

Planar Matchings for Weighted Straight Skeletons*

Therese Biedl[†]

*David R. Cheriton School of Computer Science
University of Waterloo, Waterloo
Ontario N2L 1A2, Canada
biedl@uwaterloo.ca*

Stefan Huber

*Institute of Science and Technology Austria
3400 Klosterneuburg, Austria
stefan.huber@ist.ac.at*

Peter Palfrader[‡]

*Universität Salzburg, FB Computerwissenschaften
5020 Salzburg, Austria
palfrader@cosy.sbg.ac.at*

Received 6 February 2015

Revised 18 January 2017

Published 13 April 2017

Communicated by Christian Knauer, Guest Editor

We introduce planar matchings on directed pseudo-line arrangements, which yield a planar set of pseudo-line segments such that only matching-partners are adjacent. By translating the planar matching problem into a corresponding stable roommates problem we show that such matchings always exist.

Using our new framework, we establish, for the first time, a complete, rigorous definition of weighted straight skeletons, which are based on a so-called wavefront propagation process. We present a generalized and unified approach to treat structural changes in the wavefront that focuses on the restoration of weak planarity by finding planar matchings.

Keywords: Planar matchings; pseudo-line arrangements; stable roommates; weighted straight skeletons.

*A preliminary version appeared at ISAAC'14.

[†]Supported by NSERC and the Ross and Muriel Cheriton Fellowship.

[‡]Research supported by Austrian Science Fund (FWF): P25816-N15.

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

1. Introduction

The straight skeleton is a skeletal structure of a polygon P , similar to the Voronoi diagram. It was introduced to computational geometry almost two decades ago by Aichholzer *et al.*,¹ and its definition is based on a so-called wavefront propagation process, see Fig. 1: Each edge of P emits a wavefront edge that moves towards the interior of P at unit speed in a self-parallel manner. The polygons formed by these wavefront edges at any given time $t \geq 0$ are the wavefront, denoted by $\mathcal{W}_P(t)$, and take the form of a mitered offset of P . Over time, the wavefront undergoes two different kinds of topological changes, so-called events, due to self-interference: roughly speaking, an *edge event* happens when a wavefront edge collapses, and a *split event* happens when the wavefront splits into parts. The straight skeleton $\mathcal{S}(P)$ of P is then defined as the geometric graph whose edges are the traces of the vertices of \mathcal{W}_P . Similar to Voronoi diagrams and the medial axis, straight skeletons have become a versatile tool for applications in various domains of science and industry.²

The *weighted straight skeleton*, where wavefront edges do not necessarily move at unit speed, was first mentioned by Aichholzer and Aurenhammer³ and Eppstein and Erickson⁴ and has since been used in a variety of different applications.^{5–8} Weighted straight skeletons, with both positive and negative weights, also constitute a theoretical tool to generalize straight skeletons to 3D.⁹ Even though weighted straight skeletons have already been applied in both theory and practice, their properties have been investigated only recently. Biedl *et al.*¹⁰ showed that basic properties of unweighted straight skeletons do not carry over to weighted straight skeletons in general. They also proposed solutions for an ambiguity in the definition of straight skeletons caused by certain edge events. This ambiguity was first mentioned by Kelly and Wonka⁷ and Huber.²

In this paper, we discuss another open problem in the definition of weighted straight skeletons caused by split events. In general, an event happens due to a topological change in the wavefront and the event handling was so far guided by




Fig. 1. The straight skeleton $\mathcal{S}(P)$ (blue) of a polygon P (bold) is defined as the traces of wavefront vertices over time. Instances of the wavefront $\mathcal{W}_P(t)$ at different times t are shown in gray (color online).

one fundamental principle: *Between events, the wavefront is a planar collection of wavefront polygons.* This is easily achieved when handling edge events and “simple” split events where at most four edges are in the wavefront afterwards. (Biedl *et al.*¹⁰ only studied polygons where split events were simple.) However, is it always possible to handle multiple simultaneous, co-located split events in a fashion that respects this fundamental principle?

We will show that it is necessary to weaken the requirement of strict planarity in the fundamental principle. After that, we can answer the question to the affirmative, that is, we show how to safely define weighted straight skeletons in the presence of multiple simultaneous, co-located split events. (Note that due to the discontinuous character of straight skeletons, it is not possible to tackle this problem by means of simulation of simplicity.⁴) We first rephrase this problem as a *planar matching* problem of directed pseudo-lines and show how to transform the planar matching problem into a *stable roommates* problem. For the main result, we prove that our particular stable roommates problem always possesses a solution. Such a solution tells us how to do the event handling of the wavefront in order to maintain planarity.

1.1. *Stable roommates problems*

The *stable marriage* problem was proposed over 50 years ago by Gale and Shapley¹¹ and concerns finding a pairing — say, between a set of n men and a set of n women — such that no pair of a man and a woman would prefer each other to their respective matching partners. Algorithms for solving this problem have found many real-world applications, including in economics and in college admission. Many variants, e.g., with restrictions on the preference lists, have been studied since (see, for example, Fleiner *et al.*¹² and the references therein).

In the stable marriage problem, pairings are possible only between items that are in different parts of the input. A natural generalization is to allow pairings between any two items; this is the so-called *stable roommates problem*. (We give a formal definition in Sec. 3.) While any instance of the stable marriage problem has a solution,¹¹ not all instances of the stable roommates problem do. Irving¹³ was the first to give a polynomial-time algorithm to test whether a stable roommates problem instance has a solution. Later, Tan¹⁴ and Tan and Hsueh¹⁵ gave another algorithm for this problem, which also provided much insight into the structure of solutions. We will review their main results (and use them to prove the existence of a solution in our special case) in Sec. 3.

2. Weighted Straight Skeletons

In this section, we clarify the details of the definition of the weighted straight skeleton, events, and how to create a new pairing of edges after such an event.

2.1. The wavefront

Let P denote a polygon, possibly with holes. We denote by $\sigma(e) \in \mathbb{R} \setminus \{0\}$ the *weight* of the edge e of P and call σ the *weight function*.^a For every edge e of P , let $n(e)$ denote the normal vector of e that points to the interior of P . Initially, every edge of P sends out a wavefront edge with fixed speed $\sigma(e)$. That is, the segments of the wavefront $\mathcal{W}(t)$ at time t that originate from edge e are contained in $\bar{e} + t \cdot \sigma(e) \cdot n(e)$, where \bar{e} denotes the supporting line of e . If $\sigma(e)$ is negative, the wavefront edge that emanated from e moves to the exterior of P .

2.2. Events, pairings, and some examples

Intuitively, an event happens when a wavefront vertex meets another wavefront edge or, in particular, another wavefront vertex. For unweighted straight skeletons (i.e., with all weights set to 1) the wavefront is planar between events. Here “planar” means that its edges do not intersect except at common endpoints. We can interpret events as the incidences where planarity is violated. To restore planarity, we must change the *pairing* of edges, i.e., which of the edges that are involved in the event are made to be consecutive in the polygons that constitute the wavefront after the event. For unweighted straight skeletons, we can always restore planarity by considering the cyclic order of wavefront edges meeting at p and by pairing each edge with the cyclically neighboring edge, see Fig. 2. We call this the *standard pairing technique*.

The situation becomes more complicated for the weighted straight skeleton, especially when two or more such events are co-located at the same time t . We give a few examples that demonstrate the difficulties:

- For the polygon in Fig. 3, the standard pairing yields a crossing after the event.
- For the polygon in Fig. 4, there exist multiple pairings of edges that lead to a planar wavefront after the event.




Fig. 2. A wavefront before (dotted), at (solid), and after an event (bold), with blue arrows showing movement direction of wavefront vertices. The *standard pairing technique* for handling a split event pairs each edge with its neighbor in the cyclic order with which it was not paired before (color online).

^aAs in Biedl *et al.*¹⁰ we exclude weight 0, because it leads to complications in the definition of skeleton via the wavefront. If one used an alternate definition via the roof model, a weight of 0 could be permitted.

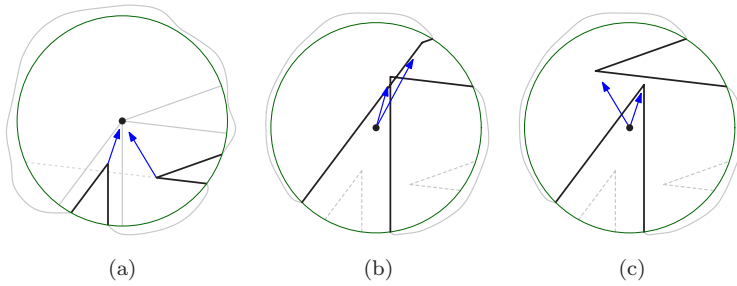


Fig. 3. (a) Two vertices move towards an event location. The immediate area of interest is shown within the green ball while the wavefront polygon outside is sketched only. At the event time, the supporting line depicted as a gray dashed line overtakes the left wavefront vertex. (b) Using the standard pairing to update the wavefront at an event can yield a non-planar wavefront. (c) Instead, the original pairing has to be preserved in this particular instance. (In both (b) and (c) the pre-event wavefront is shown dotted and in gray for reference.) (color online).

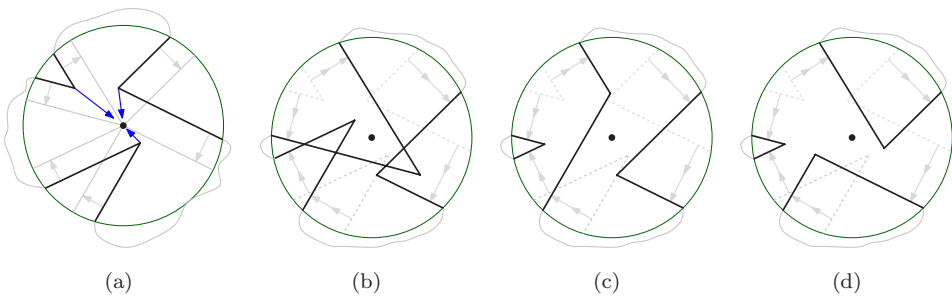


Fig. 4. Three vertices meet at an event (a), requiring a reconfiguration of the wavefront as planarity gets violated (b). There are multiple ways to restore planarity (c–d). Also the number of components of the post-event wavefront differs. (Again, the pre-event wavefront is shown dotted in the post-event figures.)

- For the polygon in Fig. 5, no pairing of edges exists such that the wavefront after the event is (strictly) planar.

This last example is the most troublesome since for the unweighted straight skeleton the main principle had been that the wavefront is planar between events, and this example shows that for weighted straight skeletons this is impossible to achieve. Because of its importance, we therefore argue the correctness of this example in more detail.

Lemma 1. *The weighted straight-skeleton wavefront propagation process of certain inputs will result in a non-planar wavefront. No pairing of wavefront edges is able to restore planarity.*

Proof. In Fig. 5 two vertices meet simultaneously at a point. By construction, one vertex lies on the supporting line of an edge that is incident to the other vertex at all

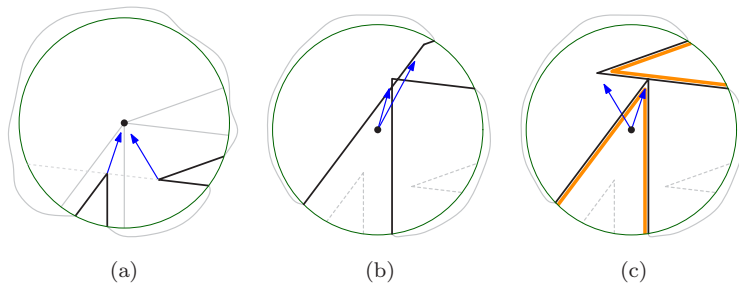


Fig. 5. (a) Two wavefront vertices meet at a point. One vertex is at all times on the supporting line (dotted) of another vertex’s edge. There are two possibilities, (b) and (c), to pair up the edges such that the wavefront remains planar in a weak sense, but not in a strict sense. (The pairing is highlighted using bold orange lines.) (color online).

times. We have three combinatorial possibilities to pair up the wavefront edges. One of them leads to a crossing. The other two possibilities are illustrated in Figs. 5(b) and 5(c). Both remaining possibilities are not (strictly) planar because some edge contains the point of a non-incident edge. \square

While in the above example (Lemma 1) the wavefront is not planar in a strict sense, there are no crossings — instead, edges only touch.

2.3. Weak planarity and the fundamental principle

Let P and Q denote two drawings (i.e., continuous maps) of a graph $H = (V, E)$ into the plane. Define the *Fréchet-distance* $d_{\mathcal{F}}(P, Q)$ to be $\inf_{\Phi: H \rightarrow H} \max_{x \in H} \times d(P(\Phi(x)), Q(x))$, where Φ is an automorphism of H and $d(\cdot, \cdot)$ denotes the standard Euclidean distance. See Sec. D.2 of Chang *et al.*¹⁶ for more details.

Define a drawing P of H to be *weakly planar* if, for any $\varepsilon > 0$, there exists a drawing Q of H that is planar (i.e., injective) such that $d_{\mathcal{F}}(P, Q) < \varepsilon$.^b

This definition implies that every planar drawing is weakly planar as well. In addition, for every weakly planar drawing P and for every $\varepsilon > 0$, the corresponding planar drawing Q in particular satisfies for all vertices v that $d_{\mathcal{F}}(Q(v), P(v)) < \varepsilon$. This definition now enables us to rephrase the fundamental principle as follows:

The wavefront is a collection of polygons that is weakly planar at all times.

2.4. Formal definition of events

The wavefront \mathcal{W} is initially planar and therefore also weakly planar. Informally, an event occurs when the wavefront is about to cease being weakly planar and event handling needs to restructure the wavefront locally such that it can continue propagating in a weakly planar fashion.

^bChang *et al.* used “weakly simple”, but we prefer the term “weakly planar” as to not use the two terms “simple drawings” and “simple graphs” for entirely unrelated concepts.

Assume that $\mathcal{W}(t')$ is weakly planar for all $t' \in [t - \delta, t]$ and some $\delta > 0$. For this time interval, we can consider \mathcal{W} to be a kinetic planar straight-line graph with a fixed set of kinetic vertices and edges. For Definition 1, we fix the vertex set and the edge set of \mathcal{W} , including the velocities of the vertices and temporarily ignore event handling. Furthermore, we denote by $B(p, r)$ the closed disk of radius r centered at p . By $\mathcal{W}(t') \cap B(p, r)$ we mean the planar straight-line graph $\mathcal{W}(t')$ with all edges truncated to fit into $B(p, r)$ or removed if they entirely reside outside $B(p, r)$.

Definition 1. Given a location p and a time t , we say that an *event* happens if at least two vertices meet for the first time at time t at p , or if $\exists \varepsilon_0 > 0 \forall \varepsilon \in (0, \varepsilon_0) \exists \delta > 0$ such that

- (i) $\mathcal{W}(t') \cap B(p, \varepsilon)$ is non-empty and weakly planar for $t' \in [t - \delta, t]$ and
- (ii) $\mathcal{W}(t') \cap B(p, \varepsilon)$ is non-empty and not weakly planar for $t' \in (t, t + \delta]$.

We say that the edges that meet in p at time t are *involved* in the event.

As this definition defines events localized at some point p , we can also talk about multiple events occurring at the same time t at different locations. If an event happens at location p and time t then, typically, weak planarity of \mathcal{W} is violated locally around p after time t . However, weak planarity is not violated if, for instance, a wavefront polygon collapses to a point. Figure 5 gives another example where weak planarity is not violated after the event. The goal of event handling is to guarantee that weak planarity holds after the event, by locally adapting the wavefront structure if needed.

Definition 2. We call the event at location p and time t *elementary* if exactly three edges are involved. We call it an *edge event* if $B(p, \varepsilon) \setminus \mathcal{W}(t - \delta)$ consists of two connected components and a *split event* otherwise. Non-elementary edge and split events are called multi-edge and multi-split events respectively.

It is known how to handle edge events and elementary split events.¹⁰ In the following, we present one unified approach that is able to correctly handle any type of event, including, in particular, multi-split events. Consequently, one side effect of our definition of weighted straight skeletons is that the distinction between edge events and split events becomes unnecessary.

2.5. Global and local weak planarity

The definition of weak planarity is global, i.e., considers the entire wavefront. In contrast, the definition of event considers only *local weak planarity*, i.e., the wavefront within a ball $B(p, \varepsilon)$ around point p . Unlike (strict) planarity, weak planarity can be violated globally while still holding everywhere locally, see Fig. 6(d). In fact, a (globally) weakly planar graph can be continuously deformed into a non-weakly-planar graph without vertices meeting and without local weak planarity being violated anywhere at any time, see Fig. 6.




Fig. 6. The continuous deformation of the dotted (blue) polygon causes a strictly planar graph (a) to become not globally weakly planar (color online).

However, as we show now, in the special case of a wavefront propagation, a non-weakly-planar wavefront cannot be obtained from a weakly planar wavefront without causing an event. To show this, fix some time interval $[t, t']$ in which no event occurs, and assume $\mathcal{W}(t)$ was weakly planar. To show that $\mathcal{W}(t')$ is weakly planar, we must find for any $\varepsilon' > 0$ a planar drawing Q' that has Fréchet distance at most ε' from $\mathcal{W}(t')$. Let $0 < \varepsilon \leq \varepsilon'$ be sufficiently small (we determine bounds for it later) and let Q be a planar drawing with $d_{\mathcal{F}}(\mathcal{W}(t), Q) \leq \varepsilon$. The goal is to transform Q into the desired drawing Q' .

Let v be a vertex, say it is at p in $\mathcal{W}(t)$ and at p' in $\mathcal{W}(t')$, and consider the balls $B(p, \varepsilon_1)$ and $B(p', \varepsilon_1)$. Here ε_1 is small enough that in both $\mathcal{W}(t)$ and $\mathcal{W}(t')$ any two vertices have distance at least $3\varepsilon_1$ (except for vertices occupying the same point) and vertex-balls intersect only edges that contain that vertex. The crucial observation is that $B(p, \varepsilon_1) \cap \mathcal{W}(t)$ equals $B(p', \varepsilon_1) \cap \mathcal{W}(t')$ after translating by $p - p'$. For if, say, edge e intersected $B(p', \varepsilon_1)$ in $\mathcal{W}(t')$, but not $B(p, \varepsilon_1)$ in $\mathcal{W}(t)$, then at some point in time interval $[t, t']$ vertex v started residing on e . Because this was not the case at time t , the supporting lines of v and e do not have a common point throughout, which means that continuing the movement of v results in a non-planarity, and hence an event, a contradiction. Similarly if some vertex v' belongs to $B(p', \varepsilon_1)$ in $\mathcal{W}(t')$, but not to $B(p, \varepsilon_1)$ in $\mathcal{W}(t)$, then at some point v and v' coincided, causing an event.

We can therefore define drawing Q' inside ball $B(p', \varepsilon_1)$ to consist of $B(p, \varepsilon_1) \cap Q$, translated by $p' - p$. With this, the Fréchet-distance within $B(p, \varepsilon_1)$ is at most ε .

If we choose $\varepsilon < \varepsilon_1$, then outside the balls at vertices, drawing Q must remain near edge-segments (i.e., maximal segments inside edges that have no vertex on them). It seems intuitive that we can complete Q' by connecting appropriately along edge-segments, but to describe this precisely we need a few notations, see also Fig. 7.

For each edge-segment s define the *tube* $T(s, \varepsilon)$ to be all points of distance at most ε from s . Further define the *core* $R(s, \varepsilon) \subset T(s, \varepsilon)$ to be the points in the tube for which the perpendicular onto s does not intersect the balls placed at the endpoints of s . This core is a rectangle that has a “shorter” side of length 2ε and a “longer” side of length $|s| - 2\varepsilon_1 \geq \varepsilon_1$. Since edge-vertex incidences do not change in time interval $[t, t']$, segment s corresponds to a segment s' in $\mathcal{W}(t')$, which also has such a tube and core. We assume that ε is small enough that in both $\mathcal{W}(t)$ and $\mathcal{W}(t')$ cores of edge-segments are disjoint unless the edge-segments coincide. (Any




Fig. 7. Balls, tubes and cores. Planar drawing Q' is obtained by translating Q near vertices and stretching Q within cores.

two edge-segments either coincide in both $\mathcal{W}(t)$ and $\mathcal{W}(t')$ or in neither, else there would be an event.)

Now define drawing Q' as follows. Near vertices, we simply translate. Namely, for any edge-segment s with ends p, q , we have $T(s, \varepsilon) = E_p \cup R(s, \varepsilon) \cup E_q$, where E_p, E_q are the maximal connected regions outside the core, containing p and q respectively. Translate $E_p \cap Q$ by $p' - p$ and $E_q \cap Q$ by $q' - q$ to obtain the parts of Q' within the corresponding regions inside the tube of s' . Multiple such regions may intersect near p , but the translation is the same for all of them, so this is consistent. Within cores we both translate and scale. Namely, we can obtain $R(s', \varepsilon)$ from $R(s, \varepsilon)$ by translation and scaling the longer side by a factor of $(|s'| - 2\varepsilon_1)/(|s| - 2\varepsilon_1) \leq |s'|/\varepsilon_1$. Apply the same transformation to $R(s, \varepsilon) \cap Q$ to obtain a drawing that fits inside $R(s', \varepsilon)$ and intersects the shorter side of $R(s', \varepsilon)$ in exactly the same way that Q intersected the shorter side of $R(s, \varepsilon)$.

Since any point of Q is within distance ε from some point in $\mathcal{W}(t)$, drawing Q resides inside the tubes. Therefore all parts of Q have been transformed into some part of Q' . The various pieces of Q' fit together, because at the boundary to the cores the pieces of Q' intersect at the same points as Q did. So Q' is a planar drawing of the wavefront. Use the same automorphism that defined $d_{\mathcal{F}}(Q, \mathcal{W}(t))$ to define one for Q' and $\mathcal{W}(t')$, then the Fréchet-distance of Q' to $\mathcal{W}(t')$ is at most $S \cdot \varepsilon$, where $S \leq \max_{s'} |s'|/\varepsilon_1$ is the maximum scale-factor used for some tube. Choosing $\varepsilon < \varepsilon'/S$ therefore shows $d_{\mathcal{F}}(Q', \mathcal{W}(t')) \leq \varepsilon'$ and we have:

Lemma 2. *Weak planarity (in a global sense) cannot get violated during the wavefront propagation without an event happening.*

2.6. Pairing edges and planar matching

Assume an event happens at time t at location p . Up until time t the wavefront \mathcal{W} is weakly planar. In case weak planarity is violated after t , we must transform the wavefront structure to restore weak planarity. This involves changing the pairing of wavefront edges.

We reduce the problem of pairing up wavefront edges during event handling to a particular matching problem, discussed in Sec. 3. This problem, which we study

independently of straight skeletons, takes a directed pseudo-line arrangement in general position as input and outputs a matching between the pseudo-lines. The matching is planar with respect to the so-called “tails”. This provides us with a means to construct a weakly planar wavefront again. In the following, we describe how to transform a weakly planar wavefront into a suitable pseudo-line arrangement for the matching problem.

The pseudo-lines stem from the supporting lines of wavefront edges and are required to be in general position. By *general position* we mean that any pair of lines intersects in exactly one unique point. In particular, this implies that no two lines are parallel, no two lines are identical, and no three lines intersect in a common point. Note that it is only the pseudo-line arrangement which is in general position. The original wavefront is not so restricted and can be in an arbitrary position.

At time t , several edges of the wavefront \mathcal{W} are incident at location p . For each such edge, either zero, one, or both endpoints approach p at time t . We construct a simplified version of the wavefront, denoted by \mathcal{W}' , as follows. First, we drop any edge both of whose endpoints reach p . For each such edge we join its two endpoints, making its two neighbors now neighbors of each other. Furthermore, any edge where no endpoint reaches p is split into two edges by a new wavefront vertex that also reaches p at time t . Thus, in \mathcal{W}' an even number of wavefront edges have exactly one endpoint at point p at time t , see Fig. 8. All changes have been of a combinatorial nature, and not changed the geometry of the wavefront at time t , so \mathcal{W}' is weakly planar at time t .

We now want to create a set \mathcal{L} of directed pseudo-lines that will be the input to the matching problem, which allows us to create a suitable pairing for the new wavefront. The idea is to use for any edge e in \mathcal{W}' the supporting line \bar{e} at time e in \mathcal{W}' and to direct it from the end where e intersected the boundary of $B(p, \varepsilon)$. However, \mathcal{L} is required to be in general position whereas some supporting lines may not be. Specifically, the subdivision of edges that reach p results in two edges with the same supporting line (but directed in opposite directions). Moreover, \mathcal{W}' may not have been strictly planar, and there may have been multiple edges with the same directed supporting line all reaching p .




Fig. 8. A multi-split event occurs at location p . The involved edges form four chains. The simplified wavefront \mathcal{W}' contains exactly eight edges, shown in bold, stemming from those chains.

We resolve such overlaps based on a strictly planar drawing \mathcal{W}^s that is very close to $\mathcal{W}'(t)$. Fix a small $\nu > 0$, and let \mathcal{W}^s be such that $d_{\mathcal{F}}(\mathcal{W}^s, \mathcal{W}'(t)) < \nu$. We can extrapolate \mathcal{W}^s to time $t + \delta$ as follows: For any edge e in $\mathcal{W}'(t)$, there exists a curve $\mathcal{C}(e)$ in \mathcal{W}^s with $d_{\mathcal{F}}(e, \mathcal{C}(e)) < \nu$. We can translate $\mathcal{C}(e)$ by $\delta \cdot \sigma(e) \cdot n(e)$, i.e., with the same speed and in the same direction as e . For sufficiently small δ , this will not result in crossings outside $B(p, \varepsilon)$.

Now we explain how to obtain \mathcal{L} . We start with the directed supporting lines of $\mathcal{W}'(t + \delta)$, and use in the following $\ell(e)$ to denote the pseudo-line representing e . For any edge e involved in the event where no other involved edge has the same supporting line \bar{e} , we can use $\bar{e} \cap B(p, \varepsilon)$ as $\ell(e)$.

For a set of two or more edges of $\mathcal{W}'(t + \delta)$ that all have the same supporting line \bar{e} , we do the following in order to arrive at a pseudo-line arrangement in general position. (Note that all these edges must be moving with the same speed since they were all incident at p at time t and at time $t + \delta$ still share a supporting line.) Prune \bar{e} to consider only the part between where it intersects the boundary of $B(p, \varepsilon)$, say this happens at points p_s and p_t . Let e_1, \dots, e_k be all edges of $\mathcal{W}'(t + \delta)$ that had supporting line \bar{e} and contained p_s , and let $e'_1, \dots, e'_{k'}$ be all edges of $\mathcal{W}'(t + \delta)$ that had supporting line \bar{e} and contained p_t .

All of e_1, \dots, e_k were overlapping in $\mathcal{W}'(t)$, and so must be replaced by curves in \mathcal{W}^s . The (translated) curves $\mathcal{C}(e_1), \dots, \mathcal{C}(e_k)$ intersect the boundary of $B(p, \varepsilon)$ within distance ν from p_s , say this was at points p_1, \dots, p_k . Similarly let $p'_1, \dots, p'_{k'}$ be the points where $\mathcal{C}(e'_1), \dots, \mathcal{C}(e'_{k'})$ intersect the boundary of $B(p, \varepsilon)$. Insert points q_1, \dots, q_k just after $p'_{k'}$ (on the boundary of $B(p, \varepsilon)$ in clockwise order), and insert points $q'_1, \dots, q'_{k'}$ just before p_1 . Now let $\ell(e_i)$ be the line segment from p_i to q_i , and let $\ell(e'_i)$ be the line segment from p'_i to q'_i . If ν is sufficiently small and we choose the points q_i close enough, then any of the new lines is within distance ν of \bar{e} , and intersects all lines of edges not supported by \bar{e} in the same order as \bar{e} did. Also none of the new lines overlap, and, if the points q_i are chosen to not trigger degeneracies, then they all intersect each other in unique points. Repeating this for all groups of overlapping edges results in a pseudo-line arrangement \mathcal{L} in general position as desired.

We use \mathcal{L} as input to the matching algorithm in the next section, and obtain as output a pairing of the pseudo-lines such that there are no crossings within $B(p, \varepsilon)$ when drawing the tails (see Corollary 1 in Sec. 3.3). If $\ell(e)$ and $\ell(e')$ were thus paired, then we declare the new wavefront \mathcal{W}^*_ν to be the one where e and e' share a common vertex after time t . Note that \mathcal{W}^*_ν depends on the parameter $\nu > 0$ that we used to find a planar wavefront \mathcal{W}^s approximating $\mathcal{W}'(t)$.

If several multi-split events happen at the same time, then this procedure is repeated for every such event independently, except that the same strictly-planar approximation \mathcal{W}^s should be used to determine the order of pseudo-lines in all event-regions.

Lemma 3. *There exists a post-event wavefront \mathcal{W}^* that is weakly planar.*

Proof. For any $\nu > 0$, let \mathcal{W}_ν^* be a post-event-wavefront constructed if we fix this value ν . This post-event-wavefront is not necessarily unique, and we cannot exclude that for different values of ν , the different post-event-wavefronts \mathcal{W}_ν^* might have different combinatorics also. However, there is only a finite number of pairings of event-edges, and so a finite number of possible post-event-wavefronts. Therefore, in the infinite sequence $(\mathcal{W}_\nu^*)_{\nu \rightarrow 0}$ of post-event wavefronts, there exists an infinite subsequence such that the post-event wavefront all have the same combinatorics. We set \mathcal{W}^* to be that post-event wavefront.

To show that \mathcal{W}^* is weakly planar, fix an $\varepsilon > 0$. There exists some $\nu \leq \varepsilon$ that belongs to the subsequence, i.e., $\mathcal{W}_\nu^* = \mathcal{W}^*$. Construct a drawing \mathcal{W}^ν that approximates \mathcal{W}^* as follows: Inside ball $B(p, \varepsilon)$, draw edge e along $\ell(e)$ as dictated by the pairing; this gives no crossing and is within distance ν of e by construction of $\ell(e)$. Outside ball $B(p, \varepsilon)$, use the (translated) curve $\mathcal{C}(e)$ for e ; this again has no crossing and is within distance ν of e . These two curves meet exactly at the boundary of $B(p, \varepsilon)$, yielding a single curve to represent e . Thus $d_{\mathcal{F}}(\mathcal{W}^\nu, \mathcal{W}^*) < \nu \leq \varepsilon$ as desired. \square

3. Matchings and Roommates

We now turn to the problem of finding a matching between pseudo-lines such that the “tails” do not cross. Formally, for an even number N , let $\mathcal{L} = \{\ell_1, \dots, \ell_N\}$ be an oriented pseudo-line arrangement in general position, i.e., a set of directed simple curves that begin and end at infinity such that any two curves intersect each other in a single point and no three curves intersect in the same point. For any two such pseudo-lines ℓ_1, ℓ_2 , let $\ell_1 \cap \ell_2$ denote the point of their intersection. Let \mathcal{C} be a closed simple curve that encloses all intersections of pseudo-lines and that intersects each (directed) pseudo-line ℓ exactly twice. Walking along the direction of ℓ , let the *begin-point* $b(\ell)$ be the first intersection point of ℓ and \mathcal{C} , and let its *end-point* be the second intersection point.

A *matching* M in \mathcal{L} is a grouping of ℓ_1, \dots, ℓ_N into pairs; we use $M(\ell_i)$ to denote the matching partner of ℓ_i . The *matching tail* of ℓ_i is the sub-curve of ℓ_i from $b(\ell_i)$ to $\ell_i \cap M(\ell_i)$.




Fig. 9. A directed pseudo-line arrangement. (Left) The matching tail. (Right) A planar matching.

Definition 3. A matching in \mathcal{L} is called *planar* if the union of the matching tails gives a planar drawing.

The *planar matching problem* is the problem of finding a planar matching M for a given pseudo-line arrangement \mathcal{L} in general position. In the following we translate the planar matching problem into a stable roommates problem.

3.1. Transformation to the stable roommates problem

We first give the definition of the stable roommates problem. Assume that we have an even number N of elements $\mathcal{A} = \{a_1, \dots, a_N\}$. Each element has a ranking of elements, which is *complete* and *strict*, i.e., all elements are ranked and no two elements are ranked the same. Let M be a matching of a_1, \dots, a_N . A pair $\{a_i, a_j\}$ is a *blocking pair* for M if a_i prefers a_j over $M(a_i)$ and a_j prefers a_i over $M(a_j)$. A matching is *stable* if there is no blocking pair. The *stable roommates problem* asks for a stable matching in \mathcal{A} .

Now let us again consider the directed pseudo-line arrangement \mathcal{L} . As we walk along a pseudo-line ℓ_i from its begin-point to its end-point, we encounter all other pseudo-lines in \mathcal{L} . This order naturally gives us a complete and strict ranking for ℓ_i if we attach ℓ_i itself to the end of the list. Thus, \mathcal{L} defines an instance of the stable roommates problem.

Lemma 4. *A directed pseudo-line arrangement has a planar matching if and only if the corresponding stable roommates instance has a stable matching.*

Proof. For a matching M , the matching tails of two pseudo-lines ℓ_i and ℓ_j cross if and only if ℓ_i prefers ℓ_j over $M(\ell_i)$ and ℓ_j prefers ℓ_i over $M(\ell_j)$. Hence, the matching is non-planar if and only if there is a blocking pair. \square

3.2. Stable partitions

In order to solve our particular stable roommates problem, we review some results on so-called stable partitions, mostly based on a paper by Tan and Hsueh.¹⁵

Let \mathcal{A} be an instance of a stable roommates problem, and let π be a permutation on \mathcal{A} , i.e., a bijective map $\mathcal{A} \rightarrow \mathcal{A}$. This map partitions \mathcal{A} into one or more *cycles*, i.e., sequences a'_0, \dots, a'_{k-1} in \mathcal{A} with $a'_0 \xrightarrow{\pi} a'_1 \xrightarrow{\pi} \dots \xrightarrow{\pi} a'_{k-1} \xrightarrow{\pi} a'_0$. A cycle with length $k \geq 3$ is called a *semi-party cycle* if a'_i prefers $\pi(a'_i)$ over $\pi^{-1}(a'_i)$ for all i with $0 \leq i < k$. A *semi-party partition* of \mathcal{A} is a permutation of \mathcal{A} where all cycles of length at least 3 are semi-party cycles.

Given a semi-party partition π , a pair $\{a_i, a_j\}$ is called a *party-blocking pair* if a_i prefers a_j over $\pi^{-1}(a_i)$ and a_j prefers a_i over $\pi^{-1}(a_j)$. A *stable partition* is a semi-party partition that has no party-blocking pairs. The cycles of a stable partition are called *parties*. An odd (even) party is a party of odd (even) cardinality. Furthermore, a_i, a_j are *party-partners* if $a_i = \pi(a_j)$ or $a_j = \pi(a_i)$.




Fig. 10. (Left) A stable partition with two pair-parties and one cycle-party of length four. (Right) The party-tail of pseudo-line ℓ (thick orange line) (color online).

Theorem 1 (Refs. 14 and 15). *For any instance \mathcal{A} of the stable roommates problem the following statements hold:*

- (1) \mathcal{A} has a stable partition, and it can be found in polynomial time.
- (2) Any stable partition of \mathcal{A} has the same number of odd parties.
- (3) \mathcal{A} has a stable matching if and only if it has a stable partition with no odd parties.

3.3. Existence of planar matchings

Now we consider parties that occur in stable roommates instances defined by a directed pseudo-line arrangement \mathcal{L} . Theorem 1(1) gives us a stable partition π for \mathcal{L} . Let a *singleton-party*, a *pair-party*, and a *cycle-party* be a party consisting of one, two, and at least three pseudo-lines, respectively. For each pseudo-line ℓ that is not in a singleton-party, let its *party-tail* be the part between $b(\ell)$ and $\ell \cap \pi^{-1}(\ell)$. For any pseudo-line ℓ that is in a singleton-party, let its *party-tail* be the part of ℓ between begin-point and end-point.

Lemma 5. *The party-tails of two pseudo-lines ℓ and ℓ' do not intersect unless ℓ and ℓ' are party-partners.*

Proof. Assume that $\ell \cap \ell'$ belongs to both party-tails, but ℓ and ℓ' are not party-partners. We first show that ℓ prefers ℓ' over $\pi^{-1}(\ell)$. This holds automatically if ℓ is in a singleton-party, because then $\pi^{-1}(\ell) = \ell$ and any pseudo-line ranks itself lowest. If ℓ is not in a singleton-party, then the party-tail of ℓ consists of the sub-curve between $b(\ell)$ and $\ell \cap \pi^{-1}(\ell)$. Since $\ell' \neq \pi^{-1}(\ell)$ by assumption, and since no three pseudo-lines intersect in a point, $\ell \cap \ell'$ comes strictly earlier than $\ell \cap \pi^{-1}(\ell)$ when walking along ℓ . By definition of the ranking for directed pseudo-lines, hence ℓ prefers ℓ' over $\pi^{-1}(\ell)$.

Similarly one shows that ℓ' prefers ℓ over $\pi^{-1}(\ell')$. Hence, $\{\ell, \ell'\}$ is a party-blocking pair and π is not a stable partition, a contradiction. \square

Lemma 6. *There cannot be two singleton-parties.*

Proof. Assume that P and P' are two singleton-parties, with $P = \{\ell\}$ and $P' = \{\ell'\}$. Since ℓ and ℓ' are in singleton-parties, their party-tails extend from their begin-points to their end-points. Since all pseudo-lines intersect within \mathcal{C} , so do ℓ and ℓ' . But ℓ and ℓ' are not party-partners, in contradiction to Lemma 5. \square

Lemma 7. *There cannot be two cycle-parties.*

Proof. Assume we have two cycle-parties $P_1 = \{\ell_0, \ell_1, \dots, \ell_{a-1}\}$ and $P_2 = \{\ell'_0, \ell'_1, \dots, \ell'_{b-1}\}$, with $\pi(\ell_i) = \ell_{i+1}$, addition modulo a , and $\pi(\ell'_i) = \ell'_{i+1}$ with addition modulo b .

Let $G(P_1)$ be the graph formed by the party tails of P_1 as follows: The vertex set comprises $b(\ell)$ and $\ell \cap \pi(\ell)$ for every pseudo-line ℓ in P_1 . We add each party-tail as two edges $(b(\ell), \ell \cap \pi(\ell))$ and $(\ell \cap \pi(\ell), \pi^{-1}(\ell) \cap \ell)$, see Fig. 11.

Note that $G(P_1)$ has the following structure: It consists of a cycle \mathcal{C}_1 of edges of the form $(\ell \cap \pi(\ell), \pi^{-1}(\ell) \cap \ell)$, where $\ell \in \{\ell_0, \dots, \ell_{a-1}\}$, and a set of edges of the form $(b(\ell), \ell \cap \pi(\ell))$ each of which is attached to one vertex of \mathcal{C}_1 . By Lemma 5, $G(P_1)$ is planar, and we consider it to have the fixed planar embedding and faces inherited from \mathcal{L} . Note that the vertices $b(\ell_0), b(\ell_1), \dots, b(\ell_{a-1})$ all lie on the simple closed curve \mathcal{C} and are ordered clockwise or counterclockwise by planarity. Therefore, $G(P_1)$ tessellates the area enclosed by \mathcal{C} into $a + 1$ regions. Note that exactly a of those regions are partially bounded by \mathcal{C} . The remaining region is the one bounded by \mathcal{C}_1 . Similarly, we define $G(P_2)$ and \mathcal{C}_2 .

Again by Lemma 5, $G(P_1) \cup G(P_2)$ is planar, and it follows that $G(P_2)$ is entirely contained in one region of $G(P_1)$. This region is not the region bounded by \mathcal{C}_1 since it must contain points on \mathcal{C} . We denote by \mathcal{R}_1 the union of all regions of $G(P_1)$ that do not contain $G(P_2)$. Likewise, we denote by \mathcal{R}_2 the union of all regions of $G(P_2)$ that do not contain $G(P_1)$. We observe that \mathcal{R}_1 and \mathcal{R}_2 are disjoint.

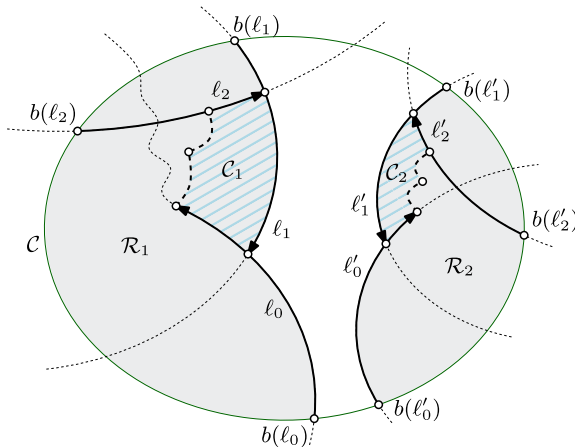


Fig. 11. The two pseudo-lines ℓ_0 and ℓ'_0 cannot intersect.

Without loss of generality, the boundary of \mathcal{R}_1 consists of parts of \mathcal{C} as well as the path $b(\ell_1), \ell_1 \cap \ell_2, \ell_0 \cap \ell_1, b(\ell_0)$. Likewise, \mathcal{R}_2 is bounded by parts of \mathcal{C} and edges stemming from ℓ'_1 and ℓ'_0 .

In the following, we will show that ℓ_0 cannot intersect \mathcal{R}_2 , and, conversely, ℓ'_0 cannot intersect \mathcal{R}_1 . Consequently, ℓ_0 does not intersect ℓ'_0 within \mathcal{C} , which is a contraction as we require each pair of pseudo-lines to intersect exactly once in the area enclosed by \mathcal{C} . This concludes the proof.

Note that ℓ_0 starts at $b(\ell_0)$, then makes up parts of the boundary of \mathcal{R}_1 until it reaches $\ell_0 \cap \ell_1$. Then, ℓ_0 moves into the interior of \mathcal{R}_1 as it makes up an edge of \mathcal{C}_1 , but not the one that is part of the boundary of \mathcal{R}_1 . Once ℓ_0 enters \mathcal{R}_1 , it can leave only by intersecting \mathcal{C} at its end-point as it is not allowed to self-intersect or to intersect ℓ_1 a second time. After ℓ_0 has left the area enclosed by \mathcal{C} , it cannot enter again, as it intersects \mathcal{C} exactly twice. So ℓ_0 cannot intersect \mathcal{R}_2 . Likewise, ℓ'_0 will exit the area enclosed by \mathcal{C} through its end-point in \mathcal{R}_2 and cannot intersect \mathcal{R}_1 . □

Lemma 8. *There cannot be a singleton-party and a cycle-party.*

Proof. We follow the same idea as in the previous proof, and use P_1 as the cycle-party and P_2 as the singleton-party. Let ℓ_0 be defined as previously, and use ℓ'_0 as the single line in P_2 . Since the tail of ℓ'_0 consists of all points between the begin-point and the end-point of ℓ'_0 , again no intersection between ℓ_0 and ℓ'_0 is possible. □

Theorem 2. *No instance of a stable roommates problem defined by a directed pseudo-line arrangement \mathcal{L} can have a stable partition with an odd party.*

Proof. Assume to the contrary that some stable partition π has an odd party P . As \mathcal{L} comprises an even number of pseudo-lines, there needs to be another odd party P' . This can only happen if there are two singleton-parties, two (odd) cycle parties, or a singleton-party and an (odd) cycle-party. These are ruled out by Lemma 6, Lemma 7, and Lemma 8, respectively. □

Theorem 3. *Every directed pseudo-line arrangement has a planar matching, and it can be found in polynomial time.*

Proof. This is a direct result of Lemma 4, Theorems 1 and 2.

By Theorem 1 we can find a stable partition in polynomial time. By Theorem 2 and Lemma 7, it consists of pair-parties, except for at most one cycle-party P that has even length. If there is no cycle-party, then the stable partition is in fact a stable matching. Otherwise, if say ℓ_1, \dots, ℓ_a is the even cycle-party, then we can find a stable matching M easily, and there are two choices: Either set $M(\ell_{2i}) = \pi(\ell_{2i})$ and $M(\ell_{2i+1}) = \pi^{-1}(\ell_{2i+1})$, or do the same after shifting all indices by one. □

3.4. Application to straight skeletons

The matching tails of the pseudo-lines play the role of wavefront edges after the event. Therefore any planar matching tells us how to pair up the wavefront edges in order to restore planarity locally at p .

Corollary 1. *There exists a weakly planar wavefront after the event if there is a planar matching for \mathcal{L} .*

Using Theorem 3, we have found a stable matching and with it a weakly planar post-event wavefront, in polynomial time.

3.5. Ambiguities in the wavefronts

Notice that if a cycle-party exists, then there are two possible post-event wavefronts to be obtained from this stable partition. Even worse, a directed pseudo-line arrangement may have many stable partitions and hence there might be many possible post-event wavefronts. Figure 12 shows an example of a directed line arrangement

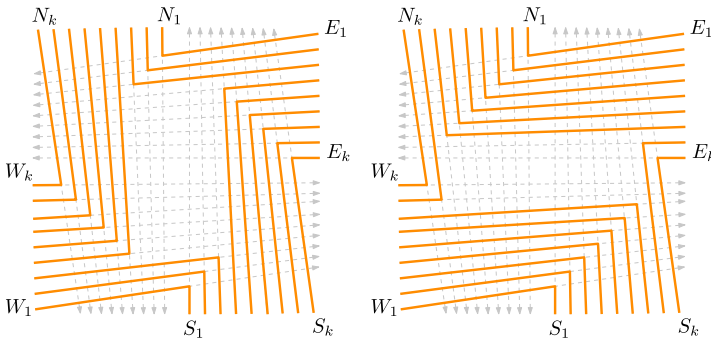


Fig. 12. An arrangement of $N = 4k$ directed lines. For any $0 \leq i \leq k$, we can find a planar matching by pairing $S_j - W_j$ and $N_j - E_j$ for $j \leq i$ and $S_j - E_j$ and $N_j - W_j$ for $j > i$. Shown here are the matchings obtained for $i = 3$ and $i = 7$.

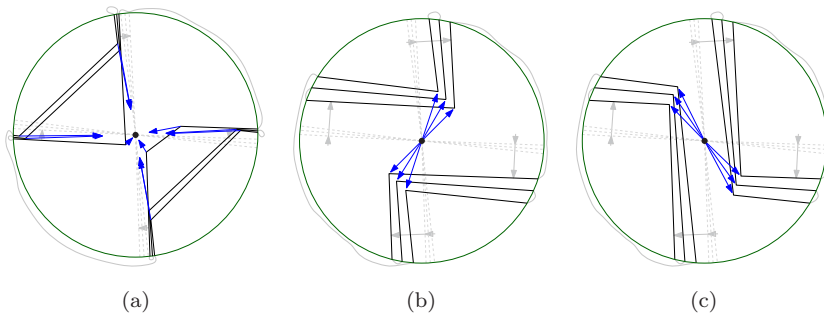


Fig. 13. A wavefront consisting of six chains approaching an event location (a). There are four possible pairings after the event, two are shown in (b) and (c). Extending this example to $2 \cdot k$ chains allows for $k + 1$ possible after-event wavefront configurations.




Fig. 14. If $\Omega(n)$ vertices move on a line towards $\Omega(n)$ other vertices, then there could be $\Omega(n^2)$ events.

that has $\Theta(N)$ many possible planar matchings. Figure 13 shows an example of a wavefront where linearly many post-event configurations are possible. Consequently, ambiguities in the development of the wavefront may be caused not only by edge events between parallel edges¹⁰ but also by multi-split events.

4. Conclusion

Although algorithms and even rudimentary implementations to construct the weighted straight skeleton have previously been presented, and even though several applications are suggested in the literature, this paper is the first to provide a concrete, constructive proof that a well-defined weighted straight skeleton actually exists in all cases. This result is based on two main ingredients: First, we have introduced and studied planar matchings on a directed pseudo-line arrangement as a generic tool independent of straight skeletons. In particular, we have shown that planar matchings always exist. Second, our interpretation of an event as a potential violation of (weak) planarity unifies the classification of edge and split events in 2D and promises to simplify the description and study of straight skeletons in dimensions higher than two, where the number of types and complexity of events would significantly increase otherwise.

The complexity of finding the weighted straight skeleton remains open. Computing a solution to the stable roommates problem can be done in $O(N^2)$ time.¹³ We potentially have $N \in \Omega(n)$ at an event. Even the number of events is not necessarily small: for unweighted polygons at most n events happen, but for weighted polygons there could be $\Omega(n^2)$ events (see Fig. 14). The trivial upper bound is $O(n^3)$ events since three supporting lines of edges must meet and this happens at most $\binom{n}{3}$ times. Hence we do not claim a running time better than $O(n^5)$ if the polygon is not in general position; surely this could be improved.

Acknowledgment

We would like to thank David Eppstein for the idea of interpreting the edge-pairing problem as a stable roommates problem.

References

1. O. Aichholzer, F. Aurenhammer, D. Alberts and B. Gärtner, A novel type of skeleton for polygons, *J. Universal Comp. Sci.* **1** (1995) 752–761.
2. S. Huber, *Computing Straight Skeletons and Motorcycle Graphs: Theory and Practice* (Shaker Verlag, 2012).
3. O. Aichholzer and F. Aurenhammer, Straight skeletons for general polygonal figures in the plane, *Voronoi's Impact on Modern Sciences II*, Vol. 21 (1998), pp. 7–21.
4. D. Eppstein and J. Erickson, Raising roofs, crashing cycles, and playing pool: Applications of a data structure for finding pairwise interactions, *Discr. Comput. Geom.* **22** (1999) 569–592.
5. F. Aurenhammer, Weighted skeletons and fixed-share decomposition, *Comput. Geom. Theor. Appl.* **40** (2008) 93–101.
6. J.-H. Haunert and M. Sester, Area collapse and road centerlines based on straight skeletons, *GeoInformatica* **12** (2008) 169–191.
7. T. Kelly and P. Wonka, Interactive architectural modeling with procedural extrusions, *ACM Trans. Graphics* **30** (2011) 14:1–14:15.
8. R. G. Laycock and A. M. Day, Automatically generating large urban environments based on the footprint data of buildings, *Proc. 8th ACM Symp. Solid Modeling and Applications* (2003), pp. 346–351.
9. G. Barequet, D. Eppstein, M. T. Goodrich and A. Vaxman, Straight skeletons of three-dimensional polyhedra, *Proc. 16th Annual European Symp. Algorithms* (2008), pp. 148–160.
10. T. Biedl, M. Held, S. Huber, D. Kaaser and P. Palfrader, Weighted straight skeletons in the plane, *Comput. Geom. Theor. Appl.* **48** (2015) 120–133.
11. D. Gale and L. S. Shapley, College admissions and the stability of marriage, *Amer. Math. Monthly* **69** (1962) 9–15.
12. T. Fleiner, R. W. Irving and D. F. Manlove, Efficient algorithms for generalized stable marriage and roommates problems, *Theor. Comp. Sci.* **381** (2007) 162–176.
13. R. W. Irving, An efficient algorithm for the “stable roommates” problem, *J. Algorithms* **6** (1985) 577–595.
14. J. J. Tan, A necessary and sufficient condition for the existence of a complete stable matching, *J. Algorithms* **12** (1991) 154–178.
15. J. J. Tan and Y.-C. Hsueh, A generalization of the stable matching problem, *Discr. Appl. Math.* **59** (1995) 87–102.
16. H.-C. Chang, J. Erickson and C. Xu, Detecting weakly simple polygons, *Proc. 26th Symp. Discrete Algorithms* (2015), pp. 1655–1670, arXiv:1407.3340.