

UC Riverside

UC Riverside Previously Published Works

Title

On a Linear Program for Minimum-Weight Triangulation

Permalink

<https://escholarship.org/uc/item/4f88b6fh>

Journal

SIAM Journal on Computing, 43(1)

ISSN

0097-5397

Authors

Yousefi, Arman

Young, Neal E

Publication Date

2014

DOI

10.1137/120887928

Peer reviewed

ON A LINEAR PROGRAM FOR MINIMUM-WEIGHT TRIANGULATION*

ARMAN YOUSEFI[†] AND 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, Hoffman, and Hu [*Math. Programming*, 31 (1985), pp. 1–14], (i) the integrality gap is constant, and (ii) given any instance, if the aforementioned heuristics find the MWT, then so does the LP.

Key words. linear programming, minimum-weight triangulation, approximation algorithm, integrality gap

AMS subject classifications. 68W25, 68U05, 90C59, 68Q25, 90C10, 90C35, 68W40, 90C05, 90C27, 90C90, 90C08

DOI. 10.1137/120887928

1. Introduction. In 1979, Garey and Johnson [16] listed minimum-weight triangulation (MWT) as one of a dozen important problems known to be neither in P nor NP-hard. In 2006 the problem was finally shown to be NP-hard [29]. The problem has a subexponential time exact algorithm [33], as well as a polynomial-time approximation scheme (PTAS) for random inputs [19]. Currently it is not known whether, for some constant $c > 1$, finding a c -approximation is NP-hard, but this is unlikely, as a quasi-polynomial-time approximation scheme exists [32]. MWT has an $O(\log n)$ -approximation algorithm [31] and, most important here, an $O(1)$ -approximation algorithm called QUASIGREEDY [25]. The constant in the big-O upper bound from [25] is astronomically large.

If restricted to simple polygons, MWT has a well-known $O(n^3)$ -time dynamic-programming algorithm [18, 23]. 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 region R to be triangulated [17, sect. 2.5.1]; also see [20, 5, 34, 24].

Linear program (LP) of Dantzig, Hoffman, and Hu for MWT. Linear-programming methods are a primary paradigm 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 relaxations. Studying their

*Received by the editors August 13, 2012; accepted for publication (in revised form) September 25, 2013; published electronically January 21, 2014. An extended abstract of this paper appeared in the proceedings of SODA 2012 [39].

<http://www.siam.org/journals/sicomp/43-1/88792.html>

[†]Computer Science Department, University of California, Los Angeles, Los Angeles, CA 90095 (yousefia@cs.ucla.edu). The research of this author was partially funded by GAANN fellowship.

[‡]Department of Computer Science and Engineering, University of California, Riverside, Riverside, CA 92521 (neal.young@ucr.edu). The research of this author was partially funded by NSF grants 0729071 and 1117954.

integrality gaps may lead to better approximation algorithms, or may widen our understanding of general methods and their limitations (as standard randomized rounding and primal-dual approaches may be insufficient for MWT).

Dantzig, Hoffman, and Hu [8] introduce the following LP (presented here as reformulated by [10]). Below, Δ denotes the set of empty triangles.¹ R denotes the region to be triangulated, minus the sides of triangles in Δ . The LP asks to assign a nonnegative weight X_t to each triangle $t \in \Delta$ so that, for each point p in the region, the triangles containing it are assigned total weight 1:

$$(1.1) \quad \text{minimize } c(X) = \sum_{t \in \Delta} c(t)X_t \text{ subject to } X \in \mathbb{R}_{\geq 0}^{\Delta} \text{ and } (\forall p \in R) \sum_{t \ni p} X_t = 1.$$

Above, $\mathbb{R}_{\geq 0}^{\Delta}$ is the set of vectors of nonnegative reals, indexed by the triangles in Δ . The cost $c(t)$ of triangle t is the sum over the edges e in t of the cost $c(e)$ of the edge, defined to be $|e|/2$ (the length of e), unless e is on the boundary of R , in which case the cost is $|e|$. (Internal edges are discounted by 1/2 since any internal edge occurs in either zero or two triangles in any triangulation.) R 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 , each on one side of e and 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, Hoffman, and Hu [8, Thm. 7], then (apparently independently) by De Loera et al. [10, Thm. 4.1(i)] and Kirsanov [22, Cor. 3.6.2]. For summaries of these results, see [11, Chap. 8] and [36]. 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 heuristics described later.

Other authors have considered *edge-based* LPs, mainly for use in branch-and-bound [26, 27, 30, 35, 3]. These edge-based LPs have unbounded integrality gaps.

LPs for maximal independent sets, which are well-studied, are closely related to all of 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 whose edge set is a subset of a given set S [28]. For a given set S , if we change the cost function in the above LP to $c(X) = \sum_{e \notin S} \sum_{t \ni e} X_t$, the LP will have a zero-cost integer solution iff there is such a triangulation. Unless $P = NP$, this implies that the LP with that cost function has unbounded integrality gap.² Thus, any bound on the integrality gap of the LP with the MWT cost function must rely intrinsically on that cost function. Similarly, given an arbitrary fractional solution X , it is NP-hard

¹That 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.

²If the LP has bounded integrality gap, it has a zero-cost fractional solution iff it has a zero-cost integer solution.

to determine whether there is an integer solution in the support of X .³ These are obstacles to standard randomized-rounding methods.

First new result: Integrality gap is constant. We show that LP (1.1) has constant integrality gap. This is the first nontrivial upper bound on the integrality gap of any MWT LP. To show it, we revisit the analysis of QUASIGREEDY [25], which shows that QUASIGREEDY produces a triangulation of cost $O(|\text{MWT}(G)|)$, where $|\text{MWT}(G)|$ is the length of the MWT of the given instance G (and also the cost of the optimal integer solution to the LP). We generalize their arguments to show that there exists a triangulation of cost $O(c(X^*))$, where $c(X^*)$ is the cost of the optimal *fractional* solution to the LP.

Our analysis also reduces the approximation ratio in their analysis by an order of magnitude, but the approximation ratio 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 excluded from) any MWT. Here is a summary. Gilbert [18] observes that the *shortest potential edge* is in every MWT. Yang, Xu, and You [38] extend this result by proving that an edge xy is in every MWT if, for any edge pq that intersects xy , $|xy| \leq \min\{|px|, |py|, |qx|, |qy|\}$. (We refer to the edges satisfying this property as the *YXY-subgraph*.) This subgraph includes every edge connecting two *mutual nearest neighbors*. Keil [21] defines another heuristic called β -skeleton as follows. An edge pq is in the β -skeleton iff there does not exist a point z in the point set such that $\angle pqz \geq \arcsin(1/\beta)$. Thus, an edge pq is in the β -skeleton iff the interior of the two circles of diameter $\beta \cdot |pq|$ passing through p and q do not contain any points. Keil [21] then shows that for $\beta \geq \sqrt{2}$, an edge that is in β -skeleton is in every MWT. Cheng and Xu [7] strengthen this to $\beta \geq 1/\sin k$, where $k \approx \pi/3.1$. Das and Joseph [9] 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. Drysdale, McElfresh, and Snoeyink [15] strengthen this to angle $\pi/4.6$. This property of e is called the *diamond property*. Dickerson, Keil, and Montague [12] 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.

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 span 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 [12], most random instances with 40,000 points are solvable in this way.

Second new result: The LP generalizes heuristics. We show that LP (1.1) generalizes these heuristics in that *if the heuristics solve a given instance as described above,*

³Given an arbitrary subset S of the edges, the problem of determining whether S contains a triangulation reduces to the problem of determining whether there is an integer solution in the support of a given fractional solution to the LP, as follows. Let set S' consist of the empty triangles whose edges are in S , so S contains a triangulation (by edges) iff S' contains a triangulation (by triangles). For each triangle $t \in S'$, solve the LP with the cost function that gives t cost zero, every other triangle in S' cost one, and all triangles not in S' cost infinity. If (for any $t \in S'$) the LP for t has no finite-cost feasible solution, then S' contains no triangulation. Otherwise, for each $t \in S'$, let X^t denote an optimal fractional solution to the LP for t . Let $\tilde{X} = \sum_{t \in S'} X^t / |S'|$ be the average of these fractional solutions. Because of the choice of the cost function, if a given X^t does not give positive weight to t , then no (integer) triangulation in S' contains t . Thus, S' contains a triangulation iff there is a triangulation in the support of \tilde{X} .

then so does the LP (that is, the extreme points of the LP are integer solutions—incidence vectors of optimal triangulations). In this sense, the LP, whose formulation requires little explicit geometry, generalizes all of these varied and generally incomparable heuristics. (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 [22, sect. 3.5].⁴) This is the first connection we know of between the heuristics and any MWT LP.

Roughly speaking, 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 an 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, Chin, and Xu [37] 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., [10, Thm. 1.1(i), Prop. 2.5], [36], or [22, Thm. 3.4.1]) and are useful for reasoning about fractional triangulations. For any fractional triangulation X and edge e ,

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

Here $\text{left}(e)$ contains the triangles that contain e and lie on one side of e , while $\text{right}(e)$ contains the triangles that contain e and lie on the other side of e . (If e is on the boundary, take $\text{right}(e) = \emptyset$.) 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., nonzeros 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 [14]. 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 excluded from every MWT, and add a constraint $\sum_{t \in \text{left}(e)} X_t = \sum_{t \in \text{right}(e)} X_t = 1$ if they prove that an interior edge e is 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 [12, 13]. In our ad hoc experiments on “typical” instances with 10^4 to 10^5 points, only a small number of variables were left undetermined by the heuristics. This allowed us to use standard linear-programming solvers to quickly solve the remaining LP. (This is in keeping with the experiments of [12], which found that most random instances on 40,000 points were solvable by heuristics.) Similarly, this preprocessing should help integer linear-programming solvers to quickly find the MWT (for instances for which the optimal solution is fractional). It is known that, asymptotically, for n random points, the expected number of remaining variables is $\Omega(n)$, but the leading constant is apparently astronomically small [6].

Remarks. The results here suggest that the LP of [8] captures much of the structure of MWT. This suggests a line of attack for improving the approximation ratio:

⁴Where G contains the center of a unit circle and $n - 1$ random points on the circle.

use systematic linear-programming methods such as randomized rounding, the primal-dual method, and lift-and-project [4] to study the integrality gap of the LP. Success would yield a better approximation (conceivably, even a PTAS, using lift-and-project). Failure would increase our understanding of the limitations of these techniques.

Implicit in our bound on the integrality gap is a polynomial-time algorithm with matching approximation ratio. Actually, there are two. Both algorithms first compute Levcopoulos and Krznaric's [25] convex partition LK of the point set (see our Lemma 2.15), then extend LK by triangulating each face f of LK. The triangulation of each f can be done either (a) using the standard dynamic program to find an MWT of f , or (b) as follows: compute the fractional solution X to the LP, then, for each face f of LK, *transpose* X into a fractional triangulation X^f of f (as described in section 2.1), and then use the cheapest triangulation of f implicit in X^f .

That the first algorithm above is an $O(1)$ -approximation algorithm follows from Levcopoulos and Krznaric's previous work [25]. However, the bound we show here— $54(\lambda + 1)$, where λ is a large constant per Lemma 2.15—is substantially smaller than their previous bound. Roughly speaking, we obtain a better bound by analyzing the transposal operation at the level of *triangles* instead of edges.

Open problems. The integrality gap is constant, but there is still a huge gap between the best lower bound known (barely above 1.0) and the upper bound shown here (astronomically large). The next step in improving our upper bound would be to reduce the value of λ in Lemma 2.15. We suspect that a primal-dual analysis is implicit in the analysis here; making the dual solution explicit might be a step in this direction.

Many different cost functions (other than the total edge length) for triangulations are studied in the literature. The MWT LP extends naturally by modifying the cost function or restricting the set of allowed triangles. (For example, the integrality of the extreme points of the LP for the simple-polygon case implies that the simple-polygon result generalizes to any linear cost function.) We conjecture that results similar to those in this paper can be obtained for other cost functions.

If MWT heuristics can solve a given instance of MWT, then so can the LP. However, the heuristics are also useful for instances that they do not completely solve: in such instances, the heuristics can still identify some edges that are in (or excluded from) every MWT, even if these do not completely determine the triangulation. Can some analogous property be shown for the LP? That is, is there some condition (e.g., based on the optimal primal-dual solution to the LP) such that, if the condition holds for an edge e , that edge must be in (or excluded from) every MWT?

DEFINITION 1.1. *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 nonempty. The (Euclidean) length of line segment pq is denoted $|pq|$. For any set E of segments, $\|E\|_2$ denotes 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.⁵ A face or polygon is empty if its interior contains no vertex.

⁵Where 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.

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 (noncrossing) 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 $\text{MCP}(G)$. A triangulation of G is a partition whose faces are empty triangles. A fractional triangulation X is a feasible solution to the LP. For any potential edge e , the weight of edge e in X , denoted X_e , is $\sum_{t \ni e} X_t$ if e is on the boundary of the region to be triangulated, and otherwise half this amount.

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 particular graph by its edge set.

2. Integrality gap is constant. This section proves our first new result.

THEOREM 2.1. *Given any instance $G = (V, E)$ of MWT, for any fractional triangulation X , there exists an integer solution of value $O(c(X))$. That is, LP (1.1) has constant integrality gap.*

Proof. Fix the MWT instance G and an arbitrary fractional triangulation X . Fix a convex partition CP of G . (Later, we will fix CP to be a particular convex partition LK with some particular properties.)

The idea of the proof is to define a “rounding” procedure that converts X into the desired integer solution. The procedure fractures X into a separate fractional triangulation X^f for each face f of CP (where X^f covers exactly f). Then, independently within each face f of CP , 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 CP), of total cost at most $\sum_{f \in \text{CP}} c(X^f)$ (and, hopefully, $O(c(X))$).

In the second step, since each f is a simple polygon, it follows from known results (see, e.g., [8, Thm. 7]; also see the introduction) that the cost of the optimal integer triangulation of f is at most the cost of X^f . Thus, the integrality gap will be $O(1)$ as long as the first step triangulates the faces so that $\sum_f c(X^f) = O(c(X))$.

The proof is divided into two parts: (i) describing a correct rounding procedure that fractures X into a fractional triangulation X^f for each face f of CP (we call this *transposing* X into f), and (ii) bounding the cost $\sum_f c(X^f)$ by $O(c(X))$.

2.1. Part (i)—fracturing X into the faces of CP . Fix any face f of CP of the convex partition CP . Our goal is to convert X into a fractional triangulation X^f of f .

We start with the observation that X , restricted to triangles that cross f , can be separated into independent layers, where each layer is a set of triangles that uniformly covers f (and possibly some points outside f). We say such a layer *blankets* f .

DEFINITION 2.2 (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 polygons overlap within f (they may overlap outside f).*

In this subsection, the polygons in blankets are always triangles.

The next lemma describes how to decompose X (over f) into blankets.

LEMMA 2.3. *There exists a set \mathcal{B} of blankets (each containing only triangles)*




FIG. 1.

and weights $\epsilon_B > 0$ 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 that “[$t \in B$]” is 1 if t is in B , else 0.

Proof. Recall that, for MWT instances consisting of a simple polygon, the LP gives optimal 0/1 solutions (see, e.g., [8, Thm. 7]). We adapt a proof of that property.

Choose any triangle t that crosses f and has $X_t > 0$. If t completely covers f , then stop and take $B = \{t\}$. Otherwise, some edge e of triangle t crosses the interior of f . Since e has positive weight, there must be a positive-weight triangle t' that has e as an edge and lies on e 's opposite side (this is implied by constraint (1.2)). Glue t and t' 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 boundary edge that crosses f . The triangles glued together in this way form the blanket B .

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.) \square

Fix the set \mathcal{B} of blankets of f from Lemma 2.3 and the corresponding weights ϵ_B .

We next describe how to convert any single blanket $B \in \mathcal{B}$ into a true triangulation \widehat{B}^f of f . The final fractional triangulation X^f will be the convex combination of these triangulations, where the triangulation of B is given weight ϵ_B .

Recall that any blanket $B \in \mathcal{B}$ consists of triangles that together uniformly cover the convex face f (and may extend outside of f). To define the triangulation \widehat{B}^f , we start with *edge transposals* (see, e.g., Lemma 4.2 of [25]). For any edge e that crosses f , *transposing* e in f slides e to its *transposal*, denoted e^f , a diagonal of f that has minimum length among four or fewer diagonals that are “near” e . We give the formal definition next, and then extend that to define transposals of triangles and blankets $B \in \mathcal{B}$.

DEFINITION 2.4 (transposing an edge [25]; see Figure 1). *Fix any triangle edge $e = x_1x_2$ that crosses f (that is, that intersects the boundary of f in two points or along an edge). The transposal of e in f , denoted e^f , is defined by the following operation: Clip the edge x_1x_2 to chord $x'_1x'_2 = (x_1x_2) \cap f$ of f . For each endpoint x'_i of $x'_1x'_2$, if the endpoint lies in the interior of an edge e of f (as opposed to being a vertex of f), then slide x'_i along e to one of the endpoints of e , called the destination of x_i . Otherwise (the endpoint is a vertex of f), take that vertex as the destination. Choose the destinations (for those where there is a choice) to minimize the length of the diagonal that connects the destinations. (Break ties consistently.) The resulting diagonal is e^f .*




FIG. 2.

Next we define what it means to transpose *triangles* and *blankets*. We give a somewhat uninformative formal definition, then describe the important properties.

DEFINITION 2.5. For any triangle t , the transposal of t in f , denoted t^f , is the convex hull of the endpoints of the transposals of the edges of t that cross f .

For any blanket $B \in \mathcal{B}$, 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, $B^f = \{t^f \mid t \in B\}$.

Consider a blanket $B \in \mathcal{B}$. By definition, the edges of triangles in B do not cross within f . But, a priori, their transposals might. We next argue that this is not the case. In fact, we prove more: roughly speaking, that transposing preserves the topology of the partition that B induces on f . More precisely, consider that partition, which comes from clipping the edges of the triangles in B into f as shown in Figure 2. Consider any edge YZ . Focus on just those chords that have an endpoint in the interior of YZ . Order these chords, as shown in the left picture, according to the order of their endpoints on YZ going from Y to Z . For chords sharing an endpoint on YZ , break ties in favor of chords that lean closer to Y .

LEMMA 2.6 (transposing preserves order). *In the above ordering of chords along YZ , all chords whose endpoints have transposal destination Y precede all chords whose endpoints have destination Z . (Informally, when transposing the edges, when we slide the endpoints to their destinations, the endpoints that slide to Y precede the endpoints that slide to Z , so no crossings are introduced.)*

Proof. Without loss of generality, assume that Z lies (one vertex) clockwise of Y . Focus on the chords xw , where x is in the interior of YZ . Let C_Y contain those chords whose endpoint x has destination Y . Let C_Z contain those whose endpoint x has destination Z . We show that, if we leave Y and travel *counterclockwise* around the boundary to Z , we encounter the chords in C_Y before we encounter the chords in C_Z . This proves the claim, because, as chords in C do not cross, as we travel counterclockwise we must encounter chords in the same order that we would if traveling clockwise from Y to Z .

Consider the perpendicular bisector of YZ . Since f is convex, the bisector intersects the boundary of f at a single point across from YZ . Suppose that this intersection point is in the interior of some edge b as shown in Figure 3. (If the intersection is a vertex of f , take b to be that point and follow similar reasoning.)

As we travel counterclockwise from Y to Z , until we pass the first endpoint of b , every chord endpoint w that we encounter is in some edge e of f that lies entirely on the Y -side of the bisector. Since both endpoints of e are closer to Y than to Z , no matter which endpoint of e is the destination of w , the destination of x will definitely be Y . Thus, until we pass the first endpoint of b , we encounter only chords in C_Y .

As we travel through the interior of edge b (or, if b is a point, through b) for all chord endpoints w that we encounter, their chords xw will have the *same* transposal (since x is in the interior of YZ and w is in the interior of b , and transposing breaks ties consistently). Thus, traveling through b , either we encounter only chords in C_Y , or we encounter only chords in C_Z .




FIG. 3.




FIG. 4.

Once we pass the other endpoint of b , until we reach Z , every chord that we encounter is an edge of f that lies entirely on the Z -side of the bisector, so, reasoning as before, we encounter only chords in C_Z . \square

Because transposing preserves order, the topological structure of the *transposal* of any blanket B is inherited from the partition that B induces on f . See the example in Figure 4.

Figure 4 shows five copies of a face f (with gray background). Copy 1 shows the face blanketed by six triangles. In copy 2, the triangle edges are clipped to their chords in the face, giving the partition that B induces on f . In copies 3 through 5, each chord is shifted to its edge transposal by sliding each endpoint to its destination. Copy 5 shows the resulting edge transposals and the transposal of B in f .

Clearly, in the partition that B induces on f (copy 2) each region is of the form $t \cap f$ for some $t \in B$. Because transposing preserves order, moving the edges of that partition to their transposals preserves the topological structure of the partition: the transposal B^f of B (copy 5) is a convex partition of f whose edges are the transposals of the edges of B , and whose regions are the transposals of the triangles in B . Also, for each triangle $t \in B$, the boundary of its transposal t^f consists of the transposals of the edges of t , together with up to three edges of f .

The transposal of a blanket is a convex partition of the face, but not quite a triangulation, because each of its regions may have up to six sides. To get a triangulation, we simply triangulate each of its regions.

DEFINITION 2.7. *The triangulated transposal of a triangle t in f , denoted \widehat{t}^f , is the MWT of the transposal t^f , except that, if t^f has no area, then t^f is the empty set. The triangulated transposal of a blanket B in f , denoted \widehat{B}^f , is the union of the triangulated transposals of the triangles in the blanket.*

The transposal B^f is a convex partition of f whose regions are the transposals of the triangles in B , so the triangulated transposal of B indeed triangulates f .

Finally, we define the fractional triangulation X^f of f . We start with the fractional triangulation X . We restrict X to triangles crossing f . We decompose this restriction of X into a convex combination of blankets of f (per Lemma 2.3). Then, in this convex combination, we replace each blanket B by its triangulated transposal \widehat{B}^f , a triangulation of f .⁶ Here is the formal definition.

DEFINITION 2.8. *Define the transposal of X in f , denoted X^f , to be the fractional triangulation of f formed by the convex combination of the transposals of the blankets in \mathcal{B} so that*

$$X_t^f = \sum_{B \in \mathcal{B}} [t \in \widehat{B}^f] \epsilon_B.$$

We now complete Part (i) of the proof of Theorem 2.1.

LEMMA 2.9. *Fix any fractional triangulation X and any convex face f . The transposal X^f of X in f defined above is a feasible fractional triangulation of f . That is, it covers the points in f uniformly with weight 1.*

Proof. As discussed, this holds because X^f is a convex combination of triangulations of f . Indeed, it covers each point p in f with total weight

$$\sum_t [p \in t] X_t^f = \sum_t [p \in t] \sum_{B \in \mathcal{B}} [t \in \widehat{B}^f] \epsilon_B = \sum_{B \in \mathcal{B}} \epsilon_B \sum_{t \in \widehat{B}^f} [p \in t] = \sum_{B \in \mathcal{B}} \epsilon_B = 1.$$

The first equality is by definition of X^f . The second just exchanges the order of summation. The third holds because \widehat{B}^f triangulates f (so exactly one $t \in \widehat{B}^f$ contains p). The last follows by Lemma 2.3. \square

2.2. Part (ii)—bounding the cost. Fix the convex partition CP and fractional solution X . By Lemma 2.9, for each face f of CP, the transposal X^f as defined in Part (i) is a fractional triangulation of f . To complete the proof of Theorem 2.1, we bound the sum of the costs of these fractional triangulations.

We start by observing that we can view X^f as taking the weight of each triangle t in X and transferring that weight to (every triangle in) the triangulated transposal \widehat{t}^f of t in f .

FACT 2.9.1.

$$X_{t'}^f = \sum_{t: t' \in \widehat{t}^f} X_t.$$

Proof.

$$X_{t'}^f = \sum_{B \in \mathcal{B}} [t' \in \widehat{B}^f] \epsilon_B = \sum_{B \in \mathcal{B}} \sum_{t \in B} [t' \in \widehat{t}^f] \epsilon_B = \sum_{t: t' \in \widehat{t}^f} \sum_{B \in \mathcal{B}} [t' \in B] \epsilon_B = \sum_{t: t' \in \widehat{t}^f} X_t.$$

The first equality is the definition of X^f . The second holds by definition of \widehat{B}^f (namely, $t' \in \widehat{B}^f$ iff $t' \in \widehat{t}^f$ for some $t \in B$). The third just exchanges the order of summation. The last follows from Lemma 2.3. \square

So far, we have considered how a blanket of triangles transposes into a single f . Next we consider how a single triangle t transposes across multiple faces. Of course, a given triangle t can cross many faces, but *in all but two* its transposal will have no area (and thus will play no part in the triangulated transposal of X in f).

⁶Alternatively, we could take X^f to be the *cheapest* triangulation \widehat{B}^f over all blankets $B \in \mathcal{B}$.




FIG. 5.

LEMMA 2.10. *Any given triangle t crosses at most two faces f in CP in which its transposal t^f has positive area. Thus, for a given t , only two faces f have $c(\widehat{t^f}) > 0$.*

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

In the two examples in Figure 5, each dashed edge is an 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* . (A face “touches” an edge if the intersection of the face and the edge, including boundaries and endpoints, is nonempty. For example, the accommodating face 2 on the left of the figure, and 2 and 5 on the right, touch all three edges of t . Each other face is nonaccommodating 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 , the third edge of t lies outside of f , so the two edges cross the interior of a single edge of f . Thus, the two edges of t that touch f must have identical transposals, 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 Figure 5, left). Since the faces are nonoverlapping and convex, 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.

Assume that no face touches the interiors 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 XY , of t whose interior f avoids. Thus, f must touch XY at an endpoint, say, Y . (For example, consider Figure 5, right. Faces 2, 3, 4, and 5 touch all three edges of t , but not all three interiors.) Since f touches XY at Y but does not touch the interior of XY , there must be an edge wY of f that extends through the interior of t . Since w is not inside t , wY must cut across t to the interior of the edge XZ . Thus,

- (2.1) *in this case, 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 Y , for otherwise the extending edges would cross inside t .

Now consider all edges in CP that extend from Y across the interior of t . Let these edges be w_1Y, w_2Y, \dots, w_kY , rotating in order around Y . (In Figure 5, $k = 3$.) CP has $k + 1$ corresponding faces f_0, f_1, \dots, f_k , also in order rotating around Y , where f_{i-1} and f_i share edge w_iY . By conclusion (2.1), only these $k + 1$ faces might be accommodating.

To finish, we observe that f_i is not accommodating unless $i \in \{0, k\}$ (the first or last face). Indeed, for $i \notin \{0, k\}$ edges $w_{i-1}Y$ and w_iY of f_i extend from Y across t




FIG. 6.

to XZ . Since these edges touch at Y , the transposal of t in f_i is thus *just the point* Y . Thus, the transposal of t in f_i has no area. \square

Our goal is to show that transposing X across the faces increases the cost of X by at most a constant factor. For any triangle t , by the lemma and Fact 2.9.1, transposing X transfers the weight X_t to the triangulated transposals \widehat{t}^f of t in *at most two* faces f .

To proceed we bound the cost of each \widehat{t}^f in terms of the cost of t . Recall that \widehat{t}^f is the MWT of its (nontriangulated) transposal t^f , which has at most six sides (up to three edges of f and up to three transposals of edges of t). We start by bounding the cost of t^f . Our bound depends on the *sensitivity* of the edges of the convex partition CP, defined as follows.

DEFINITION 2.11 (sensitivity). *An edge e is σ -sensitive if, for any potential edge e' that crosses e , for each endpoint x of e' , the distance from x to the closest endpoint of e is at most $\sigma|e'|$.*

In other words, the circle of radius $\sigma|e'|$ around each endpoint of e' contains an endpoint of e .

For the rest of this section, fix σ such that all edges of CP are σ -sensitive.

LEMMA 2.12. *For any face f of CP and triangle t , the total length of the edges in t 's transposal t^f that are not also edges of CP is at most 2σ times the length of t 's edges.*

Proof. Let f be any face of CP 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 2σ times the length of e* . This claim implies the lemma, because each edge of the transposal of t (but not of CP) is the edge transposal e^f of a unique edge e of t . To finish, we prove the claim.

For an edge e that crosses f , one of the following three cases holds: (1) e is incident to one vertex of f and properly intersects one s side of f (as in the left of Figure 6), (2) e properly intersects two sides s and s' of f (as in the right of Figure 6), or (3) e is incident to two vertices of f .

In case (1), let W be the vertex that f shares with e (and e^f). Since s is σ -sensitive, and e crosses s , the endpoint W of e is at most $\sigma|e|$ from some endpoint of s . Since $|e^f|$ is the minimum distance between W and any endpoint of s , this implies $|e^f| \leq \sigma|e|$.

In case (2), let Y be an endpoint of e and let Z and Z' , respectively, be the closest endpoints of s and s' to Y . Because e^f is the shortest segment from an endpoint of s to an endpoint of s' , $|e^f| \leq |ZZ'|$. By the triangle inequality, $|ZZ'| \leq |YZ| + |YZ'|$. Because s and s' are σ -sensitive, $|YZ|$ and $|YZ'|$ are each at most $\sigma|e|$.

In case (3), the transposal e^f of e is the same as e , so the claim holds. \square

It is straightforward to extend the bound to the *triangulated* transposal \widehat{t}^f of t . Recall that the cost of a triangle t is the sum of the costs of its edges, where the cost of an edge is half its length, unless the edge is on the boundary of the entire region, in which case the cost of the edge is its length. The cost $c(\widehat{t}^f)$ of a triangulation \widehat{t}^f is the sum of the costs of the triangles in the triangulation.

LEMMA 2.13. *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 (nontriangulated) transposal of t in f .*

Proof. Recall that t^f has at most six vertices, say, v_1, v_2, \dots, v_k , ordered clockwise. Triangulate t^f by adding up to three interior diagonals connecting the odd vertices (e.g., v_1v_3, v_3v_5, v_5v_1). The total length of the added diagonals is at most the total length of the boundary. Likewise, the sum of costs of the added diagonals is at most the sum of the costs of the edges on the boundary of t^f .

Each added edge occurs in two triangles in this triangulation, whereas each boundary edge occurs in just one triangle. Thus, adding the diagonals gives a triangulation of cost at most three times the cost of t^f . The lemma follows, as $c(\widehat{t}^f)$ is the minimum cost of any triangulation of t^f . \square

Next we gather the bounds in the previous lemmas to bound the total cost across all the faces. We are not finished, as the bound depends on not only the cost of the fractional triangulation, but also the total length of the edges in the convex partition CP and the sensitivity σ of those edges.

LEMMA 2.14. *The total cost $\sum_f c(X^f)$ is at most $3\|\text{CP}\|_2 + 12\sigma c(X)$.*

Proof. The total cost is

$$\begin{aligned} \sum_f c(X^f) &= \sum_{f,t} X_t c(\widehat{t}^f) && \text{by Fact 2.9.1,} \\ &\leq 3 \sum_{f,t} [c(\widehat{t}^f) > 0] X_t c(t^f) && \text{by Lemma 2.13,} \\ &\leq 3\|\text{CP}\|_2 + 6\sigma \sum_{t,f} [c(\widehat{t}^f) > 0] X_t c(t) && \text{by Lemma 2.12,} \\ &\leq 3\|\text{CP}\|_2 + 12\sigma \sum_t X_t c(t) && \text{by Lemma 2.10,} \\ &= 3\|\text{CP}\|_2 + 12\sigma c(X) && \text{by definition of } c(X). \quad \square \end{aligned}$$

To proceed further, we need a convex partition whose edges have constant sensitivity and total length $O(c(X))$. Levkopoulos and Krzrnaric have shown the existence of something close: a convex partition LK whose edges are 4.45-sensitive and have total length $O(\text{MCP}(G))$ (recall that $\text{MCP}(G)$ is the minimum-length convex partition of G).

LEMMA 2.15 (see [25]). *For some constant $\lambda > 0$, for any MWT instance G , there exists a convex partition LK of G , whose edges are 4.45-sensitive, having total length $\|\text{LK}\|_2 \leq \lambda \|\text{MCP}(G)\|_2$.*

Proof. Levkopoulos and Krzrnaric show that what they call the *quasi-greedy convex partition* has the above properties: for the first property, see their Lemma 5.4 and the discussion before it; for the second property, see their Corollary 5.3 [25]. \square

This convex partition will work for us: we prove next that $\|\text{MCP}(G)\|_2 \leq 18c(X)$. (Note that $\|\text{MCP}(G)\|_2$ is trivially at most the cost of any *integer* triangulation, but the bound here concerns the *fractional* triangulation, so requires proof.)




FIG. 7.

The proof uses the constraints on X and leverages a previous analysis of $\text{MCP}(G)$ due to Plaisted and Hong [31, Lemma 10].

LEMMA 2.16. $\|\text{MCP}(G)\|_2 \leq 18c(X)$.

Proof. For every vertex v in the interior of the convex hull of the vertex set 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 the convex hull 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 $\|\text{MCP}(G)\|_2 \leq 6 \sum_{v \in V} S_{\min}(v)$ [31, Lemma 10].

We claim $S_{\min}(v) \leq (3/2) \sum_{e \ni v} X_e |e|$, where X_e is $\sum_{t \ni e} X_t$ if e is on the boundary of the convex hull; otherwise X_e is half this amount. As $\sum_v \sum_{e \ni v} X_e |e| = 2 \sum_e X_e |e| = 2c(X)$, the claim implies the lemma. We prove the claim.

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

Because X satisfies constraint (1.2), rotating clockwise around v , there is a sequence v_1, v_2, \dots, v_k of distinct vertices such that for each $i = 1, \dots, k$, the triangle $v v_i v_{(i+1) \bmod k}$ has positive weight in X . (To find the sequence, take any positive-weight triangle that has v as a vertex. Let v_1 and v_2 be the other vertices, in clockwise order. By constraint (1.2), there is a positive-weight triangle that shares edge vv_2 and lies clockwise to that edge. Let v_3 be the other vertex of that triangle. By constraint (1.2), there is a positive-weight triangle that shares edge vv_3 and lies clockwise to that edge. Let v_4 be the other vertex of that triangle. Continue, stopping when the next vertex v_i that would be added is v_1 —this must happen by constraint (1.2).)

Let e_i and t_i denote edge vv_i and triangle $v v_i v_{(i+1) \bmod k}$, respectively. Note that each edge, and each triangle, is distinct. Call the sequence of edges $h = (e_1, e_2, \dots, e_k)$ a *helix*. Let $\text{wrap}(h)$ denote the number of times h wraps around v . By a standard construction the X_e 's can be expressed as a linear combination of incidence vectors of helices. (Similar to Lemma 2.3's proof, repeatedly find a helix h , choose weight ϵ_h , and subtract $\epsilon_h/\text{wrap}(h)$ from each triangle X_{t_i} in the helix, reducing coverage near v by ϵ_h .) This gives a probability distribution ϵ on helices h such that each $X_e = \sum_h [e \in h] \epsilon_h / \text{wrap}(h)$.

Now choose a helix h at random from the probability distribution ϵ . Partition h greedily into contiguous subsequences of edges such that each group g is maximal subject to the constraint that the total clockwise angle around v swept by the group's edges is at most 360° . (In Figure 7, white triangles separate groups.) Each group contains a star, and (as neighboring groups are separated by at most 180°) there are at least $\lceil 360 \text{wrap}(h) / (360 + 180) \rceil = \lceil 2 \text{wrap}(h) / 3 \rceil$ groups.

From the randomly chosen h , choose a group g uniformly at random from h 's

first $\lceil 2\text{wrap}(h)/3 \rceil$ groups. For any given edge e , the probability that e is in g is at most $\sum_h [e \in h] \epsilon_h / (2 \text{wrap}(h)/3) = (3/2)X_e$. Thus, by linearity of expectation, the expected total length $E[\|g\|_2]$ of edges in g is at most $(3/2) \sum_{e \ni v} X_e |e|$. On the other hand, every g contains a star, so $S_{\min}(v) \leq E[\|g\|_2]$. This proves Lemma 2.16. \square

For the rounding procedure, fix the (previously arbitrary) convex partition CP to be the partition LK from Lemma 2.15. The cost of the final triangulation is at most

$$\begin{aligned} & \sum_f c(X^f) && \text{because each face } f \text{ is simple,} \\ & \leq 3\|\text{LK}\|_2 + 12 \sigma c(X) && \text{by Lemma 2.14,} \\ & \leq 3\lambda\|\text{MCP}(G)\|_2 + 54 c(X) && \text{by Lemma 2.15,} \\ & \leq 3\lambda \cdot 18 c(X) + 54 c(X) && \text{by Lemma 2.16,} \\ & = 54(\lambda + 1) c(X). \end{aligned}$$

Hence, the integrality gap is at most $54(\lambda + 1)$, completing the proof of Theorem 2.1. \square

3. The LP generalizes heuristics. This section proves our second new result (the LP generalizes MWT heuristics). Here is a summary of heuristics for determining that a given potential edge $e = xy$ of G is in every MWT of a given MWT instance $G = (V, E)$:

β -skeleton: For $\beta = 1/\sin k$ where $k \approx \pi/3.1$, there does not exist a point z in the point set such that $\angle xzy \geq \arcsin(1/\beta) \approx \pi/3.1$. Equivalently, the two disks of diameter $\beta|e|$ having e as a chord are empty of points. If this condition holds, then e is in every MWT of G [21, 7].

YXY-subgraph: For every potential edge pq that crosses $e = xy$, its size $|e|$ is at most $\min\{|px|, |py|, |qx|, |qy|\}$. If this condition holds, then e is in every MWT of G [38, 18].

maximality: For every potential edge that crosses e , that edge is known to be *excluded from* every MWT. If so, then e is in every MWT of G (see, e.g., [12]).

Here is a summary of heuristics for determining that a given potential edge e of G (not on the boundary of the region to be triangulated) is excluded from every MWT of G :

independence: Some potential edge that crosses e is known to be in every MWT. If this condition holds, then e is not in any MWT of G (see e.g., [12]).

diamond: Neither of the two triangles with base e and base angle $\pi/4.6$ is empty. If this condition holds, then e is not in any MWT of G [9, 15].

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 excluded from every MWT. If this condition holds, then e is not in any MWT of G [12]. (Edge 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 nonconvex, or e is not longer than 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 status of edges other than e . For example, if one of the conditions implies that some edge e' is excluded from every MWT, then the LMT-skeleton condition may then in turn imply that some new edge e is excluded from every MWT, because e' lies on one of two triangles t or t' in the pair 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 [12], most random instances with as many as 40,000 points are solvable via the heuristics.⁷

Here is our second new result. If an instance is solvable via the heuristics, then LP (1.1) solves the instance also.

THEOREM 3.1. *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 this section gives the proof. The first step is to show that each condition above that ensures that an edge is in (or excluded from) 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 3.2. *If any of the following conditions hold, the LP forces potential edge e of G to 1:*

1. β -skeleton: *The β -skeleton condition above holds for e .*
2. YXY-subgraph: *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 [21, 7, 38, 18]. Here we extend those arguments to any optimal *fractional* triangulation X^* : if the condition holds and X^* does not give e weight 1, then a polygon P' covering e whose triangles have 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 our setting for the following reason: in the MWT setting, the polygon P identified for retriangulation is the union of noncrossing triangles, whereas here, in the fractional setting, the polygon P' is the union of triangles that *can* cross (much as in Lemma 2.3). If the triangles in P' do not 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 e does not occur in the MWT. Previous works [21, 7] show that there must be a sequence t_1, t_2, \dots, t_m of empty triangles in the MWT whose union P covers e , as shown on the left in Figure 8. Using the β -skeleton condition, [21, 7] show that this union has a triangulation that costs less than does t_1, \dots, t_m , 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, \dots, t_m of empty triangles with positive weight in X^* covering e ,

⁷Reference [12] defines 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 for the effectiveness of the LMT-skeleton.




FIG. 8.

but these triangles can *cross* (see the example in Figure 8). We extend the arguments of [21, 7] to show that, even if such a crossing occurs, a triangulation of lower cost can still be found.

Full proof. Here are the details of the proofs for part 1 (β -skeleton) and part 2 (YXY-subgraph). Part 3 (maximality) has been discussed already above, in the proof idea.

Part 1 (β -skeleton): The original β -skeleton heuristics are shown to be valid for MWT by local-improvement arguments: if an edge e is in the β -skeleton (for $\beta \approx 1/\sin(\pi/3.1)$) but *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 [21, 7]. We briefly sketch this argument and then extend it to any optimal *fractional* triangulation X^* .

Assume for the remainder of this section that e goes horizontally from the point a on the left to the point b on the right. If ab is not in the MWT, there exists a set of MWT edges that intersect e . Let e_1, \dots, e_n , be the set of edges indexed in nondecreasing order of their length. If the edges are removed from the MWT, an empty polygonal region P results. In [21, 7] it is shown that P can be retriangulated at lesser cost by a set of edges that contains ab . The idea is to generate a sequence of triangulated polygons P_0, \dots, P_n such that P_0 is the degenerate polygon ab , P_n is a triangulation of P , and $P_{j-1} \subseteq P_j$. To obtain P_j , P_{j-1} is expanded to include the endpoints v_j and v'_j of e_j . Assume v_j is above the line through ab and v'_j is below it. If both v_j and v'_j already lie on the boundary of P_{j-1} , then $P_j = P_{j-1}$. Otherwise, at least one of them will not be on the boundary of P_{j-1} . Assume without loss of generality that v_j is not on the boundary of P_{j-1} (if v'_j is also outside P_{j-1} , it will be dealt with similarly).

Since v_j is not on the boundary of P_{j-1} , edge e_j intersects a boundary edge $v_i v_k$ of P_{j-1} . Consider the sequence δ of vertices on the path from a to b on the boundary of P (there are two such paths, but the one above the line through ab is intended). On the sequence δ , vertex v_i is the last vertex before v_j that belongs to P_{j-1} , and v_k is the first vertex after v_j that belongs to P_{j-1} . This observation allows us to clearly define $v_i v_k$ in the fractional setting because in that setting, polygon P may be self-intersecting and e_j may intersect more than one boundary edge of P_{j-1} in the half-space above ab . In general, the triangle $\Delta v_i v_j v_k$ contains a subsequence δ_1 of vertices on δ from v_i to v_j and another subsequence δ_2 from v_j to v_k . The polygon $v_i \delta_1 v_j \delta_2 v_k$ is then triangulated arbitrarily, and P_j is the union of P_{j-1} and the triangulated polygon $v_i \delta_1 v_j \delta_2 v_k$ —and possibly another triangulated polygon—to include v'_j if v'_j is not on the boundary of P_{j-1} .

This construction is shown in Figure 9. The polygon with dashed boundary is P , and P_{j-1} is a triangulation of the dark gray polygon. The union of P_{j-1} and arbitrary




FIG. 9.

triangulations of the light gray polygons is P_j , which includes the endpoints v_j and v_j' of e_j . The light gray polygon above ab is $v_i\delta_1v_j\delta_2v_k$. The white vertices inside triangle $\Delta v_iv_jv_k$ are δ_1 , and the black vertices inside the triangle are δ_2 .

In [21, 7] it is shown by induction on j that for every j all edges in the triangulation of P_j are shorter than e_j . By induction all the edges of P_{j-1} are shorter than e_{j-1} and $|e_{j-1}| \leq |e_j|$. Thus it remains to show that all the new edges that are added to P_{j-1} to form P_j are shorter than e_j . The new edges triangulate the polygon $v_i\delta_1v_j\delta_2v_k$ and are all contained in triangle $\Delta v_iv_jv_k$, so each new edge is shorter than $\max\{|v_iv_j|, |v_jv_k|, |v_iv_k|\}$. Since $v_iv_k \in P_{j-1}$, $|v_iv_k| < |e_{j-1}| \leq |e_j|$. It remains to show that $|v_iv_j| < |e_j|$. The argument for v_jv_k is similar. The following two facts are used for this part of the argument.

FACT 3.2.1 (see [21, Lemma 2]). *Let ab be an edge in the β -skeleton (where $\beta \approx 1/\sin(\pi/3.1)$). For any edge pq , if pq intersects ab , then it has length $|pq|$ greater than $\max\{|ab|, |ap|, |aq|, |bp|, |bq|\}$.*

FACT 3.2.2 (see [7, Remote Length Lemma]). *Let ab be an edge in the β -skeleton (where $\beta \approx 1/\sin(\pi/3.1)$). Let p, q, r , and s be four other distinct points of the point set such that pq intersects ab , rs intersects ab , pq does not intersect rs , and p and s lie on the same side of the line through ab . Then $|qr| < \max\{|pq|, |rs|\}$.*

The argument to show $|v_iv_j| < |e_j|$ is as follows. If v_i lies in triangle Δav_jb , then $|v_iv_j| \leq \max\{|av_j|, |v_jb|, |ab|\} < |e_j|$. The second inequality holds based on Fact 3.2.1. If v_i is outside Δav_jb , consider the convex hull of the path from a to v_j on P_j . Vertex v_i must lie in some triangle $\Delta v_cv_dv_j$, where v_c and v_d are hull vertices. Thus, $|v_iv_j| \leq \max\{|v_cv_j|, |v_cv_d|, |v_dv_j|\}$. Since v_c and v_d are hull vertices, they were added in the growth process in the past. Thus, the edges e_c and e_d with endpoints v_c and v_d were processed before e_j , so neither e_c nor e_d is longer than e_j . Combining this observation with Fact 3.2.2 gives $|v_cv_j| < \max\{|e_c|, |e_j|\} \leq |e_j|$. Using a similar argument, one can show that v_cv_d and v_dv_j are both shorter than e_j . Thus, $|v_iv_j| \leq \max\{|v_cv_j|, |v_cv_d|, |v_dv_j|\} < |e_j|$. This completes the proof that all the edges of P_j are shorter than e_j , and thus the new triangulation of P costs less.

Next we extend the above arguments to any optimal *fractional* triangulation X^* . If X^* does not give ab weight 1, there is a triangle t with positive X_t^* that properly intersects ab . For any side d of t that intersects ab there must be a triangle t' with positive $X_{t'}^*$ that has d as a side and lies on the other side of d from t . The existence of a triangle t' is a consequence of constraints (1.2). Repeating the same argument for the new triangle(s) gives a set of triangles that cover ab .




FIG. 10.

Let $\Gamma = (t_1, t_2, \dots, t_m)$ be the sequence of triangles in the order they intersect ab in the direction from a to b . Triangle t_1 is incident on a and t_m is incident on b . All triangles have a positive weight in X^* . The triangles in the sequence may or may not cross each other, as shown in Figure 10. Let P (the shaded area in the figure) be the polygon formed by the boundary edges of triangles in the sequence. As shown in the right of the figure, polygon P is self-intersecting if some of the triangles in the sequence cross. Next, we consider both cases and derive a contradiction in each case.

Case 1—no triangles in Γ cross. For this case we apply directly the technique used in [21] and [7] to retriangulate the interior of P at lower cost. Lowering the weight of those triangles in Γ by $\epsilon > 0$ and raising the weight of new triangles by ϵ gives a fractional triangulation of cost less than $c(X^*)$.

Case 2—some triangles in Γ cross. In this case the technique of [21] and [7] cannot be directly applied because, in that setting, the polygon P identified for retriangulation is the union of noncrossing triangles, whereas in this case, P is the union of triangles that cross.

Let d_1, \dots, d_n be the set of edges of triangles in Γ that intersect ab indexed in the order they intersect ab (in the direction from a to b). The only part of the argument used in [21] and [7] that does not directly generalize to our case concerns Fact 3.2.2. Fact 3.2.2 holds for any pair of edges pq and rs that intersect ab and do not intersect each other. Recall that in the MWT setting the edges intersecting ab do not intersect each other, so Fact 3.2.2 holds for any pair of those edges. However, in the current case some edges in d_1, \dots, d_n may intersect each other. Thus, Fact 3.2.2 does not automatically hold in this case. This issue is resolved by the following technical lemma.

FACT 3.2.3. *Let $d_i = pq$ and $d_j = rs$ be two edges of triangles t_1, t_2, \dots, t_m such that pq intersects ab , rs intersects ab , pq intersects rs , and p and s lie on the same side of the line through ab . Then ps and qr are both shorter than $\max\{|pq|, |rs|\}$.*

Proof. Assume that endpoints p and s are above the line through ab . Let ℓ be the intersection point of d_i and d_j , and assume without loss of generality that point ℓ is also above the line through ab . Let a_i and a_j be the intersection points of d_i and d_j with ab . Assume $i < j$ so that $|aa_i| < |aa_j|$.

We first show that the triangle bounded by d_i , d_j , and ab contains a vertex from the boundary of polygon P (see Figure 11). Consider a maximal contiguous subsequence of edges of d_1, \dots, d_n that starts with d_i and in which every edge in the subsequence intersects d_j . Let $d_h = uv$ be the last edge in this subsequence. Note that such d_h is well-defined. Let u and v be the endpoints of d_h below and above ab ,




FIG. 11.




FIG. 12.

respectively, and let a_h be the intersection point of d_h and ab . Clearly a_h is between a_i and a_j . Based on the maximality of the subsequence d_i, \dots, d_h , edge d_h is the last edge that intersects d_j , so d_{h+1} does not intersect d_j . Also d_{h+1} intersects ab at a point between a_h and a_j and shares an endpoint with d_h . Edge d_{h+1} cannot be incident on v because any edge connecting v to a point on line segment $a_h a_j$ intersects d_j . Therefore, d_{h+1} is incident on u . Also, d_{h+1} intersects $a_h a_j$, and does not intersect d_j or d_h . Thus, the other endpoint of d_{h+1} must be in the triangle $\Delta a_i l a_j$. Let x be this endpoint.

The argument so far shows the existence of a point x in the triangle bounded by d_i , d_j , and ab (see Figure 12). Thus, $\angle a_i l a_j \leq \angle axb < \pi/3$. The first inequality holds because x is in triangle $\Delta a_i l a_j$, and the second inequality holds for the following reason. From the definition of β -skeleton, edge ab is in the β -skeleton iff there does not exist a point z in the point set such that $\angle azb \geq \arcsin(1/\beta)$. For $\beta = 1/\sin(\pi/3.1)$, this implies that $\angle awb \leq \pi/3.1$ for any point w in the point set.

In triangle Δqlr the angle $\angle qlr$ is less than $\pi/3$, so $|qr| < \max\{|\ell q|, |\ell r|\} < \max\{|pq|, |rs|\}$, and a similar argument on triangle Δpls shows that $|ps| < \max\{|pq|, |rs|\}$. This completes the proof of Fact 3.2.3. \square

The remainder of the argument is similar to Case 1, where triangles do not cross. The interior of polygon P can be retriangulated at a lesser cost using the techniques in the original β -skeleton arguments. Finally, lowering the weight of triangles in t_1, \dots, t_m by $\epsilon > 0$ and raising the weight of new triangles by ϵ gives a fractional triangulation that costs less than $c(X^*)$.




FIG. 13.

Part 2 (YXY-subgraph): The argument used above for the β -skeleton works for the YXY -subgraph as well. The only parts of the argument for the β -skeleton that use geometric properties of edges in the β -skeleton are Facts 3.2.1 and 3.2.2. Thus, it suffices to show that these two facts hold for the edges of the YXY -subgraph too.

Let ab be an edge of the YXY -subgraph and pq be any edge that intersects ab . By definition of the YXY -subgraph, $|ab| \leq \min\{|pa|, |pb|, |qa|, |qb|\}$, so the union of two disks centered at a and b with radius $|ab|$ does not contain p or q (see Figure 13). If the angle $\angle paq$ in triangle $\triangle apq$ and the angle $\angle pbq$ in triangle $\triangle bpq$ are both greater than 90° , we have $|pq| > \max\{|pa|, |pb|, |qa|, |qb|, |ab|\}$, and the proof of Fact 3.2.1 is complete. To show that $\angle pbq > 90^\circ$, let s and t be the intersections of pq with the circle centered at a . We have $\angle pbq > \angle sbt > 90^\circ$. The second inequality holds because $\angle sbt$ is an inscribed angle (i.e., is an angle formed by two chords with a common endpoint in the circle centered at a) and its intercepted arc (i.e., the part of the circle which is “inside” the angle) is greater than 180° . The argument to show $\angle paq > 90^\circ$ is similar.

We next show that Fact 3.2.2 also holds for any edge ab of the YXY -subgraph. Let p, q, r , and s be four other distinct points of the point set such that pq and rs both intersect ab . We first consider the case that pq and rs do not intersect. Let $h(q)$ and $h(r)$, respectively, denote the distance of q and r from the line through ab . Lemma 2 in [38] states that if $h(q) \leq h(r)$, then $|qr| < |rs|$. Similarly, if $h(r) \leq h(q)$, then $|qr| < |pq|$. Hence $|qr| < \max\{|pq|, |rs|\}$, which proves Fact 3.2.2 in this case. If, on the other hand, pq and rs intersect, we use Fact 3.2.3. The proof of Fact 3.2.3 is almost the same whether ab is an edge of the YXY -subgraph or an edge of the β -skeleton. \square

LEMMA 3.3. *If any of the following conditions hold 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. diamond: *The diamond condition 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




FIG. 14.

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 3.2.

Full proof. As discussed in the proof idea, part 1 (independence) and part 3 (LMT-skeleton) are straightforward. We give the detailed proof of part 2 (diamond).

Part 2 (diamond): Like the β -skeleton and the YXY -subgraph, the original diamond heuristic for MWT is proved by local-improvement arguments: if the condition holds for an edge e that is in the MWT, then a polygon covering e within the MWT can be retriangulated at lesser cost, contradicting the optimality of the MWT [9, 15]. We first give a summary of the results in [15] and then explain how to extend those results to any optimal *fractional* triangulation X^* .

Suppose e is horizontal, p and q are its endpoints, and p is on the left of q . Let Δpqw and Δpqu be the two isosceles triangles with base angle $\pi/4.6$ above and below pq and C be the disk with diameter e , as shown in Figure 14. Suppose that Δpqw contains a point a' and Δpqu contains a point b' . If e is in the MWT, $a'b'$ is not in the MWT, and there is a set of triangles in the MWT that intersect $a'b'$. Consider the sequence of triangles encountered when tracing $a'b'$ toward a' , starting from edge e and stopping with the first triangle that has a vertex inside disk C , and let P_1 be the polygon formed by the boundary edges of triangles in the sequence. Let a be the vertex found inside C —if all else fails, then $a = a'$. In Figure 14, P_1 is the dark gray area above e . Similarly, consider the sequence of triangles encountered when tracing $a'b'$ toward b' , starting from edge e and stopping with the first triangle that has a vertex inside disk C , and let P_2 be the polygon formed by the boundary edges of triangles in the sequence. In Figure 14, P_2 is the dark gray area below e . The boundary edges of P_1 are grouped naturally into two chains, one from p to a and one from q to a . Vertex a does not belong to any of the two chains. P_1 is a *fan* on p (or q) if all triangles in P_1 are incident on p (or q). Similarly P_2 can be a fan on p (or q).

Drysdale, McElfresh, and Snoeyink [15] prove the following two facts to show how P_1 or P_2 or their union can be triangulated at lesser cost.⁸

FACT 3.3.1 (following [15, Lemma 8]). *If P_1 (or P_2) is not a fan, it can be retriangulated at lower cost.*

⁸In summarizing the result of [15], we use the same notation and names. The only exception is that [15] uses A and B instead of P_1 and P_2 .

FACT 3.3.2 (see [15, Lemma 9]). *When both P_1 and P_2 are fans, then their union can be retriangulated at lower cost.*

The above facts contradict the optimality of the MWT. Thus, e cannot be in any MWT. Next we extend the above arguments to any optimal *fractional* triangulation X^* . We find polygons P'_1 and P'_2 corresponding to P_1 and P_2 in the above argument and show that if X^* does not give e weight zero, then $P'_1 \cup P'_2$ can be retriangulated at lesser cost. Lowering the weight of those triangles by $\epsilon > 0$ and raising the weight of other triangles by ϵ gives a fractional triangulation that costs less than X^* . The details of the argument follow.

If X^* does not give pq weight 0, there is a triangle t with positive X_t^* that has pq as a side. Triangle t intersects $a'b'$. For any side d of t that intersects $a'b'$ there must be a triangle t' with positive $X_{t'}^*$ that has d as a side and lies on the other side of d from t . By repeating the same argument for the new triangle(s), a sequence Γ of triangles can be obtained that completely covers $a'b'$. As shown in Figure 15 triangles in the sequence Γ may or may not cross each other. The left figure shows the case that no triangles cross, while the right figure shows the case that some triangles cross. If triangles in the sequence do not cross, the original arguments from [15] apply. However, if triangles in the sequence cross, additional analysis is needed.




FIG. 15.

Consider the set of triangles in Γ encountered when tracing $a'b'$ toward a' , starting from edge e and stopping with the first triangle that has a vertex inside disk C , and let P'_1 be the polygon formed by the boundary edges of these triangles. Also let a be the vertex found inside C —if all else fails, then $a = a'$. Similarly, define P'_2 to be the polygon formed by the boundary edges of the set of triangles encountered when tracing $a'b'$ from e toward b' until a vertex is inside C . Figure 16 shows P'_1 and P'_2 (the dark shaded regions) in two cases. In the left figure there are no crossing triangles, while in the right figure some triangles in P'_1 cross.

We consider the following cases and show how in each case $P'_1 \cup P'_2$ can be retriangulated at lower cost.

Case 1— P'_1 is not a fan and some triangles in P'_1 overlap. This case is shown in Figure 17 and its magnified version in Figure 18. In the sequence of triangles in P'_1 encountered when tracing $a'b'$ toward a' , let t' be the first triangle that crosses some of the previous triangles in the sequence (the light gray triangle labeled in Figure 17). Let P'' be the polygon that covers the set of all triangles before t' (the dark gray polygon labeled P'' in Figure 17). Polygon P'' and triangle t' share an edge. Let $p'q'$ be this edge such that p' is to left of q' (see Figure 18). The boundary edges of P'' can be grouped naturally into two chains, one from p to p' and one from q to q' . We use the following additional facts from [15] to show that P'' can be retriangulated




FIG. 16.




FIG. 17.




FIG. 18.

at lower cost. Note that the following facts can be applied to P'' because P'' is the union of triangles that do not cross.

FACT 3.3.3 (see [15, Lemma 7]). *If the chain of boundary vertices from p to p' has three consecutive vertices x , y , and z with $|yq| \geq |zq|$ and the internal angle $\angle xyz$ is less than π , then P'' can be retriangulated to decrease its cost. (The same is true with q and q' exchanging roles with p and p' , respectively.)*

We say the *clockwise limit*⁹ on the directions of the boundary edges on the chain (from p to p') is perpendicular to qw if for every two consecutive vertices p_i and p_{i+1} on the chain, p_{i+1} is above or on the line through p_i that is perpendicular to qw .

FACT 3.3.4 (see [15, Lemma 6]). *If the chain of boundary vertices from p to p' has no three consecutive vertices x , y , and z with $|yq| \geq |zq|$ that form an internal angle ($\angle xyz$) of less than π , then the clockwise limit on the directions of the boundary*

⁹We use the term *clockwise limit* to be consistent with the terminology of the original paper [15]. The word *limit* in this context has nothing to do with the well-known mathematical limit.

edges is perpendicular to qw . (The same is true with q and q' exchanging roles with p and p' , respectively, and “counterclockwise” replacing “clockwise”).

Facts 3.3.4 implies that if the clockwise limit on the direction of boundary edges (on the chain from p to p') is not perpendicular to the line through qw , then there are three consecutive vertices x , y , and z with $|yq| > |zq|$ that form an internal angle ($\angle xyz$) of less than π , and the existence of such vertices by Fact 3.3.3 implies that P'' can be retriangulated to decrease its cost. The remainder of the proof for Case 1 shows that one of the edges on the boundary of P'' violates the mentioned limit on the direction of boundary edges, and thus P'' can be retriangulated at lower cost.

In Figure 18, the darker shaded area is P'' . Triangles in P'' do not overlap. Triangle $t' = \triangle p'q'v$ is the first triangle that overlaps some of the previous triangles. Triangle t' and polygon P'' share edge $p'q'$, and t' crosses some triangles covered by P'' . Thus, t' crosses some boundary edges of P'' either on the chain from p to p' or on the chain from q to q' . Assume without loss of generality that triangle t' crosses an edge on the chain from p to p' . Since edge $p'q'$ is a boundary edge of P'' , the other two sides of triangle t' ($p'v$ and $q'v$) should intersect P'' . Let $p_i p_{i+1}$ be the first edge on the chain from p to p' that is intersected by $q'v$ when moving from q' to v .

The line through $q'v$ divides the plane into a half-space above it (the half-space containing point w) and a half-space below it. It is easy to see that p_i is in the half-space above $q'v$ and p_{i+1} is below it, so p_{i+1} cannot be above the line through p_i that is perpendicular to qw . This means that the clockwise limit on the direction of the boundary edges is not perpendicular to qw . This, combined with Facts 3.3.3 and 3.3.4, implies that the interior of P'' can be retriangulated at lower cost.

Case 2— P'_1 is not a fan and no two triangles in P'_1 cross. Since triangles in P'_1 do not overlap, Fact 3.3.1 directly implies that P'_1 can be retriangulated at a lesser cost.

The previous two cases show that if P'_1 is not a fan, we can retriangulate some triangles in P'_1 to reduce the cost of triangulation. A similar argument applies to P'_2 , so there remains the case where both P'_1 and P'_2 are fans. We consider this case next.

Case 3— P'_1 and P'_2 are both fans. P'_1 is a fan on p (or q) if all triangles in P'_1 are incident on p (or q), and thus no two triangles in P'_1 overlap. Similarly if P'_2 is a fan, no two triangles in P'_2 overlap. Additionally, when both P'_1 and P'_2 are fans, triangles in P'_1 are all above the line through pq , triangles in P'_2 are below the line through pq , and no triangle in P'_1 intersects a triangle in P'_2 . Hence, no two triangles in $P'_1 \cup P'_2$ can cross, and by Fact 3.3.2 $P'_1 \cup P'_2$ can be retriangulated at lower cost.

In all of the above cases, by lowering the weight of some triangles in $P'_1 \cup P'_2$ by $\epsilon > 0$ and raising the weight of some other triangles in the new triangulation by ϵ , a fractional triangulation costing less than X^* can be obtained, which contradicts the optimality of X^* . This completes the proof of part 2 (diamond property) and Lemma 3.3. \square

Assume (as in the statement of Theorem 3.1) 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 3.2 and 3.3 (by a simple inductive proof) that every edge that can be deduced to be excluded from 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., [8, Thm. 7], [10, Thm. 4.1(i)], [22, Cor. 3.6.2])

that, for any simple polygon f , each basic 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 Theorem 3.1.

Acknowledgments. We thank the two anonymous referees for suggestions on improving the presentation of the results.

REFERENCES

- [1] O. AICHHOLZER, F. AURENHAMMER, S. W. CHENG, N. KATOH, G. ROTE, M. TASCHWER, AND Y. F. XU, *Triangulations intersect nicely*, Discrete Comput. Geom., 16 (1996), pp. 339–359.
- [2] E. ANAGNOSTOU AND D. CORNEIL, *Polynomial-time instances of the minimum weight triangulation problem*, Comput. Geom., 3 (1993), pp. 247–259.
- [3] F. AURENHAMMER AND Y. XU, *Optimal triangulations*, in Encyclopedia of Optimization, 2nd ed., Vol. 4, Springer, New York, 2009, pp. 2757–2764.
- [4] E. BALAS AND M. PERREGAARD, *Lift-and-project for mixed 0–1 programming: Recent progress*, Discrete Appl. Math., 123 (2002), pp. 129–154.
- [5] M. G. BORGELT, C. BORGELT, C. LEVCOPOULOS, AND J. S. B. MITCHELL, *Fixed parameter algorithms for the minimum weight triangulation problem*, Int. J. Comput. Geom. Appl., 18 (2008), pp. 185–220.
- [6] P. BOSE, L. DEVROYE, AND W. EVANS, *Diamonds are not a minimum weight triangulation's best friend*, Int. J. Comput. Geom. Appl., 12 (2002), pp. 445–453.
- [7] 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, ACM, New York, 1996, pp. 196–203.
- [8] G. B. DANTZIG, A. J. HOFFMAN, AND T. C. HU, *Triangulations (tilings) and certain block triangular matrices*, Math. Programming, 31 (1985), pp. 1–14.
- [9] G. DAS AND D. JOSEPH, *Which triangulations approximate the complete graph?*, in Optimal Algorithms, Lecture Notes in Comput. Sci. 401, Springer, Berlin, 1989, pp. 168–192.
- [10] J. A. DE LOERA, S. HOSTEN, F. SANTOS, AND B. STURMFELS, *The polytope of all triangulations of a point configuration*, Doc. Math., 1 (1996), pp. 103–119.
- [11] J. A. DE LOERA, J. RAMBAU, AND F. SANTOS, *Triangulations: Structures for Algorithms and Applications*, Algorithms Comput. Math. 25, Springer-Verlag, Berlin, 2010.
- [12] M. T. DICKERSON, J. M. KEIL, AND M. H. MONTAGUE, *A large subgraph of the minimum weight triangulation*, Discrete Comput. Geom., 18 (1997), pp. 289–304.
- [13] M. T. DICKERSON, R. L. SCOT DRYSDALE, S. A. MCELDFRESH, AND E. WELZL, *Fast greedy triangulation algorithms*, Comput. Geom., 8 (1997), pp. 67–86.
- [14] D. P. DOBKIN, H. EDELSBRUNNER, AND M. H. OVERMARS, *Searching for empty convex polygons*, Algorithmica, 5 (1990), pp. 561–571.
- [15] R. L. DRYSDALE, S. MCELDFRESH, AND J. S. SNOEYINK, *On exclusion regions for optimal triangulations*, Discrete Appl. Math., 109 (2001), pp. 49–65.
- [16] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability*, W. H. Freeman & Co., San Francisco, CA, 1979.
- [17] P. GIANNOPOULOS, C. KNAUER, AND S. WHITESIDES, *Parameterized complexity of geometric problems*, Comput. J., 51 (2007), pp. 372–384.
- [18] P. D. GILBERT, *New Results on Planar Triangulations*, Master's thesis, Applied Computation Theory Group, University of Illinois at Urbana-Champaign, Champaign, IL, 1979.
- [19] 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 (SODA), ACM, New York, SIAM, Philadelphia, 1996, pp. 252–260.
- [20] M. HOFFMANN AND Y. OKAMOTO, *The minimum weight triangulation problem with few inner points*, Comput. Geom., 34 (2006), pp. 149–158.
- [21] J. M. KEIL, *Computing a subgraph of the minimum weight triangulation*, Comput. Geom., 4 (1994), pp. 18–26.
- [22] D. KIRSANOV, *Minimal Discrete Curves and Surfaces*, Ph.D. thesis, Harvard University, Cambridge, MA, 2004.
- [23] G. T. KLINCSEK, *Minimal triangulations of polygonal domains*, Combinatorics, 79 (1980), pp. 121–123.
- [24] 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, Lecture Notes in Comput. Sci. 4271, Springer, Berlin, 2006, pp. 49–57.

- [25] C. LEVCOPOULOS AND D. KRZNNARIC, *Quasi-greedy triangulations approximating the minimum weight triangulation*, J. Algorithms, 27 (1998), pp. 303–338.
- [26] Y. KYODA, *A Study of Generating Minimum Weight Triangulation within Practical Time*, Master's thesis, Dept. Inform. Sci., Univ. Tokyo, Tokyo, Japan, 1996.
- [27] Y. KYODA, K. IMAI, F. TAKEUCHI, AND A. TAJIMA, *A branch-and-cut approach for minimum weight triangulation*, in Algorithms and Computations, Lecture Notes in Comput. Sci. 1350, Springer, Berlin, 1997, pp. 384–393.
- [28] E. L. LLOYD, *On triangulations of a set of points in the plane*, in the 18th Annual IEEE Symposium on Foundations of Computer Science, IEEE, Long Beach, CA, 1977, pp. 228–240.
- [29] W. MULZER AND G. ROTE, *Minimum-weight triangulation is NP-hard*, J. ACM, 55 (2008), 11.
- [30] 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, ACM, New York, 1996, pp. 517–518.
- [31] D. A. PLAISTED AND J. HONG, *A heuristic triangulation algorithm*, J. Algorithms, 8 (1987), pp. 405–437.
- [32] J. REMY AND A. STEGER, *A quasi-polynomial time approximation scheme for minimum weight triangulation*, J. ACM, 56 (2009), 15.
- [33] W. D. SMITH, *Studies in Computational Geometry Motivated by Mesh Generation*, Ph.D. thesis, Princeton University, Princeton, NJ, 1988.
- [34] A. SPILLNER, *A faster algorithm for the minimum weight triangulation problem with few inner points*, in Proceedings of the 1st Workshop on Algorithms and Complexity in Durham (ACiD 2005, Durham, UK), 2005, pp. 135–146.
- [35] A. TAJIMA, *Optimality and integer programming formulations of triangulations in general dimension*, in Algorithms and Computation, Lecture Notes in Comput. Sci. 1533, Springer, Berlin, 1998, pp. 377–386.
- [36] F. TAKEUCHI, H. IMAI, AND K. IMAI, *Polytopes of linear programming relaxation for triangulations*, in Kyoto University Research Information Repository 1068, Kyoto University, Kyoto, 1998, pp. 121–133.
- [37] C. A. WANG, F. CHIN, AND Y. F. XU, *A new subgraph of minimum weight triangulations*, J. Combin. Optim., 1 (1997), pp. 115–127.
- [38] 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, Lecture Notes in Comput. Sci. 834, Springer, Berlin, Heidelberg, 1994, pp. 423–427.
- [39] A. YOUSEFI AND N. E. YOUNG, *On a linear program for minimum-weight triangulation*, in Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2012, pp. 811–823.