

Disjoint Compatible Geometric Matchings*

Mashhood Ishaque[†]

Diane L. Souvaine*

Csaba D. Tóth^{‡*}

Abstract

We prove that for every even set of n pairwise disjoint line segments in the plane in general position, there is another set of n segments such that the $2n$ segments form pairwise disjoint simple polygons in the plane. This settles in the affirmative the *Disjoint Compatible Matching Conjecture* by Aichholzer *et al.* [1]. The key tool in our proof is a novel subdivision of the free space around n disjoint line segments into at most $n + 1$ convex cells such that the dual graph of the subdivision contains two edge-disjoint spanning trees.

1 Introduction

A *planar straight-line graph* (PSLG, for short) is a graph $G = (V, E)$ where the vertices are distinct points in the plane, the edges are straight-line segments between vertices such that every edge can intersect vertices and other edges only at its endpoints. A simple polygon is a 2-regular connected PSLG. A **geometric matching** is a 1-regular PSLG. The edges of a geometric matching are pairwise disjoint line segments in the plane. Two PSLGs (V, E_1) and (V, E_2) on the same vertex set are **compatible** if no edge in E_1 crosses any edge in E_2 .

A simple polygon with an even number of edges is the union of two disjoint compatible geometric matchings. Conversely, the union of two disjoint compatible geometric matchings is the union of pairwise disjoint simple polygons, each of which has an even number of edges. Aichholzer *et al.* [1] conjectured that for every geometric matching M with an even number of edges, there is a disjoint compatible geometric matching (*Disjoint Compatible Matching Conjecture*). The condition that M has an even number of edges cannot be dropped: for every odd $n \geq 3$, there are geometric matchings with n edges that do not admit a disjoint compatible matching [1]. Our main result is the proof of the Disjoint Compatible Matching Conjecture.

Theorem 1 *For every geometric matching with an even number of edges, and no three collinear vertices, there is a disjoint compatible geometric matching.*

The general framework for the proof of Theorem 1. Aichholzer *et al.* [1] proposed three stronger conjectures, each of which would immediately imply Theorem 1. Even though we do not prove any of these stronger conjectures, they laid out a framework that does lead to the proof of Theorem 1.

Before presenting two of these stronger conjectures, we recall a few definitions. For a finite set M of disjoint line segments in the plane, a *convex subdivision* of the free space is a set C of interior disjoint convex cells such that every cell is disjoint from the segments in M , and the union of their closures tile the plane. In a **dual (multi-)graph** G associated with M and C , the vertices correspond to the cells in C , and each endpoint of a segment in M corresponds to an edge between two incident cells (see Fig. 1). Note that M and C do not always determine the dual graph uniquely, since a segment endpoint may be incident to

*A preliminary version of this paper has appeared in the *Proceedings of 27th Symposium on Computational Geometry (Paris, 2011)*, pp. 125–134.

[†]Department of Computer Science, Tufts University, Medford, MA, {mishaq01, dls, cdtoth}@cs.tufts.edu. Partially supported by NSF Grants #CCF-0830734 and #CBET-0941538.

[‡]Department of Mathematics and Statistics, University of Calgary, Calgary, AB, cdtoth@ucalgary.ca. Partially supported by NSERC grant RGPIN 35586.

more than two convex cells, and the corresponding dual edge connects only two of those cells. In any case, the dual multi-graph has exactly $2n$ edges, some of which may be parallel. For brevity, we use the term *dual graph* even if G has parallel edges. An *orientation* of a (multi-)graph G is an assignment of a direction to each edge of G . An orientation is *even* if the indegree of every node is even. It is well known that every connected (multi-)graph with an even number of edges has an even orientation [21].

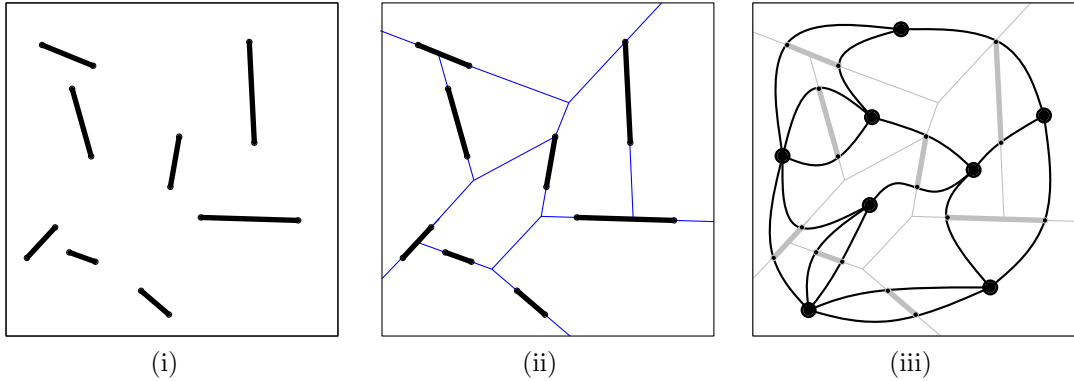


Figure 1: (i) A geometric matching with 8 edges in a bounding box. (ii) A convex subdivision with 8 cells. (iii) An associated dual graph.

Let M be a geometric matching with n edges. The *Two Trees Conjecture* [1] asserts that there is a convex subdivision with $n + 1$ cells such that an associated dual graph G is the edge-disjoint union of two spanning trees, and for every segment $uv \in M$, the edges of G corresponding to u and v are in different spanning trees. The *Extension Conjecture* [1] stipulates that if n is even, then there is a convex subdivision with $n + 1$ cells such that an associated dual graph G admits an even orientation with the property that for every node of G of indegree 2, the two incoming edges do not correspond to the two endpoints of any segment $uv \in M$. We no longer believe that the stronger conjectures hold [5], but we use several elements of these conjectures. For example, an even orientation with the above property induces a disjoint compatible geometric matching *even if* the dual graph has fewer than $n + 1$ cells (see Fig. 2).

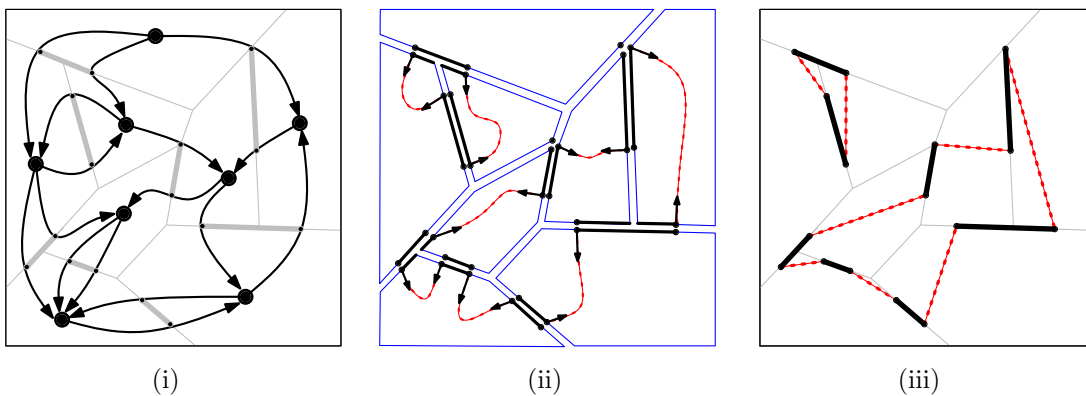


Figure 2: (i) An even orientation of the dual graph from Fig. 1(iii) such that whenever a node has indegree 2, the two incoming edges do not correspond to the two endpoints of the same segment. (ii) The even orientation maps an even number of segment endpoints to each cell. The segment endpoints assigned to each cell can be matched without crossings within that cell. (iii) The union of the matchings in all cells forms a disjoint compatible geometric matching.

We first prove a weaker form of the Two Trees Conjecture (Theorem 2), and show that there exists a convex subdivision with *at most* $n + 1$ cells whose dual graph *contains* two edge-disjoint spanning trees. Combining Theorem 2 with a purely graph theoretic lemma (Lemma 3), which is essentially a weaker form

of the Extension Conjecture, we prove the Disjoint Compatible Matching Conjecture.

Theorem 2 *For every set M of n disjoint line segments in the plane, with no three collinear segment endpoints, there is a convex subdivision of the free space into at most $n + 1$ cells such that an associated dual graph G contains two edge-disjoint spanning trees.*

Lemma 3 *Let $G = (V, E)$ be a multigraph with $4k$ edges that contains two edge-disjoint spanning trees, and let X be a collection of disjoint pairs of adjacent edges in E , called conflict pairs. Then G has an even orientation such that if the indegree of a vertex is 2, then the two incoming edges are not in X .*

Outline. In Section 2, a class of convex subdivisions $\mathcal{D}(M)$ and corresponding dual graphs are defined for a geometric matching M . Section 3 derives a sufficient conditions when a dual graph of a subdivision in $\mathcal{D}(M)$ contains two edge-disjoint spanning trees. Specifically, we show that the dual graph contains two edge-disjoint spanning trees if the subdivisions contains no “critical polygons,” a class of polygons that we shall define below. To complete the proof of Theorem 2, it remains to *construct* a convex subdivision in $\mathcal{D}(M)$ with no critical polygons. This is done algorithmically in Section 4. The algorithm starts with a simple initial convex subdivision with $n + 1$ cells in $\mathcal{D}(M)$, and modifies it recursively until all critical polygons are eliminated. The number of cells may decrease during the process, but the dual graph of the resulting convex subdivision contains two edge-disjoint spanning trees. Finally we prove Lemma 3 in Section 5, and show that the combination of Theorem 2 and Lemma 3 readily implies the Disjoint Compatible Matching Conjecture (Theorem 1).

Motivation and related work. The disjoint compatible matching conjecture is in the family of augmentation problems for planar straight-line graphs. In a typical augmentation problem, we are given a PSLG $G = (V, E)$ and a graph property P , and asked to find a (minimum) set of new edges such that $G' = (V, E \cup F)$ is a PSLG with property P . Theorem 1 shows that every 1-regular PSLG with an even number of edges and no three collinear vertices can be augmented to a 2-regular PSLG. Previously, Bose *et al.* [6] proved that every PSLG matching can be augmented to a PSLG spanning tree of maximum degree 3, and the bound 3 is the best possible. Hoffmann *et al.* [17] showed that every PSLG can be augmented to a connected PSLG such that the degree of every vertex increases by at most two. Al-Jubei *et al.* [3] proved a PSLG matching on $n \geq 4$ vertices can be augmented to a 3-connected PSLG iff no edge is a chord of the convex hull of the vertices, and the augmented PSLG has always at most $2n - 2$ edges. We refer to a recent survey by Hurtado and Tóth [18] on augmenting PSLGs.

The union of two disjoint compatible perfect geometric matchings is a collection of disjoint polygons in the plane. This simple fact is crucial for the best current upper bound on the number of simple polygons on a fixed set of n points in the plane [7, 16, 25]. The maximum number of perfect geometric matchings on a set of n points in the plane is known to be between $\Omega(3^n)$ and $O(10.05^n)$, as shown by García *et al.* [14] and Sharir and Welzl [26], respectively. Replacing a geometric matching with a compatible matching can be thought of as a reconfiguration in the configuration space of all perfect geometric matchings on a fixed point set. The reconfigurations can be represented by the *compatible matching graph* (resp., *disjoint compatible matching graph*), whose nodes correspond to the perfect geometric matchings on a fixed set of n points in the plane (n even), and two nodes are adjacent iff the corresponding matchings are compatible (resp., disjoint and compatible). Aichholzer *et al.* [1] showed that the compatible matching graph is connected, and its diameter is $O(\log n)$. Razen [24] constructed an n -element point set in general position where the diameter of this graph is $\Omega(\log n / \log \log n)$. Theorem 1 shows that the *disjoint compatible matching graph* has no isolated nodes, but it remains an open problem whether it is always connected.

The stronger conjectures suggested by Aichholzer *et al.* [1] and our proof of Theorem 1 rely on the *convex subdivision method*. A disjoint compatible matching is constructed by studying the properties of a dual graph of a convex subdivision. This method has a long history in visibility problems in computational geometry, although the definition of dual graphs is not uniform. For example, O’Rourke [23] used connectivity

properties of a dual graph of a convex subdivision to show that the free space around every set of $n \geq 5$ segments in the plane in general position can be covered by the visibility polygons of $\lfloor 2n/3 \rfloor$ point guards, and this bound is the best possible. See also [27, 28] for variants of this problem. Recently, Al-Jubei *et al.* [2] defined a general class of convex subdivisions for a set of disjoint line segments, and proved that there is a convex subdivision in this class whose dual graph is 2-edge-connected. They proved that a bridge (ie., a cut-edge) in the dual graph is equivalent to the existence of a simple polygon along the boundaries of some convex cells with certain properties (a “forbidden” pattern). Starting from an arbitrary convex subdivision in the class, they applied a sequence of local deformations until all forbidden configurations were removed. In this paper, we follow a similar strategy with a considerably more sophisticated machinery.

2 Convex Subdivisions

We define a general class of convex subdivisions for a set of disjoint line segments in the plane. Let M be a set of n disjoint line segments in the plane, with no three collinear segment endpoints. A *convex subdivision* of the free space is a set C of interior disjoint convex sets (called *cells*) such that every cell is disjoint from the segments in M , and the union of their closures tiles the plane. We can interpret a convex subdivision as a cell complex, and define the vertices and edges of a subdivision. The *vertices* of the subdivision are the endpoints of the segments in M as well as any other point that is incident to three or more cells. We will refer to these two types of vertices as *segment endpoints* and *Steiner points*. The *edges* of the subdivision are line segments between two consecutive vertices along the boundary of a cell, and rays lying on the boundary of a cell incident to a vertex. Note that every segment in M is the union of one or more edges of the subdivision. Every edge of the subdivision is either a *segment edge* which lies along some line segment in M , or an *extension edge* which lies in the free space. In the following definition, we orient the extension edges, but keep all segment edges undirected.

Definition 1 For a set M of disjoint line segments, let $\mathcal{D}(M)$ be the set of all convex subdivisions C that admit an orientation of the extension edges with the following properties.

- The out-degree of every segment endpoint is 1;
- the out-degree of every Steiner point in the free space is 0 or 1;
the Steiner points in the free space whose out-degree is 0 are called *sinks*;
- the out-degree of every Steiner point lying in the relative interior of a segment is 0;
- every cycle passes through at least one segment endpoint.




Figure 3: (i) A convex subdivision formed by a geometric matching and directed extension edges, with a sink q . (ii) The extension-paths emitted by segment endpoints u and v . (iii) The extension-tree rooted at the sink q . (iv) The extension-tree rooted at r .

Let $N = N(C)$ denote the 1-skeleton of the convex subdivision $C \in \mathcal{D}(C)$, that is, the **planar network** formed by all segments and all extension edges. The network N is partially directed: extension edges are directed and segment edges are undirected.

If no three segment endpoints are collinear, then it is easy to construct a convex subdivision in $\mathcal{D}(C)$ with no sinks as follows. Consider the segment endpoints v_1, v_2, \dots, v_{2n} in an arbitrary order. For each v_i , draw a directed edge (*extension*) along the supporting line of the incident line segment $v_i v_j$ that starts from v_i in direction $\overrightarrow{v_j v_i}$ and ends when it reaches another line segment, a previous extension, or infinity. This simple algorithm does not produce sinks if no three segment endpoints are collinear. (Aichholzer *et al.* [1] considered convex subdivisions obtained by this simple algorithm only.)

Definition 2 Let C be a convex subdivision in $\mathcal{D}(M)$ with directed extension edges.

- The **extension-path** of a vertex p of outdegree 1 is the open directed path along extension edges starting from p to the first point $q \neq p$ lying on a segment or at a sink, or to infinity. The extension-path does not include the endpoints of the path, and so it is disjoint from all segments.
- An **extension-tree** rooted at a vertex v lying on a segment or at a sink is the union of all (open) extension-paths that terminate at v . Every leaf of an extension tree is a segment endpoint.

Al-Jubei *et al.* [2] considered convex subdivisions with no sinks. By allowing sinks, we have significantly more freedom in constructing and modifying convex subdivisions. Figure 4 illustrates how one can modify adjacent extension edges, create a sink, and reduce the number of cells by one. Note that we can relocate a sink q to any other Steiner vertex q' of the same extension-tree by changing the directions of the extension edges along the path between q and q' . We will employ such local deformations in Section 4 for constructing a convex subdivision with no critical polygons.




Figure 4: (i) A convex subdivision for 4 segments with 5 cells. (ii) Extensions edges are deformed so that two incoming edges at q become collinear. (iii) By removing the outgoing edge of q , it becomes a sink, and we obtain a convex subdivision with 4 cells. (iv) We can move the sink to any Steiner vertex within the same extension-tree by changing the directions of some extension edges.

Let $S = S(C)$ denote the **set of sinks** of a convex subdivision $C \in \mathcal{D}(M)$, and let $s = |S(C)|$. It is easy to verify that if there are $|M| = n$ segments, then C has $n - s + 1$ cells. Indeed, the free space of the n segments and s sinks, $\mathbb{R}^2 \setminus (M \cup S)$, is a connected region with $n + s$ **holes**. Insert the $2n$ extension-paths one by one. The insertion of each extension-path either splits a region into two regions or decreases the number of holes. Eventually, the regions are convex cells, which are simply connected. So there were exactly $n - s$ steps in which a region is split into two.

A segment endpoint v may be incident to more than two cells. For a convex subdivision $C \in \mathcal{D}(M)$, we define the **dual graph** $G = G(C)$ such that the edge of G corresponding to a segment endpoint v connects the two cells adjacent to the **unique out-going edge** of v . Double edges are possible, corresponding to two endpoints of a line segment lying on the common boundary of two cells.

3 Sufficient Conditions for Two Spanning Trees

In this section, we establish a sufficient condition for a convex subdivision $C \in \mathcal{D}(M)$ to imply that the dual graph $G(C)$ contains two edge-disjoint spanning trees. Fix a convex subdivision $C \in \mathcal{D}(M)$ with directed extension edges for a set M of n disjoint line segments. The dual graph $G = G(C)$ has exactly $2n$ edges and at most $n + 1$ nodes, so there are enough edges for two edge-disjoint spanning trees. According to a well-known result by Nash-Williams [22] and Tutte [29], a (multi-)graph $G = (V, E)$ contains two edge-disjoint spanning trees if and only if, for every partition $V = \bigcup_{i=1}^r V_i$ of the vertex set into r nonempty classes, there are at least $2(r - 1)$ edges between different vertex classes. See Fig. 5 for an example. It is enough to consider vertex sets $V_i \subset V$ that each induce connected subgraphs of G , otherwise we could break some V_i into two classes without increasing the number of edges between different classes.

For a subset $C' \subseteq C$ of convex cells, let $G(C')$ be the subgraph of G induced by C' . We will determine the number of edges of each $G(C')$. Then we find a sufficient condition that guarantees that for every partition $C = \bigcup_{i=1}^r C_i$, there are at most $2n - 2(r - 1)$ edges within the induced subgraphs $G(C_i)$, and so there are at least $2(r - 1)$ edges between different vertex classes C_1, \dots, C_r .

Proposition 4 *If every subset $C' \subseteq C$ induces at most $2(|C'| - 1)$ edges of the dual graph G , then G contains two edge-disjoint spanning trees.*

Proof. For every partition $C = \bigcup_{i=1}^r C_i$, the induced subgraphs $G(C_i)$, $i = 1, \dots, r$, jointly have at most $\sum_{i=1}^r 2(|C_i| - 1) = 2|C| - 2r \leq 2(n + 1) - 2r$ edges. The total number of edges is $2n$. Therefore, there are at least $2(r - 1)$ edges between distinct vertex classes. \square

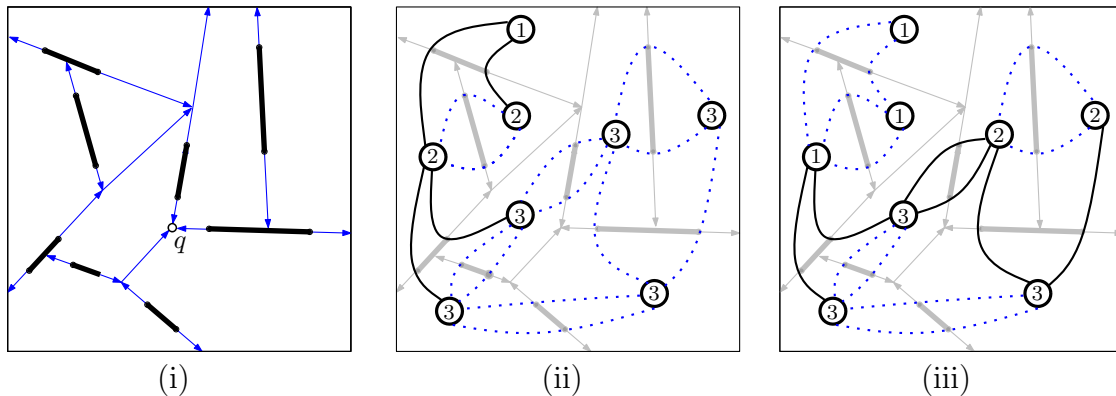


Figure 5: (i) A convex subdivision $C \in \mathcal{D}(M)$ for a set M of 8 segments. (ii) A partition of C into three classes, labeled ①, ②, and ③. In the dual graph, 4 solid edges are between distinct node classes, 12 dotted edges are within the node classes. (iii) Another partition of C into three classes. In the dual graph, 6 solid edges are between distinct classes, more than the requisite 4.

We can relax this condition noting that the total number of cells in C is only $n - s + 1$. Intuitively, every sink “decreases” the total number of cells by one. We distribute this beneficial effect among disjoint subsets of C . For every subset $C' \subseteq C$, we will assign a set of sinks $S(C')$, of cardinality $s(C')$, such that disjoint subsets of C are assigned disjoint sets of sinks, that is, for every partition $C = \bigcup_{i=1}^r C_i$ we have $\sum_{i=1}^r s(C_i) \leq s$.

Proposition 5 *If every subset $C' \subseteq C$ induces at most $2(|C'| + s(C') - 1)$ edges of the dual graph G , then G contains two edge-disjoint spanning trees.*

Proof. For every partition $C = \bigcup_{i=1}^r C_i$, the induced subgraphs $G(C_i)$, $i = 1, \dots, r$, jointly have at most $\sum_{i=1}^r 2(|C_i| + s(C_i) - 1) \leq 2|C| + 2s - 2r = 2(n - s + 1) + 2s - 2r = 2(n + 1) - 2r$ edges. The total number of edges is $2n$. Therefore, there are at least $2(r - 1)$ edges between distinct vertex classes. \square

Associating a region and sinks to sets $C' \subseteq C$. Let $C' \subseteq C$ be a set that induces a connected dual graph $G(C')$. Let $A(C')$ be the union of the closures of all cells in C' , that is, $A(C') = \bigcup_{c \in C'} \text{cl}(c)$. See Fig. 6(i)-(ii). Denote by $\partial A(C')$ the boundary of $A(C')$. Since every edge of $G(C')$ corresponds to an edge of the network $N(C)$ between cells in C' , the interior of $A(C')$ is connected.

Remark 1 Although we would like to assign to $C' \subseteq C$ every sink in the interior of $A(C')$, instead we will assign sinks to C' more carefully. Some sinks in the interior of a hole of $A(C')$ will be assigned to C' , rather than to the hole. Similarly, $A(C')$ may lie in a hole of another domain $A(C'')$ for some $C'' \subseteq C$, and some sinks in the interior of $A(C')$ may be assigned to C'' instead of C' . This will be crucial later in Lemma 8.

We define a polygonal domain $R(C') \subseteq A(C')$, which will be used for assigning sinks to C' , and it will also simplify counting the number of edges and vertices of $G(C')$ later in Lemma 7. Refer to Fig. 6. Construct region $R(C')$ from $A(C')$ by removing a thin channel along every path of $N(C)$ that

- connects a point of $\partial A(C')$ to a sink or to another point of $\partial A(C')$ via the interior of $A(C')$, and
- does not contain the outgoing extension edge of any segment endpoint.




Figure 6: (i) A convex subdivision $C \in \mathcal{D}(M)$, the cells in a subset $C' \subseteq C$ are shaded. (ii) Domain $A(C')$ is connected, and its complement $\mathbb{R}^2 \setminus A(C')$ has four connected components H_1, \dots, H_4 . The dotted paths are removed from $A(C')$ to create $R(C')$. (iii) In this instance, the boundary of $R(C')$ is composed of two weakly simple polygonal curves: γ_1 is the boundary of a hole, and γ_2 is an outer boundary. The boundaries are perturbed into *simple* polygonal curves for clarity. Sink q_2 is in $S(C')$ because it is in the interior of a hole of $R(C')$ and an extension-path goes from a Steiner vertex on the boundary of this hole to q_2 . However, sinks q_1 and q_3 are not in $S(C')$: sink q_1 is in the interior of a hole but no extension path can reach it from the boundary of the hole; and q_3 is on the outer boundary of $R(C')$.

The boundary of $R(C')$ is the union of the boundary of $A(C')$ and all paths removed from $A(C')$. Note that the interior of $R(C')$ is connected, since the edges of $N(C)$ that correspond to edges of $G(C')$ have not been removed. When we remove a path between two points of $\partial A(C')$, we effectively merge two components of $\partial A(C')$ into one component of $\partial R(C')$. When we remove a path between a boundary point and a sink q , we effectively move q to the boundary of $\partial R(C')$. As a result, no path of segment edges traverses the interior of $R(C')$ between two points of $\partial R(C')$; and no extension-path traverses the interior of $R(C')$ from a Steiner point of $\partial R(C')$ to another point of $\partial R(C')$.

We say that a closed polygonal curve (p_1, p_2, \dots, p_n) is *weakly simple* if either $n \in \{1, 2\}$, or $n \geq 3$ and for every $\varepsilon > 0$ and $1 \leq i \leq n$, there are points p'_i in the ε -neighborhood of p_i such that $(p'_1, p'_2, \dots, p'_n)$ is a simple closed polygonal curve. Similarly, an open polygonal curve (p_1, p_2, \dots, p_n) , where p_1 and p_n are points at infinity, is *weakly simple* if $n \geq 2$ and for every $\varepsilon > 0$ and $1 < i < n$, there are points p'_i in the ε -neighborhood of p_i such that $(p_1, p'_2, \dots, p'_{n-1}, p_n)$ is a simple open polygonal curve. Note that every (closed or open) weakly simple polygonal curve with $n \geq 2$ vertices subdivides the plane into a left and a right side (one of which may be empty).

The boundary of $R(C')$ may have several connected components. Each connected component is a weakly simple polygonal curve. We distinguish two types of boundary curves: *boundaries of holes* (closed curve whose interior is disjoint from $R(C')$); and *outer boundaries* (either a single closed curve whose interior contains $R(C')$, or one or more open curves).

Sink Assignment Rule. We can now assign sinks to each subset $C' \subseteq C$ as follows (refer to Fig. 6(iii)). Let $S(C') \subseteq S(C)$ be the set of sinks $q \in S(C)$ that satisfy one of the following conditions.

- q lies in the interior of $R(C')$;
- q lies on the boundary of a hole of $R(C')$;
- q lies in the interior of a hole of $R(C')$ and there is an extension-path from a Steiner point on the boundary of this hole to q .

Proposition 6 *If $C_1, C_2 \subset C$ such that $C_1 \cap C_2 = \emptyset$ and $G(C_1)$ and $G(C_2)$ are connected subgraphs, then $S(C_1) \cap S(C_2) = \emptyset$.*

Proof. Suppose, to the contrary, that there is a sink $q \in S(C_1) \cap S(C_2)$. Since C_1 and C_2 are disjoint, the corresponding domains $R(C_1)$ and $R(C_2)$ are interior disjoint, and the boundaries of their holes are also disjoint. Therefore q has to lie in the interior of a hole of $R(C_1)$ or $R(C_2)$. Assume without loss of generality that q lies in the interior of a hole of $R(C_1)$ such that there is an extension-path from a Steiner point on the boundary of that hole to q . Since $R(C_1)$ is connected, the entire domain $R(C_2)$ lies in (the closure of) a hole of $R(C_1)$. The extension-path from $\partial R(C_1)$ to q would have to traverse $R(C_2)$ from a Steiner point on the outer boundary of $R(C_2)$ to the boundary of a hole of $R(C_2)$. However, such an extension path cannot traverse $R(C_2)$ by the construction of $R(C_2)$. The contradiction shows that there is no sink $q \in S(C_1) \cap S(C_2)$. \square

Counting the number of nodes and edges of $G(C')$. We next give a formula for the number of edges in $G(C')$ for $C' \subseteq C$. Let Γ denote the set of all (open or closed) oriented weakly simple polygonal curves in $N(C)$ that subdivide the plane into a left and a right side. The orientation of a curve $\gamma \in \Gamma$ is independent from the directions of any extension edges along γ . In fact, every curve $\gamma \in \Gamma$ can be considered twice, once in each direction. Each curve in Γ partitions the plane into three parts: curve γ , an open *left*, and an open *right* side (depending the orientation of γ). We define *local features* for a curve $\gamma \in \Gamma$. A local feature of γ is an edge of $N(C)$ incident to a vertex of γ , of one of the following types (refer to Fig. 7).

- (a) edge e lies on the *left* side of γ , and either e is a segment edge or the extension edge emitted by a segment endpoint;
- (b) edge e lies on the *left* side of γ , and it is the outgoing extension edge of a Steiner vertex in γ ;
- (c) edge e lies on the *right* side of γ , and either e is a segment edge or the extension edge emitted by a segment endpoint;
- (d) edge e lies on the *right* side of γ , and it is the outgoing extension edge of a Steiner vertex in γ .

Intuitively, an (a)- or (b)-feature is an opportunity to enter into the *left* side of γ along a directed extension edge or an (undirected) segment edge: an (a)-feature offers an entry into the left side of γ along a segment edge or an extension edge that leaves a segment endpoint; a (b)-feature offers an entry into the left side of γ entirely in the free space of M . The (c)- and (d)-features have analogous meaning, with an edge entering into the *right* side of γ .

Let $a(\gamma)$, $b(\gamma)$, $c(\gamma)$, and $d(\gamma)$ denote the number of (a)-, (b)-, (c)-, and (d)-features of γ , respectively. Note that if $-\gamma$ and γ denote the same curve in Γ taken with opposite orientations, then we have $a(\gamma) = c(-\gamma)$, and $b(\gamma) = d(-\gamma)$.




Figure 7: (i) A convex subdivision, the polygonal curve is the boundary between the white and the shaded cells, where the left side of γ is shaded. The local features of γ are marked. For example, the segment edge pu is an (a)-feature, lying on the left side of γ , with vertices $p \in \gamma$ and segment endpoint u . (ii) A simplified figure, which shows only the edges of γ and its local features.

Lemma 7 Consider a subset $C' \subseteq C$ that induces a connected subgraph $G(C')$. Let h be the number of holes of $R(C')$. Let s_{in} denote the number of all sinks lying in the interior of $R(C')$. Orient each component of the boundary $\partial R(C')$ such that $R(C')$ lies on their left side, and let a (resp., b) be the total number of (a)-features (resp., (b)-features) along $\partial R(C')$. Then graph $G(C')$ has $2(|C'| - 1) - a - 2b + 2h + 2s_{\text{in}}$ edges.

Proof. Since C' induces a connected subgraph of $G(C')$, the domain $R(C')$ is connected. Let x be the number of line segments in M lying in the interior of $R(C')$. Now $R(C') \setminus (M \cup S)$ is a connected polygonal domain with exactly $h + x + s_{\text{in}}$ holes.

The number of extension-paths that enter the interior of $R(C') \setminus M$ from its boundaries is exactly $a + b + 2x$. Indeed, every segment in the interior of $R(C')$ emits two extension-paths. Every (b)-feature emits an extension-path into the interior of $R(C')$. Every (a)-feature either emits an extension-path into the interior of $R(C') \setminus M$, or is a segment edge. If an (a)-feature is a segment edge, then it is contained in some segment $uv \in M$ that enters the interior of $R(C')$ along e . If segment uv enters the interior of $R(C')$ along e , then the remaining portion of uv remains in the interior of $R(C')$, since no path of segment edges traverses the interior of $R(C')$ by construction. The endpoint of uv corresponding to the (a)-feature and lying in the interior of $R(C')$ emits an extension-path into the interior of $R(C') \setminus M$.

Insert the $a + b + 2x$ extension-paths one after the other into $R(C')$. The insertion of each extension-path either splits a region into two regions or decreases the number of holes. Eventually, we obtain convex cells, which are simply connected. So there were exactly $(a + b + 2x) - (h + x + s_{\text{in}}) = a + b + x - h - s_{\text{in}}$ steps in which a region is split into two regions. Thus the number of cells is $|C'| = 1 + a + b + x - h - s_{\text{in}}$.

A segment endpoint corresponds to an edge between two cells in C' if its outgoing extension edge lies in the interior of $R(C')$. There are exactly $a + 2x$ such segment endpoints. Hence, $G(C')$ has

$$\begin{aligned} a + 2x &= 2(a + b + x - h - s_{\text{in}}) - a - 2b + 2h + 2s_{\text{in}} \\ &= 2(|C'| - 1) - a - 2b + 2h + 2s_{\text{in}} \end{aligned}$$

edges, as required. \square

Critical polygons. We show next that if $G(C')$ has more than $2(|C'| + s(C') - 1)$ edges for some $C' \subseteq C$, then a single component of the boundary $\partial R(C')$ is responsible for this. Let $\Gamma_+ \subset \Gamma$ be the set of all *closed* weakly simple polygonal curves in Γ oriented *counterclockwise*. For every $\gamma \in \Gamma_+$, let $s(\gamma)$ denote the number of sinks q such that either $q \in \gamma$, or q lies in the interior of γ and there is an extension-path from a Steiner point of γ to q .

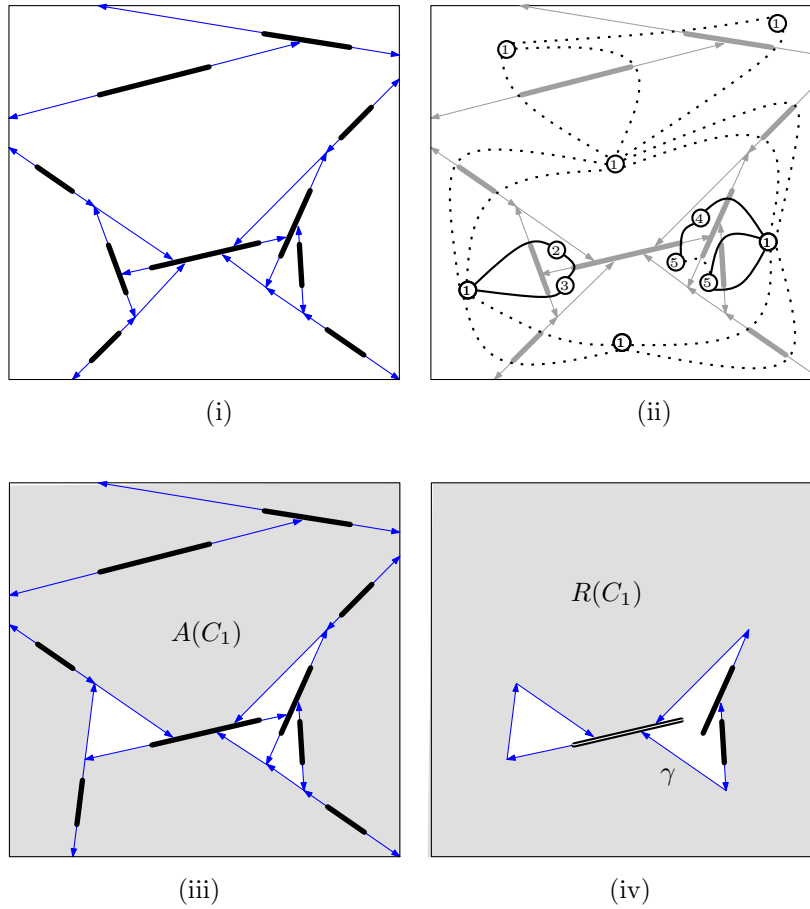


Figure 8: (i) A convex subdivision $C \in \mathcal{D}(M)$. (ii) The dual graph does not contain two edge-disjoint spanning trees because there are only 7 edges between the 5 vertex classes labeled ①...⑤. Subgraph $G(C_1)$ induced by the 6 nodes labeled ① has 12 dotted edges, but $2 \cdot (6 - 1) = 10$, so C_1 violates the condition in Propositions 4 and 5. (iii) The cells ① are shaded, their union is region $A(C_1)$. (iv) Region $R(C_1)$ has one hole, bounded by a weakly simple polygon γ , such that $c(\gamma) = d(\gamma) = s(\gamma) = 0$. So γ is a 0-critical polygon.

Lemma 8 *If $C \in \mathcal{D}(M)$ is a convex subdivision such that*

$$c(\gamma) + 2d(\gamma) + 2s(\gamma) \geq 2$$

for every $\gamma \in \Gamma_+$, then the dual graph $G(C)$ contains two edge-disjoint spanning trees.

Proof. By Proposition 5, it is enough to show that if $G(C')$ has more than $2(|C'| + s(C') - 1)$ edges for some $C' \subseteq C$, then there is a curve $\gamma \in \Gamma_+$ such that $c(\gamma) + 2d(\gamma) + 2s(\gamma) < 2$.

Consider a subset $C' \subseteq C$ such that $G(C')$ has more than $2(|C'| + s(C') - 1)$ edges. (See an example in Fig. 8). By Lemma 7, we have $2(|C'| + s(C') - 1) < 2(|C'| - 1) - a - 2b + 2h + 2s_{\text{in}}(C')$, that is, $a + 2b + 2(s(C') - s_{\text{in}}(C')) < 2h$. Let $\gamma_1, \dots, \gamma_h \in \Gamma_+$ denote the weakly simple polygons on the boundaries of the h holes of $R(C')$. By discarding features on the outer boundaries of $R(C')$, it is clear that

$$a \geq \sum_{i=1}^h c(\gamma_i) \quad \text{and} \quad b \geq \sum_{i=1}^h d(\gamma_i).$$

Thanks to the careful definition of the set $S(C')$ of sinks assigned to C' , we now have $s(C') - s_{\text{in}}(C') = \sum_{i=1}^h s(\gamma_i)$. This yields

$$\sum_{i=1}^h (c(\gamma_i) + 2d(\gamma_i) + 2s(\gamma_i)) < 2h.$$

For at least one of the holes γ_i , we have $c(\gamma_i) + 2d(\gamma_i) + 2s(\gamma_i) < 2$. □

A closed counterclockwise polygonal curve $\gamma \in \Gamma_+$ is called *critical* if $c(\gamma) + 2d(\gamma) + 2s(\gamma) \leq 1$. In particular, it is *0-critical* if $c(\gamma) + 2d(\gamma) + 2s(\gamma) = 0$, and *1-critical* if $c(\gamma) + 2d(\gamma) + 2s(\gamma) = 1$. See Fig. 9 for a few examples. Observe that if $\gamma \in \Gamma_+$ is critical, then it has no (d)-features, it does not pass through any sink, and there is no extension-path from γ to a sink in the interior of γ . A reformulation of Lemma 8 using this terminology is the following.

Corollary 9 *If $C \in \mathcal{D}(M)$ is a convex subdivision with no critical polygons, then the dual graph $G(C)$ contains two edge-disjoint spanning trees.*

Intuitive interpretation of critical polygons. Recall that $N = N(C)$ denotes the planar network formed by all segments and extensions of a convex subdivision $C \in \mathcal{D}(M)$. The line segments in M are undirected, and the extensions are directed. If a polygon γ is 0-critical then one cannot travel from γ to the exterior of γ respecting the directions of extension edges. Intuitively, “there is no way out” from γ . If γ is 1-critical, then every path to the exterior of γ that respects the directions of extension edges has to pass through the single (c)-feature of γ .

D-paths. We can relax the condition that the directions of *all* extension edges should be respected, since we know that there is no extension-path that starts from the exterior of a critical polygon γ and terminates at a sink on or in the interior of γ . We say that a directed path in $N(C)$ is a *D-path*, if it respects the directions of all extension edges that are not part of any extension-tree rooted at a sink. That is, a D-path may traverse segment edges and edges of extension-trees rooted at sinks in an arbitrary direction.

Proposition 10

- If $\gamma \in \Gamma_+$ is 0-critical, then there is no D-path from a point in γ to the exterior of γ .
- If $\gamma \in \Gamma_+$ is 1-critical, then every such D-path enters the exterior of γ along the same edge, which is incident to a vertex lying on a line segment in M .

Proof. Suppose that there is a D-path from $\gamma \in \Gamma_+$ to its exterior, and let e be the first edge that enters the exterior of γ . If e is a segment edge (undirected) or an extension edge directed away from γ , then it is (c)- or (d)-feature of γ . If e is an extension edge directed towards γ , then it lies on an extension-path that terminates at a sink q . If this extension-path leaves γ along some extension edge e' before it reaches q , then

e' is a (d)-feature of γ . If, however, sink q lies on or in the interior of γ , then $s(\gamma) \geq 1$. In all cases, we have $c(\gamma) + 2d(\gamma) + 2s(\gamma) > 0$.

The only possibility for $c(\gamma) + 2d(\gamma) + 2s(\gamma) = 1$ is that $c(\gamma) = 1$. In this case, any such D-path enters the exterior of γ along the same segment edge, which is incident to a vertex lying on a line segment. \square




Figure 9: Schematic drawing of some critical polygons. Only the critical polygons and their features are shown: segment edges are bold, extension-paths are dotted (and curved). Upper row: 0-critical polygons with no (c)- or (d)-features, that is, they have no D-paths to their exteriors. Lower row: 1-critical polygons with exactly one (c)-feature each, that is, all D-paths to the exterior pass through the same segment edge.

Maximal critical polygons. For every closed polygon $\gamma \in \Gamma_+$, let $C(\gamma)$ denote the set of cells in the interior of γ ; and let $R(\gamma) = R(C(\gamma))$. We define the *size* of γ to be the cardinality of $C(\gamma)$. We say that a critical polygon γ is *maximal* if there is no other critical polygon γ' such that $C(\gamma) \subset C(\gamma')$. Similarly, we define maximal 0-critical and maximal 1-critical polygons.

In the next section, we will construct a convex subdivision recursively, by modifying the boundary of a maximal critical polygon (defined below). The disjointness of maximal critical polygons is essential for this approach.

Proposition 11

- (i) If $\gamma_1, \gamma_2 \in \Gamma_+$ are maximal 0-critical, then $R(\gamma_1)$ and $R(\gamma_2)$ are disjoint.
- (ii) If there is no 0-critical polygon in Γ_+ and $\gamma_1, \gamma_2 \in \Gamma_+$ are maximal 1-critical, then $R(\gamma_1)$ and $R(\gamma_2)$ are disjoint.

Proof. Suppose that $\gamma_1, \gamma_2 \in \Gamma_+$ are maximal 0-critical polygons but they intersect (Fig. 10(i)-(ii)). By definition, they have no (c)- or (d)-features. Hence there is no such feature along the boundary of $R(\gamma_1) \cup R(\gamma_2)$. Let R be the union of $R(\gamma_1)$, $R(\gamma_2)$, and all holes of $R(\gamma_1) \cup R(\gamma_2)$. Then the boundary of R is a weakly simple polygon in Γ_+ with no (c)- or (d)-feature. So it is a 0-critical polygon containing all cells of $C(\gamma_1)$ and $C(\gamma_2)$, contradicting the maximality of γ_1 and γ_2 .

Now suppose that there is no 0-critical polygon in Γ_+ , and $\gamma_1, \gamma_2 \in \Gamma_+$ are maximal 1-critical polygons (Fig. 10(iii)-(iv)). That is, γ_1 and γ_2 have no (d)-features, and they each have exactly one (c)-feature. Let R be the union of $R(\gamma_1)$, $R(\gamma_2)$, and all holes of $R(\gamma_1) \cup R(\gamma_2)$. Then the boundary of R contains no (d)-feature, and at most two (c)-features. Any (c)-feature of ∂R corresponds to a (c)-feature of either γ_1 or γ_2 . However, ∂R must have at least one (c)-feature, since we assumed that there is no 0-critical polygon.




Figure 10: (i)–(ii) Schematic drawings of two intersecting 0-critical polygons γ_1 and γ_2 in the same convex subdivision. (iii)–(iv) Two intersecting 1-critical polygons γ_3 and γ_4 in the same convex subdivision.

If ∂R has exactly one (c)-feature (Fig. 10(iii)–(iv)), then its outer boundary is a 1-critical polygon that is strictly larger than γ_1 and γ_2 , contradicting the maximality of γ_1 and γ_2 . So we may assume that R has exactly two (c)-features, each of which must be the same edge in N as the corresponding (c)-feature of γ_1 or γ_2 . It follows that the (c)-feature of γ_1 (resp., γ_2) is in the exterior of $R(\gamma_2)$ (resp., $R(\gamma_1)$).




Figure 11: (i)–(ii) Two intersecting 1-critical polygons γ_1 and γ_2 in the same convex subdivision, where $R(\gamma_1)$ and $R(\gamma_2)$ are interior disjoint, but their common boundary $\gamma_1 \cap \gamma_2$ contains an extension edge e . (iii)–(iv) Two intersecting 1-critical polygons γ_3 and γ_4 in the same convex subdivision, where $R(\gamma_1)$ and $R(\gamma_2)$ overlap, and $R(\gamma_1) \cap R(\gamma_2)$ has two connected components, each of which is a 0-critical polygon.

We first show that the regions $R(\gamma_1)$ and $R(\gamma_2)$ overlap. Suppose, to the contrary, that $R(\gamma_1)$ and $R(\gamma_2)$ are interior-disjoint, but they have some common edges. Assume first that their common boundary $\gamma_1 \cap \gamma_2$ contains an extension edge e (Fig. 11(i)–(ii)). Then e is part of an extension tree rooted at some point r . The root r also lies in $\gamma_1 \cap \gamma_2$, otherwise the extension-path enters the exterior of γ_1 or γ_2 , creating a (d)-feature. Since neither γ_1 nor γ_2 contains any sink, the root r lies on some line segment $uv \in M$. Since $R(\gamma_1)$ and $R(\gamma_2)$ are interior-disjoint, r is a (c)-feature of both γ_1 and γ_2 , hence $R(\gamma_1) \cup R(\gamma_2)$ has no (c)-feature, contradicting the above assumption that it has at least two (c)-features. Assume now that $\gamma_1 \cap \gamma_2$ contains no extension edge, but it contains a segment edge along some segment $uv \in M$. Then the two endpoints of $(\gamma_1 \cap \gamma_2) \cap uv$ are (c)-features of both γ_1 and γ_2 , and hence neither γ_1 nor γ_2 is critical.

We have shown that $R(\gamma_1)$ and $R(\gamma_2)$ overlap, that is, $C(\gamma_1) \cap C(\gamma_2) \neq \emptyset$ (Fig. 11(iii)–(iv)). Let R' be a connected component of $\text{int}(R(\gamma_1)) \cap \text{int}(R(\gamma_2))$. As noted above, R' is disjoint of the (c)-features of γ_1 and γ_2 . We conclude that $\partial R'$ has neither (c)- nor (d)-features, since it would be a local feature of the same type for γ_1 or γ_2 . This contradicts our assumption that there are no 0-critical polygons.

□

4 Constructing a Convex Subdivision

Let M be a set of n disjoint line segments with no three collinear endpoints. In this section we construct a convex subdivision $C \in \mathcal{D}(M)$ with no critical polygon. We start with an initial convex subdivision $C_0 \in \mathcal{D}(M)$, which has no sinks, and then apply a sequence of *modifications* in order to eliminate critical polygons. For a convex subdivision $C \in \mathcal{D}(M)$, let $s(C)$ denote the total number of sinks, and let M_C be the set of segments that intersect or lie in the interior of some critical polygons. A modification step may have two possible outcomes: (1) We create a new sink and $s(C)$ strictly increases. In this case, we redraw the convex subdivision by a simple “left-right” algorithm described below that preserves all existing sinks. (2) There are no new sinks but we destroy a maximal critical polygon and M_C strictly decreases. In this case we consider the next maximal critical polygon and recurse. In each modification step, the function $f(C) = n \cdot s(C) - |M_C|$ strictly increases. Since the value of $f(C)$ is always between $-n$ and n^2 , after fewer than $n^2 + n$ modification steps, we obtain a convex subdivision $C \in \mathcal{D}(M)$ with $M_C = \emptyset$, that is, with no critical polygons.

Left-right subdivisions. Our initial convex subdivision $C_0 \in \mathcal{D}(M)$ is a “left-right” subdivision defined below. If a modification step creates a new sink, we redraw a left-right subdivision that preserves all extension-trees rooted at sinks, and draws new extensions for all other segment endpoints.

Let M be a set of n disjoint line segments in the plane such that no three endpoints are collinear, and no two segment endpoints have the same x -coordinate. Assume that we are given a set T of $s \geq 0$ pairwise disjoint extension trees (initially $s = 0$), each of which terminates at a sink, and have convex angles at their vertices. We construct extensions for all remaining segment endpoints in two phases.

Sort the segments in M according to the x -coordinates of their left endpoints in increasing order. Let u_i and v_i be the left and right endpoint, respectively, of the i th segment. In the first phase, consider the left endpoints u_1, \dots, u_n , successively. If u_i is not incident to an extension-tree in T , then draw a directed segment (*left extension*) along $\overrightarrow{v_i u_i}$ from u_i until it reaches another segment, a previous extension, or infinity. In the second phase, consider the right endpoints v_1, \dots, v_n . If v_i is not incident to an extension tree in T , then draw a directed segment (*right extension*) along $\overrightarrow{u_i v_i}$ from v_i until it reaches another segment, a previous extension, or infinity. The algorithm extends all segments beyond both of their endpoints, all angles are convex, and so it produces a convex subdivision in $\mathcal{D}(M)$ with exactly s sinks. The convex subdivision constructed by this algorithm, for any set T , is called a *left-right subdivision* of M .




Figure 12: (i) Disjoint line segments with an extension tree rooted at a sink. (ii) Left extensions drawn in the first phase. (iii) Right extensions drawn in the second phase.

We show that a left-right subdivision contains no 0-critical polygon. Recall that $N = N(C)$ is the network of all segment and extension edges of C (see Fig. 12(iii)). For a left-right subdivision $C \in \mathcal{D}(M)$,

let $N_1 \subset N$ be the union of all segments edges, all extension-trees rooted at sinks, and all *left* extensions (see Fig. 12(ii)). That is, N_1 is the union of edges created at the end of phase 1 of the above subdivision algorithm. Recall that a D-path is a directed path in N that respects the directions of extension edges that are not part of any extension-tree rooted at a sink (intuitively, a sink “neutralizes” the directions of an extension-tree).

Proposition 12 *If $C \in \mathcal{D}(M)$ is a left-right subdivision, then N_1 contains a D-path from any point $p \in N_1$ to infinity.*

Proof. We construct a D-path starting from a point $p \in N_1$ by repeating the following three steps: (1) If p lies on a segment, then walk to its left endpoint. (2) If p is a point incident to an extension-tree rooted at a sink, then walk to the leftmost leaf of the extension-tree. (3) If p lies on a left extension in N_1 , then follow the left extension to the end. To prove that this walk goes to infinity, it is enough to show that it has no loops. Each left segment endpoint u_i either emits a left extension or is incident to an extension-tree rooted at a sink and the leftmost vertex of the extension-tree is strictly to the left of u_i . Therefore the left-endpoints of the segments along the walk have decreasing x -coordinates, and so the walk visits every left segment endpoint at most once. \square

Corollary 13 *Left-right subdivisions have no 0-critical polygons.*

Proof. Combine Propositions 10 and 12 \square

Subdivisions with no 0-critical polygons. Our algorithm performs a sequence of local modifications starting from a left-right subdivision that will either eliminate all critical polygons from the resulting subdivisions or else introduce a new sink. It maintains the invariant that the subdivision always remains in the class $\mathcal{D}^*(M) \subseteq \mathcal{D}(M)$ defined here. Refer to Fig. 13.

Definition 3 *A subdivision $C \in \mathcal{D}(M)$ is in the class $\mathcal{D}^*(M)$ if it has the following properties.*

1. $N(C)$ contains a D-path from every point $p \in N$ to infinity.
2. If a segment endpoint v lies on or in the interior of a critical polygon, or if v is incident to any extension-tree rooted at a sink, then its outgoing edge is collinear with the line segment.
3. If a Steiner point p of outdegree 1 lies on or in the interior of a critical polygon, but it is not incident to any extension-tree rooted at a sink, then its outgoing edge is collinear with an incoming edge.

Proposition 14 *Every left-right subdivision is in $\mathcal{D}^*(M)$.*

Proof. Property 1 follows from Proposition 12. For property 2, consider an endpoint v of a segment $uv \in M$. If v is incident to an extension tree in T , rooted at a sink, then v is a leaf of this tree. Since every tree in T has convex angles at their vertices, the outgoing edge of v is collinear with segment uv . If v is not incident to any tree in T , then the outgoing edge is an extension of uv by construction. For property 3, let p be a Steiner vertex of outdegree 1, but not part of any tree in T . By construction, the outgoing edge of p is collinear with an incoming edge. \square

In fact, the vertices of a left-right subdivision satisfy all these conditions in the exterior of critical polygons, too. The intuition for the definition of class $\mathcal{D}^*(M)$ is that our local modifications successively destroy maximal critical polygons, but their effect remains largely outside of the surviving critical polygons. That is, on the boundary and inside the *current* critical polygons, we always find a structure similar to left-right subdivisions. The special conditions on the extension-trees rooted at sinks ensure that at every incident




Figure 13: A subdivision $C \in \mathcal{D}^*(M)$. The weakly simple polygon γ with white interior is maximal 1-critical. A sink q lies in the interior of γ . The outgoing edges of the segment endpoints v_1 and v_2 are not collinear with the incident segment in M ; these endpoints are in the exterior of γ and not incident to any extension-tree rooted at a sink. The outgoing edges of the Steiner points p_1, \dots, p_5 are not collinear with any incoming edges; these Steiner points are either in the exterior of γ or incident to an extension-tree rooted at q .

segment endpoint v , the incident extension edge of the tree and the segment edge already determine convex angles (that is, two straight angles), independently of any incoming extension edges incident to v .

As an immediate consequence of Definition 3, the subdivisions in $\mathcal{D}^*(M)$ have no 0-critical polygons.

Proposition 15 *Let $C \in \mathcal{D}^*(M)$. If $\gamma \in \Gamma_+(C)$ is a 1-critical polygon, then for every $p \in N(C)$ in the closed polygonal domain bounded by γ , there is a D-path $\delta(p) \subset N(C)$ from p to the (c)-feature of γ .*

Proof. By Definition 3, N contains a D-path $\delta(p)$ from p to infinity. By Proposition 10, $\delta(p)$ has to pass through the (c)-feature of γ . \square

We can now focus on 1-critical polygons. For our local modification steps, it will be convenient to group together consecutive collinear extension edges of N . Let a *straight-extension* be a maximal path along consecutive collinear extension edges in N . For example, in a left-right subdivision, every left extension (resp., right extension) is a straight-extension, even though they may consist of several collinear extension edges of N (see Fig. 12(iii)). Definition 3(2) requires that if a straight-extension is not part of an extension-tree rooted at a sink and its endpoint p is a Steiner point on or in the interior of a critical polygon, then p must be in the relative interior of some other straight-extension.

Testing whether a segment is in M_C . While the number of sinks remains the same, we measure the progress of our algorithm with $|M_C|$, the number of segments that intersect or lie in the interior of some critical polygon. If $C \in \mathcal{D}^*(M)$, there is a simple equivalent condition for $m \in M_C$. Two D-paths in N are *endpoint-disjoint* if no segment endpoint lies in the relative interior of both paths.

Proposition 16 *Let $C \in \mathcal{D}^*(M)$ and $m \in M$. We have $m \notin M_C$ if and only if $N(C)$ contains two endpoint-disjoint D-paths from any point $p \in m$ to infinity.*

Proof. First consider a segment $m \in M_C$. Let γ be the maximal 1-critical polygon that intersects m or contains m in its interior. If m contains the (c)-feature of γ , then all D-paths from any point $p \in m$ to infinity must pass through the endpoint of m lying in the exterior of γ . For all other segments $m \in M_C$, all D-paths from any point $p \in m$ to infinity must pass through the (c)-feature of γ , and then follow the line segment containing the (c)-feature to its endpoint in the exterior of γ . Hence any two D-paths to infinity have to pass through the same segment endpoint.

Next consider a segment $m \notin M_C$. By Definition 3(1), there is a D-path from any $p \in m$ to infinity in N . Suppose that every D-path in N from $p \in m$ to infinity passes through the same segment endpoint w . Let $W \subset N$ be the set of all points $q \in N$ such that all D-paths from q to infinity go through w . The set W is bounded, and its outer boundary is a 1-critical polygon $\gamma \in \Gamma_+$. Segment m lies on or in the interior of γ , contradicting our assumption that $m \notin M_C$. \square

One step of the recursion. It remains to show that if a convex subdivision in $\mathcal{D}^*(M)$ has a critical polygon, then we can perform a local modification step that strictly increases $f(C) = n \cdot s(C) - |M_C|$.

Lemma 17 *If $C \in \mathcal{D}^*(M)$ and C has a 1-critical polygon, then there is a convex subdivision $C' \in \mathcal{D}^*(M)$ such that $n \cdot s(C) - |M_C| < n \cdot s(C') - |M_{C'}|$.*




Figure 14: A left-right subdivision. The interior of a 1-critical polygon γ is shaded, its unique (c)-feature is the edge $p_0 q_0$. Dotted lines indicate the boundary of the geodesic hull H_γ . The convex vertices of H_γ are (q_0, \dots, q_5) , where q_1 is counterclockwise, and q_2, q_3, q_4, q_5 are clockwise.

Proof. Let $C \in \mathcal{D}^*(M)$ and let $\gamma \in \Gamma_+$ be a maximal 1-critical polygon (refer to Fig. 14). Let e_0 be the unique (c)-feature of γ , let $p_0 = e_0 \cap \gamma$, and let q_0 be the endpoint of the line segment containing e_0 that lies either on or in the exterior of γ . Let H_γ denote the geodesic hull of $\gamma \cup \{q_0\}$ with respect to the line segments of M lying in the exterior of γ . The geodesic hull H_γ is a closed polygonal domain, with boundary ∂H_γ . Note that every convex vertex of H_γ is either q_0 or a convex vertex of γ ; and every reflex vertex of H_γ is a segment endpoint in the exterior of γ . Let (q_0, q_1, \dots, q_k) denote the convex vertices of H_γ in counterclockwise order.

Every convex vertex q_i , $i = 1, \dots, k$, is a Steiner point in the free space: q_i cannot be at a (c)- or (d)-feature of γ by our assumptions. It cannot be at an (a)-feature, either, because an incident line segment would create a (c)-feature of γ by Def. 3(2). Every q_i , $i = 1, \dots, k$, is incident to two edges of γ , which are extension edges.

We define an *orientation* for these vertices as follows. Vertex q_i is *counterclockwise* if the edge preceding q_i in a counterclockwise order along γ is directed into q_i . Vertex q_i is *clockwise* if the edge following q_i

in a counterclockwise order along γ is directed into q_i . If both edges are directed into q_i , then it is both clockwise and counterclockwise.

A brief preview. A very brief preview of the remainder of the proof of Lemma 17 follows. Choose a convex vertex q_i , $1 \leq i \leq k$, modify the outgoing edge of q_i , and make some additional adjustments to obtain a convex subdivision $C' \in \mathcal{D}^*(M)$. One of the following two strategies will suffice. (1) Replace the outgoing edge of q_i with $q_i q_{i+1}$, and introduce a new sink at q_{i+1} . In this case, recompute a left-right subdivision (including the new sink) C' , and the number of sinks increases by one. (2) Replace the outgoing edge of q_i with an edge of ∂H_γ , and recompute some of the extension edges lying in the geodesic hull H_γ . In this case, the number of sinks remains the same but M_C strictly decreases.

We proceed with the details. Distinguish two cases.

Case 1: There are two consecutive convex vertices, q_i and q_{i+1} , $1 \leq i \leq k-2$, such that q_i is oriented counterclockwise and q_{i+1} is oriented clockwise (e.g., q_1 and q_2 in Fig. 14 or q_2 and q_3 in Fig. 15(ii)). We distinguish two subcases.

Subcase 1a: $q_i q_{i+1}$ is an edge of H_γ (see $q_2 q_3$ in Fig. 15). Replace the outgoing extension edge of q_i with the directed edge $q_i q_{i+1}$, thereby introducing a new sink at q_{i+1} . Build a valid extension-tree rooted at q_{i+1} (described below), which is disjoint from any extension-trees rooted at existing sinks of C . Then compute a left-right subdivision C' with these $s(C) + 1$ sinks (the $s(C)$ sinks of C and a new sink at q_{i+1}). Subdivision C' is in $\mathcal{D}^*(M)$ because it is a left-right subdivision, and it has one more sink than C .




Figure 15: Subcase 1a. (i) A 1-critical polygon γ with shaded interior. q_2 is counterclockwise and q_3 is clockwise, and $q_2 q_3$ is an edge of the geodesic ∂H_γ . (ii) Introduce a sink at q_3 , which is the root of the directed tree T_0 formed by edge $q_2 q_3$ and all extension-paths leading to q_2 and q_3 . Extend T_0 recursively with the extension-trees rooted at v_1 and v_2 to obtain a tree which has convex angles at all vertices. We perturb the tree at v_1 and v_2 so that they include very short outgoing edges collinear with the incident line segments, and obtain an extension-tree rooted at q_3 .

It remains to build an extension-tree rooted at q_{i+1} . Let T_0 be the union of $q_i q_{i+1}$, all extension-paths of N leading to q_i , and all extension paths of N leading to q_{i+1} . To show that T_0 is a crossing-free directed tree rooted at q_{i+1} , we need the following claim. Let $\gamma(q_i, q_{i+1})$ denote the portion of γ from q_i to q_{i+1} in counterclockwise order.

Claim 1 *There is no extension-path in $N(C)$ to q_i (resp., q_{i+1}) from any other vertex of $\gamma(q_i, q_{i+1})$.*

Proof. Suppose, to the contrary, that there is such an extension-path δ . Then the composition of δ with (part of) $\gamma(q_i, q_{i+1})$ is a weakly simple polygon with nonempty interior and without (c)- or (d)-features. This contradicts the fact that $N(C)$ has no 0-critical polygons, completing the proof. \square

Since the original outgoing edges of q_i and q_{i+1} are not part of T_0 , and any other vertex has outdegree at most one, T_0 is a tree. The tree T_0 does not contain any extension edge that crosses $q_i q_{i+1}$, because every such extension edge leads to an interior vertex of $\gamma(q_i, q_{i+1})$, which are not in T_0 by Claim 1. Therefore T_0 is a plane directed tree rooted at q_{i+1} . Its leaves are adjacent to segments in M by definition. At the leaves of T_0 lying on the boundary or in the interior of γ , the adjacent segment is collinear with the leaf by Def. 3(2). But this does not necessarily hold for leaves in the exterior of γ (see, e.g., segment endpoint v_1 in Fig. 15(i)).

We augment T_0 to a tree in which every leaf is collinear with the incident segment edge. Set $T = T_0$, and augment it by recursively appending to it any extension tree rooted at a segment endpoint v , which is at a leaf of T but the leaf is not collinear with the incident line segment. The resulting directed tree T consists of extension edges, it is rooted at q_{i+1} , but is it not necessarily an extension-tree because some of its non-leaf vertices may be segment endpoints. Perturb T at each such segment endpoint v , by creating a sufficiently short outgoing edge vv' collinear with the segment incident to v (see Fig. 15(ii)). After this perturbation, we obtain an extension-tree rooted at q_{i+1} , as required.

Subcase 1b: $q_i q_{i+1}$ is not an edge of H_γ (see edge $q_2 q_3$ in Fig. 16). We replace the outgoing edge of q_i or q_{i+1} with the incident edge of ∂H_γ . This provides a new edge from γ to the exterior of γ , and ensures that some segment in M_C has two endpoint-disjoint D-paths to infinity (c.f. Claim 5 below). After modifying the outgoing edge of q_i or q_{i+1} , carefully rebuilding nearby extensions yields a subdivision $C' \in \mathcal{D}^*(M)$ with $M_{C'} \subset M_C$. To control the extent of the modifications, we will define a polygonal domain $P \subset H_\gamma$: all extension edges in the exterior of P remain unchanged, and the extension edges inside P are redrawn. Polygon P is defined in the next two paragraphs.




Figure 16: Subcase 1b. (i) A 1-critical polygon γ with shaded interior. q_2 is counterclockwise and q_3 is clockwise. The geodesic H_γ between q_2 and q_3 wraps around the segment endpoints r_2 and r_3 . The ray $q_2 \vec{r}_2$ hits the boundary of γ at point t . (ii) Modification in Subcase 1b: the outgoing edge of q_2 is replaced by $q_2 r_2$, and the outgoing edge of r_2 is replaced by $r_2 t$. The extensions originating in the shaded polygon P are redrawn. After the modification, segment uv has two endpoint-disjoint D-paths to infinity.

By Proposition 15, $N(C)$ contains D-paths $\delta(q_i)$ and $\delta(q_{i+1})$ from q_i and q_{i+1} , respectively, to the (c)-feature of γ . The D-paths $\delta(q_i)$ and $\delta(q_{i+1})$ meet at some point \hat{q} , and we denote by $\delta(q_i, q_{i+1})$ the union of the initial portions of $\delta(q_i)$ and $\delta(q_{i+1})$ up to \hat{q} (see Fig. 17(i)). Let $q_i r_i$ and $r_{i+1} q_{i+1}$ be the edges of ∂H_γ incident to q_i and q_{i+1} , respectively (possibly $r_i = r_{i+1}$). Since r_i and r_{i+1} are reflex vertices of the

geodesic hull H_γ , they are endpoints of some line segments lying in the exterior of γ . The extension-paths emitted by r_i and r_{i+1} reach $\gamma(q_i, q_{i+1})$, so $N(C)$ contains D-paths $\delta(r_i)$ and $\delta(r_{i+1})$ from r_i and r_{i+1} , respectively, to the (c)-feature of γ . Both $\delta(r_i)$ and $\delta(r_{i+1})$ reach $\delta(q_i, q_{i+1})$ at some points \hat{r}_i and \hat{r}_{i+1} , respectively (possibly $\hat{r}_i = \hat{r}_{i+1}$). Since $\delta(r_i)$ and $\delta(r_{i+1})$ do not cross, the points $(q_i, \hat{r}_i, \hat{r}_{i+1}, q_{i+1})$ appear in this order along path $\delta(q_i, q_{i+1})$. Assume, by applying a reflection if necessary, that \hat{q} does not lie between q_i and \hat{r}_i along $\delta(q_i, q_{i+1})$.

The ray $\overrightarrow{q_i r_i}$ hits $\gamma(q_i, q_{i+1})$ at some point t (see Fig. 17(ii)). Note that any extension edge that crosses $q_i t$ is directed from the right to the left side of $\overrightarrow{q_i t}$, because no extension leaves $\gamma(q_i, q_{i+1})$. Let $\alpha(r_i, q_{i+1})$ be the directed path¹ from r_i to q_{i+1} that starts with the directed segment segment $r_i t$ and then follows γ counterclockwise from t to q_{i+1} . Let x be the first point on $\alpha(r_i, q_{i+1})$ that lies on a straight-extension that contains an edge of $\delta(q_i, q_{i+1})$. Point $x \in \alpha(r_i, q_{i+1})$ exists, because the outgoing edge of q_{i+1} is part of $\delta(q_i, q_{i+1})$. Let $\delta(x)$ be a D-path from x to the (c)-feature of γ . Now let P be the polygonal domain enclosed by segment $q_i r_i$, the initial portion of path $\alpha(r_i, q_{i+1})$ from r_i to x , and paths $\delta(q_i)$ and $\delta(x)$.




Figure 17: (i) The portions of δ and ∂H_γ between a counterclockwise vertex q_i and a clockwise vertex q_{i+1} . The paths $\delta(q_i)$, $\delta(r_i)$, $\delta(r_{i+1})$, and $\delta(q_{i+1})$ go to the (c)-feature of γ (which is outside of the figure). The interior of γ is shaded light gray. (ii) Three straight-extensions cross segment $r_i t$, one of them contains an edge of $\delta(q_i)$. Polygon P is shaded dark gray. (iii) A straight-extension crosses $q_i r_i$, and it contains an edge of $\delta(q_i)$.

Let e_0 be the straight-extension emitted by q_i (see Fig. 17(iii)). We will replace e_0 with $q_i r_i$. Any extension edge that crosses $q_i r_i$ in $N(C)$ will be truncated. Let E be the set of all edges of polygon P that are contained in some straight-extension of $N(C)$ that intersects $q_i r_i$. Clearly, we have $e_0 \in E$. Let every edge $e \in E$ inherit its direction from the straight-extension it is contained in.

Claim 2 *The interior of P lies on the right side of every edge $e \in E$.*

Proof. Let P_0 be the polygonal domain enclosed by $q_i r_i$, $\delta(q_i)$, and $\delta(r_i)$. Consider an edge $e \in E$. Since e crosses neither $\delta(q_i)$ nor $\delta(r_i)$, we have $e \subset P_0$. Recall that \hat{q} does not lie between q_i and \hat{r}_i along path $\delta(q_i, q_{i+1})$. This implies that $P_0 \subset P$. Hence edge $e \in E$ lies on $\delta(q_i)$, and the interiors of both P_0 and P are on the right side of e . \square

Every straight-extension that contains an edge $e \in E$ crosses segment $q_i r_i$ (see Fig. 18(i)). Order the segments in E by e_0, \dots, e_ℓ , according to their intersections with $q_i r_i$. The straight-extensions containing the segments in E decompose polygon P into *sectors*, each of which is adjacent to a unique edge in E by Claim 2. For $j = 0, \dots, \ell - 1$, denote by S_j the sector adjacent to $e_j \in E$.

¹Path $\alpha(r_i, q_{i+1})$ is not part of $N(C)$, and does not have to follow the directions of extension edges along γ .

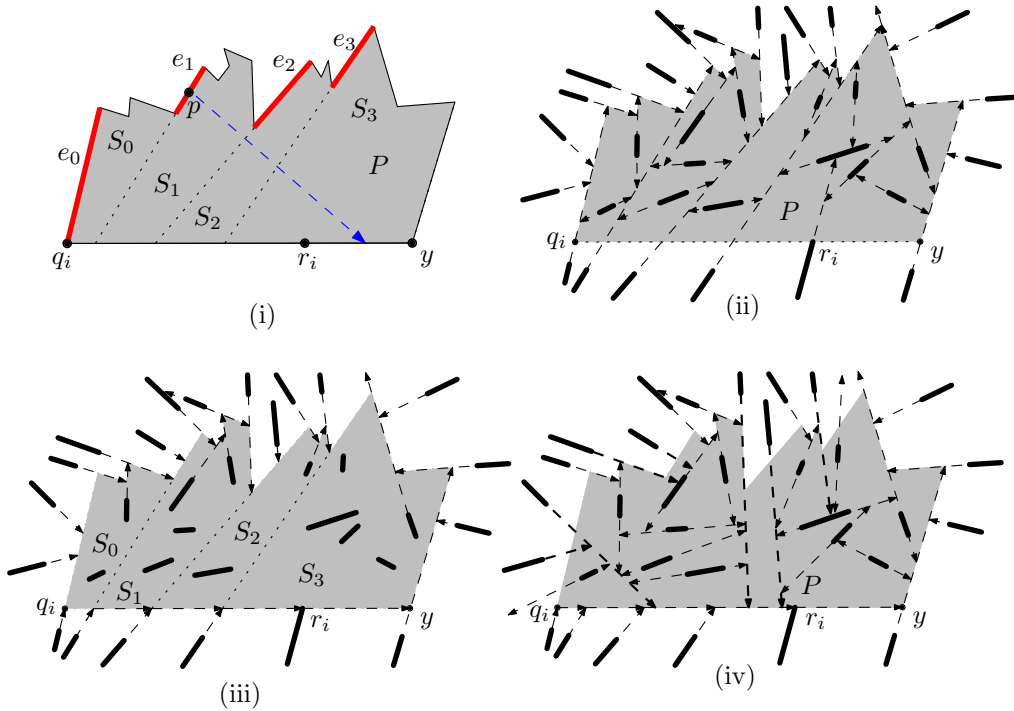


Figure 18: (i) Polygon P with edges $E = \{e_0, e_1, e_2, e_3\}$ is decomposed into 4 sectors. A ray enters the interior of P at point p of edge e_1 . (ii) The edges of the subdivision $C \in \mathcal{D}^*(M)$ that intersect P . (iii) The outgoing edge of q_i is replaced by $q_i r_i$, and the outgoing edge of r_i is replaced by $r_i y$. The straight-extensions inside polygon P are erased unless they overlap with the boundary of P . (iv) New straight-extensions are created in polygon P .

Claim 3 *If p is a point on an edge $e \in E$, then any ray shot from p into the interior of P cannot hit any other edge in E .*

Proof. If a ray is emitted from a point $p \in e$ of some edge $e \in E$, then it cannot hit another edge of E in the same sector, and it can only cross the boundary between any two adjacent sectors from left to right. \square

Let $r_i y = r_i t \cap \partial P$ be the portion of $r_i t$ on the boundary of P (that is $y = x$ if $x \in r_i t$ as in Fig. 17(ii), and $y = t$ otherwise as in Fig. 16(ii)).

We are now ready to construct a new convex subdivision C' , by modifying some extensions of C . Refer to Fig. 18. All extensions in the exterior of P remain the same (but some new extensions starting from P may enter the exterior of P). All extension-trees rooted at sinks remain the same (even if they are in the polygon P). All straight-line extensions that start from the interior of P and overlap with some edge of P are also preserved. Replace the straight-extension e_0 emitted by q_i by the directed segments $q_i r_i$ and $r_i y$ (these become the outgoing edges of q_i and r_i , respectively). The edges of $N(C)$ that cross $q_i y$ (from the exterior of P into the interior of P) are truncated such that they now end at $q_i y$. In particular, all edges in E are erased. Extend the edges of $N(C)$ that hit some $e \in E$ from the exterior of P successively into the interior of P until they hit a line segment, the boundary of ∂P , or another extension (by Claim 3, they do not hit any other edge in E). See Fig. 18(iii).

It remains to replace the straight-extensions of $N(C)$ that start from a segment endpoint u in the interior of P but do not overlap with any edge of P and are not part of any extension-tree rooted at a sink. Denote by U the set of segments endpoints in the interior of P whose straight-extensions are still missing. Consider the sectors $S_0, \dots, S_{\ell-1}$ in this order. Recall that no segment in M crosses the boundary between sectors. For each sector S_j , we draw extensions for the endpoints $U \cap S_j$ in two phases. In the first phase, successively draw the extensions whose directions point towards the right of edge e_j . In the second phase, draw

successively the extensions whose directions point towards the left of e_j . The extensions drawn in the first phase remain in polygon P (they may enter sectors $S_{j'}$, $j' > j$); the extensions drawn in the second phase may exit polygon P through the segment e_j . Denote by $C' \in \mathcal{D}(M)$ the resulting convex subdivision of M (Fig. 18(iv)).

Claim 4 $N(C')$ contains a D-path from every point $p \in N(C')$ to infinity.

Proof. The extensions of N lying outside of polygon P have not changed. Suppose that $p \in N(C)$, and its D-path $\delta(p)$ to infinity traversed polygon P . If $\delta(p)$ reached ∂P and overlaps with an edge $e \in E$, then it now enters the interior of P , reaches a point $p' \in \partial P$ (along extensions drawn in the first phase in each sector), and then goes to infinity along $\delta(p')$. If $\delta(p)$ reaches ∂P from the exterior at the edge $q_i y$, then it now follows $q_i y$ to y , and goes to infinity along $\delta(y)$. Finally, if p lies in the interior of P , then $N(C')$ now contains a path from p to a point $p' \in \partial P$ (along the extensions drawn in the first phase in each sector), and then path $\delta(p')$ continues in the exterior of P to infinity. \square

Claim 5 We have $M_{C'} \subsetneq M_C$.

Proof. We use Proposition 16 to check whether a segment $m \in M$ is in M_C . First we show that every segment $m \in M \setminus M_C$ continues to have two endpoint-disjoint D-paths to infinity. Let $m \in M \setminus M_C$. It lies in the exterior of γ , and it had two endpoint-disjoint D-paths to infinity in $N(C)$. The modification affects only those paths that intersect polygon P . Every D-path intersecting P reaches γ , and so it has to pass through the (c)-feature q_0 . Out of two endpoint-disjoint D-paths from m , at most one passes through q_0 , and so at most one intersects polygon P . Let $p \in \partial P$ be the point where this D-path reaches ∂P . It is enough to show that $N(C')$ contains a D-path from p to q_0 . It was shown in Claim 4 that there is a D-path from p to q_0 , which remains in the polygonal domain $R(\gamma) \cup P$.

Next we show that at least one segment $m \in M_C$ has two endpoint-disjoint D-paths to infinity in $N(C')$. Let $u \in \gamma$ be the segment endpoint whose extension-path reaches q_i , and let $uv \in M$ be the segment incident to u (see Fig. 16(ii)). We have $uv \in M_C$, since $u \in \gamma$. After the modification, the extension-path of u ends at r_i , which is the endpoint of a segment in the exterior of γ . From the other endpoint v , however, there is still a path to q_0 within the polygonal domain $R(\gamma) \cup P$. In $N(C)$, r_i had two endpoint-disjoint D-paths to infinity: one goes to q_0 followed by some D-path α_1 , and a D-path α_2 entirely in the exterior of γ . After the modification, $N(C')$ contains two endpoint-disjoint D-paths from uv to infinity: one going through q_0 followed by α_1 , and another one through r_i followed by α_2 . Hence $uv \in M_C$ but $uv \notin M_{C'}$. \square

Claim 6

- If a segment endpoint $v \in N(C')$ lies on or in the interior of a critical polygon, or if v is incident to an extension-tree rooted at a sink, then its outgoing edge is collinear with the line segment.
- If a Steiner point $p \in N(C')$ of outdegree 1 lies on or in the interior of a critical polygon, but it is not part of any extension-tree rooted at a sink, then its outgoing edge is collinear with an incoming edge.

Proof. As a result of the modifications in polygon P , the first condition is violated only at the segment endpoint r_i (the new outgoing edge of r_i is contained in $r_i y$, which is not collinear with the incident segments in M). However, r_i was in the exterior of all critical polygons in $N(C)$ by the maximality of $\gamma \in \Gamma_+$, and by Claim 5 it is also in the exterior of all critical polygons in $N(C')$. The second condition is violated only for the Steiner vertex q_i . However, q_i now has an extension-path to r_i , which lies in the exterior of all critical polygons of $N(C')$, and so q_i is also in the exterior of all critical polygons of $N(C')$. \square

By Claims 4, 5, and 6, the modification in subcase 1b produces a subdivision $C' \in \mathcal{D}^*(M)$ such that $s(C') = s(C)$ and $|M_{C'}| < |M_C|$.

Case 2: Vertex q_1 is oriented clockwise or q_k is oriented counterclockwise (e.g. q_1 in Fig. 19(i)). We may assume w.l.o.g. that q_1 is clockwise (by applying a reflection if necessary). We distinguish two subcases.




Figure 19: Subcase 2a. (i) A 1-critical polygon γ with shaded interior. q_1 is oriented clockwise. q_0q_1 is an edge of the geodesic ∂H_γ . (ii) We replace the outgoing edge of q_1 with a directed edge $q_1q'_0$, which creates a new D-path from q_1 to the exterior of γ that does not pass through the bottleneck (c)-feature q_0 .

Subcase 2a: the line segment q_0q_1 lies in the geodesic hull H_γ (see q_0q_1 in Fig. 19(i)). Replace the outgoing edge of q_1 with the directed edge $q_1q'_0$, where q'_0 is a point on the outgoing edge of q_0 in a small neighborhood of q_0 (Fig. 19). Let P be the polygonal domain bounded by $q_1q'_0, q_0q'_0$, and the D-path $\delta(q_1)$ (note that the D-path $\delta(q_0)$ degenerates to a simple point). Recompute the extensions within P in the same way as it was done in Subcase 1b above. We obtain a new convex subdivision $C' \in \mathcal{D}^*(M)$, analogously to subcase 1b. All segments in $M \setminus M_C$ still have two endpoint-disjoint D-paths to infinity. Let u be the segment endpoint along γ whose extension-path reaches q_1 . It is an endpoint of some segment $uv \in M_C$. We show that $uv \notin M_{C'}$, and so $|M_{C'}| \leq |M_C|$, as required.

We use Proposition 16 to verify that $uv \notin M_{C'}$. Let α_v be the D-path from v that starts with the extension-path emitted by v and goes to infinity. Since $uv \in M_C$, the D-path α_v passes through q_0 , and q'_0 . After the modification, there is still a D-path from v to q_0 in the polygonal domain $R(\gamma) \cup P$. On the other hand, u now has a D-path to q'_0 along an extension-path that uses the new edge $q_1q'_0$; this D-path bypasses the bottleneck q_0 . So there are now two endpoint-disjoint D-paths from any point in M to q'_0 . Note that q'_0 is in the exterior of γ , and by the maximality of the 1-critical polygon γ , it has two endpoint-disjoint D-paths to infinity. The composition of these D-paths (from m to q'_0 and from q'_0 to infinity) provides two endpoint-disjoint D-paths from m to infinity.

Subcase 2b: q_0q_1 does not lie in the geodesic hull H_γ (see q_0q_1 in Fig. 20). We can repeat the same argument as in subcase 1b, with the only difference (in fact, simplification) that path $\delta(q_0)$ degenerates to a simple point. This implies that $\hat{q} = q_0$ and so $\delta(q_0, q_1) = \delta(q_1)$.

Let q_1r_1 be the edge of ∂H_γ incident to q_1 . Since r_1 is a reflex vertex of the geodesic hull H_γ , it is an endpoint of some line segments lying in the exterior of γ . The ray $\overrightarrow{q_1r_1}$ hits $\gamma(q_0, q_1)$ at some point t (Fig. 20(i)). Note that any extension edge that crosses q_1t is directed from the left to the right side of $\overrightarrow{q_1t}$, because no extension leaves $\gamma(q_0, q_1)$. Let $\alpha(r_1, q_0)$ be the directed path from r_1 to q_0 that starts with the directed segment r_1t and then follows γ clockwise from t to q_0 . Let x be the first point on $\alpha(r_1, q_0)$ such that either x lies on a straight-extension that contains an edge of $\delta(q_0, q_1)$, or $x = q_0$. Let $\delta(x)$ be a D-path from x to q_0 . Now let P be the polygonal domain enclosed by segment q_1r_1 , the initial portion of path $\alpha(r_1, q_0)$ from r_1 to x , and paths $\delta(q_1)$ and $\delta(x)$. Analogously to subcase 1b, we construct a subdivision

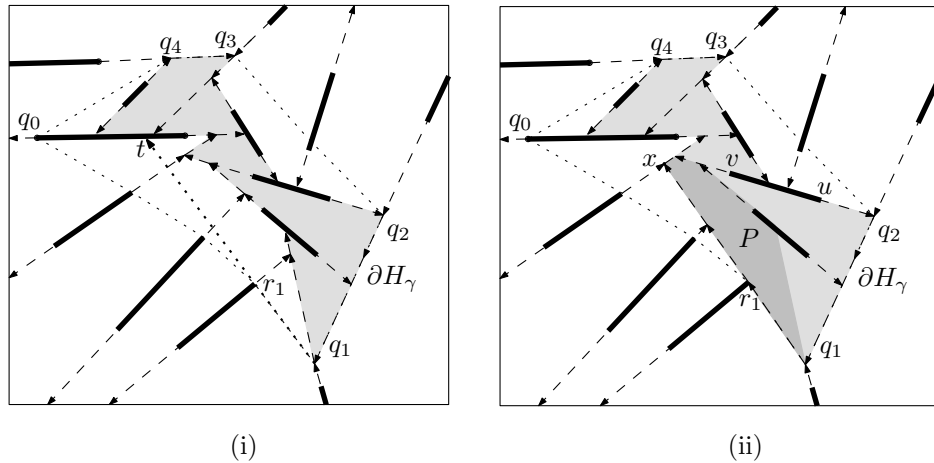


Figure 20: Subcase 2b. (i) A 1-critical polygon γ with shaded interior. q_1 is oriented clockwise. The geodesic H_γ between q_0 and q_1 wraps around the segment endpoint r_1 . The ray $q_2\vec{r}_2$ hits the boundary of γ at point t . (ii) Modification in Subcase 2b: x is the first intersection point of ray $q_1\vec{r}_1$ with a straight-extension that contains an edge of $\delta(q_i, q_{i+1})$. The outgoing edge of q_1 is replaced by q_1r_1 , and the outgoing edge of r_1 is replaced by r_1x . The extension edges in the shaded polygon P are redrawn. After the modification, segment uv has two endpoint-disjoint D-paths to infinity.

$C' \in \mathcal{D}^*(C)$ by replacing the outgoing edge of q_1 by q_1r_1 , and redrawing the extensions inside polygon P (Fig. 20(ii)). This completes Subcase 2b, and the proof of Lemma 17. \square

5 From Two Edge-Disjoint Spanning Trees to Disjoint Compatible Matchings

In this section, we prove the disjoint compatible matching conjecture. We first prove Lemma 3, which involves purely graph theoretic arguments. Then we show that Theorem 2 and Lemma 3 readily imply the existence of a compatible disjoint matching for any matching M .

We mention two related results. Farber *et. al.* [11] and Cordovil and Moreira [10] proved independently that if $G = (V, E)$ is the union of two edge-disjoint spanning trees on n vertices, then one can exchange all edges of the two trees in $n - 1$ steps such that each step exchanges one pair of edges and maintains a partition of E into two spanning trees. These exchange steps, however, do not seem to be sufficient to reach a partition of E into two *conflict-free* spanning trees. Cannon *et al.* [8] have recently designed a polynomial-time algorithm for testing whether a multi-graph with given conflict pairs has an even orientation such that if the indegree of a vertex is 2, then the two incoming edges are not in conflict.

5.1 The union of two edge-disjoint spanning trees

We start with a brief review of well-known properties of multigraphs formed by the edge-disjoint union of two spanning trees. We use the following a shorthand notation throughout this section: for a set S and an element t , we write $S + t = S \cup \{t\}$ and $S - t = S \setminus \{t\}$. Let $G = (V, E)$ be a multigraph on n vertices which is the edge-disjoint union of a *blue* spanning tree (V, B) and a *red* spanning tree (V, R) . Since G has n vertices and $2n - 2$ edges, the sum of the vertex degrees is $4n - 4$. Since each vertex is incident to both spanning trees, the minimum vertex degree is 2. It follows that G has at least two vertices of degree 2 or 3.

Vertex pruning. Let $v_1 \in V$ be a vertex of degree 2, incident to edges $e_1, e_2 \in E$. Then $\{e_1, e_2\}$ is a cut set of G , and so each spanning tree has at least one edge in $\{e_1, e_2\}$. Assume without loss of generality

that $e_1 \in B$ and $e_2 \in R$. Then v is a leaf in each spanning tree. It follows that $(V - v_1, B - e_1)$ and $(V - v_1, R - e_2)$ are two edge-disjoint spanning trees on the vertex set $V - v_1$. *Pruning vertex v_1* is the operation that replaces $(V, B \cup R)$ by $(V - v_1, (B - e_1) \cup (R - e_2))$.

Now let $v_2 \in V$ be a vertex of degree 3, incident to edges $e_1 = v_2u_1$, $e_2 = v_2u_2$, and $e_3 = v_2u_3$. Then two of these edges are in the same tree, and the third edge is in the other tree. Assume without loss of generality that $e_1, e_2 \in B$ and $e_3 \in R$. Note that if we remove v_2 and all three incident edges from G , then the blue tree breaks into exactly two connected components, where u_2 and u_3 are in distinct components. The two components can be reconnected by adding a new edge between u_2 and u_3 . Consequently, $(V - v_2, B - e_1 - e_2 + u_1u_2)$ and $(V - v_2, R - e_3)$ are two edge-disjoint spanning trees on the vertex set $V - v_2$. *Pruning vertex v_2* is the operation that replaces $(V, B \cup R)$ by $(V - v_2, (B - e_1 - e_2 + u_1u_2) \cup (R - e_3))$.




Figure 21: Multigraph G_1 , which is the edge-disjoint union of a dotted red tree and a dashed blue tree. Algorithm 1 successively prunes vertices v_1, \dots, v_7 of degree 2 or 3, and maintains the two edge-disjoint spanning trees.

We can reduce G to a single vertex by successively pruning vertices of degree 2 or 3 (see Fig. 21).

Algorithm 1. Successive Pruning

1. Input: a multigraph $G_1 = (V_1, E_1)$ such that $E_1 = B_1 \cup R_1$, where (V_1, B_1) and (V_1, R_1) are edge-disjoint spanning trees.
2. Put $i = 1$. While $|E_i| > 0$, do:
 - (a) pick an arbitrary vertex $v_i \in V_i$ which has degree 2 or 3 in G_i ;
 - (b) prune v_i from $G_i = (V_i, B_i \cup R_i)$ to obtain $G_{i+1} = (V_{i+1}, B_{i+1} \cup R_{i+1})$;
 - (c) put $i = i + 1$.

Edge flip at a vertex of degree 3. Consider again a vertex v_2 of degree 3 in G , incident to edges $e_1 = v_2u_1$, $e_2 = v_2u_2$, and $e_3 = v_2u_3$. Assume, without loss of generality, that $e_1, e_2 \in B$ and $e_3 \in R$. If we delete e_1 and e_2 , the blue tree breaks into three components, containing the vertices u_1 , u_2 and v_2 , respectively, where $\{v_2\}$ is a one-vertex component of $(V, B - e_1 - e_2)$. Suppose that u_3 is in the same component as u_2 . Then $(V, B - e_2 + e_3)$ and $(V, R - e_3 + e_2)$ are also two edge-disjoint spanning trees of G . An *edge flip* at v_2 is the operation that replaces the edge partition $E = B \cup R$ by $E = (E - e_2 + e_3) \cup (E - e_3 + e_2)$.

Contracting double edges. Suppose that $G = (V, E)$ contains two parallel edges, that is, edges $e_1, e_2 \in E$ with two common endpoints $u, v \in V$. Since no spanning tree has double edges, e_1 and e_2 are in distinct spanning trees, say $e_1 \in B$ and $e_2 \in R$. The *contraction of e_1 and e_2* is the operation of deleting edges e_1 and e_2 and identifying vertices u and v . We obtain a multigraph G' with $n - 1$ vertices and $2(n - 2)$ edges. Exactly one edge is contracted in each of (V, B) and (V, R) , hence G' is the edge-disjoint union of two spanning trees.

5.2 Conflict-free tree representations

Our proof of Lemma 3 is based on the concept of **conflict-free tree representation**, defined below. Recall that a *vertex splitting* in a (multi-)graph $G = (V, E)$ is the operation that replaces a vertex $v \in V$ with two vertices v_1 and v_2 , and replaces each edge $vu \in E$ incident to v with a new edge incident to either v_1 or v_2 (see Fig. 22(i)–(ii)). We assume that the replacement edge is identified with the original edge, and in particular it preserves the same conflicts.




Figure 22: (i) A graph G with an even number of edges. Conflicting pairs of edges are joined by small circular arcs. (ii) A conflict-free tree representation with a leaf at u after some vertex splitting steps. (iii) An even orientation on the tree. (iv) The induced even orientation of the original graph G .

Definition 4 Let $G = (V, E)$ be a multigraph and let $X \subseteq \binom{E}{2}$ be a collection of disjoint pairs of adjacent edges in E (called *conflict pairs*).

- A subset of edges $R \subseteq E$ has a conflict-free tree representation if, after some appropriate vertex splitting operations, the edges in R form a tree in which adjacent edges are not in conflict.
- Let $U \subseteq V$ be a subset of vertices. A subset of edges $R \subseteq E$ has a conflict-free tree representation with leaves at U if, after some appropriate vertex splitting operations, the edges in R form a tree in which adjacent edges are not in conflict, and all copies of each vertex $u \in U$ are leaves.

Proposition 18 Let $G = (V, E)$ be a multigraph and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. If $R \subseteq E$ has a conflict-free tree representation and $|R|$ is even, then (V, R) has an even orientation such that if the indegree of a vertex is 2, then the two inbound edges are not in conflict.

Proof. It is well-known that every connected graph with an even number of edges has an even orientation [21]. After some vertex splitting operations, (V, R) becomes a spanning tree (V', R') with no two adjacent edges in conflict. (as in Fig. 22). Fix an arbitrary even orientation for the tree (V', R') . Since all indegrees are even, for every directed edge $\vec{e} \in R'$, there is another edge $\vec{f} \in R'$ directed into the same vertex such that e and f are not in conflict. This orientation is an even orientation of (V, R) with the same property. Hence, if a vertex has in-degree 2, then the two inbound edges are not in conflict. \square

In the proof of Lemma 3, we find an even orientation for a multigraph $G = (V, E)$ which *contains* two edge-disjoint spanning trees. Our technique for constructing an even orientation focuses on the subgraph of G formed by two edge-disjoint spanning tree, and all other edges will have only supporting roles. To avoid conflicts between the extra edges and the edges of the two spanning trees, we choose two spanning trees with special properties, established in the following proposition.

Proposition 19 Let $G = (V, E)$ be a multigraph that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. There is a partition $E = B \cup R \cup Y$ such that (V, B) and

(V, R) are spanning trees, and if two parallel edges of E are in conflict, then either both of them are in $B \cup R$ or neither of them is in $B \cup R$. (See Fig. 23(i)–(ii).)

Proof. We proceed by induction on $n = |V|$. If $n = 1$, then both spanning trees are empty, and the claim trivially holds. Let $n > 1$ and assume that the claim holds for all smaller multigraphs. If no two parallel edges are in conflict, then any two edge-disjoint spanning trees have the required property. Suppose that edges $e_1, e_2 \in E$ are parallel and in conflict. Let $u, v \in V$ be the endpoints of e_1 and e_2 . Contract edges e_1 and e_2 (that is, identify u and v , and remove all edges between u and v), and denote by G' the resulting multi-graph on $n - 1$ vertices. Since G contains two edge-disjoint spanning trees, but no spanning tree has parallel edges, the resulting graph G' also contains two edge-disjoint spanning trees. By induction, G' contains two edge-disjoint spanning trees, say B' and R' , such that conflicting parallel edges are either both in $B' \cup R'$, or neither of them is in $B' \cup R'$. Let B'_0 and R'_0 be the subgraphs of G corresponding to B' and R' , respectively. Note that B'_0 and R'_0 each have two components such that u and v lie in distinct components. Now $B'_0 \cup \{e_1\}$ and $R'_0 \cup \{e_2\}$ are edge-disjoint spanning trees in G with the required property. \square




Figure 23: (i) The dual (multi-)graph G from Fig. 1(iii). It contains two edge-disjoint spanning trees: the dotted red tree and the dashed blue tree. Two additional edges are solid black. Conflicting pairs of edges are joined by small circular arcs. (ii) G also contains two edge-disjoint spanning trees, (V, B) and (V, R) , with the property that if two parallel edges are in conflict then either both of them are contained in $B \cup R$ or neither of them is in $B \cup R$. (iii) After some vertex splitting operations, we obtain a multigraph H in which $(V, B \cup R)$ is a subgraph, and all other edges have exactly one endpoint in V and are not in conflict with any adjacent edge.

Corollary 20 Let $G = (V, E)$ be a multigraph that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Then there is a sequence of vertex splitting operations on G that produces a multigraph $H = (V \cup W, B \cup R \cup Y)$ where where (V, B) and (V, R) are spanning trees on V , each edge in Y connects a vertex in V to a leaf in W , and no edge in Y is in conflict with any adjacent edge. (See Fig. 23(iii).)

Proof. Let $E = B \cup R \cup Y$ be the edge partition from Proposition 19. We now use vertex splitting operation to turn one endpoint of each edge $e \in Y$ into a leaf as follows. Recall that every edge $e \in E$ is in conflict with at most one other edge. Assign successively every edge $uv \in Y$ to either u or v such that (1) if $e = uv$ is in conflict with a parallel edge $f = uv$ (hence f is also in Y), then assign e and f to different endpoints; (2) if uv is in conflict with a nonparallel but adjacent edge vw , then assign uv to u (the vertex not incident to the conflicting edge); (3) if uv is not in conflict with any adjacent edge, then assign it to u or v arbitrarily. Apply a vertex splitting operation successively for every $e \in Y$: if an edge $e = uv \in Y$ is assigned to u , then split v into two vertices, one of which becomes a leaf incident to e only. \square

Edge partition algorithm. We next present an algorithm that partitions the edges of $G = (V, E)$ into nonempty subsets such that each subset has a conflict-free tree representation with a leaf at a common

vertex $v \in V$. In general we cannot hope for a single conflict-free tree representation for *all* edges in E . For example, if $G = (V, E)$ consists two parallel edges, and the two edges are in conflict, then E does not have any conflict-free tree representation.

Let $G = (V, E)$ be a multigraph on n vertices that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be set of disjoint conflict pairs. Let $H = (V \cup W, B \cup R \cup Y)$ be the multigraph from Corollary 20. We apply a sequence of vertex splitting operations on H to obtain a partition the edge set E .

The edge partition algorithm is guided by a successive pruning algorithm, with possible edge flips between consecutive prune steps. Put $G_1 = (V, B \cup R)$. Now G_1 is the edge-disjoint union of a *blue* spanning tree $(V_1, B_1) = (V, B)$ and a *red* spanning tree $(V_1, R_1) = (V, R)$. The successive pruning algorithm will produce a sequence of multigraphs $G = G_1, G_2, \dots, G_n = (\{v_n\}, \emptyset)$, where $G_i = (V_i, E_i)$ is an edge-disjoint union of two spanning trees: a *blue tree* (V_i, B_i) and a *red tree* (V_i, R_i) . Simultaneously with the pruning steps on G_i , we maintain a dynamic data structure that consists of a partition \mathcal{E}_i of E into subsets, each of which has a conflict-free tree representation. Each set in \mathcal{E}_i is associated with a vertex or an edge of G_i , as described below. When we prune a vertex v_i from G_i , the sets associated with v_i and its incident edges are “rearranged” to form new sets associated with the vertices and edges of G_{i+1} . Eventually, we have $G_n = (\{v_n\}, \emptyset)$, and all sets in \mathcal{E}_n are associated with v_n . The edge partition \mathcal{E}_n will be the output of the algorithm.

We maintain the following invariants for $i = 1, 2, \dots, n$. Each edge $e \in E_i$ between vertices $u, v \in V_i$ corresponds to a unique nonempty edge set $E_i(e) \in \mathcal{E}_i$, which has a conflict-free tree representation with leaves at both u and v . In the conflict-free tree representation of $E_i(e)$, two edges incident to the leaves u and v are called the *designated edges* of $E_i(e)$ (the two designated edge may coincide if $E(e) = \{e\}$). Any other set $A_i \in \mathcal{E}_i$ is associated with some vertex $v \in V_i$, such that A_i has a conflict-free tree representation with a leaf at v . The edge in A_i incident to the leaf v is called the *designated edge* of A_i . Several sets may be associated to the same vertex $v \in V_i$, or none at all. For every vertex $v \in V_i$, denote by $\mathcal{E}_i(v) \subset \mathcal{E}_i$ the edge sets associated with v .

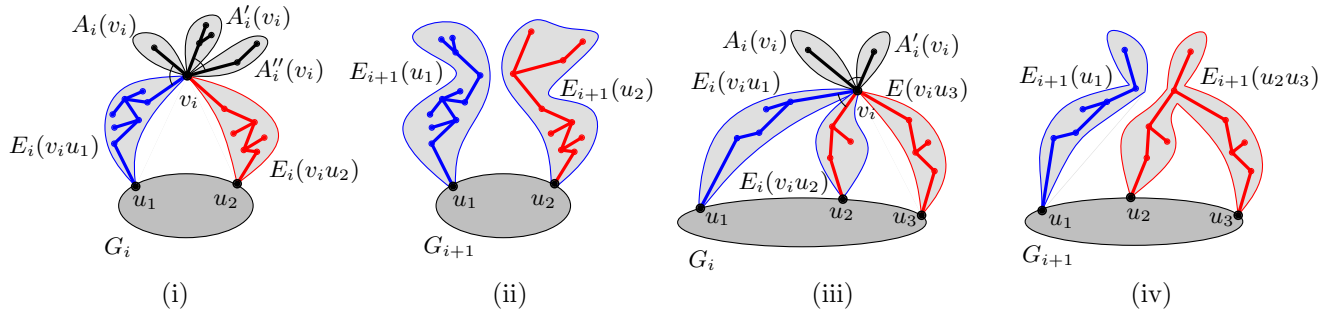


Figure 24: (i) Graph G_i with a vertex v_i of degree 2. (ii) Graph G_{i+1} after pruning v_i (iii) Graph G_i with a vertex v_i of degree 3. (iv) Graph G_{i+1} after pruning v_i .

Initialization and a general step i . For every edge $e \in B \cup R$, create a set $E_1(e) = \{e\}$ corresponding to $e \in E_1$. For every edge $e \in Y$, incident to $v \in V$ and a leaf $w \in W$, create a set $A_1(v)$, associated with vertex v . Each edge set in \mathcal{E}_1 contains a single edge, and so they each have a conflict-free tree representation in which each vertex is a leaf. Consider now one step of the successive pruning algorithm where we are given $G_i = (V_i, E_i)$ and $E_i = B_i \cup R_i$. Let $v_i \in V_i$ be a vertex of degree 2 or 3 in G_i that we wish to prune. We distinguish between two cases based on the degree of v_i .

Case 1: v_i has degree 2 in G_i . Refer to Fig. 24(i)–(ii). Assume without loss of generality that v_i is incident to the edges $v_i u_1 \in B_i$ and $v_i u_2 \in R_i$ (possibly $u_1 = u_2$). Vertex v_i and the edges $v_i u_1, v_i u_2$ are removed from G_i to obtain G_{i+1} . We need to take care of the associated edge sets $E_i(v_i u_1), E_i(v_i u_2)$ and all sets in $\mathcal{E}_i(v_i)$. Put $E_{i+1}(u_1) = E_i(v_i u_1)$; and then successively augment $E_{i+1}(u_1)$ with sets $A_i(v_i) \in \mathcal{E}_i(v_i)$ such that the designated edge of $A_i(v_i)$ is in conflict with neither the designated edge of $E_i(v_i u_1)$ at v_i , nor the designated edge of any other set $A'_i(v_i) \in \mathcal{E}_i(v_i)$ that has already been added to $E_{i+1}(u_1)$. Let $E_{i+1}(u_2)$ be the union of $E_i(v_i u_2)$ and all remaining sets from $\mathcal{E}_i(v_i)$ (whose designated edges are in conflict with neither the designated edge of $E_i(v_i u_2)$ incident to v_i , nor each other). Note that both $E_{i+1}(u_1)$ and $E_{i+1}(u_2)$ has a conflict-free tree representation, obtained by simply identifying the copies of vertex v_i in their conflict-free tree representations. We construct \mathcal{E}_{i+1} from \mathcal{E}_i by removing $E_i(v_i u_1), E_i(v_i u_2)$ and all sets in $\mathcal{E}_i(v_i)$; and adding $E_{i+1}(u_1)$ and $E_{i+1}(u_2)$. Let the designated edge of $E_{i+1}(u_1)$ (resp., $E_{i+1}(u_2)$) be the designated edge of $E_i(v_i u_1)$ (resp., $E_i(v_i u_2)$) incident to u_1 (resp., u_2).

Case 2: v_i has degree 3 in G_i . Refer to Fig. 24(iii)–(iv). Assume without loss of generality that v_i is incident to the edges $v_i u_1 \in B_i$ and $v_i u_2, v_i u_3 \in R_i$. If the designated edges of $E_i(v_i u_2)$ and $E_i(v_i u_3)$ are in conflict, then perform an edge flip at v_i in the graph G_i (see G_2 in Fig. 25, for an example). We may now assume that v_i is incident to the edges $v_i u_1 \in B_i$ and $v_i u_2, v_i u_3 \in R_i$; and the designated edges of $E_i(v_i u_2)$ and $E_i(v_i u_3)$ are not in conflict. Put $E_{i+1}(u_1) = E_i(v_i u_1)$; and then successively augment $E_{i+1}(u_1)$ with sets $A_i(v_i) \in \mathcal{E}_i(v_i)$ such that the designated edge of $A_i(v_i)$ is in conflict with neither the designated edge of $E_i(v_i u_1)$ at v_i , nor the designated edge of any other set $A'_i(v_i) \in \mathcal{E}_i(v_i)$ that has already been added to $E_{i+1}(u_1)$. Let $E_{i+1}(u_2 u_3)$ be the union of $E_i(v_i u_2), E_i(v_i u_3)$, and all remaining sets from $\mathcal{E}_i(v_i)$ (whose designated edges are in conflict with neither the designated edges of $E_i(v_i u_2)$ and $E_i(v_i u_3)$ incident to v_i , nor each other). Now $E_{i+1}(u_1)$ and $E_{i+1}(u_2 u_3)$ each have a conflict-free tree representation. We construct \mathcal{E}_{i+1} from \mathcal{E}_i by removing $E_i(v_i u_1), E_i(v_i u_2), E_i(v_i u_3)$ and all sets in $\mathcal{E}_i(v_i)$; and adding $E_{i+1}(u_1)$ and $E_{i+1}(u_2 u_3)$. Let the designated edge of $E_{i+1}(u_1)$ be the designated edge of $E_i(v_i u_1)$ incident to u_1 ; and let the designated edges of $E_{i+1}(u_2 u_3)$ be the designated edges of $E_i(v_i u_2)$ and $E_i(v_i u_3)$ incident to u_2 and u_3 , respectively.

After $n - 1$ prune steps, we obtain $G_n = (\{v_n\}, \emptyset)$. That is, G_n has only one vertex, and no edges. It follows that $\mathcal{E}_n = \mathcal{E}_n(v_n)$. This completes the description of our edge partition algorithms. See Fig. 25 for an example.

As noted above, the edge partition algorithm is guided by the successive pruning algorithm. Since G_i , for $i = 1, \dots, n - 1$, has at least two vertices of degree 2 or 3, and we always have at least two possible choices for the vertex $v_i \in V_i$ we prune in step i . These choices are crucial for the proof of Lemma 3. We now point out several important properties of the edge partition algorithm, which hold independently of the choices for the vertices v_1, \dots, v_{n-1} .

Proposition 21 *Let $G = (V, E)$ be a multigraph on n vertices that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Then the above algorithm returns: (1) a vertex $v_n \in V$, (2) a partition of E into a collection \mathcal{E}_n of nonempty edge sets, and (3) a conflict-free tree representation with a leaf at v_n for each set $A \in \mathcal{E}_n$. A vertex $v \in V$ may have several copies (due to vertex splitting) over the conflict-free tree representations of the sets in \mathcal{E}_n , but at most two copies of each $v \in V$ has degree 2 or higher.*

Proof. The first claim follows directly from the invariants maintained for the sets in $E_i, i = 1, 2, \dots, n$. For the second claim, notice that $H = (V \cup W, B \cup R \cup Y)$ (cf. Corollary 20) was obtained from $G = (V, E)$ by a sequence of vertex splitting operations, where exactly one copy of each vertex has degree 2 or higher. In the remainder of the algorithm, every vertex v_i is split only once, when pruning v_i . Hence at most two

copies of each vertex $v \in V$ has degree 2 or higher over the conflict-free tree representations of all sets in \mathcal{E}_n . \square

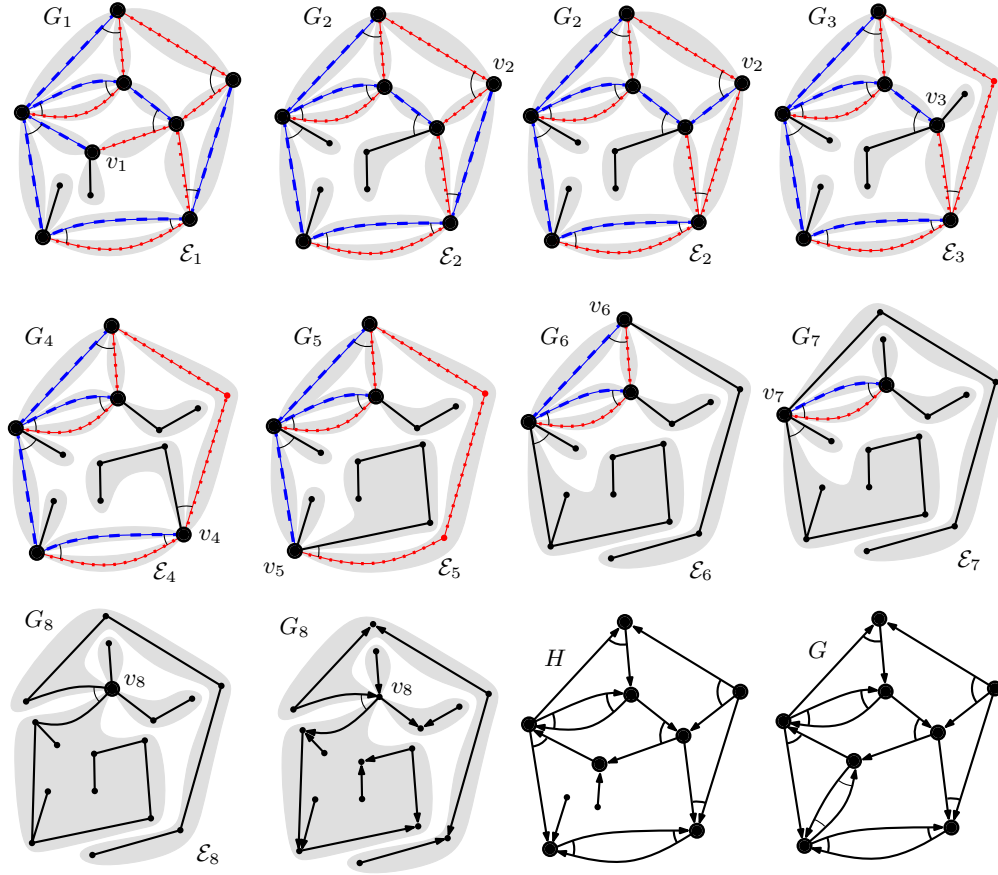


Figure 25: A sequence of graphs G_1, \dots, G_8 is obtained by successively pruning vertices v_1, \dots, v_7 . Each G_i , $i = 1, \dots, 8$, is the edge-disjoint union of a dotted red spanning tree and a dashed blue spanning tree. Conflicts are indicated by small circular arcs. Before v_2 is pruned, an edge flip is applied at v_2 to ensure that the two incident edges of the same color are not in conflict. Each G_i , $i = 1, \dots, 8$, corresponds to an edge partition \mathcal{E}_i , where each set has a conflict-free tree representation, as shown in the figure. There are four sets in \mathcal{E}_8 , each of which contains a designated edge incident to v_8 . Each even set has a conflict-free even orientation; and each odd set has an orientation where the indegree of v_8 is odd. In this instance, the two edges oriented into v_8 are not in conflict, hence we obtain a conflict-free even orientations for H and hence for G .

If all edge sets in \mathcal{E}_n have even size, then they each have an even orientation. In this case, by Proposition 18, graph G has an even orientation such that whenever the indegree of a vertex is 2, then the two incoming edges are not in conflict. If some sets in \mathcal{E}_n have odd size, we can still construct a desired orientation for G under certain conditions.

Proposition 22 Let $G = (V, E)$ be a multigraph with n vertices and an even number of edges that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Suppose that there is a vertex v_0 and a partition of E into a collection \mathcal{E} of subsets with the following properties:

- every even set $A \in \mathcal{E}$ has a conflict-free tree representation;
- every odd set $A \in \mathcal{E}$ has a conflict-free tree representation with a leaf at v_0 .
- if \mathcal{E} contains exactly two odd sets, then their edges incident to v_0 are not in conflict.

Then G has an even orientation such that whenever the indegree of a vertex is 2, then the two incoming edges are not in conflict.

Proof. We show that E can be partitioned into *even* sets each of which has a conflict-free tree representation. Every such edge set has a conflict-free even orientation and so Proposition 18 completes the proof. If \mathcal{E} contains no odd sets, then our proof is complete. Since $|E|$ is even, then number of odd sets in \mathcal{E} is even. If \mathcal{E} contains exactly two odd sets but their edges incident to v_0 are not in conflict, then the union of the two odd sets is even, and it has a conflict-free tree representation. If \mathcal{E} contains four or more odd sets, then they can be paired up so that the edges incident to v_0 in each pair are not in conflict. The union of each pair has even cardinality and has a conflict-free tree representation. \square

We can now establish Lemma 3 in several important special cases.

Corollary 23 *Let $G = (V, E)$ be a multigraph with n vertices and an even number of edges that contains two edge-disjoint spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Assume that*

- G contains more than $2n - 2$ edges, or
- G has a vertex such that no two incident edges are in conflict.

Then G has an even orientation such that whenever the indegree of a vertex is 2, then the two incoming edges are not in conflict.

Proof. First assume that G has more than $2n - 2$ edges. Consider the multigraph $H = (V \cup W, B \cup R \cup Y)$ from Corollary 20. Since B and R are spanning trees, we have $|B| = |R| = n - 1$ and so $Y \neq \emptyset$. Let $v_0 w_0 = e_0$ be an arbitrary edge in Y with $v_0 \in V$ and $w_0 \in W$. Recall that e_0 is not in conflict with any adjacent edge in H . Perform the edge partition algorithm such that vertex v_0 is never pruned (that is $v_i \neq v_0$ for $i = 1, \dots, n - 1$). This is possible, since G_i has at least two vertices of degree 2 or 3 for $i = 1, \dots, n - 1$. The algorithm returns an edge partition $\mathcal{E}_n = \mathcal{E}_n(v_0)$, with $\{e_0\} \in \mathcal{E}_n$. If \mathcal{E}_n contains exactly two odd sets, then one of them is $\{e_0\}$, and so the designated edges of these two sets are not in conflict. Proposition 22 completes the proof.

Assume now that G has a vertex v_0 such that no two incident edges are in conflict. Perform the edge partition algorithm such that vertex v_0 is never pruned (that is $v_i \neq v_0$ for $i = 1, \dots, n - 1$). This is possible, since G_i has at least two vertices of degree 2 or 3 for $i = 1, \dots, n - 1$. The algorithm returns an edge partition $\mathcal{E}_n = \mathcal{E}_n(v_0)$. The designated edges of no two sets in \mathcal{E}_n are in conflict, and Proposition 22 completes the proof. \square

5.3 Contracting double edges in conflict

By Corollary 23, we may assume that $G = (V, E)$ is the edge-disjoint union of two spanning trees. Then G has $|V| = n$ vertices and $|E| = 2n - 2$ edges, and so there are at most $n - 1$ conflict pairs in X . If any two conflicting edges have at most one common endpoint, then there is a vertex $v_0 \in V$ such that no two incident edges are in conflict, and Corollary 23 yields a desired even orientation for G . Therefore, we may assume that G contains at least one pair of parallel edges in conflict.

Let $G = (V, E)$ be a multigraph on n vertices which is the edge-disjoint union of two spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Successively contract any pair of parallel edges that are in conflict, and denote the resulting graph by $\widehat{G} = (\widehat{V}, \widehat{E})$. (Refer to Fig. 26(i)–(ii)). Note that $\widehat{E} \subset E$. Each vertex $\hat{v} \in \widehat{V}$ is the result of identifying a set $U_{\hat{v}} \subseteq V$ of vertices. Let $G[\hat{v}]$ denote the subgraph of G induced by $U_{\hat{v}}$. By construction, $G[\hat{v}]$ is the edge-disjoint union of two spanning trees, a red tree $(U_{\hat{v}}, R_{\hat{v}})$




Figure 26: (i) G is the edge-disjoint union of two spanning trees: a dotted red tree and a dashed blue tree. Edge pairs in conflict are marked with small circular arcs. (ii) \widehat{G} is obtained from G by successively contracting parallel edges in conflict. Odd (even) vertices are marked with empty (full) dots. (iii) \widehat{H} is obtained by adding two leaves in conflict to each odd vertex of \widehat{G} . (iv) An even orientation for \widehat{H} such that whenever the indegree of a vertex is 2, the two incoming edges are not in conflict.

and a blue tree $(U_{\hat{v}}, B_{\hat{v}})$, each of which is conflict-free. We say that a vertex $\hat{v} \in \widehat{V}$ is *odd* if both $R_{\hat{v}}$ and $B_{\hat{v}}$ are odd; and $\hat{v} \in \widehat{V}$ is *even* if both $R_{\hat{v}}$ and $B_{\hat{v}}$ are even.

We shall construct an even orientation for \widehat{G} such that (1) the indegree of every *odd* vertex $\hat{v} \in \widehat{V}$ is at least 2; and (2) if the indegree of an even vertex $\hat{v} \in \widehat{V}$ is 2, then the two incoming edges are not in conflict. See Fig. 26(iii)–(iv). The minimum degree requirement is enforced by augmenting \widehat{G} with pairs of leaves in conflict: For each odd vertex \hat{v} , introduce two new vertices \hat{w}_1, \hat{w}_2 , two new edges $\hat{v}\hat{w}_1, \hat{v}\hat{w}_2$, and a new conflict pair $\{\hat{v}\hat{w}_1, \hat{v}\hat{w}_2\}$. Denote the augmented multigraph by $\widehat{H} = (\widehat{V} \cup \widehat{W}, \widehat{E} \cup \widehat{S})$, and the augmented conflict pairs by \widehat{X} . We refer to the new edges in \widehat{H} as *special edges*. Consider an even orientation of \widehat{H} such that if the indegree of a vertex is 2, then the two incoming edges are not in conflict. At each odd vertex $\hat{v} \in \widehat{V}$, the two special edges must be oriented into \hat{v} . However, the two special edges are in conflict, so at least two edges of \mathcal{E} must also be oriented into \hat{v} .

Lemma 24 *Let $G = (V, E)$ be a multigraph on n vertices which is the edge-disjoint union of two spanning trees and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Assume that not all vertices of \widehat{G} are odd. Then G has an even orientation such that whenever the indegree of a vertex is 2, then the two incoming edges are not in conflict.*

Proof. Compute \widehat{H} and \widehat{X} as described above. We perform the edge partition algorithm on \widehat{H} , with certain additional constraints to be specified below. The algorithm maintains the edge partitions $\mathcal{E}_1, \dots, \mathcal{E}_n$ for the edges of \widehat{H} (that is, including the special edges). Recall that the algorithm is guided by a sequence of multigraphs $\widehat{G} = G_1, G_2, \dots, G_n$. Perform the edge partition algorithm on \widehat{H} until the first step ℓ , $1 \leq \ell < n$, where there exists an *even* vertex $\hat{x} \in V_\ell$ of degree 2 or 3. For all steps $i \leq \ell$, the following invariants are automatically maintained by the algorithm:

- For every edge $uv \in E_i$, the edge set $E_i(uv) \in \mathcal{E}_i$ has odd cardinality
- Every set $A_i(v) \in \mathcal{E}_i(v)$ for every $v \in V_i$, either consists of a single special edge or has even cardinality.

Both invariants clearly hold for $i = 1$, where every set $E_1(uv)$ associated with an edge $uv \in E_1$ contains exactly one edge; and every $A_1(v) \in \mathcal{E}_1(v)$ associated with a vertex $v \in V_1$ contains exactly one edge, which is a special edge. For maintaining the invariants, consider cases 1 and 2 of the edge partition algorithm. Note that in every step $i \leq \ell$, the vertex v_i is odd, and so it is incident to two special edges (leaves) in conflict. In case 1, two new edge sets are created, $E_{i+1}(u_1)$ and $E_{i+1}(u_2)$, each of which is associated with a vertex. They are each formed by the union of a set $E_i(v_i u_1)$ or $E_i(v_i u_2)$ associated with an edge, one special edge

of v_i , and some other sets associated with vertex v_i . Due to the invariants, both $E_{i+1}(u_1)$ and $E_{i+1}(u_2)$ are even. In case 2, two new edge sets are created: $E_{i+1}(u_1)$ associated with vertex u_1 , and $E_{i+1}(u_2u_3)$ associated with edge $u_2u_3 \in E_{i+1}$. $E_{i+1}(u_1)$ is the union of $E_i(v_iu_1)$, one special edge of v_i , and some other sets associated with vertex v_i ; and so it is even. $E_{i+1}(u_2u_3)$ is the union of $E_i(v_iu_2)$, $E_i(v_iu_3)$, one special edge of v_i , and some other sets associated with vertex v_i ; and so it is odd. We have established that both invariants hold for the edge partition \mathcal{E}_ℓ .

Now consider step ℓ of the algorithm, where $\hat{x} \in V_\ell$ is an even vertex of degree 2 or 3. We distinguish between two possible scenarios.

Case A: For any two edges of G_ℓ incident to \hat{x} , the two designated edges incident to \hat{x} are not in conflict. In this case, complete the edge partition algorithm so that \hat{x} is never pruned (that is, $v_i \neq \hat{x}$ for $i = \ell, \dots, n-1$). The algorithm returns the edge partition $\mathcal{E}_n = \mathcal{E}_n(\hat{x})$. Note that $\mathcal{E}_n(\hat{x})$ contains all sets of $\mathcal{E}_\ell(\hat{x})$, which are even by invariants 2. The designated edges of the other sets in $\mathcal{E}_n(\hat{x})$ are not in conflict with each other.

Based on the edge partition \mathcal{E}_n of \hat{H} , we now construct an edge partition \mathcal{E} for the original graph $G = (V, E)$. The edge partition algorithm for \hat{H} has split each vertex $\hat{v} \in \hat{V}$, $\hat{v} \neq \hat{x}$, into exactly two copies (in the step when \hat{v} is pruned). If \hat{v} is an odd vertex, then its two special edges are incident to distinct copies of \hat{v} . Modify the sets in $\mathcal{E}_n(\hat{x})$ as follows: replace the two copies of each vertex $\hat{v} \in \hat{V}$, $\hat{v} \neq \hat{x}$, by the spanning trees $(U_{\hat{v}}, R_{\hat{v}})$ and $(U_{\hat{v}}, B_{\hat{v}})$; and remove any special edges incident to \hat{v} (in case \hat{v} is odd). Denote the resulting collection of sets by \mathcal{E}' . Note that this operation preserves the parity of every set in \mathcal{E}_n . Each edge set in \mathcal{E}' has a conflict-free tree representation with a leaf in $U_{\hat{x}}$. So far, we have not replaced \hat{x} with the two corresponding even trees $(U_{\hat{x}}, R_{\hat{x}})$ and $(U_{\hat{x}}, B_{\hat{x}})$. We can now define the edge partition \mathcal{E} of G . Let \mathcal{E} contain all even sets of \mathcal{E}' , the set $B_{\hat{x}}$, and the union of $R_{\hat{x}}$ with all odd sets of \mathcal{E}' . Every set in \mathcal{E} is even, and has a conflict-free tree representation. Proposition 22 completes the proof.




Figure 27: (i) \hat{H} from Figure 26(iii) is the edge-disjoint union of two spanning trees: a dotted red tree and a dashed blue tree. Edge pairs in conflict are marked with small circular arcs. (ii) The edge partition algorithm returns sets \mathcal{E}_4 , each with a conflict-free tree representation, and each incident to a common vertex \hat{v}_4 . Vertices $\hat{v}_1, \hat{v}_2, \hat{v}_3$ have been each split into two vertices. (iii) Each copy of $\hat{v}_1, \hat{v}_2, \hat{v}_3$ is replaced by a conflict-free copy of $G[\hat{v}_i]$. The resulting graph now has an even orientation. (iv) This gives a required even orientation for graph G from Figure 26(i).

Case B: There are two edges in E_i incident to \hat{x} such that their designated edges incident to \hat{x} are in conflict. See Fig. 27. In this case, prune \hat{x} in step ℓ , that is put $v_\ell = x$. Vertex $v_\ell = \hat{x}$ is incident to 2 or 3 edges in G_ℓ , say $v_\ell u_1, v_\ell u_2$, and possibly $v_\ell u_3$. We may assume without loss of generality that the designated edges of $v_\ell u_1$ and $v_\ell u_2$ are in conflict. Let $\hat{u} = u_1$, and complete the edge partition algorithm so that \hat{u} is never pruned, that is, $v_i \neq \hat{u}$ for $i = \ell+1, \dots, n-1$. Invariant 1 implies that $E_\ell(\hat{x}\hat{u})$ is odd; and invariant 2 implies that every set in $\mathcal{E}_\ell(\hat{u})$ is even. In step ℓ , the union of $\mathcal{E}_\ell(\hat{x}\hat{u})$ and some of the even sets in $\mathcal{E}_\ell(\hat{u})$ forms an odd set, say $A(\hat{u}) \in \mathcal{E}_{\ell+1}(\hat{u})$. This set $A(\hat{u})$ remains in the edge partition in steps $\ell+1, \dots, n-1$ of the algorithm, and so $A(\hat{u}) \in \mathcal{E}_n = \mathcal{E}_n(\hat{u})$.

Based on the edge partition \mathcal{E}_n of \hat{H} , we now construct an edge partition \mathcal{E} for the original graph $G = (V, E)$. Similarly to Case A, replace the two copies of each vertex $\hat{v} \in \hat{V}$, $\hat{v} \neq \hat{u}$, by the spanning trees

$(U_{\hat{v}}, R_{\hat{v}})$ and $(U_{\hat{v}}, B_{\hat{v}})$; and remove any special edges incident to \hat{v} (in case \hat{v} is odd). Denote the resulting collection of sets by \mathcal{E}' , and let $A \in \mathcal{E}$ be the set that contains $A(\hat{u})$. Since this operation preserves the parity of all edge sets in \mathcal{E}_n , the set A is odd. Each edge set in \mathcal{E}' has a conflict-free tree representation with a leaf in $U_{\hat{u}}$. Let $v_0 \in U_{\hat{u}}$ be the designated leaf of $A \in \mathcal{E}'$. Partition \mathcal{E}' into three sets $\mathcal{E}' = \mathcal{E}^0 \cup \mathcal{E}^- \cup \mathcal{E}^+$ as follows: let \mathcal{E}^0 contain every set in \mathcal{E}' with a leaf at v_0 ; the remaining sets in \mathcal{E}' are partitioned between the collections \mathcal{E}^- and \mathcal{E}^+ such that if their leaves incident to $U_{\hat{u}}$ are in conflict then they are in different collections. We can now define the edge partition \mathcal{E} of G . Let \mathcal{E} contain all sets of \mathcal{E}^0 , the union of $B_{\hat{u}}$ and all sets in \mathcal{E}^- , and the union of $R_{\hat{x}}$ and all sets in \mathcal{E}^+ . Every set in \mathcal{E} has a conflict-free tree representation with a leaf at v_0 . Proposition 22 completes the proof unless \mathcal{E} contains exactly two odd edge sets and their edges incident to v_0 are in conflict.

Suppose now that \mathcal{E} contains exactly two odd edge sets, and their edges incident to v_0 are in conflict. By construction, one of the two odd edge sets in \mathcal{E} is A . Recall that $E_\ell(\hat{x}\hat{u}) \subset A$, and the edge in A incident to $v_0 \in U_{\hat{u}}$ is the designated edge of $E_\ell(\hat{x}\hat{u})$ incident to \hat{u} . Consider step ℓ of the edge partition algorithm above. In Case B, we assumed that the designated edge of $E_\ell(\hat{x}\hat{u})$ is in conflict with some other edge incident to \hat{x} . Since \widehat{G} has no two parallel edges in conflict, the edge set $E_\ell(\hat{x}\hat{u})$ has two distinct designated edges, incident to \hat{x} and \hat{u} , respectively. In particular, $E_\ell(\hat{x}\hat{u})$ contains at least two edges. Let $j < \ell$ be the largest index such that some edge in $E_\ell(\hat{x}\hat{u})$ is incident to v_j . Since $j < \ell$, v_j is an odd vertex. In step j , vertex v_j was pruned. That is, v_j was split into two copies: one is incident to $E_\ell(\hat{u}\hat{x})$, and the other is incident to some component $B_{j+1} \in \mathcal{E}_{j+1}$, which becomes part of some set $B \in \mathcal{E}$, $B \neq A$. Merge A into B by splitting \hat{u} into two copies and identifying the two copies of v_j (effectively, reversing the vertex split operation performed in step j). Proposition 22 completes the proof. \square

We are now ready to prove Lemma 3.

Lemma 3. *Let $G = (V, E)$ be a multigraph with $4k$ edges that contains two edge-disjoint spanning trees, and let $X \subseteq \binom{E}{2}$ be a collection of disjoint conflict pairs. Then G has an even orientation such that if the indegree of a vertex is 2, then the two incoming edges are not in conflict.*

Proof. Let $n = |V|$ be the number of vertices in G . If $4k > 2n - 2$, then Corollary 23 completes the proof. Suppose that $4k = 2n - 2$, and so G is the edge-disjoint union of two spanning trees. Compute the multigraph $\widehat{G} = (\widehat{V}, \widehat{E})$. Since \widehat{G} is also the edge-disjoint union of two spanning trees, it has $|\widehat{E}| = 2|\widehat{V}| - 2$ edges. If a vertex $\hat{v} \in \widehat{V}$ is odd, then the subgraph $G[\hat{v}]$ of G has $4k_{\hat{v}} - 2$ edges for some $k_{\hat{v}} \in \mathbb{N}$. If all vertices $\hat{v} \in \widehat{V}$ are odd, then the total number of edges in G is

$$|E| = |\widehat{E}| + \sum_{\hat{v} \in \widehat{V}} (4k_{\hat{v}} - 2) = (2|\widehat{V}| - 2) - 2|\widehat{V}| + 4 \sum_{\hat{v} \in \widehat{V}} k_{\hat{v}} = -2 + 4 \sum_{\hat{v} \in \widehat{V}} k_{\hat{v}},$$

which is not a multiple of 4. Since G has $4k$ edges, not all vertices of \widehat{G} are odd, and Lemma 24 completes the proof. \square

5.4 Disjoint Compatible Matchings

We recall a result by Aichholzer *et al.* [1] [Lemma 2] for points in convex position.

Proposition 25 [1] *Let P be a set of points in convex position. Let M be a (not necessarily perfect) matching on P such that every segment in M is on the boundary of the convex hull of P . Then there is a perfect matching of P that is disjoint and compatible with M if and only if $|P|$ is even and $|P| = 2$ implies $M = \emptyset$.*

We can now prove Theorem 1 and settle the Disjoint Compatible Matching Conjecture.

Theorem 1. *For every perfect straight-line matching M with an even number of edges, and no three collinear vertices, there is a disjoint compatible perfect straight-line matching.*

Proof. Let M be a set of disjoint line segments in the plane, with no three collinear segment endpoints. Refer to Fig. 1 and 2. By Theorem 2, there is a convex subdivision $C \in \mathcal{D}(M)$ with no critical polygons. In the dual graph, the nodes correspond to cells in C , and the edges correspond to segment endpoints. Two edges of the dual graph are *in conflict* if they are adjacent and correspond to the two endpoints of the same segment. By Lemma 3, the dual graph has an even orientation such that if a vertex has indegree 2, then the two incoming edges are not in conflict. Fix such an orientation for the remainder of the proof.

For every cell $c \in C$, assign incident segment endpoints to c if the corresponding edge is oriented into c . In every cell $c \in C$, independently, we construct a perfect matching on the segment endpoints assigned to c . Since the orientation is even, every cell is assigned to an even number of segment endpoints. Furthermore, if a cell is assigned to exactly two segment endpoints, then these endpoints are not connected by a segment in M . By Proposition 25, there is a perfect matching on the segment endpoints assigned to each c . The union of these perfect matchings is a perfect matching on all segment endpoints. It is disjoint from M and compatible with M , as required. \square

Acknowledgements. The Disjoint Compatible Matchings Conjecture was posed at the *U.P.C. Workshop on Combinatorial Geometry (Caldes de Malavella, Catalunya, Spain, May 2006)*. The third author heard of the problem independently, started working on it with a collaborator, and raised the problem at a meeting of the *MIT-Tufts Combinatorial Geometry Research Group* in 2007. We thank all our collaborators in papers [1], [2], [5], and [8].

References

- [1] O. Aichholzer, S. Bereg, A. Dumitrescu, A. García, C. Huemer, F. Hurtado, M. Kano, A. Márquez, D. Rappaport, S. Smorodinsky, D. Souvaine, J. Urrutia, D. Wood, Compatible geometric matchings, *Comput. Geom.* **42** (2009), 617–626.
- [2] M. Al-Jubeih, M. Ishaque, M. Hoffmann, D. L. Souvaine, and C. D. Tóth, Convex partitions with 2-edge connected dual graphs, *J. Combin. Opt.* **22** (3) (2010), 409–425.
- [3] Marwan Al-Jubeih, Gill Barequet, Mashhood Ishaque, Diane L. Souvaine, Csaba D. Tth, and Andrew Winslow, Constrained tri-connected planar straight line graphs, in *Thirty Essays on Geometric Graph Theory (J. Pach, ed.)*, vol. 29 of Algorithms and Combinatorics, Springer, 2012, to appear.
- [4] T. F. Banchoff, Global geometry of polygons I: The theorem of Fabricius-Bjerre, *Proc. AMS* **45** (1974), 237–241.
- [5] N. M. Benbernou, E. D. Demaine, M. L. Demaine, M. Hoffmann, M. Ishaque, D. L. Souvaine, and C. D. Tóth, Disjoint segments have a convex partition with a 2-edge connected dual graph, in *Proc. 19th Canadian Conf. Comp. Geom.*, Ottawa, ON, 2007, pp. 13–16. Erratum in *Proc. 20th Canadian Conf. Comp. Geom.*, Montreal, QC, 2008, p. 223.
- [6] P. Bose, M. E. Houle, and G. T. Toussaint, Every set of disjoint line segments admits a binary tree, *Discrete Comput. Geom.* **26** (3) (2001), 387–410.
- [7] K. Buchin, C. Knauer, K. Kriegel, A. Schulz, and R. Seidel, On the number of cycles in planar graphs, in *Proc. 13th COCOON*, vol. 4598 of LNCS, Springer, 2007, pp. 97–107.
- [8] S. Cannon, M. Ishaque, and C. D. Tóth, Even orientations with forbidden pairs and demands, in *Proc. 6th Conf. on Fun with Algorithms*, vol. 7288 of LNCS, Springer, 2012, to appear.

- [9] B. Chazelle and D. P. Dobkin, Optimal convex decompositions, *Comput. Geom.* **2** (1985), 63–133.
- [10] R. Cordovil and M. L. Moreira, Bases-cobases graphs and polytopes of matroids, *Combinatorica* **13** (2) (1993), 157–165.
- [11] M. Farber, B. Richter, and H. Shank, Edge-disjoint spanning trees: a connectedness theorem, *J. Graph Theory* **8** (1985), 319–324.
- [12] A. Frank, T. Jordán, and Z. Szigeti, An orientation theorem with parity conditions, *Discrete Appl. Math.* **115** (1-3) (2001), 37–47.
- [13] A. Frank and Z. Király, Graph orientations with edge-connection and parity constraints *Combinatorica* **22** (1) (2002), 47–70.
- [14] A. García, M. Noy, and J. Tejel, Lower bounds on the number of crossing-free subgraphs of K_N , *Comput. Geom. Theory Appl.* **16** (2000), 211–221.
- [15] S. K. Ghosh, *Visibility Algorithms in the Plane*, Cambridge Univ. Press, 2007.
- [16] M. Hoffmann, M. Sharir, A. Sheffer, Cs. D. Tóth, and E. Welzl, Counting plane graphs: flippability and its applications, in *Proc. WADS*, vol. 6844 of LNCS, Springer, 2011, pp. 524–535.
- [17] M. Hoffmann, B. Speckmann, and C. D. Tóth, Pointed binary encompassing trees: simple and optimal, *Comput. Geom. Theory Appl.* **43** (1) (2010), 35–41.
- [18] F. Hurtado and C. D. Tóth, Plane geometric graph augmentation: a generic perspective, in *Thirty Essays on Geometric Graph Theory (J. Pach, ed.)*, vol. 29 of Algorithms and Combinatorics, Springer, 2012, to appear.
- [19] M. Keil, Polygon decomposition, in: *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds. Elsevier, Amsterdam, 2000, pp. 491–518.
- [20] M. Keil and J. Snoeyink, On the time bound for convex decomposition of simple polygons, *Internat. J. Comput. Geom. Appl.* **12** (2002), 181–192.
- [21] L. Lovász and M. D. Plummer, *Matching Theory*, AMS Chelsea, 2009.
- [22] C. St.-J. A. Nash-Williams, Edge-disjoint spanning trees in finite graphs, *J. London Math. Soc.* **36** (1961), 445–450.
- [23] J. O’Rourke, *Art Gallery Theorems and Algorithms*, Oxford Univ. Press, 1987.
- [24] A. Razen, A lower bound for the transformation of compatible perfect matchings, in: *Abstracts of the 24th European Workshop on Comput. Geom.*, 2008, pp. 115–118.
- [25] M. Sharir, A. Sheffer, and E. Welzl, Counting plane graphs: perfect matchings, spanning cycles, and Kasteleyn’s technique, in *Proc. 28th Sympos. Comput. Geom.*, ACM Press, 2012, to appear. See also [arXiv:1109.5596v1](https://arxiv.org/abs/1109.5596v1)
- [26] M. Sharir and E. Welzl, On the number of crossing-free matchings, cycles, and partitions, *SIAM J. Comput.* **36** (2006), 695–720.
- [27] C. D. Tóth, Illuminating both sides of line segments, in: *Discrete and Computational Geometry (J. Akiyama, M. Kano, M. Urabe, eds.)*, vol. 2098 of LNCS, Springer, Berlin, 2001, pp. 370–380.
- [28] C. D. Tóth, Illumination in the presence of opaque line segments in the plane, *Comput. Geom. Theory Appl.* **21** (3) (2002), 193–204.
- [29] W. T. Tutte, On the problem of decomposing a graph into n connected factors, *J. London Math. Soc.* **36** (1961), 221–230.