A Stronger Lower Bound on Parametric Minimum Spanning Trees

David Eppstein

17th Algorithms and Data Structures Symp. (WADS 2021)
August 2015

Minimum spanning tree problem

Input: Graph, each edge has weight

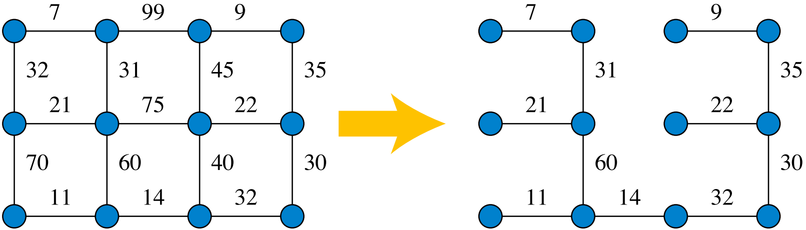

Output: Tree minimizing sum of edge weights

Many applications!

Parametric minimum spanning tree problem

Input: Graph, each edge has linear function $\lambda \mapsto$ weight

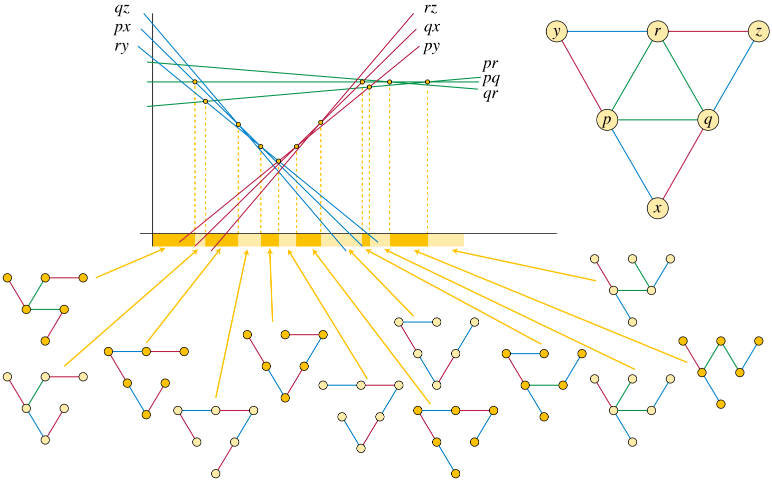

Output: Minimum spanning trees at each value of the parameter λ

Why should we study this problem?

Because it helps us trade off two different optimization criteria!

For a graph with quantities a_{uv} and b_{uv} at each edge uv

Given any spanning tree T, consider their sums:

$$A_T = \sum_{uv \in T} a_{uv} \qquad B_T = \sum_{uv \in T} b_{uv}$$

Search for tree optimizing nonlinear combination $f(A_T, B_T)$

Quasiconvex bicriterion problems

Special case of bicriterion minimum spanning tree seeking to maximize $f(A_T, B_T)$, where f is *quasiconvex*: all lower sets $\{(x, y) \mid f(x, y) \leq \theta\}$ are convex

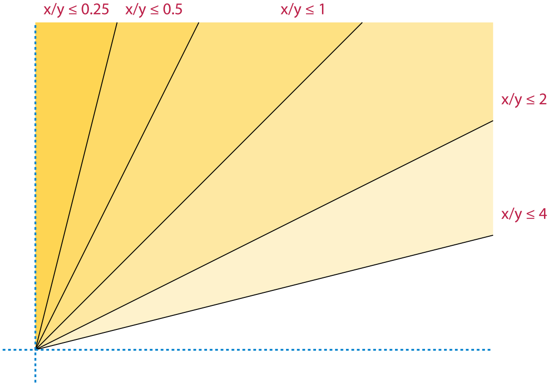

Example: return on investment $f(x, y) = x/y$, with $x, y > 0$

Converting bicriterion to parametric problems

Each spanning tree T of the graph
$\Rightarrow$ a point (A_T, B_T)

Max of quasiconvex function is a vertex of their *convex hull*

Exponentially many spanning trees, but convex hull is much smaller!

Its vertices are parametric minimum and maximum spanning trees weighted by $\lambda a_{uv} + b_{uv}$

Optimize by finding all these trees and testing which one is best

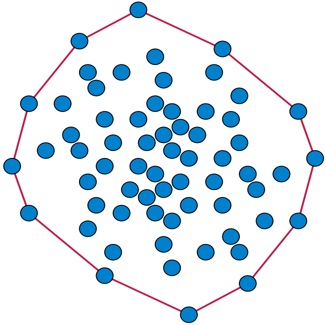

Some bicriterion-optimal tree problems

- Each edge has an eventual profit and investment cost; goal is to maximize return on investment

  Use two-criterion combination A/B

- Each edge has a normal distribution on weights, A and B are its parameters, and the goal is to find a tree minimizing $A + \sqrt{B}$ with high probability of having weight below some threshold

  Use two-criterion combination $A + \sqrt{B}$

- Each edge has a cost and a failure probability (represented by its log-likelihood), and we wish to minimize the cost-reliability ratio

  Use two-criterion combination Ae^B

What is known about parametric MST algorithms?

Two current fastest algorithms:

- $O(mn \log n)$ [Fernández-Baca et al. 1996]
- $O(n^{2/3} \log^{O(1)} n)$ per tree [Agarwal et al. 1998]

To know which of these two algorithms is faster, we need to know how many different trees a single instance can have

See also faster algorithms for planar graphs [Fernández-Baca and Slutzki 1997] or for some problems of finding a single optimal tree in a parametric problem [Katoh and Tokuyama 2001; Chan 2005])

How many trees can there be?

In this example: $n = 6$, $m = 9$, $\#$ trees $= 12$

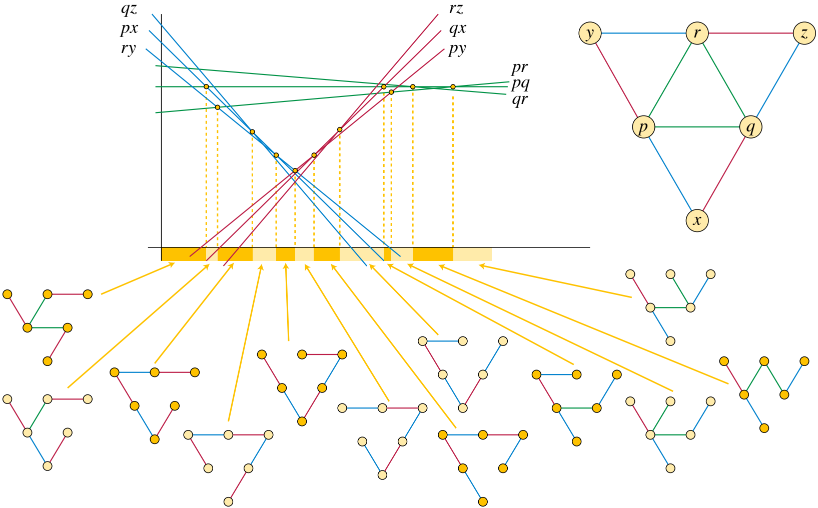

What was known about the number of trees?

$O(mn^{1/3})$ [Dey 1998]

Based on bounds on halving lines of point sets and their generalization to matroids

$\Omega\big(m\alpha(n)\big)$ [Eppstein 1998]

From conversion of lower envelope of line segments to graph with many 3-edge paths

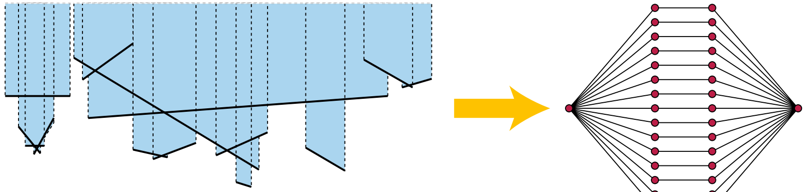

New result

Some graphs have $\Omega(m \log n)$ parametric MSTs

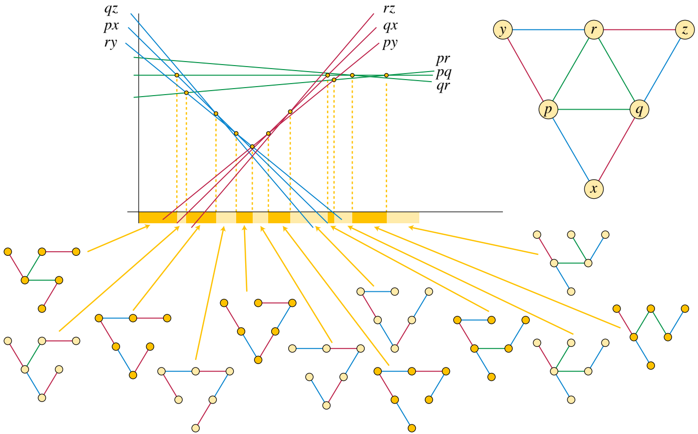

This example graph is an instance of our general construction

The graphs

Recursive construction

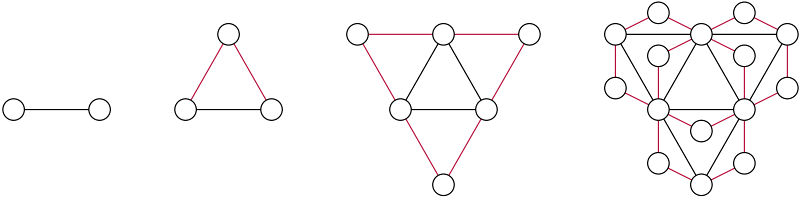

Repeatedly replace all edges by triangles

Replacements act like nonlinearly weighted edges

Within the larger graph, each triangle acts like an edge with weight equal to the *bottleneck path weight*, smallest w such that $\exists$ path of max weight $\leq w$ between the replaced edge's endpoints

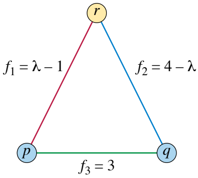

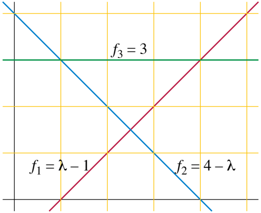

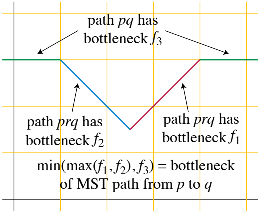

Global structure of replacement weights

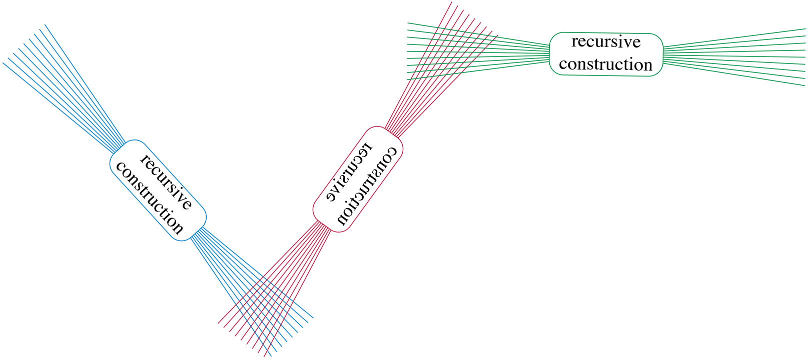

Three copies of smaller construction (one for each edge of replacement triangles), plus $\Theta(n)$ breakpoints where copies interact

$$3T(n/3) + \Theta(n) \Rightarrow \Theta(n \log n)$$

From $n \log n$ to $m \log n$ (graph view)

Subdivide recursive-triangle graph so each edge $\Rightarrow$ four-edge path

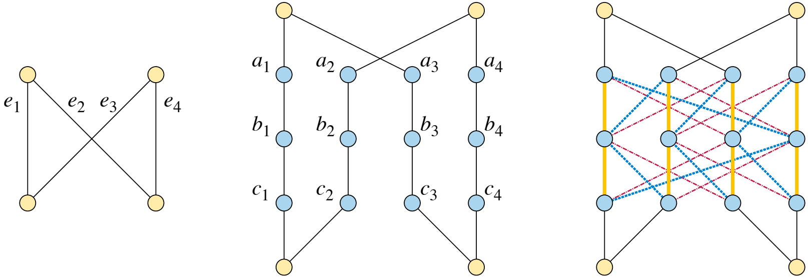

Pack many copies of subdivided graph into a single larger graph, using disjoint sets of edges in interiors of paths

From $n \log n$ to $m \log n$ (weight view)

Shared parts of paths have small weight, always in MST

Interiors of subdivided paths form disjoint $n \log n$ constructions

Arrange so each subdivided path has a range of parameters within which its bottlenecks control the MST

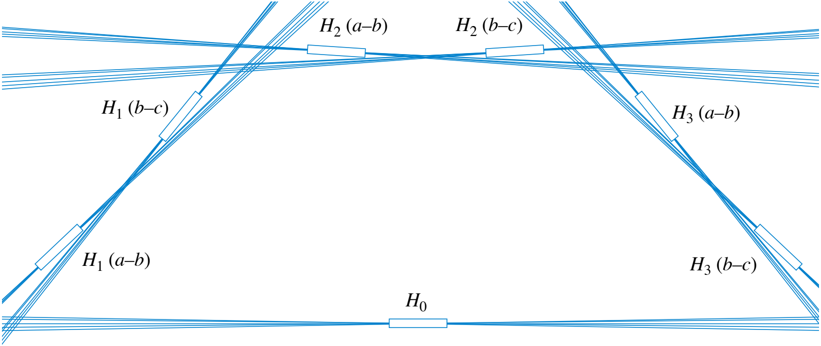

A problem I didn't solve

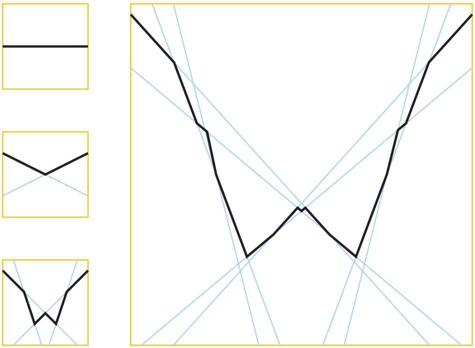

How many breakpoints for a min/max formula over n linear functions?

This is a *bottleneck shortest path* on a series-parallel graph

Parametric MST upper bounds apply, but new lower bounds do not

Known: $\Omega\big(n\alpha(n)\big)$
[Eppstein 1998]

Conclusions and open problems

New $\Omega(m \log n)$ bound on combinatorial complexity of the parametric MST problem

Still far from the known $O(mn^{1/3})$ upper bound,
but progress seems difficult
(closely related to notorious k-set problem in discrete geometry)

Previous $\Omega(m\alpha(n))$ bound also applied to parametric bottleneck shortest path, but the new bound does not – is it also $\Omega(m \log n)$?

References I

Pankaj K. Agarwal, David Eppstein, Leonidas J. Guibas, and Monika R. Henzinger. Parametric and kinetic minimum spanning trees. In *Proc. 39th IEEE Symp. Foundations of Computer Science (FOCS '98)*, pages 596–605, 1998. doi: 10.1109/SFCS.1998.743510.

Timothy M. Chan. Finding the shortest bottleneck edge in a parametric minimum spanning tree. In *Proc. 16th ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 917–918. SIAM, 2005. URL https://dl.acm.org/citation.cfm?id=1070432.1070561.

Tamal K. Dey. Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.*, 19(3):373–382, 1998. doi: 10.1007/PL00009354.

References II

David Eppstein. Geometric lower bounds for parametric matroid optimization. *Discrete Comput. Geom.*, 20(4):463–476, 1998. doi: 10.1007/PL00009396.

David Fernández-Baca and Giora Slutzki. Linear-time algorithms for parametric minimum spanning tree problems on planar graphs. *Theoret. Comput. Sci.*, 181(1):57–74, 1997. doi: 10.1016/S0304-3975(96)00262-9.

David Fernández-Baca, Giora Slutzki, and David Eppstein. Using sparsification for parametric minimum spanning tree problems. *Nordic J. Comput.*, 3(4):352–366, 1996.

Naoki Katoh and Takeshi Tokuyama. Notes on computing peaks in *k*-levels and parametric spanning trees. In Diane L. Souvaine, editor, *Proc. 17th Symposium on Computational Geometry (SoCG 2001)*, pages 241–248. ACM, 2001. doi: 10.1145/378583.378675.