Computational Geometry and Parametric Matroid Optimization

David Eppstein

Dept. of Information and Computer Science
University of California, Irvine

Outline

I. Parametric minimum spanning trees and parametric matroids

II. $\Omega(mn^{1/3})$ lower bound for general matroids

III. $O(mn^{1/3})$ upper bound

IV. $O(mn \log n)$ time algorithm for spanning trees

V. Constructive solid geometry, series parallel graphs, and an $\Omega(m\alpha(n))$ lower bound for spanning trees

I. Parametric minimum spanning trees and parametric matroids

Parametric Edge Weights

Given a graph

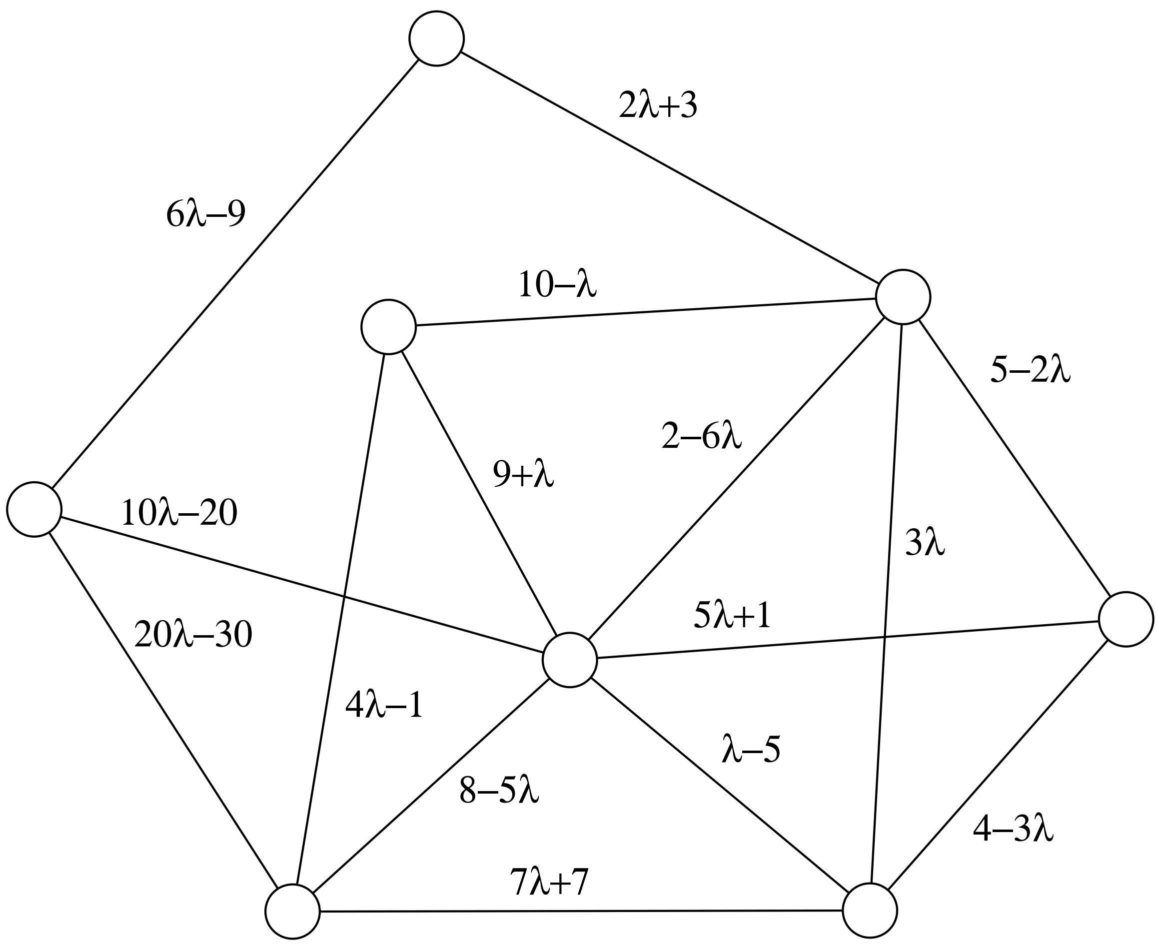

Edges labeled with linear functions $f(e) = a_e\lambda + b_e$

Parametric Minimum Spanning Tree

Plug in a real number as $\lambda \rightarrow$ real edge weights

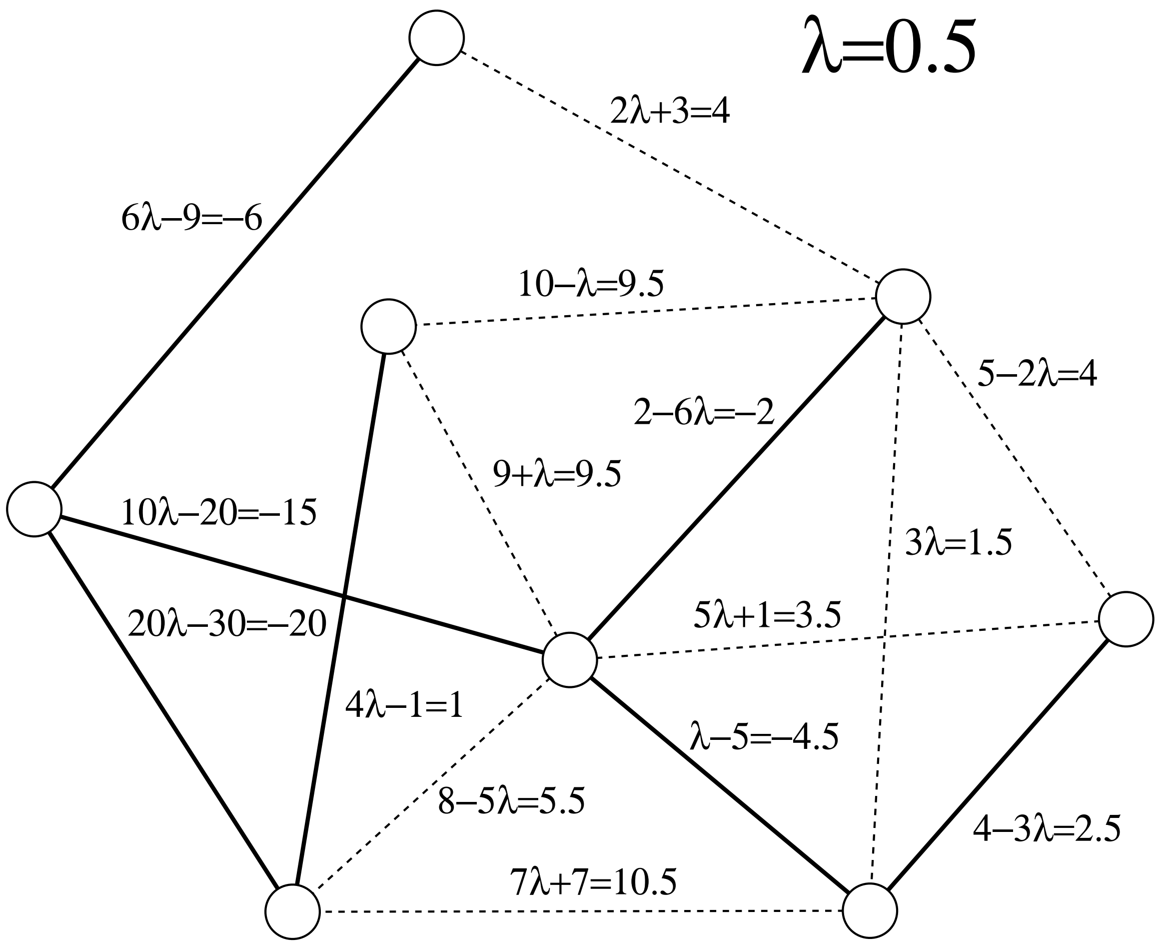

Different λ give different minimum spanning trees

What questions are we interested in?

As λ varies continuously, how does the minimum spanning tree change?

How many different spanning trees does one get?

How quickly can we compute this sequence of spanning trees?

How quickly can we find an "optimal" value of λ (e.g. minimum-ratio spanning tree)

Generalization to matroids

Matroid = set of objects, certain subsets designated as *independent sets*

Any subset of an independent set is independent

If A and B are independent, with $|B| > |A|$, then for some $x \in B - A$, $A \cup \{x\}$ is independent

Basis = maximal independent set

If objects have weights, minimum weight basis can be found by a greedy algorithm

Important special classes of matroid

Graphical matroid: Objects are edges in a graph, independent sets are forests, minimum weight basis is minimum spanning tree

Uniform matroid: Independent sets have $|S| \leq k$, minimum weight basis = k smallest values

Transversal matroid: Objects are vertices from one side of a bipartite graph, independent sets are sets of endpoints of matchings

Why computational geometry?

Linear functions = lines in the plane

If we can abstract away the matroid combinatorics, leaving only a computational geometry problem (involving arrangements of lines or related objects) we can use many tools from the geometry literature

(topological sweeping, ϵ-cuttings, many-face bounds, projective duality, Davenport-Schinzel sequences, . . .)

Parametric optimization and k-sets

How many ways to split k points from n by a line?

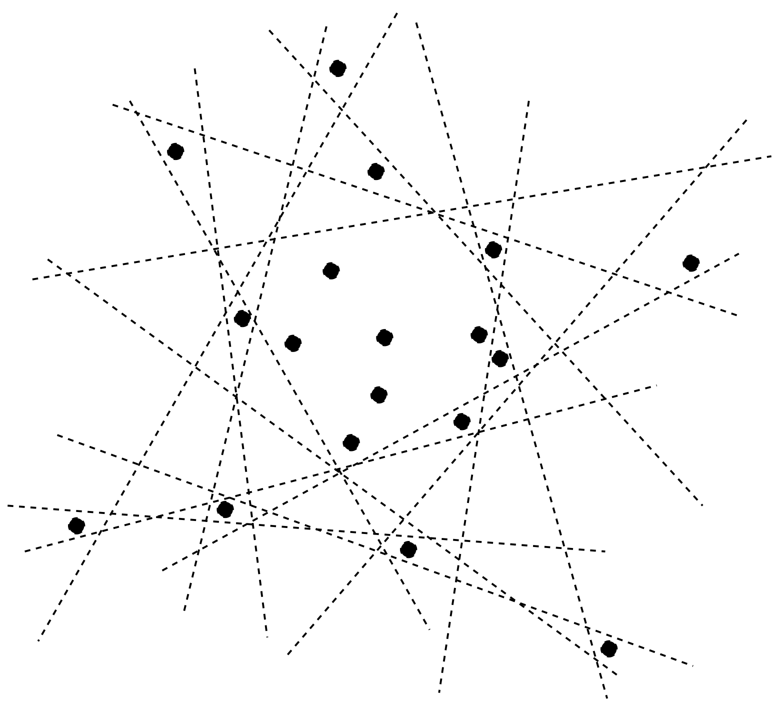

Equivalently (by projective duality):
how many ways to separate k lines from n by a point (i.e. k lines below the point, $n-k$ above)

Known bounds (until recently): $\Omega(n \log k)$, $O(nk^{1/2})$

Equivalent to parametric uniform matroid

II. $\Omega(mn^{1/3})$ lower bound

(from D. Eppstein, “Geometric lower bounds for parametric matroid optimization”, STOC 1995 and *Discrete & Computational Geometry* to appear)

Line arrangement from parametric matroid

Graph $w(e) = a_e\lambda + b_e$ as line in the (λ, w) plane

Element values at λ_0 given by crossings with vertical line $\lambda = \lambda_0$

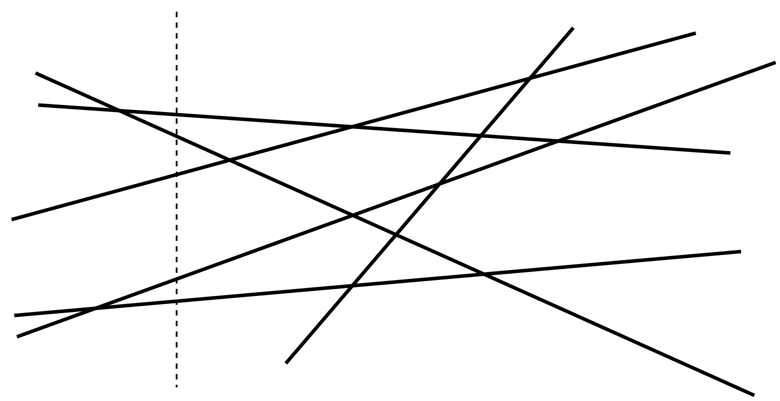

As vertical line sweeps left→right, element order stays fixed except at arrangement vertices

Assume *general position*: each arrangement vertex only involves two lines (else perturbation increases complexity)

Corollary: adjacent bases in sequence of minimum weight bases differ by *swaps* of two elements

Convex chains from parametric matroid

Imagine the following process:

- Sweep a vertical line across the arrangement
- Place a *token* at each crossing of the sweep line with a line in the minimum weight basis

tokens move along straight lines, except at swaps

Each token path is a convex chain

So, a rank-k matroid gives us a family of k disjoint convex chains in an arrangement

Parametric matroid from convex chains

Suppose we have k convex chains involving n lines

This configuration "looks like" it comes from the token-passing process of a parametric matroid

The objects and weights are obvious, but what are the independent sets of the matroid?

Define a bipartite graph (X, Y, E), where
$X =$ the set of lines in the arrangement
$Y =$ the set of convex chains
$(x, y) \in E$ if line x participates in chain y

Transversal matroid bases = sets of lines that can be matched one-for-one with the convex chains

Lemma: the token-passing process on this transversal matroid gives back the same set of convex chains we started with

Convex chains = parametric matroids

So, any parametric matroid problem leads to a set of convex chains

Any set of convex chains forms a parametric transversal matroid problem

Number of base changes in the matroid =
total number of corners in the convex chains

Result: parametric matroid problem complexity =
convex chain problem complexity

Second result: parametric transversal matroids are the worst case among parametric matroids

So, how to find chains w/ many corners?

Lemma (Erdös):
There exist configurations of n points and n lines with $\Omega(n^{4/3})$ point-line incidences

Proof:
Form a $\sqrt{n}$ by $\sqrt{n}$ grid of points; choose lines greedily to contain as many points as possible

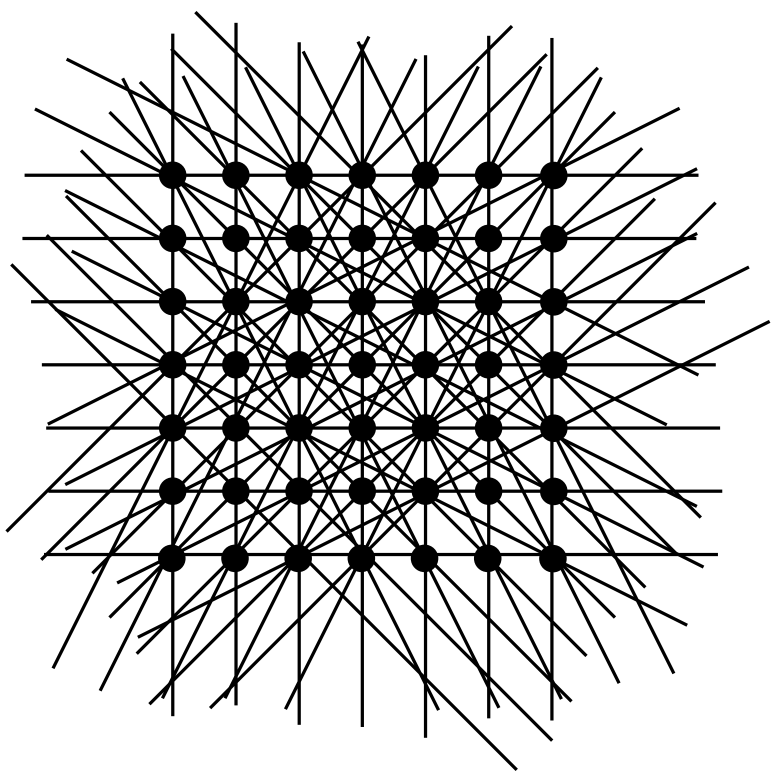

Chains w/ many corners continued

Lemma:
There exist sets of $n/3$ disjoint convex chains, formed by n lines, having a total of $\Omega(n^{4/3})$ corners

Proof:

- Form the configuration above with $n/3$ points and lines (no two points sharing an x-coordinate)

- "Lift" lines to form chains around each point

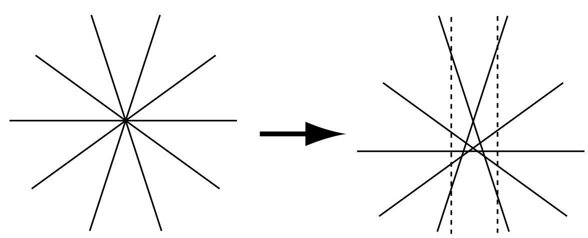

- Add two extra lines dropping steeply from the left and right side of each convex chain to make the chains disjoint

What if rank is much smaller than n?

Form a $3k$-line arrangement of k convex chains with $\Omega(k^{4/3})$ corners

"Flatten" by compressing the vertical dimension

Connect $O(n/k)$ flattened arrangements into one large convex chain

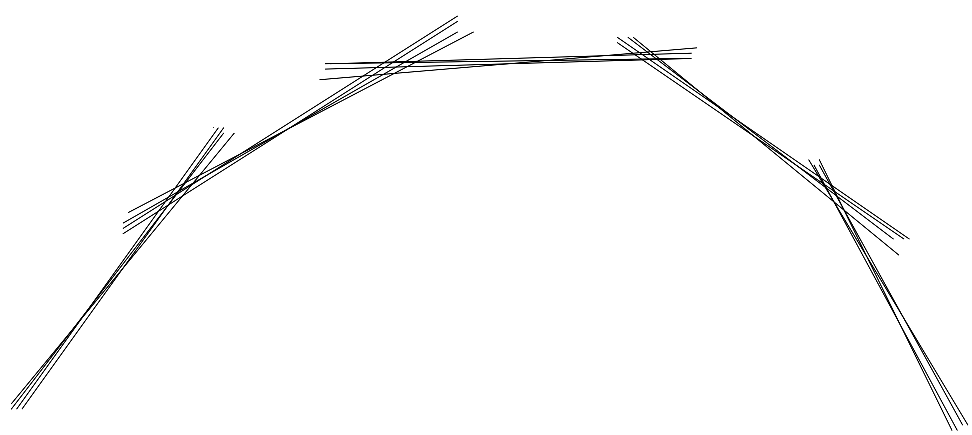

Glue together chains from each flattened arrangement

Result: rank-k matroids can have $\Omega(nk^{1/3})$ basis changes

III. $O(mn^{1/3})$ upper bound

(from T. Dey, "Improved bounds for *k*-sets and *k*th levels", FOCS 1997)

Projective duality

Correspondence from

(x, y) plane to (a, b) plane
point (x, y) to line $b = (-x)a + y$
line $y = ax + b$ to point (a, b)

Preserving all point-line incidence combinatorics

Any incidence-based statement has a dual form

Used as an aid to intuition:
take statement you really want to solve
translate mechanically to dual form
using a “dictionary” of primal-dual correspondences
often dual is more understandable than primal

Duality and convex chains

Convex chain
= intersection of halfspaces below lines
= set of points that are below all lines

Dual of convex chain
= set of lines that are above all points
= upper chain of convex hull

Point on upper hull dualizes to line tangent to chain

Vertex of upper hull dualizes to segment of chain

Duality and parametric matroids

We want to show that disjoint convex chains in an arrangement of n lines can't have many corners

Dually: a set of convex polygons sharing n vertices (satisfying some property dual to disjointness) can't have very many edges.

But what happens if we have polygons forming a n-vertex graph with many edges?

Lemma: any drawing of an n-vertex m-edge graph has $\Omega(m^3/n^2)$ crossings

(used previously to prove many-face bounds in arrangements, bounds on 3-dimensional k-sets, etc.)

The upper bound

Suppose there are m basis changes $\to$ m corners of convex chains $\to$ m edges in an n-vertex graph.

Then (lemma) there are $\Omega(m^3/n^2)$ crossings.

Reversing the duality, a crossing (point shared by two upper hulls) corresponds to a bitangent (line shared by two chains).

But, two chains have only as many bitangents as they have crossings!

And the number of chain crossings is at most $2nk$ (each line crosses each chain at most twice)

So $\Omega(m^3/n^2) \le 2nk$ and $m = O(nk^{1/3})$

IV. $O(mn \log n)$ time algorithm

(from D. Fernández-Baca, G. Slutzki, and D. Eppstein, “Using sparsification for parametric minimum spanning tree problems”, SWAT 1996 and *Nordic J. Computing* 1996)

How to compute all parametric minimum spanning trees?

First attempt at an algorithm:

- Construct arrangement of lines

- When does a vertex correspond to a swap?

  When only the edge corresponding to the higher-slope line is in the MST prior to the vertex, and the other edge induces a cycle in the MST which contains the first edge

- Sweep the arrangement by a vertical line

- As each vertex is swept, test if it gives a swap ($O(\log n)$ using Sleator-Tarjan dynamic trees)

Total time $O(m^2 \log n)$

How to speed this up?

Sparsification!

Divide-and-conquer technique for speeding up dynamic graph algorithms
[Eppstein et al, J. ACM to appear]

Divide graph into $m/2$-edge subgraphs G and H

Compute sequence of MST's for each subgraph

Merge the results

Key property used:
$\mathrm{MST}(G \cup H) = \mathrm{MST}(\mathrm{MST}(G) \cup \mathrm{MST}(H))$

Merge step

Given parametric MST solution on two subgraphs

How to merge?

Each solution is formed by a collection of disjoint convex chains

View chains as sets of line segments

Portion of a line not belonging to one of these segments can not participate in the overall parametric MST (from the key property)

Sweep arrangement of all segments by a vertical line, test whether each swept vertex gives a swap

Analysis

How many crossings in the arrangement of line segments?

Formed by m lines in at most $2n$ convex chains
Each line can only cross each chain twice
So at most $4mn$

Arrangement can be swept in time proportional to number of crossings, so the bottleneck is the Sleator-Tarjan swap test

Overall recurrence

$$\begin{aligned} T(m,n) &= 2T(m/2,n) + O(mn\log n) \\ &= O(mn\log n \log \frac{m}{n}) \end{aligned}$$

With "Improved Sparsification" can remove the $\log(m/n)$ term by reducing n at each level of the recursion

V. Constructive solid geometry, series parallel graphs, and an $\Omega(m\alpha(n))$ lower bound

Constructive solid geometry (CSG)

Given some set of simple base shapes

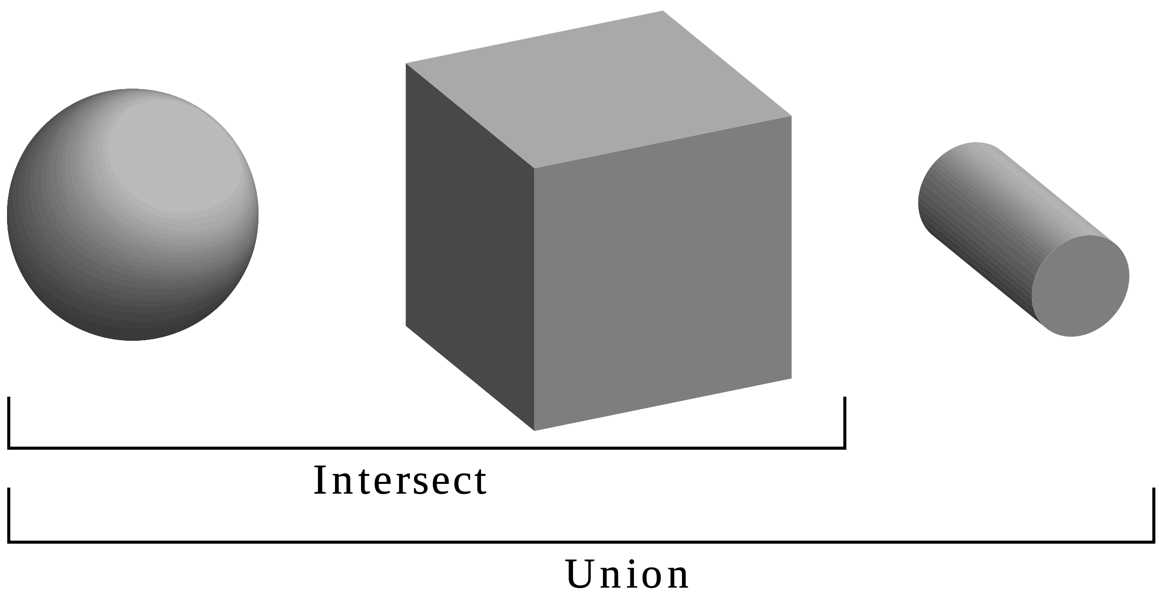

Form complex shapes by unions and intersections

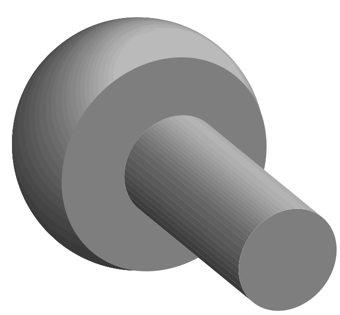

CSG of monotone paths

Base shapes: halfspaces bounded above by lines

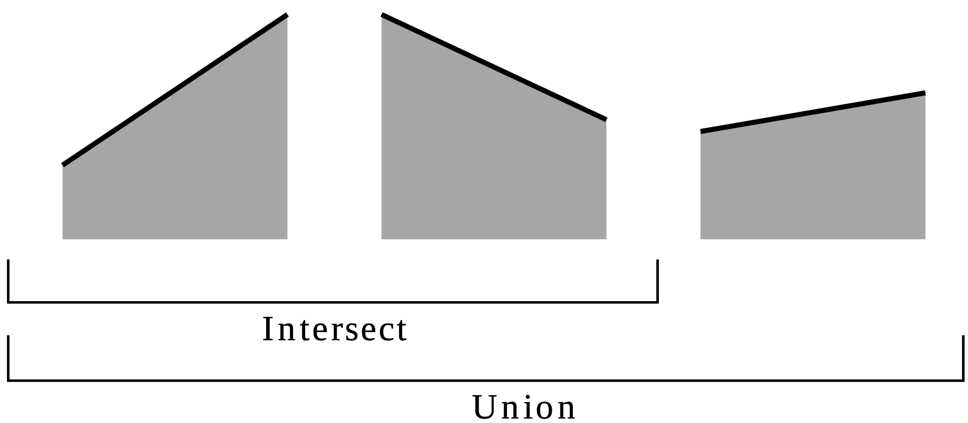

Unions and intersections:
regions bounded above by monotone paths

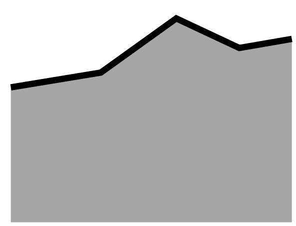

How many vertices can a path formed
from *n* halfspaces have?

Series-parallel graphs

Graphs with two designated terminals s, t

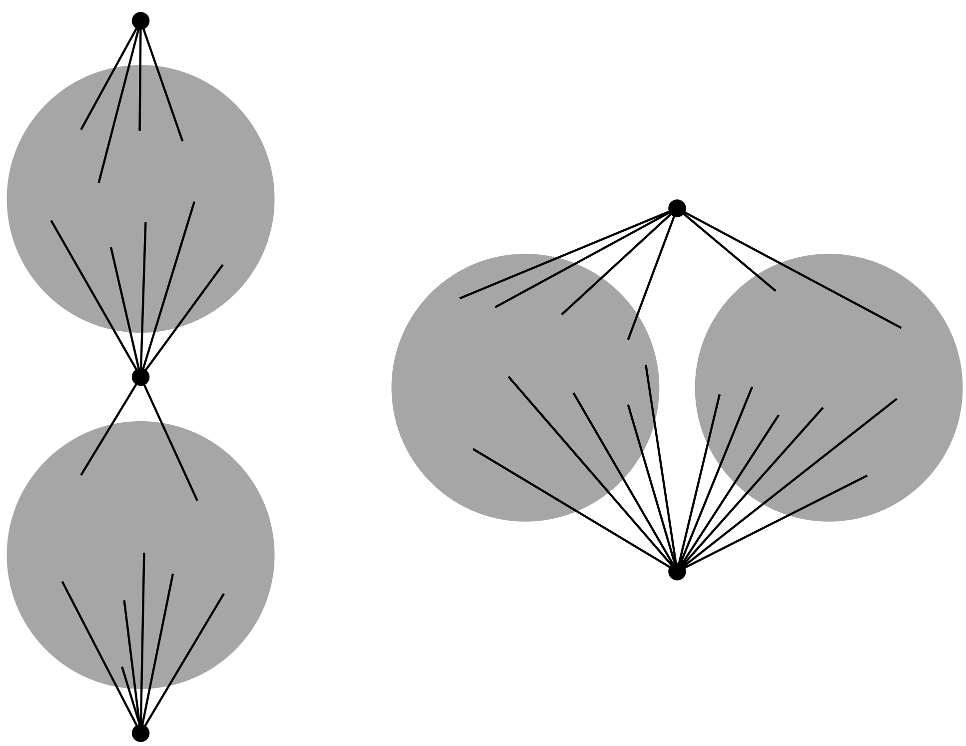

Formed from simple graphs (edges) by operations:
Series connection – identify $t_G = s_H$
Parallel connection – identify $s_G = s_H$ and $t_G = t_H$

Monotone-path CSG $\leq$ parametric MST

Given a parametric series parallel graph
MST(λ) has a unique path from *s* to *t*
How heavy is the heaviest edge on that path?

Equivalent to monotone path CSG!
Halfspace = single edge
Union = series connection
Intersection = parallel connection

Convex corner = MST edge swap
Concave corner = two equally heavy edges on path

Without loss of generality, convex $\geq$ corners/2
(otherwise, turn the picture upside down)

So monotone path CSG gives us lower bounds on parametric graph MST

Best current construction

Union of regions below line segments
(aka upper envelope)

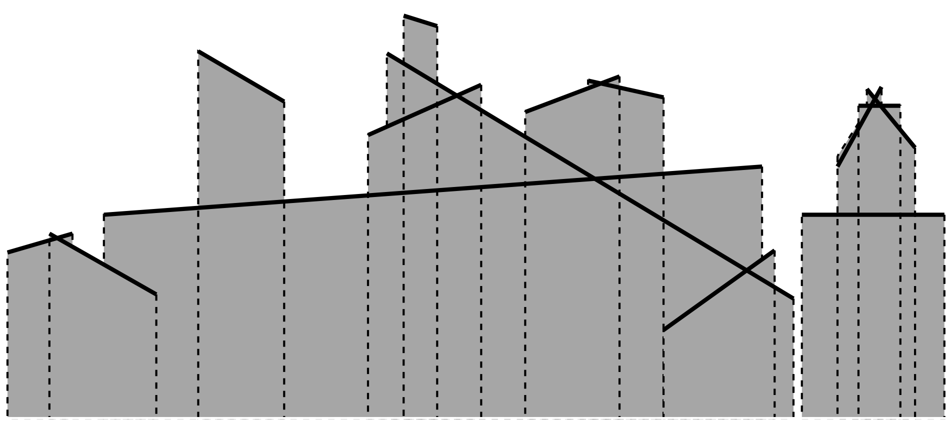

Known to have $\Theta(n\alpha(n))$ complexity
[Wiernik and Sharir, Disc. & Comput. Geom. 1988]

Form regions below segments
by intersecting three halfspaces

Resulting series parallel graph

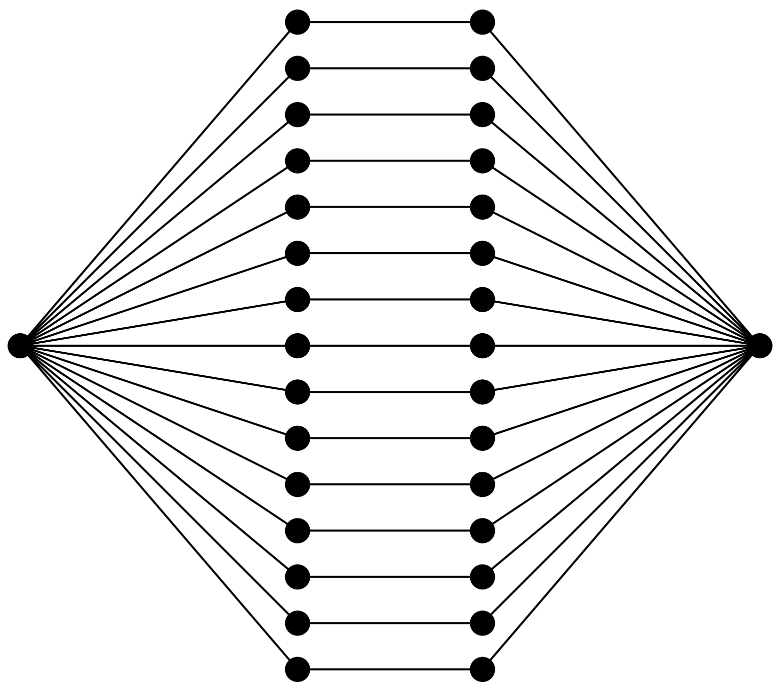

By combining several graphs on the same vertex set, get $\Omega(m\alpha(n))$ lower bound

Conclusions

- Many parametric matroid problems have simple reformulations as geometry problems
- Computational geometry tools help solve them
- Can matroid theory return the favor?
- What happens in higher dimensions?
  (i.e., more than one parameter)
- What about nonlinear weights?
  (moving points in $\mathbb{R}^d \rightarrow$ quadratic distances)
- Of special interest in geometry:
  How many swaps in uniform matroids?
  (Aka the k-set problem)