the entire arrangement, and each one can only be part of at most two polygons. \square

We now complete the lower bound for matroid optimization by constructing a collection of nonoverlapping polygons with many vertices. Known results on faces in arrangements give the bound $\Omega(n^{2/3}r^{2/3})$ on $AP(n, r)$, however, the bound in the theorem below is stronger for the case of interest in which $r < n$.

Theorem 7. *For any n and $r = O(n^2)$ there is a collection of r nonoverlapping polygons in an arrangement of n lines, having $\Omega(nr^{1/3})$ vertices.*

Proof. We first form an arrangement A' with $O(r)$ lines, such that some r faces of the arrangement have total complexity $\Omega(r^{4/3})$ [6]. We then add steep left and right tangents to these faces, as in the proof above, so that we have $O(r)$ convex chains with the same asymptotic complexity. To complete the construction we flatten A' by an affine transformation and connect $O(n/r)$ copies of the flattened arrangement in a large convex chain shape, as in the proof of Theorem 3 and as depicted in Fig. 4. The sets of r chains in adjacent copies of A' can be connected where the copies cross, so we get r chains overall and $\Omega(nr^{1/3})$ total complexity. As in Theorem 4, we can add additional lines to the base of each of these chains, to form them back into convex polygons. \square

Corollary 2. *There exist parametric matroid optimization problems with complexity $MO(n, r) = \Omega(nr^{1/3})$.*

5. Conclusions

We have described three different problems from combinatorial geometry and related them all to a common nongeometric problem, of parametric matroid optimization. We then used these relations to prove lower bounds on the matroid problem.

These results also give new hope for results on the geometric side of the problem, on the longstanding open problem of bounds on k -sets. The matroid upper and lower bounds are both of the form nr^c , giving rise to the possibility that a similar lower bound could be proved for k -sets. The fact that our problem of polygons in arrangements is very similar to other problems with $O(n^{4/3})$ bounds gives reason to believe that similar bounds might hold for matroid optimization, and its special case the k -set problem.

An interesting related problem concerns similar parametric optimization questions for nonmatroidal problems. Fernández-Baca and Slutzki [14] show that many such problems, when restricted to certain classes of graphs, produce a sequence of solutions of polynomial length. One such question of particular interest is that of *parametric shortest*

paths [26], [38]: in a graph with linearly varying edge weights, how many times can the shortest path between two vertices change? Parametric shortest-path problems have been extensively studied for their application to minimum cost flow [17], but a solution to this question would also help clarify the processor bounds needed for fast parallel linear programming with two variables per inequality [30]. An argument similar to the proof of Savitch's theorem can be used to show an $n^{O(\log n)}$ bound on the number of shortest paths (G. S. Lueker and N. Megiddo, personal communications; Carstensen attributes this result to Gusfield [19]). Carstensen proved a matching $n^{\Omega(\log n)}$ lower bound for this problem [4]. Do other nonmatroidal parametric optimization problems give similar quasi-polynomial bounds?

Another direction for generalization is in the functions defining the matroid element weights as functions of the parameter t . Work of Gusfield *et al.* on parametric sequence alignment [20] can be interpreted as a shortest-path problem with weights depending linearly on two parameters, and, similarly to the situation in two dimensions, multiparameter uniform matroid optimization problems are essentially equivalent to levels in higher-dimensional arrangements [1], [2], [8], [11]. Tamaki and Tokuyama [34] have also investigated parametric matroid problems with quadratic element weight functions, but much more work remains to be done in this direction.

A final open problem concerns algorithms for parametric matroid optimization. It is known that a k -level in a line arrangement (essentially equivalent to the set of base changes in a parametric uniform matroid) can be constructed in time $O(n \log n + x \log^2 n)$, where x denotes the output complexity [7], [10]. Recently, we used some of the ideas from this paper (in particular the equivalence between matroid optimization and convex polygons in an arrangement) as part of an algorithm for finding all minimum spanning trees in a graph with linearly varying edge weights, in time $O(mn \log n)$ [16]. However, there may still be room for improvement in this bound, and known results for nonuniform and nongraphic matroids are even further from the output complexity bounds proven here.

References

1. B. Aronov, B. M. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and R. Wenger. Points and triangles in the plane and halving planes in space. *Discrete & Computational Geometry* 6:435–442, 1991.
2. I. Bárány, Z. Füredi, and L. Lovász. On the number of halving planes. *Combinatorica* 10:175–183, 1990.
3. P. M. Camerini, F. Maffioli, and C. Vercellis. Multi-constrained matroidal knapsack problems. *Mathematical Programming* 45:211–231, 1989.
4. P. J. Carstensen. Parametric cost shortest path problems. Unpublished Bellcore memo, available from the Literaturstelle at the Inst. für Ökonometrie und Operations Research, Univ. Bonn, 8 May 1984.
5. R. Chandrasekaran. Minimal ratio spanning trees. *Networks* 7:335–342, 1977.
6. K. L. Clarkson, H. Edelsbrunner, L. J. Guibas, M. Sharir, and E. Welzl. Combinatorial complexity bounds for arrangements of curves and surfaces. *Discrete & Computational Geometry* 5:99–160, 1990.
7. R. Cole, M. Sharir, and C. K. Yap. On k -hulls and related problems. *SIAM Journal on Computing* 16:61–77, 1987.
8. T. K. Dey and H. Edelsbrunner. Counting triangle crossings and halving planes. *Discrete & Computational Geometry* 12:281–289, 1994.
9. H. Edelsbrunner, L. J. Guibas, and M. Sharir. The complexity and construction of many faces in arrangements of lines and of segments. *Discrete & Computational Geometry* 5:161–196, 1990.
10. H. Edelsbrunner and E. Welzl. Constructing belts in two-dimensional arrangements with applications. *SIAM Journal on Computing* 15:271–284, 1986.

11. D. Eppstein. Improved bounds for intersecting triangles and halving planes. *Journal of Combinatorial Theory, Series A* 62:176–182, 1993.
12. D. Eppstein and D. S. Hirschberg. Choosing subsets with maximum weighted average. *Journal of Algorithms* 24:177–193, 1997.
13. P. Erdős, L. Lovász, A. Simmons, and E. G. Straus. Dissection graphs of planar point sets. *A Survey of Combinatorial Theory*, pp. 139–149. North-Holland, Amsterdam, 1973.
14. D. Fernández-Baca and G. Slutzki. Parametric problems on graphs of bounded tree-width. *Journal of Algorithms* 16(3):408–430, 1994.
15. D. Fernández-Baca and G. Slutzki. Linear-time algorithms for parametric minimum spanning tree problems on planar graphs. *Proc. 2nd Latin Amer. Symp. Theoretical Informatics*, pp. 257–271. Lecture Notes in Computer Science, Vol. 911. Springer-Verlag, Berlin, 1995.
16. D. Fernández-Baca, G. Slutzki, and D. Eppstein. Using sparsification for parametric minimum spanning tree problems. *Nordic Journal on Computing* 3(4):352–366, 1996.
17. A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by cancelling negative cycles. *Journal of the Association for Computing Machinery* 36:873–886, 1989.
18. D. Gusfield. Bounds for the parametric minimum spanning tree problem. *Proc. Humboldt Conf. Graph Theory, Combinatorics and Computing*, pp. 173–183. Congressus Numerantium XXVI. Utilitas Mathematica, Winnipeg, Manitoba, 1979.
19. D. Gusfield. Sensitivity analysis for combinatorial optimization. Tech. Rep. UCB/ERL M80/22, Univ. of California, Berkeley, May 1980.
20. D. Gusfield, K. Balasubramanian, and D. Naor. Parametric optimization of sequence alignment. *Algorithmica* 12:312–326, 1994.
21. D. Halperin and M. Sharir. On disjoint concave chains in arrangements of (pseudo) lines. *Information Processing Letters* 51:53–56, 1994.
22. S. Hart and M. Sharir. Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes. *Combinatorica* 6:151–177, 1986.
23. R. Hassin and A. Tamir. Maximizing classes of two-parametric objectives over matroids. *Mathematics of Operations Research* 14:362–375, 1989.
24. J. E. Hersherberger and J. Snoeyink. Convex polygons made from few lines and convex decompositions of polyhedra. *Proc. 3rd Scand. Workshop Algorithm Theory*, pp. 376–387. Lecture Notes in Computer Science, Vol. 621. Springer-Verlag, Berlin, 1992.
25. H. Ishii, S. Shiode, and T. Nishida. Stochastic spanning tree problem. *Discrete Applied Mathematics* 3:263–273, 1981.
26. R. M. Karp and J. B. Orlin. Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Applied Mathematics* 3:37–45, 1981.
27. N. Katoh and T. Ibaraki. On the total number of pivots required for certain parametric combinatorial optimization problems. Working Paper 71, Inst. Econ. Res., Kobe Univ. Commerce, 1983.
28. N. Katoh, T. Tokuyama, and K. Iwano. On minimum and maximum spanning trees of linearly moving points. *Discrete & Computational Geometry* 13(2):161–176, March 1995.
29. L. Lovász. On the number of halving lines. *Annales Universitatis Scientiarum Budapestinensis de Rolando Eötvös Nominatae, Sect. Math.* 14:107–108, 1971.
30. G. S. Lueker, N. Megiddo, and V. Ramachandran. Linear programming with two variables per inequality in poly-log time. *SIAM Journal on Computing* 19:1000–1010, 1990.
31. N. Megiddo. Applying parallel computation algorithms in the design of serial algorithms. *Journal of the Association for Computing Machinery* 30(4):852–865, 1983.
32. J. Pach, W. Steiger, and E. Szemerédi. An upper bound on the number of planar k -sets. *Discrete & Computational Geometry* 7:109–123, 1992.
33. R. Ravi and M. X. Goemans. The constrained minimum spanning tree problem. *Proc. 5th Scand. Workshop Algorithm Theory*, pp. 66–75. Lecture Notes in Computer Science, Vol. 1097. Springer-Verlag, Berlin, 1996.
34. H. Tamaki and T. Tokuyama. How to cut pseudo-parabolas into segments. *Proc. 11th Symp. Computational Geometry*, pp. 230–237. ACM, New York, 1995.
35. S. Toledo. Maximizing non-linear concave functions in fixed dimension. In *Complexity in Numerical Computations* (P. M. Pardalos, ed.), pp. 429–447. World Scientific, Singapore, 1993.
36. D. J. A. Welsh. *Matroid Theory*. Academic Press, New York, 1976.

37. A. Wiernik and M. Sharir. Planar realizations of nonlinear Davenport–Schinzel sequences by segments. *Discrete & Computational Geometry* 3:15–47, 1988.
38. N. E. Young, R. E. Tarjan, and J. B. Orlin. Faster parametric shortest path and minimum-balance algorithms. *Networks* 21:205–221, 1991.

Received November 12, 1996, and in revised form February 19, 1997.

Note added in proof. Very recently, Tamal Dey has proven an $O(nr^{1/3})$ upper bound on the number of base changes in a parametric matroid optimization problem, matching the $\Omega(nr^{1/3})$ lower bound given here. The same upper bound also applies to k -sets and parametric minimum spanning trees. Dey’s results will appear at the 38th IEEE Symposium on Foundations of Computer Science.