On a Linear Program for Minimum-Weight Triangulation

EXTENDED ABSTRACT[*]

Arman Yousefi[†] Neal E. Young[‡]

Abstract

Minimum-weight triangulation (MWT) is NP-hard. It has a polynomial-time constant-factor approximation algorithm, and a variety of effective polynomial-time heuristics that, for many instances, can find the exact MWT. Linear programs (LPs) for MWT are well-studied, but previously no connection was known between any LP and any approximation algorithm or heuristic for MWT. Here we show the first such connections: for an LP formulation due to Dantzig et al. (1985): (i) the integrality gap is bounded by a constant; (ii) given any instance, if the aforementioned heuristics find the MWT, then so does the LP.

1 Introduction

In 1979, Garey and Johnson listed minimum-weight triangulation (MWT) as one of a dozen important problems not known to be in P nor NP-hard [15]. In 2006 the problem was finally shown to be NP-hard [28]. The problem has a sub-exponential time exact algorithm [32], as well as a polynomial-time approximation scheme (PTAS) for random inputs [18]. It is still not known whether, for some $\lambda > 1$, finding a λ-approximation is NP-hard, but this is unlikely as a quasi-polynomial-time approximation scheme exists [31]. MWT has an $O(\log n)$-approximation algorithm [30], and, most important here, an $O(1)$-approximation algorithm called QuasiGreedy [24]. The constant in the big-O upper bound from [24] is large (we estimate 100,000 or more).

If restricted to simple polygons, MWT has a well-known $O(n^3)$-time dynamic-programming algorithm [17, 22]. Polynomial-time algorithms also exist for instances with a constant number of "shells" [2] and for instances with only a constant number of vertices in the interior of the convex hull of V [16, §2.5.1], [19, 4, 33, 23].

Linear program of Dantzig et al. for MWT. LP methods are one of the primary emerging paradigms for the design of approximation algorithms. For many hard combinatorial optimization problems, especially so-called packing and covering problems, the polynomial-time approximation algorithm with the best approximation ratio is based on linear programming, either via randomized rounding or the primal-dual method. The design of a good approximation algorithm is often synonymous with bounding the integrality gap of an underlying LP.

MWT has several straightforward linear programming (LP) relaxations. Studying their integrality gaps may lead to better approximation algorithms, or may widen our understanding of general methods and their limitations. (Standard randomized-rounding and primal-dual approaches may be insufficient for MWT approximation algorithms.)

Dantzig et al. (1985) introduce the following LP (presented here as reformulated by [9]). Below $\triangle$ denotes the set of empty triangles.[1] R denotes the region to be triangulated minus the sides of triangles in $\triangle$. The LP asks to assign a non-negative weight X_t to each triangle $t \in \triangle$ so that, for each point in the region, the triangles containing it have total weight 1:

$$(1.1) \quad \text{minimize } c(X) = \sum_{t\in\triangle} c(t)X_t, \text{ subject to}$$

$$X \in \mathbb{R}^{\triangle}_{\geq 0} \text{ and } (\forall p \in R) \sum_{t \ni p} X_t = 1.$$

Above, the cost $c(t)$ of triangle t is the sum over the edges in t of the cost $c(e)$ of the edge, where $c(e)$ is $|e|/2$ (the length of e), unless e is on the boundary of R, in which case $c(e) = |e|$. (Internal edges are discounted by 1/2 since any internal edge occurs in either zero or two triangles in any triangulation.) R

[*] A full version of this paper appears on arxiv.org.

[†] Research partially funded by GAANN fellowship.

[‡] Research partially funded by NSF grant 0729071. Both authors are at the University of California, Riverside, USA.

[1] That is, triangles lying in the region to be triangulated, whose vertices are in the given set of points, but otherwise contain none of the given points.

doi:10.1137/1.9781611973099.65

Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, 811-823 (2012)

as specified is infinite, but can easily be restricted to a polynomial-size set of points without weakening the LP. (E.g., let R contain, for each possible edge e, two points p and q, one on each side of e, very near e.)

For the simple-polygon case, the above LP finds the exact MWT (every extreme point has 0/1 coordinates, and so corresponds to a triangulation). This was shown by Dantzig et al. (1985) [7, Thm. 7], then (apparently independently) by De Loera et al. (1996) [9, Thm. 4.1(i)] and Kirsanov (2004) [21, Cor. 3.6.2]. For summaries of these results, see [10, Ch. 8] and [35]. Kirsanov describes an instance (a 13-gon with a point at the center) for which this LP has integrality gap just above 1, as well as instances (50 random points equidistant from a center point) that are solved by the LP but not by the LMT-skeleton heuristic.

Other authors have considered *edge-based* LP's, mainly for use in branch-and-bound [25, 26, 29, 34, 3]. These edge-based LPs have unbounded integrality gaps. LPs for *maximal independent sets*, which are well studied, are closely related to all the above LPs, as triangulations can be defined as maximal independent sets of triangles (or of edges). The above LPs enforce some, but not all, well-studied inequalities for maximal independent sets.

It is known to be NP-hard to determine whether there exists a triangulation that uses only edges in a given subset S [27]. If we change the cost function in the above LP to $c(X) = \sum_{e\in S}\sum_{t\ni e} X_t$, the LP will have a zero-cost integer solution iff there is such a triangulation. Thus, the LP with that cost function has unbounded integrality gap (unless P=NP). Thus, any analysis of the LP for MWT must rely intrinsically on the cost function. (This is also an obstacle for many randomized-rounding methods.)

First new result. We show that LP (1.1) *has constant integrality gap.*

This is the first non-trivial integrality-gap upper bound for any MWT LP. To show it, we revisit the analysis of QuasiGreedy [24], which shows that QuasiGreedy produces a triangulation of cost $O(|\text{MWT}(G)|)$, where $|\text{MWT}(G)|$ is the length of the MWT of G (and also the cost of the optimal *integer* solution to the LP). We generalize and strengthen their arguments to show that there exists a triangulation of cost at most $O(c(X^*))$, where $c(X^*)$ is the cost of the optimal fractional triangulation (i.e., solution to the LP).

Our analysis reduces the approximation ratio in their analysis by an order of magnitude, but it remains a large constant.

MWT heuristics. Much of the MWT literature concerns polynomial-time heuristics that, given an instance, find edges that must be in (or out of) any MWT. Here is a summary. Gilbert observe that the *shortest potential edge* is in every MWT [17]. Yang et al. extend this result by proving that an edge xy is in every MWT if, for any edge pq that intersects xy, $|xy| \le \min\{|px|, |py|, |qx|, |qy|\}$ [37]. (We refer to the edges satisfying this property as the *YXY subgraph.*) This subgraph includes every edge connecting two *mutual nearest neighbors*. Keil et al. show (for some $\beta > 1$) that, if, for an edge pq, the two circles of diameter $\beta \cdot |pq|$ passing through p and q are empty (of other vertices), then pq is in every MWT [20]. Cheng et al. strengthen this to $\beta \approx 1.17682$ [6]. The set of such edges is called the *β-skeleton.* Das and Joseph show that an edge e cannot be in any MWT if both of the two triangles with base e and base angle $\pi/8$ contain other vertices [8]. Drysdale et al. strengthen this to angle $\pi/4.6$ [14]. This property of e is called the *diamond property.* Dickerson et al. describe a simple local-minimality property such that, if an edge e lacks the property, the edge cannot be in any MWT. Using this, they show that the so-called *LMT skeleton* must be *in* the MWT [11].

A primary use of the heuristics is to solve some instances of MWT exactly in polynomial time, as follows: *Given an instance, use the heuristics to identify edges that are in the MWT. If the regions left untriangulated by these edges are simple polygons (equivalently, if the edges connect the given points) then find the MWT of each region independently using the standard dynamic programming algorithm.* (The MWT will be the union of the MWT's of the regions.) According to [11] (1997), most random instances with 40,000 points are solvable in this way.

Second new result. We show that LP (1.1) *generalizes these heuristics* in that *if the heuristics solve a given instance as described above, then so does the LP* (that is, the extreme points of the LP are incidence vectors of optimal triangulations).

In fact the LP appears to be stronger than the heuristics, in that some natural instances are solved by the LP, but not by the heuristics [21, §3.5].[2] In this sense, the LP, whose formulation requires little explicit geometry, generalizes all of these varied and generally incomparable heuristics. This is the first connection we know of between the heuristics and any MWT LP.

[2] Where C contains the center of a unit circle and $n-1$ random points on the circle.

Roughly, the heuristics are based on a combination of (i) local-improvement arguments about the MWT and (ii) logical closure (once the heuristic determines the status of one edge with respect to the MWT, this in turn determines the status of other edges, and so on). We extend these arguments to apply to the optimal fractional triangulation X^*. This is possible because (i) X^* looks "locally" like a MWT and (ii) the LP enforces logical closure of linear constraints on X^*.

After we finished the body of this work, we became aware of and examined additional heuristics by Wang et al. [36] and Aichholzer et al. [1]. We conjecture that the LP generalizes them as well.

An equivalent formulation of the LP. The following constraints are equivalent to the last constraints in LP (1.1) (see e.g. [9, Thm. 1.1(i), Prop. 2.5], [35], or [21, Thm. 3.4.1]) and are useful for reasoning about fractional triangulations. For any fractional triangulation X and edge e,

$$(1.2) \quad \sum_{t \in S(e)} X_t \; - \sum_{t \in T(e)} X_t \; = \; [e \in \text{boundary}(R)].$$

Here $S(e)$ contains the triangles that contain e and lie on one side of e, while $T(e)$ contains the triangles that contain e and lie on the other side of e. (If e is on the boundary, take $T(e) = \emptyset$.) The notation $[x \in S]$ denotes 1 if $x \in S$ and 0 otherwise.

Practical considerations. Using the $O(n^2)$ constraints (1.2) instead of the constraints in (1.1) gives an equivalent LP with total size (i.e., non-zeros in the constraint matrix) proportional to the number of empty triangles. The empty triangles can be identified, and the LP constructed, in time proportional to their number [13]. Their number is always $O(n^3)$, but often smaller (e.g. $O(n^2)$ in expectation for randomly distributed points).

The time to construct and solve the LP can be further reduced by a preprocessing step based on the heuristics — remove any variable X_t if the heuristics prove any edge of t to be out of every MWT, and add a constraint $\sum_{t \in S(e)} X_t = \sum_{t \in T(e)} X_t = 1$ if they prove an interior edge e to be in every MWT. For randomly distributed points, only $O(n)$ edges (in expectation) have the diamond property, forming $O(n^2)$ possible empty triangles, from which the modified LMT skeleton can be computed in $O(n^2)$ time [11, 12]. On "typical" instances with $10^4 - 10^5$ points, only a very small number of variables are left undetermined by the heuristics. (For n random points, the expected number is $\Omega(n)$, but with an apparently astronomically small leading constant [5].) This allows standard LP software to quickly solve the LP, and integer-LP solvers to quickly find the MWT.

Remarks. We do not give an algorithm per se, and the integrality-gap bound, though constant, is large. But both results suggest that the LP of Dantzig et al. captures much of the structure of MWT. This suggests a clear line of attack for finding an approximation algorithm with reduced approximation ratio: study the integrality gap of the LP, trying systematic LP methods such as the primal-dual method. If constantly many rounds of lift-and-project (applied to the LP) yield an LP with integrality gap $1 + \epsilon$. Or, if randomized-rounding, primal-dual, and similar approaches fail, this may help us better understand their limitations.

2 Definitions

The *interior* of a segment pq is $pq - \{p, q\}$. The *interior* of a polygon P consists of P minus its boundary. Two sets *properly intersect* (or *overlap*, or *cross*) if the intersection of their interiors is non-empty. The (Euclidean) length of line segment pq is $|pq|$. For any set E of segments, $|E|$ is the total length of segments in E.

A *planar straight-line graph* (PSLG) is an undirected graph $G = (V, E)$ along with a planar embedding that identifies each vertex with a planar point and each edge with the line segment connecting its endpoints, so that each edge intersects other edges (and V) only at its endpoints. The *length* of G is the sum of the Euclidean lengths of its edges. G partitions the plane into polygonal *faces*.[3] A face or polygon is *empty* if its interior contains no vertex.

A *diagonal*, or *potential edge*, of G is any segment $pq \notin E$ connecting two vertices of a face, and contained in that face, so that $G' = (V, E \cup \{pq\})$ is still a PSLG. A *partition* of G is a PSLG that extends G by adding (non-crossing) diagonals; equivalently, the faces of the partition refine the faces of G. A *convex partition* of G is a partition whose faces are empty and strictly convex. The minimum-length convex partition of G is denoted MCP(G). A *triangulation* of G is a partition whose faces are empty triangles. A *fractional triangulation* X is a feasible solution to the LP.

[3] Where two points are in the same face if there is a path between them that intersects no edge, with the caveat that the term *face* excludes the single such unbounded region.

Formally, an instance of MWT is specified by a planar point set V, implicitly defining a PSLG $G = (V, E)$ where E contains the edges on the boundary of the convex hull of V. A solution is a minimum-length triangulation of G.

Throughout, we fix an instance $G = (V, E)$ of MWT specified by a given point set V. Unless stated otherwise, every graph considered is a partition of G. Since the vertex set V is the same for all such graphs, we identify each graph with its edge set.

3 LP (1.1) has constant integrality gap

PROPOSITION 3.1. *Given any instance $G = (V, E)$ of MWT, for any fractional triangulation X, there exists an integer solution of value $O(c(X))$. Thus, LP (1.1) has constant integrality gap.*

The rest of this section proves the proposition.

Fix the MWT instance G and an arbitrary fractional triangulation X. For now, also fix an arbitrary convex partition C. (Later, we will specify how to choose C.)

Summary of proof. The idea is to define a sort of "rounding" procedure that converts X into the desired integer solution. The main step of the procedure converts X into a separate fractional triangulation X^f for each face f of C (covering just f). Next, independently within each face f of C, the procedure replaces the fractional triangulation X^f by the optimal integer triangulation of f. The final "rounded" solution is then the union of these integer triangulations (one for each face f of C), of total cost at most $\sum_{f \in C} c(X^f)$ (and, hopefully, $O(c(X))$).

In the second step, since each f is a simple polygon, it follows from known results (e.g. [7, Thm. 7]; see the introduction) that the cost of the integer triangulation of f is at most the cost of X^f. Thus, the integrality gap will be $O(1)$ as long as the main step triangulates the faces so that $\sum_f c(X^f) = O(c(X))$.

The proof divides into two parts: (1) defining the rounding procedure and showing that it produces a feasible fractional triangulation X^f of each face f, and (2) showing that $\sum_f c(X^f) = c(X)$.

Proofs of the main theorems are postponed to subsequent sections.

The rounding procedure. In addition to the given G, X, and C, let f be an arbitrary face of the convex partition C.

To convert X into a fractional triangulation X^f of f, start by focusing on just the triangles that cross f and have positive weight in X. We "break" each such triangle t into a set $\widehat{t^f}$ of triangles within t. Then, in X^f, we give each triangle in $\widehat{t^f}$ weight X_t.

To break each triangle t into a set $\widehat{t^f}$ of triangles in f, we leverage the concept of edge transposals from [24, (see e.g. Lemma 4.2)]. The reader may skim the details of the definition on first reading.

DEFINITION 3.1. (TRANSPOSALS OF TRIANGLES) *Given a triangle t crossing face f, the* triangulated transposal $\widehat{t^f}$ of t (in f) *is obtained as follows. First, orient*[4] *each edge e of t so that t lies to the right of e. Next, for each edge e of t independently, obtain its* edge transposal *in f, denoted e^f, as follows:*

1. Clip e to $\tilde{e} = e \cap f$.
2. Obtain e^f by sliding each endpoint p of $\tilde{e}$ to an "adjacent" vertex of f: if p is a vertex of f, leave it there, otherwise p lies in one edge YZ of f, slide it to Y or Z, choosing the destinations of the endpoints to minimize $|e^f|$ (and breaking ties consistently).
3. Let e^f inherit e's orientation in the natural way.

Now define the (non-triangulated) transposal t^f *of triangle t to be the polygon containing those points in f that, for every each edge e of t, lie to the right of its transposal e^f.*

Define the triangulated transposal, $\widehat{t^f}$, *of t to be a minimum-length triangulation of the transposal t^f. If t^f has no area, then $\widehat{t^f} = \emptyset$.*

The transposal t^f has at most six sides. It might have no area.

Below is a convex face f (with vertices A, C, D, F, G) blanketed by a collection of triangles (numbered 1-6, with thin edges). To the right of that each triangle is clipped into the face. To the far right are the transposals of the triangles. The only triangles whose transposals have positive area are triangles 3 and 4; their respective transposals are the 4-gon labeled 3 and the 3-gon labeled 4 (shown to the far right). (The dashed, gray arrow extending from each edge points towards its edge transposal. The only edge transposal that does not lie on the boundary of f is $(Ae)^f = AD$.)

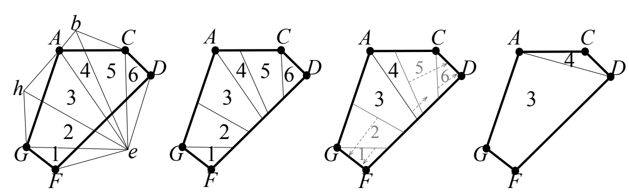

[4] I.e., order the two endpoints.

The fractional triangulation X^f of face f is then obtained (from X) simply by giving weight X_t to each triangle s in the triangulated transposal $\widehat{t}^f$ of t:

DEFINITION 3.2. (TRANSPOSAL X^f OF X IN f) *The transposal of X in f, denoted X^f, assigns to each potential triangle in s in f the weight $X^f_s = \sum_{t:s\in\widehat{t}^f} X_t$. (Here t ranges over triangles that cross f. More than one t may contribute to X^f_s.)*

That X^f is a fractional triangulation of f is not immediate from the definition. This requires proof:

THEOREM 3.1. *The transposal X^f of X in f is a fractional triangulation of f. That is, it covers the points in f uniformly with weight 1.*

Section 3.1 gives the proof. The proof uses two observations — (1) for any given f, the fractional triangulation X, restricted to triangles that cross f, can be decomposed into layers, each of which looks like an actual triangulation (but possibly extending outside of f), and (2) within each layer, the triangles can be collectively "morphed" into their triangulated transposals, while maintaining uniform coverage of f.

Bounding the cost. Having defined the fractional triangulation X^f of each face f, it remains to bound their total cost $\sum_f c(X^f)$ (over all faces of C).

The bound will depend on the *sensitivity* of the edges of C, defined as follows:

DEFINITION 3.3. (SENSITIVITY) *An edge e is r-*sensitive *if, for any potential edge d that crosses e, for each endpoint x of d, the distance from x to its closest endpoint of e is at most $r|d|$.*

The core of the cost bound is this theorem:

THEOREM 3.2. *If every edge in the convex partition C is r-sensitive, then $\sum_f c(X^f)$ is at most $3|C| + 12r\, c(X)$.*

Section 3.2 gives the proof, which is based on several observations. (1) Although a given triangle t can cross arbitrarily many faces, and has a transposal t^f in each of those faces, t crosses at most *two* faces in which its transposal has positive area. Thus, t can contribute to the cost of at most *two* faces. (2) the cost of any transposal t^f of t (not counting the edges in C) cannot exceed the cost of t by much. This follows from the definition of edge transposals and the sensitivity of C's edges, which imply that, for each edge e of t, the transposal of e cannot be much longer than e. (3) Each transposal t^f has at most six sides, so triangulating it to obtain the triangulated transposal $\widehat{t}^f$ increases the cost by a constant factor.

Thm. 3.2 gives an upper bound of $3|C|+12r\, c(X)$.

To use this bound we need C to have $|C| = O(c(X))$ and $r = O(1)$. Existing results by Levcopoulos and Krznaric's get us most of the way there:

THEOREM 3.3. ([24]) *For some constant $\lambda > 0$, and any MWT instance G, there exists a convex partition* LK *of G, whose edges are 4.45-sensitive, having total length $|\text{LK}| \le \lambda\,|\text{MCP}(G)|$. (Recall that* MCP$(G)$ *is the minimum-length convex partition of G.)*

Proof. Levcopoulos and Krznaric show that what they call the *quasi-greedy convex partition* has these properties: for Property (1), see their Lemma 5.4 and the discussion before it; for Property (2), see their Corollary 5.3 [24]. □

We now fix the (previously arbitrary) convex partition C to be the partition LK from Thm. 3.3. To use the bound in Thm. 3.3, we need to show that $|\text{MCP}(G)|$ is $O(c(X))$.

This is relatively easy. Using the constraints on X and a previous analysis of MCP(G) due to Plaisted and Hong [30, Lemma 10], we show the following bound:

LEMMA 3.1. $|\text{MCP}(G)| \le 18\, c(X)$

The proof is in Section 3.3.

Finally, combining the two theorems and the lemma, the cost of the final integer triangulation is at most

$\sum_f c(X^f)$	As each f is simple.
$\le 3\lvert\text{LK}\rvert + 12\, r\, c(X)$	By Thm. 3.2.
$\le 3\lambda\lvert\text{MCP}(G)\rvert + 54\, c(X)$	By Thm. 3.3.
$\le 3\lambda \cdot 18\, c(x) + 54\, c(X)$	By Lemma 3.1.
$= 54(\lambda + 1)\, c(X)$	

Above λ is the constant from Thm. 3.3. (Although λ is large, the bound above is still substantially smaller than the approximation ratio proven for QUASIGREEDY in [24].)

This (with the proofs of Thm. 3.1, Thm. 3.2, and Lemma 3.1 below), proves Proposition 3.1.

3.1 Proof of Thm. 3.1. Let G, X, C, and f be as above. We start with the observation that the fractional triangulation X, restricted to triangles that cross f, can be decomposed into a weighted sum of incidence vectors of what that we call *blankets*:

DEFINITION 3.4. (BLANKET) *A set B of empty polygons with endpoints in V* blankets *the face f if the union of the polygons contains f and no two of the polygons overlap within f (they may overlap outside f).*

If the polygons in B are triangles, the transposal *of B (in f), denoted B^f, is the set containing, for each triangle $t \in B$, the transposal t^f of t. That is, $B^f = \{t^f \mid t \in B\}$. The* triangulated transposal, *denoted $\widehat{B}^f$, of B (in f) is just the (multiset) union, over all triangles $t \in B$, of the triangulated transposal of t. That is, $\widehat{B}^f = \bigcup_{t\in B} \widehat{t}^f$.*

The next lemma says that, over f, X can be decomposed into a weighted sum of blankets.

LEMMA 3.2. *There exists a set $\mathcal{B}$ of blankets, with a weight $\epsilon_B > 0$ for each $B \in \mathcal{B}$, such that $\sum_{B\in\mathcal{B}} \epsilon_B = 1$ and, for every triangle t crossing f, $X_t = \sum_{B\in\mathcal{B}} [t \in B]\, \epsilon_B$.*

(Recall "$[t \in B]$" is 1 if $t \in B$, else 0.)

Proof. Recall that, for instances consisting of a simple polygon, the LP gives optimal 0/1 solutions (e.g., [7, Thm. 7]). We adapt a proof of that property.

Choose any triangle t that crosses f and has $X_t > 0$. If any edge e of this triangle crosses (the interior of) f, since e has positive weight, there must be a positive-weight triangle s that has e as an edge and lies on e's opposite side (see Constraint (1.2)). Glue t and s together to form a polygonal region. Continue in this way, growing the polygonal region by repeatedly gluing a new triangle to any boundary edge e that crosses f. Stop when the region has no such boundary edges. The triangles glued together in this way must form a blanket B of f.

Let ϵ_B be the minimum weight of any triangle in B. This gives the first blanket B and its weight ϵ_B. Subtract ϵ_B from each X_t for $t \in B$. This reduces X's coverage of f uniformly by ϵ_B. To generate the remaining blankets in $\mathcal{B}$ (and their weights), iterate this process as long as X still covers f with positive (and necessarily uniform) weight.

(The process does terminate, as each iteration brings some X_t to zero.) □

We will also use the following lemma, whose (long) proof we delay.

Let f' denote f minus points on potential edges.

LEMMA 3.3. *For any blanket $B \in \mathcal{B}$, the triangulated transposal $\widehat{B}^f$ of B also blankets f.*

(The lemma is essentially the theorem we are proving, restricted to the special case when the triangles t crossing f have integer weight $X_t \in \{0, 1\}$, i.e., those with $X_t = 1$ blanket f.)

Fix any point $p \in f'$. We will use the lemmas above to show that X^f covers p with weight 1.

Restrict attention to triangles t that cross f. Recall that X^f is obtained from X by "transferring" weight X_t from each triangle t to the triangulated transposal of t. So X^f covers p with weight $\sum_t \sum_{s\in\widehat{t}^f} [p \in s] X_t$.

By Lemma 3.2, each weight X_t can be split into the sum of the weights of the blankets B containing t. That is, $X_t = \sum_{B\in\mathcal{B}} [t \in B]\epsilon_B$.

Combining these two observations, X^f covers p with weight

$$\begin{aligned}
&\sum_t \sum_{s\in\widehat{t}^f} [p \in s] \sum_{B\in\mathcal{B}} [t \in B]\epsilon_B \\
&\quad = \sum_{B\in\mathcal{B}} \epsilon_B \sum_{t\in B} \sum_{s\in\widehat{t}^f} [p \in s] \\
&\quad = \sum_{B\in\mathcal{B}} \epsilon_B \sum_{s\in\widehat{B}^f} [p \in s].
\end{aligned}$$

The final sum on the right, $\sum_{s\in\widehat{B}^f} [p \in s]$, is the number of triangles that cover p in the triangulated transposal of B. By Lemma 3.3, this number is 1. Thus, each blanket B contributes ϵ_B to the coverage of p by X^f. Thus, X^f covers p with weight $\sum_{B\in\mathcal{B}} \epsilon_B$, which equals 1.

To finish proving Thm. 3.1, we prove Lemma 3.3.

The idea is to morph B continuously into its (non-triangulated) transposal $B^f = \{t^f \mid t \in B\}$. Specifically, morph the edges of triangles in B as follows: *First, for every triangle edge e, clip e to the chord $e\cap f$ of f, giving a set of chords. Next, for every side YZ of face f (in any order), do the following step: simultaneously, for every chord za having an endpoint $z \in YZ$, slide the endpoint z continuously along YZ at unit rate to the corresponding endpoint (Y or Z) of za's transposal za^f. As the endpoint z moves, move the chord za as well (as shown below).*

Below are the start, middle, and end of one step of the morphing process for a single side YZ of f. The moving chords are labeled 1-9. Chords not touching YZ don't move and aren't shown.

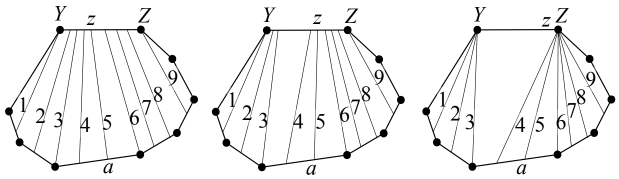

This morphing process morphs each edge e of each triangle $t \in B$ to the chord $e \cap f$, and then morphs that chord continuously until it arrives at its transposal, e^f. We will show below that as the chords move *no crossings are introduced.* Thus, the following invariant is maintained: *the regions (each one coming from a triangle $t \in B$) collectively blanket f.*

Clearly each triangle $t \in B$ (after being clipped to $t \cap f$) is morphed into its transposal t^f. Thus, the final set of regions is exactly B^f, which (by the invariant) must blanket f.

Since the triangulated transposal $\widehat{B}^f$ of B is obtained from B simply by triangulating each polygon in $\widehat{B}^f$ (preserving the exact covering of f'), the lemma follows.

To complete the proof, we show that during morphing no chords cross. Consider the step for any side YZ of f.

Observation: *As the endpoint y of a chord ya slides along YZ, the transposal of ya is invariant.*

Proof. By the definition of transposal, $ya^f = ya^f$ for any z that is (with y) in the interior of YZ. Thus, the transposal doesn't change while y stays in the interior of YZ. And, if (e.g.) y is moving towards Y, then YA is the transposal of ya, so Y must be a closest point in $\{Y, Z\}$ to A. This implies that the transposal of Ya is also YA. Thus, when the chord's endpoint y arrives at Y, the transposal of the chord does not change. □

The observation implies that the morphing process indeed maps each edge e of each triangle $t \in B$ to its transposal e^f.

For any two points $y, z \in YZ$, let $y \prec z$ denote that y comes before z when traveling from Y along YZ. Overloading notation, for any two points a and b on the boundary of f, minus YZ, let $a \prec b$ denote that a comes before b when traveling from Y to Z along the boundary minus YZ. (In the diagram, $y \prec z$ and $a \prec b$.)

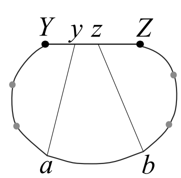

Now let ya be an arbitrary chord such that the step slides y towards Y. Let zb be an arbitrary chord such that the step slides z towards Z. To finish the proof, we will show that $y \preceq z$. Thus, the morphing process does not cause chords to cross. (On consideration, this implies that each triangle t, after being clipped to $t \cap f$, gets morphed to a corresponding region t^f with the claimed properties.)

Fix A and B such that the transposals of ya and zb are YA and ZB, respectively.

Observation: *If $b \preceq a$, then $B \preceq A$.* (Transposing preserves the order of the non-YZ endpoints.)

Proof. In the case that a and b both lie in the interior of a single side of f, it must be that the transposals of ya and zb are the same (because y and z are also both in the interior of a single side, YZ), so A equals B (so $B \preceq A$). In the remaining case (by the definition of transposal), there exist two distinct sides PQ and RS (other than YZ) of f such that $a, A \in PQ$ and $b, B \in RS$. Since $b \preceq a$, this implies $B \preceq A$. □

Observation: $A \prec B$.

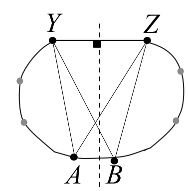

Proof. Since YA is the transposal of ya, point Y must be a closest point in $\{Y, Z\}$ to A; that is, A must lie on the Y-side of the bisector of YZ. Likewise B must lie on the Z-side of the bisector. It follows from the convexity of f that $A \preceq B$. And, since ties are broken consistently in choosing transposals, it cannot be that $A = B$. □

The last two observations imply that $a \prec b$. Assuming inductively that chords ya and zb are non-crossing at the start of the step, this implies that $y \preceq z$. Thus, as y slides towards Y and z slides towards Z, the chords remain non-crossing throughout the step.

This concludes the proof of Thm. 3.1.

3.2 Proof of Thm. 3.2. We want to bound the total cost of the fractional triangulations that X induces in all faces f of C, that is, $\sum_f c(X^f)$.

In this section, for convenience, we define $c(\widehat{t}^f) = c(t^f) = 0$ if t does not cross f or if t^f has area zero. We will prove the following lemmas:

Lemma 3.4: *Any given triangle t crosses at most two faces f in which its transposal t^f has positive area. Thus, for a given t, only two faces f have $c(t^f) > 0$.*

Lemma 3.5: *For any triangle t and face f, the cost of t^f minus the edges in C is at most $2r$ times the cost of t.* (Recall that r is the sensitivity of C's edges.)

Lemma 3.6: *The cost of the triangulated transposal $\widehat{t^f}$ is at most three times the cost of the (non-triangulated) transposal t^f.*

Before we prove the lemmas, we note that they imply the theorem as follows. The total cost is

$$\begin{aligned}
&\textstyle\sum_f c(X^f) \;=\; \sum_{f,t} X_t\, c(\widehat{t^f}) && \text{by def'n of } X^f\\
&\le 3 \textstyle\sum_{f,t} X_t\, c(t^f) && \text{by Lemma 3.6}\\
&\le 3|C| + \; 6r \textstyle\sum_{t,f}\, X_t\, c(t)\, [c(t^f) > 0] && \text{by Lemma 3.5}\\
&\le 3|C| + 12r\, \textstyle\sum_t\, X_t\, c(t) && \text{by Lemma 3.4}\\
&= 3|C| + 12r\, c(X).
\end{aligned}$$

Lemma 3.4. *Given any triangle t, there are at most two faces f of C in which t's transposal t^f has positive area.*

Proof. Fix a triangle $t = \Delta ABD$ and consider how the faces of C can overlap t. Say that a face is *accommodating* if t's transposal t^f in f has positive area.

We start with two examples.

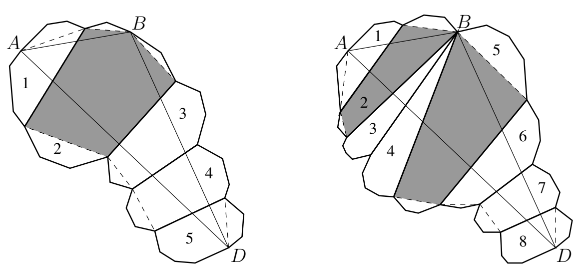

Above, each dashed edge is the edge transposal of an edge of t. Within each accommodating face, the (positive-area) transposal of t is dark.

We claim that *every accommodating face touches all three edges of t* (where touching an endpoint of an edge counts as touching the edge). (For example, the accommodating faces 2 on the left, and 2 and 5 on the right, touch all three edges of t. Each other face is non-accommodating and, except for 3 and 4 on the right, touches only two edges of t.) The claim holds because, if a face f touches only two edges of t, then $f \cap t$ lies within a "corner" of t. Then two of t's edges must cross the same two sides (or vertices) of f in the same way, and those two edges will have the same transposals (directed oppositely), forcing t^f to have no area.

Now consider the case that t has a face f that touches the *interior* of all three edges of t (as in the figure to the left, above). Since no other face f' can cross f, no face other than f can touch all three edges of t. By the claim, then, only face f might be accommodating, so the lemma holds.

So assume that no face touches the interior of all three edges of t.

By the claim, any accommodating face f still has to touch all three edges of t, but now there is at least one edge, say AB, of t whose interior f avoids. Thus, f must touch AB at an endpoint, say, B. (For example, consider the figure on the right above. Faces 2, 3, 4, and 5 touch all three edges of t, but not all three interiors.) Since f touches AB at B, but does not touch the interior of AB, there must be an edge Bx of f that extends through the interior of t. Since x is not inside t, Bx must cut across t to the interior of the edge AD. Thus, *any accommodating face f must share some vertex v with t, and an edge of the face must extend from v across the interior of t.*

If there are two accommodating faces, they must extend an edge across t from the *same* vertex v, for otherwise the extending edges would cross inside t. Let this vertex be B.

Now consider all edges in C that extend from B across the interior of t. Let these edges be $Bx_1, Bx_2, \ldots, Bx_k$, rotating in order around B. (In the picture above, $k = 3$.) C has $k+1$ corresponding faces $f_0, f_1, \ldots, f_k$, also in order rotating around B, where f_{i-1} and f_i share edge Bx_i. By the conclusion of the paragraph before last, only these $k+1$ faces might be accommodating.

To finish, we observe that f_i is not accommodating unless $i \notin \{0, k\}$ (the first or last face). Indeed, for $i \notin \{0, k\}$ edges Bx_{i-1} and Bx_i of f_i extend from B across t to AD. Since these edges touch at B, the transposal of AD in f_i is thus *just the point B*. Thus, the transposal of t in f_i has no area. □

Lemma 3.5. *Assume C's edges are r-sensitive. For any face f and triangle t, the total cost of the edges that are t's transposal in f but not in C is at most $2r$ times the cost of t.*

Proof. Let f be any face of C and e be any edge that crosses f. We claim that *the length of the edge* *transposal e^f of e in f is at most $2r$ times the length of e.* This claim implies the lemma, because each edge in the transposal of t, but not in C, is the edge transposal e^f of an edge e in t that is not on the boundary of f. We prove the claim.

For an edge e that crosses a face f one of the following cases holds: (1) e is incident to two vertices of f, or (2) e is incident to one vertex of f and properly intersects one s side of f (as in the figure on the left below) or (3) e properly intersects two sides s and s' of f (as in the figure on the right below).

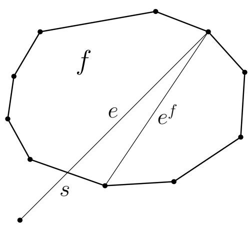

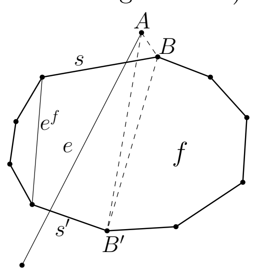

In case (1), the transposal e^f of e is the same as e, so the claim holds. In case (2), since s is r-sensitive, and e^f is the shortest segment from the endpoint of e to an endpoint of s, $|e^f| \le r|e|$. In case (3), let A be an endpoint of e and let B and B' respectively be the closest endpoints of s and s' to A. Because e^f is the shortest segment from an endpoint of s to an endpoint of s', $|e^f| \le |BB'|$.

By the triangle inequality, $|BB'| \le |AB| + |AB'|$.

Because s and s' are r-sensitive, $|AB|$ and $|AB'|$ are each at most $r\,|e|$, proving the lemma. ☐

LEMMA 3.6. *For any face f and any triangle t, the cost $c(\widehat{t^f})$ of the triangulated transposal of t in f is at most three times the cost $c(t^f)$ of the transposal of t in f.*

Proof. As observed previously, t^f is a convex polygon with at most six sides. Let set B contain the edges in t^f. Let set D contain (up to three) diagonals of t^f, connecting alternating vertices around the boundary of t^f such that $B \cup D$ partition t^f into a triangulation T (a set of triangles).

Recall that $c(e) = |e|/2$ for each $e \in D$, while $c(e) \in \{|e|/2, |e|\}$ for $e \in B$. By the choice of diagonals, $|D| \le |t^f|$, so $c(D) = \sum_{e \in D} c(e) \le c(t^f)$. Clearly $c(B) = \sum_{e \in B} c(e) = c(t^f)$.

Each edge in B occurs in one triangle in T, while each edge in D occurs in two. Thus, $c(T) = \sum_{t \in T} c(t) = c(B) + 2\,c(D) \le 3\,c(t^f)$.

The lemma follows, as $c(\widehat{t^f}) \le c(T)$. ☐

This concludes the proof of Thm. 3.2.

3.3 Proof of Lemma 3.1. For every vertex v in the interior of V, define a *star* at v to be a subset of edges incident to v such that no two successive edges (around v) are separated by an angle of 180 degrees or more. For every vertex v on the boundary of V, define the (only) star at v to consist of the two boundary edges incident to v. Let $S_{\min}(v)$ denote the minimum cost of any star at v. Plaisted and Hong show $|\textsc{mcp}(G)| \le 6 \sum_v S_{\min}(v)$ [30, Lemma 10].

We claim $\sum_v S_{\min}(v) \le (3/2) \sum_v \sum_{e \ni v} X_e|e|$.

As $\sum_v \sum_{e \ni v} X_e|e| = 2 \sum_e X_e|e| = 2\,c(X)$, this implies the lemma.

We prove the claim.

It's easy to see that, for any boundary vertex v, $S_{\min}(v) = \sum_{e \ni v} X_e|e|$, so restrict attention to just an interior vertex v and its edges.

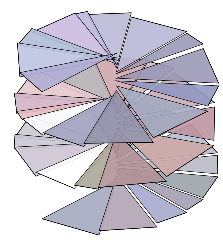

Because X satisfies constraint (1.2), rotating around v, there is a sequence $e_1, e_2, \ldots, e_k$ of edges such that each e_i forms a positive-weight triangle with its "neighboring" edge e_{i+1} (e_1 if $i = k$).

Call this sequence of edges a *helix*, h. Let $w(h)$ denote the number of times h wraps around v. Let N^h_e be the multiplicity of e in h. By a standard construction the X_e's can be expressed as a linear combination of incidence vectors of helices. (Similar to Lemma 3.2's proof, repeatedly find a helix h, choose weight ϵ_h, and subtract $\epsilon_h N^h_t / w_h$ from each X_t, reducing coverage near v by ϵ_h.) This gives a probability distribution ϵ on helices such that each $X_e = \sum_h \epsilon_h\, N^h_e / w(h)$.

Now choose a helix h at random from the probability distribution ϵ. Break (partition) h greedily into disjoint *groups* of contiguous edges such that each group g is maximal subject to the constraint that the neighboring edges' angles in g's total at most 360°. (In the figure, white triangles separate groups.) Consideration shows that each group contains a star, and (as neighboring groups are separated by at most 180°), there are at least $\lceil 360\, w(h)/(360 + 180) \rceil = \lceil 2\, w(h)/3 \rceil$ groups.

From the randomly chosen h, choose a group g uniformly at random from h's first $\lceil 2w(h)/3 \rceil$ groups.

For any given edge e, the probability that e is in g is at most $\sum_h \epsilon_h\, N^h_e / (2w(h)/3) = (3/2) X_e$. Thus, by linearity of expectation, the expected total length $E[|g|]$ of edges in g is at most $(3/2) \sum_{e \ni v} X_e|e|$. On the other hand, g contains a star, so $E[|g|] \ge S_{\min}(v)$. This proves Lemma 3.1.

4 LP (1.1) generalizes MWT heuristics

Fix any MWT instance $G = (V, E)$. It is known that any of the following conditions implies that a given potential edge e of G is in every MWT of G.

β-skeleton: For $\beta \approx 1.17682$, the two disks of diameter $\beta\,|e|$ having e as a chord are empty [20, 6].

YXY-subgraph: Every potential edge pq that crosses $e = xy$ has $|e| \le \min\{|px|, |py|, |qx|, |qy|\}$ [37, 17].

maximality: Every potential edge that crosses e is known to be *out of* every MWT (see e.g. [11]).

Similarly, any of the following conditions implies that a given potential edge e of G (not on the boundary of the region to be triangulated) is *out of* every MWT of G.

independence: Some potential edge that crosses e is known to be *in* every MWT e.g. [11].

diamond: Neither of the two triangles with base e and base angle $\pi/4.6$ are empty [8, 14].

LMT skeleton: For every two triangles t and t' for which e is *locally minimal*, one of the edges of t or t' is known to be *out of* every MWT [11].

In the LMT-skeleton condition, e is *locally minimal* for two triangles t and t' if $t \cap t' = e$ and t and t' together are a minimum-length triangulation of the quadrilateral $Q = t \cup t'$ — that is, either Q is non-convex, or e is at least as short as the other diagonal of Q.

Let E^* denote the set of edges that can be deduced to be in every MWT by applying the logical closure of the above six rules. (Logical closure is necessary because the maximality, independence, and LMT-skeleton conditions depend on the known statuses of edges other than e. For example, if one of the conditions implies that some edge e is out of every MWT, then the LMT-skeleton condition may then in turn imply that some new edge e' is out of every MWT, because e lies on one of two triangles t or t' for which e' is locally minimal.)

Ideally, the set E^* gives a partition of G in which every face is empty. If this happens, then the remaining edges in the MWT can be found by triangulating each remaining face independently using the standard dynamic-programming algorithm, and we say G is *solvable* via the heuristics. According to [11] (1997), most random instances with as many as 40,000 points are solvable via the heuristics.[5] Next we show that if an instance is solvable via the heuristics, then Linear Program (1.1) solves the instance also.

PROPOSITION 4.1. *For any instance G of MWT, let E^* be the partition of G defined above. If every face of E^* is empty, then every optimal extreme point of the LP (for G) is the incidence vector of a minimum-length triangulation.*

The remainder of the section gives the proof. The first step is to show that each condition above that ensures that an edge is in (or out of) every MWT also ensures that the LP gives the edge weight 1 (or 0) in any optimal fractional solution.

Say that LP (1.1) *forces a potential edge e to z* (where $z \in \{0, 1\}$) if, for every optimal fractional triangulation X^* of G, the weight that X^* gives to e is z.

LEMMA 4.1. *If any of the following conditions holds, the LP forces potential edge e of G to 1.*

1. *The β-skeleton condition above holds for e.*
2. *The YXY-subgraph condition above holds for e.*
3. (maximality) *The LP forces every potential edge that crosses e to 0.*

Proof idea. Part (3) is relatively straightforward: if X^* gives weight 0 to every edge that crosses e, then no triangle t that crosses e has positive X^*_t, so the only way X^* can cover points near e is with triangles that have e as a side.

The original β-skeleton and the YXY-subgraph heuristics are shown to be valid for MWT by local-improvement arguments: if the condition holds for an edge e that is *not* in the MWT, then a polygon P covering e within the MWT can be retriangulated at lesser cost, contradicting the optimality of the MWT [20, 6, 37, 17]. Here we extend those arguments to any optimal *fractional* triangulation X^*: if the condition holds and X^* does not give e fractional weight 1, then a polygon P' covering e whose triangles have

[5] [11] define the modified LMT-skeleton to be the set of edges that can be deduced to be in every MWT via (the logical closure of) just the diamond, LMT-skeleton, maximality, and independence conditions above. The use of logical closure is crucial to the effectiveness of the LMT skeleton.

positive weight in X^* can be retriangulated (lowering the weight of those triangles by $\epsilon > 0$ and raising the weight of other triangles by ϵ), giving a fractional triangulation that costs less than X^*.

The original arguments are intricate geometric case analyses, typically taking several pages. The arguments do not extend completely to this setting for the following reason: in the MWT setting, the polygon P identified for re-triangulation is the union of non-crossing triangles, whereas here, in the fractional setting, the polygon P' is the union of triangles that *can* cross (much as in Lemma 3.2). If the triangles in P' don't cross, then the original arguments apply, but in general additional analysis is needed. To illustrate, consider the β-skeleton. Suppose for contradiction that the β-skeleton condition holds for an edge $e = ab$ but it does not occur in the MWT. [20, 6] show that there must be a sequence $t_1, t_2, \ldots, t_k$ of empty triangles in the MWT whose union P covers e as shown to the left, below. Using the β-skeleton condition, they show that this union has a triangulation that costs less than does $t_1, \ldots, t_k$, contradicting the optimality of the MWT.

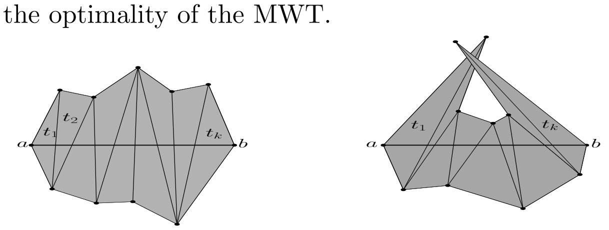

In the current context, if e has weight below 1 in X^*, then there must (similarly) exist a sequence $t_1, t_2, \ldots, t_k$ of empty triangles with positive weight in X^* covering e, but these triangles can cross as shown to the right above. We extend their arguments to show that, even if such crossing occurs, a triangulation of lower cost can still be found. ○○○

(For the full proof, see the full version of the paper.)

Lemma 4.2. *If any of the following conditions holds for a potential edge e of G (not on the boundary of the region to be triangulated), the LP forces e to 0.*

1. (independence) *The LP forces a potential edge that crosses e to 1.*
2. *The diamond condition above holds for e.*
3. (LMT skeleton) *For every two triangles t and t' for which e is locally minimal, the LP forces one of the edges of t or t' to 0.*

Proof idea. Part (1) is straightforward: if potential edges e and e' cross, then the LP covering constraint for a point near the intersection of e' and e implies that the total weight of potential triangles that have e or e' as sides is at most 1.

Part (3), the LMT skeleton, is straightforward. If an optimal fractional triangulation X^* gives e positive weight, then (by constraint (1.2) implied by the LP) there must be two triangles t and t' with positive X^*_t and $X^*_{t'}$ whose intersection is e. Edge e must be locally minimal for t and t' (otherwise X^* could be improved by reducing X^*_t and $X^*_{t'}$ and raising the weights of the other two triangles that triangulate $t \cup t'$).

Part (2), the diamond condition, is handled as the β-skeleton and YXY-subgraph are handled in the proof idea of Lemma 4.1. ○○○

(For the full proof, see the full version of the paper.)

Assume (as in the statement of Proposition 4.1) that the set E^* of edges that can be deduced to be in every MWT of G gives a partition of G in which every face is empty. It follows from Lemmas 4.1 and 4.2 (by a simple inductive proof) that every edge that can be deduced to be out of every MWT is forced to 0 by the LP, and every edge that can be deduced to be in every MWT is forced to 1. Thus, in any optimal fractional triangulation X^*, no potential triangle t that crosses an edge in E^* has positive weight X^*_t. Thus, the optimal fractional triangulations X^* are exactly those that, for each face f of the partition, induce an optimal fractional triangulation of the simple polygon f. It is known (e.g. [7, Thm. 7], [9, Thm. 4.1(i)], [21, Cor. 3.6.2]) that, for any simple polygon f, each optimal fractional triangulation is the incidence vector of an actual triangulation of f. Thus, each optimal extreme point of the LP for G is also the incidence vector of a triangulation of G, proving Proposition 4.1.

5 Remarks and open problems.

The constant factor proven here can be reduced in several places, at the expense of complicating the argument. The main challenge though is of course reducing λ.

Triangulations optimizing functions other than the total edge length are of interest. MWT LPs extend naturally to such problems, by modifying the cost function or by restricting the set of empty triangles. (For example, the integrality of the LP for the simple-polygon case implies that the simple-polygon result extends directly to such problems.) Can results similar to those in this paper be obtained for such problems?

The upper bound shown here on the integrality gap of the LP is constant but quite large. The only known lower bounds are barely above 1. We suspect that much better upper bounds can be shown, and that these should lead to an approximation algorithm with a better approximation ratio. We suspect that implicit in the analysis here is a primal-dual argument; making the dual solution explicit might be a step in this direction.

The LP studied here has an integrality gap above 1, so cannot be used directly to derive a PTAS. Applying sufficiently many rounds of lift-and-project to the LP will bring the integrality gap to $1 + \epsilon$. Are only $O(1)$ rounds required? Does this lead to a PTAS?

Does the LP generalize the heuristics in a stronger sense? Is there some condition, based on the optimal primal/dual pair, such that, if the condition holds for an edge e, then that edge must be, or cannot be, in any MWT?

References

[1] O. Aichholzer, F. Aurenhammer, S.W. Cheng, N. Katoh, G. Rote, M. Taschwer, and Y.F. Xu. Triangulations intersect nicely. *Discrete and Computational Geometry*, 16(4):339–359, 1996.

[2] E. Anagnostou and D. Corneil. Polynomial-time instances of the minimum weight triangulation problem. *Computational Geometry*, 3(5):247–259, 1993.

[3] F. Aurenhammer and Y. Xu. Optimal triangulations. In *Encyclopedia of Optimization, Second Edition*, volume 4, pages 160–166, 2000.

[4] M.G. Borgelt, C. Borgelt, C. Levcopoulos, and J.S.B. Mitchell. Fixed parameter algorithms for the minimum weight triangulation problem. *International Journal of Computational Geometry and Applications*, 18(3):185–220, 2008.

[5] P. Bose, L. Devroye, and W. Evans. Diamonds are not a minimum weight triangulation's best friend. *International Journal of Computational Geometry and Applications*, 12(6):445–453, 2002.

[6] S.W. Cheng and Y.F. Xu. Approaching the largest β-skeleton within a minimum weight triangulation. In *Proceedings of the twelfth annual symposium on Computational geometry*, pages 196–203, ACM, 1996. ACM.

[7] G.B. Dantzig, A.J. Hoffman, and T.C. Hu. Triangulations (tilings) and certain block triangular matrices. *Mathematical programming*, 31(1):1–14, 1985.

[8] G. Das and D. Joseph. Which triangulations approximate the complete graph? *Optimal Algorithms*, pages 168–192, 1989.

[9] J.A. de Loera, S. Hosten, F. Santos, and B. Sturmfels. The polytope of all triangulations of a point configuration. *Documenta Mathematica*, 1:103–119, 1996.

[10] J.A. De Loera, J. Rambau, and F. Santos. *Triangulations: Structures for algorithms and applications*, volume 25. Springer, 2010.

[11] M.T. Dickerson, J.M. Keil, and M.H. Montague. A large subgraph of the minimum weight triangulation. *Discrete and Computational Geometry*, 18(3):289–304, 1997.

[12] M.T. Dickerson, R.L. Scot Drysdale, S.A. McElfresh, and E. Welzl. Fast greedy triangulation algorithms. *Computational Geometry*, 8(2):67–86, 1997.

[13] D.P. Dobkin, H. Edelsbrunner, and M.H. Overmars. Searching for empty convex polygons. *Algorithmica*, 5(1):561–571, 1990.

[14] R.L. Drysdale, S. McElfresh, and J.S. Snoeyink. On exclusion regions for optimal triangulations. *Discrete Applied Mathematics*, 109(1-2):49–65, 2001.

[15] M.R. Garey and D.S. Johnson. *Computers and intractability*, volume 174. Freeman San Francisco, CA, 1979.

[16] P. Giannopoulos, C. Knauer, and S. Whitesides. Parameterized Complexity of Geometric Problems. *The Computer Journal*, 51(3):372–384, November 2007.

[17] P.D. Gilbert. New results on planar triangulations, 1979. Masters thesis.

[18] M.J. Golin. Limit theorems for minimum-weight triangulations, other Euclidean functionals, and probabilistic recurrence relations. In *Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 252–260, Society for Industrial and Applied Mathematics, 1996. Society for Industrial and Applied Mathematics.

[19] M. Hoffmann and Y. Okamoto. The minimum weight triangulation problem with few inner points. *Computational Geometry*, 34(3):149–158, 2006.

[20] J.M. Keil. Computing a subgraph of the minimum weight triangulation. *Computational Geom-*

etry, 4(1):18–26, 1994.

[21] D. Kirsanov. *Minimal discrete curves and surfaces*. PhD thesis, Harvard, 2004.

[22] G.T. Klincsek. Minimal triangulations of polygonal domains. *Combinatorics*, 79:121–123, 1980.

[23] C. Knauer and A. Spillner. A fixed-parameter algorithm for the minimum weight triangulation problem based on small graph separators. In *Graph-Theoretic Concepts in Computer Science*, pages 49–57, Springer, 2006. Springer.

[24] D. Krznaric and C. Levcopoulos. Quasi-greedy triangulations approximating the minimum weight triangulation. *J. of Algorithms*, 27(2):303–338, 1998.

[25] Y. Kyoda. A study of generating minimum weight triangulation within practical time. Master's thesis, Dept. Inform. Sci., Univ. Tokyo, Tokyo, Japan, 1996.

[26] Y. Kyoda, K. Imai, F. Takeuchi, and A. Tajima. A branch-and-cut approach for minimum weight triangulation. In *Proceedings of Algorithms & Computations, 8th International Symposium ISAAC*, pages 384–393, 1997.

[27] E.L. Lloyd. On triangulations of a set of points in the plane. *IEEE Symposium on Foundations of Computer Science*, pages 228–240, 1977.

[28] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *Journal of the ACM (JACM)*, 55(2):11, 2008.

[29] T. Ono, Y. Kyoda, T. Masada, K. Hayase, T. Shibuya, M. Nakade, M. Inaba, H. Imai, K. Imai, and D. Avis. A package for triangulations. In *Proceedings of the twelfth annual symposium on Computational geometry*, pages 517–518, ACM, 1996. ACM.

[30] D.A. Plaisted and J. Hong. A heuristic triangulation algorithm. *Journal of Algorithms*, 8(3):405–437, 1987.

[31] J. Remy and A. Steger. A quasi-polynomial time approximation scheme for minimum weight triangulation. *Journal of the ACM (JACM)*, 56(3):15, 2009.

[32] W. D. Smith. *Studies in computational geometry motivated by mesh generation*. PhD thesis, Princeton Univ., Princeton, NJ, 1988.

[33] A. Spillner. A faster algorithm for the minimum weight triangulation problem with few inner points. In *Proc. 1st Workshop Algorithms and Complexity in Durham (ACiD 2005, Durham, UK)*, pages 135–146, 2005.

[34] A. Tajima. Optimality and integer programming formulations of triangulations in general dimension. *Algorithms and Computation*, pages 378–387, 1998.

[35] F. Takeuchi, H. Imai, and K. Imai. Polytopes of linear programming relaxation for triangulations. *Kyoto University Research Information Repository*, 1068:121–133, 1998.

[36] C.A. Wang, F. Chin, and Y.F. Xu. A new subgraph of minimum weight triangulations. *Journal of Combinatorial Optimization*, 1(2):115–127, 1997.

[37] B. Yang, Y. Xu, and Z. You. A chain decomposition algorithm for the proof of a property on minimum weight triangulations. In *Algorithms and Computation*, volume 834 of *Lecture Notes in Computer Science*, pages 423–427. Springer Berlin / Heidelberg, 1994.