

Manhattan-Geodesic Embedding of Planar Graphs

Bastian Katz¹, Marcus Krug¹, Ignaz Rutter¹, and Alexander Wolff²

¹ Faculty of Informatics, Universität Karlsruhe (TH), KIT, Germany
{katz, krug, rutter}@iti.uka.de

² Institut für Informatik, Universität Würzburg, Germany

Abstract. In this paper, we explore a new convention for drawing graphs, the (*Manhattan-*) *geodesic* drawing convention. It requires that edges are drawn as interior-disjoint *monotone* chains of axis-parallel line segments, that is, as geodesics with respect to the Manhattan metric. First, we show that geodesic embeddability on the grid is equivalent to 1-bend embeddability on the grid. For the latter question an efficient algorithm has been proposed. Second, we consider *geodesic point-set embeddability* where the task is to decide whether a given graph can be embedded on a given point set. We show that this problem is \mathcal{NP} -hard. In contrast, we efficiently solve *geodesic polygonization*—the special case where the graph is a cycle. Third, we consider geodesic point-set embeddability where the vertex–point correspondence is given. We show that on the grid, this problem is \mathcal{NP} -hard even for perfect matchings, but without the grid restriction, we solve the matching problem efficiently.

1 Introduction

In this paper we consider a new convention for drawing graphs. One of the most popular conventions is the *orthogonal* drawing convention, which requires edges to be drawn as interior-disjoint *rectilinear* chains, that is, chains of axis-parallel line segments. Restricting the number of edge directions potentially yields very clear drawings. We go a step further and insist that, additionally, edges are drawn as *monotone* chains. Such chains are called *Manhattan paths*. The idea behind monotonicity is that following the course of a monotone curve is potentially easier than following the course of a curve that is allowed to make detours. Manhattan paths are geodesics with respect to the Manhattan metric. Therefore we name our new convention the (*Manhattan-*) *geodesic* drawing convention.

In the Euclidean plane, geodesics are straight-line segments, and the classic result of König, Fáry, and Stein says that the class of graphs that have a straight-line drawing is exactly the class of planar graphs. Since there are efficient (linear-time) planarity-testing algorithms, we can decide efficiently whether a given graph has a Euclidean-geodesic drawing. We consider the same problem, which we call (MANHATTAN-) GEODESIC EMBEDDABILITY, with respect to the Manhattan distance. As an example take K_4 , the complete graph on four vertices, which has a geodesic drawing in the Euclidean plane but not in the

Manhattan plane. To avoid problems of drawing resolution, both questions are also interesting on the grid. The Euclidean case has been solved, for example, by Schnyder [14] who can draw any planar n -vertex graph on a grid of size $(n - 2) \times (n - 2)$, which is asymptotically optimal in the worst case.

Fixed point set. Next, we consider the setting where we are given not just a graph, but also a set of points (in the plane or on the grid) to which the vertices of the graph must be brought into correspondence. We call this problem GEODESIC POINT-SET EMBEDDABILITY. Kaufmann and Wiese [7] considered point-set embeddability (PSE) with respect to the *polyline drawing convention*. They showed that it is \mathcal{NP} -hard to decide whether a graph can be embedded on a point set with at most one bend per edge and that two bends suffice for any planar graph and any point set. Cabello [1] showed that it is \mathcal{NP} -hard to decide whether a planar graph has a straight-line embedding on a given point set.

A special case of both the straight-line and the orthogonal drawing convention has also been considered. Rappaport [12] showed that it is \mathcal{NP} -hard to decide whether a set P of n points has an *orthogonal polygonization*, that is, whether the n -cycle can be realized on P using horizontal or vertical edges only. O'Rourke [9] proved that if one forbids 180° -degree angles in the vertices, then there exists at most one simple rectilinear polygon with vertex set P . He also showed how to reconstruct the polygon from P in $O(n \log n)$ time. We refer to Demaine's survey [2] about problems related to polygonization.

PSE with the same drawing convention but with respect to a different graph class—perfect matchings—was considered by Rendl and Woeginger [13]. They showed that given a set of n points in the plane, one can decide in $O(n \log n)$ optimal time whether each point can be connected to exactly one other point with an axis-parallel line segment. They also showed that the problem becomes hard if one insists that the segments do not cross. Hurtado [5] gave a simple $O(n \log n)$ -time algorithm for the same problem under the geodesic drawing convention. The idea is to alternately go up and down the occupied grid columns.

Fixed correspondence. We further restrict the placement of the vertices by making the bijection between vertices and points part of the input. We call the resulting problem LABELED GEODESIC PSE. A special case of this problem (where the graph is a perfect matching) has applications in VLSI layout. Insisting on geodesic connections makes sure that signals reach their destinations as fast as possible. For example, a popular, but more restrictive wiring technique in VLSI layout, *single-bend wiring*, uses special geodesic connections. Raghavan et al. [11] have shown that one can decide our perfect matching problem efficiently when insisting on at most one bend per edge.

For the same problem with given vertex–point correspondence but under the polyline drawing convention, Pach and Wagner [10] showed that it is possible to embed any planar graph on any set of points, but they also showed that some edges may require $\Omega(n)$ bends. Goaoc et al. [4] showed that it is \mathcal{NP} -hard to decide whether a given graph can be 1-bend embedded on a given set of points with given vertex–point correspondence.

	GEODESIC	GEODESIC POINT-SET EMBEDDABILITY		
	EMBEDDABILITY	unrestricted	labeled (on grid)	labeled (off grid)
planar graph	\mathcal{P} [Thm. 1]	hard [Thm. 2]	hard [Thm. 4]	open
matching	trivial	\mathcal{P} [5]	hard [Thm. 4]	\mathcal{P} [Thm. 5]
polygonization	trivial	\mathcal{P} [Thm. 3]	open	open

Table 1: Overview over results in geodesic embeddability; hard is short for \mathcal{NP} -hard

Our Contribution. Drawing graphs with (Manhattan) geodesics opens up a large new field of research; we have done the following first steps.

- We show that GEODESIC EMBEDDABILITY on the grid is equivalent to deciding whether the given graph has a rectilinear *one*-bend drawing on the grid, see Section 2. Liu et al. [8] proposed an algorithm to decide the latter question efficiently. It is easy to see that a rectilinear one-bend drawing of an n -vertex graph fits on the $n \times n$ grid.
- We then prove that GEODESIC PSE is \mathcal{NP} -hard on (and off) the grid, reducing (in two steps) from HAMILTONIAN CYCLE, see Section 3. In contrast, we give a complete and easy-to-check characterization of all *yes*-instances of GEODESIC POLYGONIZATION, which is the special case of GEODESIC PSE where the input graph is restricted to a cycle.
- We show that LABELED GEODESIC MATCHING on the grid is \mathcal{NP} -hard by reduction from 3-PARTITION, see Section 4. This implies hardness of LABELED GEODESIC PSE on the grid. Our proof vitally exploits the space limitation of the grid. On the other hand, we show that LABELED GEODESIC MATCHING becomes easy if we loosen or drop this limitation.

We give a list of results and open questions in geodesic embeddability in Table 1. In the remainder of the paper, by a *grid geodesic* (or, even shorter, a *geodesic*) we mean a Manhattan-geodesic connecting two grid points on the grid. A *geodesic grid embedding* (or *geodesic embedding* for short) of a graph G is a drawing of G such that the vertices of G are mapped to grid points and the edges of G are mapped to interior-disjoint grid geodesics.

2 Geodesic Embeddability

In this section we ask whether a given planar graph has a geodesic embedding on the grid, that is, we allow the vertices to be mapped to arbitrary grid points. Clearly, this question makes only sense for graphs of maximum degree 4, but K_4 , for instance, does *not* have a geodesic embedding on the grid.

In the following, we show that a graph has a geodesic embedding on the grid if and only if it has an orthogonal embedding on the grid with at most one bend per edge. Liu et al. [8] characterized planar graphs which are orthogonally 1-bend embeddable and proposed an efficient decision algorithm for this problem. Hence, we have the somewhat surprising result that we can efficiently recognize graphs that admit a geodesic embedding on the grid.

Theorem 1. *Let $G = (V, E)$ be a planar graph. Then G has a geodesic embedding on the grid if and only if G is 1-bend embeddable on the grid.*

Proof. The “if”-direction is trivially true, so we immediately turn to the “only if”-direction. Suppose that G has a geodesic embedding \mathcal{E} on the grid. We turn \mathcal{E} into an *orthogonal representation* as introduced by Tamassia [15]. Such a representation consists of lists, one for each face of the given embedding. The list for a face f has, for each edge e incident to f , an entry describing (a) the shape of e in terms of left (-90°) and right ($+90^\circ$) turns, and (b) the angle that the edge makes with its successor in the cyclic order of the edges around f .

Since \mathcal{E} is geodesic, the angles along each edge sum up to a value in $\{-90^\circ, 0^\circ, +90^\circ\}$. From the representation of \mathcal{E} we compute a new representation where we replace the shape entry of each edge by the corresponding sum. The result is a valid representation since for each face the sum of the inner angles remains the same and for each vertex the sum of the angles between consecutive incident edges also remains the same. Since the new representation is valid, Tamassia’s flow network [15] yields the corresponding (1-bend) embedding of G . \square

3 Geodesic Point-Set Embeddability

In this section, we ask whether a given planar graph can be embedded on a given set of grid points. We assume that we are not given a bijection between vertices and points.

First, we show that this problem, GEODESIC PSE, is \mathcal{NP} -hard by reduction from the problem HAMILTONIAN CYCLE COMPLETION (HCC), which is \mathcal{NP} -hard [6]. Our proof also works in the case where the (Manhattan-) geodesics are not restricted to the grid. HCC is defined as follows. Given a non-Hamiltonian cubic graph G , decide whether G has two vertices u and v such that $G + uv$ (i) is planar, (ii) has a Hamiltonian cycle H , and (iii) has an embedding such that u and v are adjacent to at most two faces on the same side of H .

Theorem 2. *GEODESIC PSE is \mathcal{NP} -hard, even for subdivisions of cubic graphs.*

Proof. Our proof is by reduction from HCC. Given an instance $G = (V, E)$ of HCC, note that $n = |V|$ is even and let $k = \frac{n}{2} + 1$. Given three non-negative integers k_0, k_1, k_2 , let $P_0 = \{(-j, 0) \mid j = 0, \dots, k_0 - 1\}$, $P_1 = \{(j, nj) \mid j = 1, \dots, k_1\}$, $P_2 = \{(j, -nj) \mid j = 1, \dots, k_2\}$, and $P(k_0, k_1, k_2) = P_0 \cup P_1 \cup P_2$, see Fig. 1a. Note that the points in $P(k_0, k_1, k_2)$ are placed such that between any two consecutive non-empty rows of the integer grid there are $n - 1$ empty rows. We now construct a graph $G' = (V', E')$ by splitting every edge of G by a vertex of degree 2. This yields $|V'| = |V| + |E| = 2n - 1 + k$. In the following, we show that G' can be embedded on $P(2n - 1, k_1, k_2)$ for some k_1, k_2 with $k_1 + k_2 = k$ if and only if G is a *yes*-instance of HCC.

Assume G is a *yes*-instance of HCC. Then there is a pair $\{u, v\}$ of vertices such that $G + uv$ contains a Hamiltonian cycle and u and v are incident to two faces on either side of this cycle. Without loss of generality, we can assume that

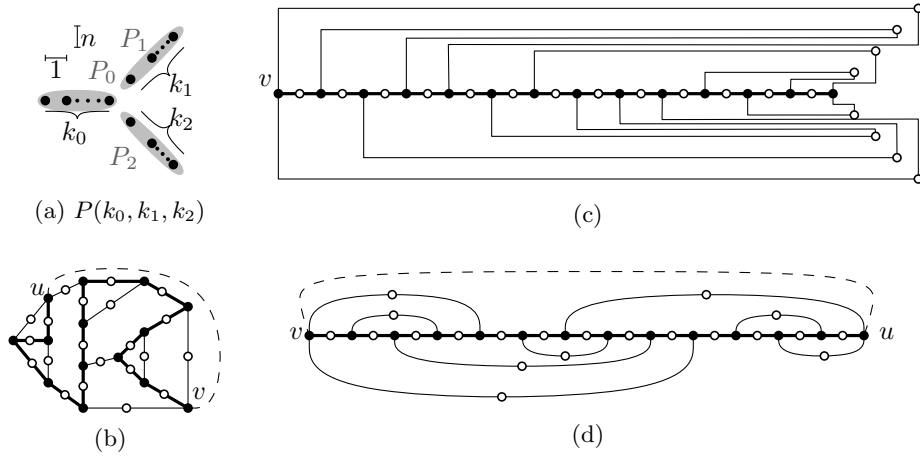


Fig. 1: Reduction of HCC to GEODESIC PSE.

uv is incident to the outer face. An example of a plane graph G' is depicted in Fig. 1b; the splitting nodes are marked with circles, the original nodes of G with black disks. Maintaining the combinatorial embedding, we can embed the Hamiltonian path connecting u and v including its splitting nodes on a set of $2n - 1$ points on a horizontal line as in Fig. 1d. We embed the faces inside the cycle above the path and the faces outside the cycle below. Since each vertex of G' has degree at most 3, each vertex has at most one edge going up or down—except u and v , which both have exactly one edge going up and one going down. Set k_1 and k_2 to the numbers of edges inside and outside the cycle, respectively. Then we can map the splitting vertices of the remaining edges to the point sets P_1 and P_2 , and route the edges as follows, see Fig. 1c. Each splitting node v that is mapped to a point in $P_1 \cup P_2$ has two neighbors, a left neighbor v^- and a right neighbor v^+ (according to their x -coordinates). We route the edge vv^- with one bend and the edge vv^+ with two bends. Note that the empty rows leave enough space for all horizontal edge segments.

Conversely, assume G' has a geodesic embedding on $P(2n - 1, k_1, k_2)$ with $k_1 + k_2 = k$. Then, the k vertices that are mapped to points in $P_1 \cup P_2$ are incident to at most $2k = n + 2$ edges. This is due to the fact that each such edge has its lexicographically larger endpoint in either P_1 or P_2 , and we claim that no point in $P_1 \cup P_2$ can be adjacent to more than two lexicographically smaller points. To see the claim, note that for any point $v \in P_1$ the set of lexicographically smaller points is contained in the third quadrant with respect to v . Clearly, at most two geodesics can go from v to points in any fixed quadrant. For points in P_2 , the argument is symmetric. Thus our claim holds.

Since G is cubic, G' has $3n$ edges. This leaves $3n - (n + 2) = 2n - 2$ edges incident to points in P_0 only. Since $|P_0| = 2n - 1$, P_0 induces a path π that alternates between vertices of degree 3 (original nodes) and degree 2 (splitting

nodes). There are two possibilities: either both endpoints—call them s and t —have degree 2 or both have degree 3. In the former case, π would contain $n - 1$ degree-3 vertices, and s and t would be adjacent to the only remaining degree-3 vertex (not in P_0). This would mean that G is Hamiltonian—contradiction.

Thus we may assume that s and t have degree 3. In this case, π witnesses a Hamiltonian path connecting s and t in G . This Hamiltonian path can be completed to a Hamiltonian cycle by an edge through the outer face of G . Since both u and v are incident to one edge pointing up and one edge pointing down from the path, they are incident to two faces on either side of the cycle in this embedding. This shows that G is indeed a *yes*-instance of HCC. \square

Now we turn to the case in which the instance consists of a simple cycle. We show that this problem, which we name GEODESIC POLYGONIZATION, can be solved efficiently. We start with a simple characterization of the *yes*-instances. To this end, we partition the grid points into two groups as follows. Let B be an axis-aligned rectangle. We say that a grid point p in B is *even* if its rectilinear distance to the lower left corner of B is even. Otherwise, we say that p is *odd*. We call a set of points *degenerate* if the set is contained in an axis-parallel line. It is clear that a degenerate point set does not have a polygonization. We now characterize all point sets that do have a polygonization. The proof, which is omitted here due to space limitations, is constructive, see our technical report [6]. It yields an efficient algorithm that computes a geodesic polygonization for any set of grid points with the given properties.

Theorem 3. *Let P be a non-degenerate set of points on the grid, let $\mathcal{B}(P)$ be the bounding box of P , and let h and w be the number of rows and columns spanned by $\mathcal{B}(P)$, respectively. Then P has a geodesic polygonization if and only if either (i) h or w is even or (ii) P does not contain all even points w.r.t. $\mathcal{B}(P)$.*

4 Labeled Geodesic Point-Set Embeddability

In Section 3, we showed that GEODESIC PSE is \mathcal{NP} -hard. In this section, we study the variant where the vertex–point correspondence is given. First, we settle the complexity of the problem on the grid.

Theorem 4. *Labeled GEODESIC PSE on the grid is \mathcal{NP} -hard, even if the given graph is a perfect matching.*

Proof. We reduce 3-PARTITION to Labeled GEODESIC MATCHING (LGM), which is a special case of Labeled GEODESIC PSE. An instance of 3-PARTITION consists of a multiset $A = \{a_1, \dots, a_{3m}\}$ of $3m$ positive integers, each in the range $(B/4, B/2)$, where $B = (\sum A)/m$, and the question is whether there exists a partition of A into m subsets A_1, \dots, A_m of A , each of cardinality three, such that the sum of the numbers in each subset is B . Since 3-PARTITION is *strongly* \mathcal{NP} -hard [3], we may assume that B is bounded by a polynomial in m .

Based on an instance A of 3-PARTITION, we now construct an instance M of LGM consisting of pairs of grid points such that M is a *yes*-instance of LGM




Fig. 2: Example of the reduction from 3-PARTITION to LGM using $A_1 = \{a_1, a_5, a_7\}$, $A_2 = \{a_2, a_3, a_8\}$, and $A_3 = \{a_4, a_6, a_9\}$ (not to scale)

if and only if A is a *yes*-instance of 3-PARTITION. Figure 2 shows an example instance M . The instance M consists of three types of point pairs.

The first type represents the numbers in A . We define $3m$ sets S_1, \dots, S_{3m} of grid points, all lying on the diagonal $\ell : y = -x$, in this order from left to right. For $1 \leq i \leq 3m$, the points in S_i occupy a_i consecutive grid points, and two consecutive sets are separated by a large gap of $L = Bm + m - 1$ grid points. The gap between the last point of S_{3m} and the origin is also L . The points in the sets T_{3m}, \dots, T_1 lie on the line $\ell' : y = -x + L$, in this order from left to right. Again, points within a set are consecutive grid points, and between consecutive sets there are large gaps of L grid points. The matching is as follows. For $1 \leq i \leq 3m$ and for $1 \leq j \leq a_i$, the j -th point in S_i (counting from the left) matches the j -th point in T_i (counting from the *right*). The a_i point pairs in $S_i \cup T_i$ represent the number a_i .

The second type of point pairs forms a sort of “dot mask”, the heart of our construction. These pairs lie on the x -axis. The geodesics between them are obviously line segments and pairwise disjoint. The leftmost segment goes from $-N$ to 0 , where $N = 3mL + mB + 2(m - 1)$. The following $m - 1$ segments have unit length and leave gaps of width B . The rightmost segment goes from L to $L + N$.

The third type of point pairs gives rise to geodesics that resemble “fences” ensuring that all geodesics that represent a number from A go through the same gap in the mask. There are $m - 1$ such pairs. Their upper endpoints are consecutive grid points on the diagonal ℓ . They lie above the points in S_1 , leaving

a gap of $m - 1$ grid points. The corresponding lower endpoints lie one unit above the, say, right endpoints of the unit-length segments on the x -axis. The matching is as follows: from left to right and for $1 \leq j \leq m - 1$, the j -th upper endpoint matches the j -th lower endpoint.

It is easy to see that any geodesic embedding of M induces a partition of A : due to the fences, all edges corresponding to the same element of A must be routed through the same gap of the dot mask, each of the m gaps has width B , and each of the mB edges must go through some gap.

Conversely, given a partition, we construct a geodesic embedding of the matching. We start by drawing the dot mask whose layout only depends on the numbers B and m . Then, we analyze the first subset of the partition, A_1 , and connect the points $S^1 = \bigcup_{a_j \in A_1} S_j$ to the corresponding points in $T^1 = \bigcup_{a_j \in A_1} T_j$, starting with the leftmost point in S^1 and the rightmost point in T^1 . For each connection, we use the bottommost geodesic that goes above all geodesics we have drawn so far. Next, we draw the first (that is, leftmost) fence. Also in this case, we use the bottommost geodesic that goes above all geodesics we have drawn so far. We repeat these two steps, connecting the points corresponding to a subset of the partition and drawing a fence. Since we left enough horizontal and vertical space, this process does not get stuck. The fences direct the next B connections into the gaps, which have exactly the right width.

Since we assumed that B is polynomial in m , the numbers L and N , which determine the grid size needed by M , are also polynomial in m . Given an embedding, the partition can be constructed from it efficiently, and vice versa. Thus our reduction is polynomial. \square

Next, we show that LGM becomes easy if we loosen or drop the space limitation of the grid. We call an instance of LGM—a set M of n pairs of grid points (we call such pairs also *edges*)—*sparse* if the minimum distance between any two occupied columns and between any two occupied rows is at least $n + 1$. In the remainder of this section, we give an efficient algorithm that solves sparse instances of LGM. Clearly, the algorithm can also be used for an instance that does not “live” on the grid, by underlaying the instance with a fine enough grid.

We say that an edge $e \in M$ is *downward* if its lexicographically larger endpoint e^+ lies below its lexicographically smaller endpoint e^- , otherwise e is *upward*. Clearly, M does not have a geodesic embedding if the bounding box of an edge *crosses* (that is, splits into two connected components) the bounding box of another edge. This can be tested easily, so from now on we assume that M is *non-crossing*, that is, no two bounding boxes of edges in M cross.

Let g and g' be any two geodesics. We say that g is *below* g' if there is a vertical line that intersects g below g' . We say that an edge $e \in M$ is *strictly below* an edge $e' \in M$ if, for any geodesic embedding γ of the two edges, $\gamma(e)$ is below $\gamma(e')$.

The *precedence graph* π_M is a directed graph whose vertex set is M and whose edges represent the strictly-below relationship. The precedence graph can be computed efficiently by a simple line sweep. It is clear that M does not have




Fig. 3: (a) Critical region $R^\uparrow(e, f)$ with $g \in M^\uparrow(e, f)$, (b) a matching M , (c) the precedence graph π_M and the extended precedence graph Π_M of M

a geodesic embedding if π_M contains a cycle. Now we construct a supergraph of π_M whose acyclicity is equivalent to the realizability of M .

For any point $a = (x_a, y_a) \in \mathbb{R}^2$, let $Q_1(a) = \{(x, y) \in \mathbb{R}^2 \mid x_a \leq x, y_a \leq y\}$ be the *first quadrant w.r.t. a* and define the other three quadrants w.r.t. a accordingly, in counterclockwise order. Let e be a downward edge and let f be any other edge in M . For such a pair (e, f) , we define the *upper critical region of e and f* as $R^\uparrow(e, f) = (Q_1(f^-) \cup Q_1(f^+)) \cap \mathcal{B}(e)$ (see Fig. 3a) and the *lower critical region of e and f* as $R^\downarrow(e, f) = (Q_3(f^-) \cup Q_3(f^+)) \cap \mathcal{B}(e)$. The critical regions for upward edges are defined by replacing Q_1 by Q_2 and Q_3 by Q_4 . Let $M^\uparrow(e, f)$ and $M^\downarrow(e, f)$ be the sets of edges in M with at least one endpoint in $R^\uparrow(e, f)$ and $R^\downarrow(e, f)$, respectively.

Let $G = (M, E)$ be a directed graph with vertex set M . We say that an edge (e, f) of G *produces* the edge (e, g) if $g \in M^\uparrow(e, f)$ and the edge (g, f) if $g \in M^\downarrow(f, e)$. Now the *extended precedence graph* Π_M is the closure of π_M with respect to production.

Lemma 2. *If Π_M contains a cycle, M does not admit a geodesic embedding.*

Proof. We claim that an edge (e, f) in Π_M means that if M has some geodesic embedding γ , then $\gamma(e)$ is below $\gamma(f)$. Clearly, the claim holds for every edge in π_M . Now suppose edge (e, g) has been produced by the edge (e, f) and M has a geodesic embedding γ . Then we know that $g \in M^\uparrow(e, f)$. Assume that e is downward. By definition, at least one of the endpoints of g —call it q —lies in $Q_1(p) \cap \mathcal{B}(e)$, where p is one of the endpoints of f . (In particular, p lies in $\mathcal{B}(e)$, otherwise $Q_1(p) \cap \mathcal{B}(e)$ would be empty.) Due to the existence of edge (e, f) , $\gamma(e)$ is below $\gamma(f)$ and thus below p . Since e is downward, $\gamma(e)$ must also be below q and hence below $\gamma(g)$. The case that e is upward and the case that (e, f) has produced an edge (g, f) can be argued symmetrically. Now induction yields the claim. \square

By Lemma 2, for M to have a geodesic embedding, it is necessary that Π_M is acyclic. We now show that this condition is also sufficient, by giving an algorithm that computes an embedding if Π_M is acyclic.

The algorithm sweeps a vertical line from left to right over the plane. Events occur only at the vertices of M . During the sweep, we partition the edges in M

into three groups. *Completed edges* have both endpoints to the left of the sweep-line. We have already embedded these edges as geodesics. *Partial edges* have one endpoint on either side of the sweep-line. A partial edge is embedded as a *partial geodesic* ending at the sweep-line. Finally, *untouched edges* have both endpoints to the right of the sweep-line. We have not started embedding these edges yet.

Let c and c' be two consecutive occupied grid columns, with c to the left of c' . Assume that we have already computed a partial geodesic embedding up to c . Let u_1, \dots, u_s be the set of upward partial edges which do not end at c , sorted from bottom to top (including the edges starting at c). We process the edges in this order. Assuming that we have already embedded u_1, \dots, u_{i-1} , we proceed depending on whether u_i ends at c' or not.

If u_i ends at c' , we embed u_i as the bottommost geodesic just above all edges u_1, \dots, u_{i-1} , that is, there is no geodesic for u_i containing a point strictly below this geodesic. Hence, u_1 has a vertical segment only on the last unoccupied column to the left of c' . By induction, u_i has vertical segments only on the last i unoccupied columns.

If u_i does not end at c' , let U_i denote the set of edges preceding u_i in Π_M . Then u_i must necessarily be embedded above all edges in U_i . We embed u_i as the bottommost geodesic above all edges u_1, \dots, u_{i-1} and above all endpoints of edges in U_i which are on c' . If there is no such restriction, we embed u_i as a straight-line segment.

We then proceed similarly with the downward edges. Let d_1, \dots, d_t be the set of partial downward edges that do not end at c , sorted from top to bottom, that is, sorted inversely to the upward edges. Let D_i denote the set of edges succeeding d_i in Π_M . We embed each edge d_i as the topmost geodesic below the geodesics d_1, \dots, d_{i-1} and below all endpoints of edges in D_i . As before, d_i has vertical segments only on the last i columns left of c' . A sample output of the algorithm is illustrated in Fig. 3b.

Since there are at most n edges by the definitions of the top- and bottommost geodesics, we need at most n unoccupied columns between c and c' . Since M is sparse, there are at least n unoccupied rows between two occupied rows on c' , so we can embed the given edges between two occupied points on c' .

Theorem 5. *Let M be a sparse non-crossing matching with n edges on the grid. Then M has a geodesic embedding if and only if Π_M is acyclic. In $\mathcal{O}(n^3)$ time, we can compute a geodesic embedding of M or prove that no such embedding exists.*

Proof. The “only if” part has been proved in Lemma 2, so we immediately turn to the “if” part. We first compute Π_M . If Π_M contains a cycle, we reject. Otherwise, we use the above embedding algorithm to compute an embedding γ of M .

Concerning running time, it is clear that π_M can be computed by a simple plane sweep in $\mathcal{O}(n^2)$ time. For computing Π_M , we need $\mathcal{O}(n^2)$ iterations, one for each edge. An iteration takes linear time since all endpoints in the corresponding two critical regions can be reported in linear total time. The embedding algorithm runs in $\mathcal{O}(n^2)$ time.




Fig. 4: Relative positions of downward/upward partial edges




Fig. 5: Impossible crossing of upward/downward partial edges

To show that γ is plane and geodesic, we maintain the following invariants during the execution of the algorithm.

1. All completed and partial edges are (partially) embedded as geodesics.
2. For every partial downward edge the partial embedding is not upward; vice versa for partial upward edges.
3. If the left endpoint of some downward edge e is above the left endpoint of an edge e' and the (partial) geodesic for e is below e' in the embedding, then Π_M contains a path from e to e' . A symmetric statement holds for upward edges.
4. The partial embedding respects all constraints corresponding to edges in Π_M .
5. No two (partial) geodesics intersect.

It is easy to see that invariants 1–4 are maintained by the algorithm. Invariant 1 yields that γ is geodesic. It remains to show that γ is plane (invariant 5).

Suppose that the algorithm introduces a crossing when going from grid column c to grid column c' and there is no crossing to the left of c . By definition of the top- and bottommost geodesic there is no intersection between two upward or two downward edges, respectively. That is, the algorithm can only introduce intersections between an upward and a downward edge. Let d_i be a downward edge and let u_j be an upward edge such that d_i and u_j intersect. Then d_i must be above u_j on c , otherwise there would be no crossing between the two edges. We now make a case distinction depending on whether or not there is an edge $e \in M$ with left endpoint e^- on c' such that d_i must be embedded below e and u_j must be embedded above e , that is, (d_i, e) and (e, u_j) are in Π_M .

First assume that there is no such edge e . Let V'_M be the points of V_M that lie to the left of or on c' . Let p_1 be the lowest point of V'_M such that d_i lies below p_1 . Similarly, let p_2 be the highest point of V'_M such that u_j lies above p_2 . Clearly, p_2 is below p_1 and by assumption there are at least n unoccupied rows between p_1 and p_2 . By definition of the top- and bottommost geodesic for d_i and u_j the two edges do not cross, see Fig. 4.

Now assume that there is an edge e with (d_i, e) and (e, u_j) in Π_M . If the left endpoints d_i^- and u_j^- are in the same column, there is an edge (u_j, d_i) in π_M , which induces a cycle in Π_M , contradicting the assumption that Π_M is acyclic. Otherwise, the endpoints d_i^- and u_j^+ are in different columns.

Assume that d_i^- is to the left of u_j^- . Since d_i is above u_j , its right endpoint d_i^+ cannot be in the critical region $R^\downarrow(u_j, e)$ since this would imply that d_i must be below u_j , which violates invariant 4 at c . Hence, d_i^+ must be to the right of u_j^+ . In this case, however, u_j^+ is in $R^\uparrow(d_i, e)$, that is, d_i must be below u_j , which again violates invariant 4 at c (see Fig. 5). The case that d_i^- is to the right of u_j^- is similar. \square

Acknowledgments. We thank the anonymous referees for their detailed and helpful comments. AW thanks Ferran Hurtado for inviting him to a workshop in 2006, Manuel Abellanas for bringing LABELED GEODESIC PSE to his attention there, and Stefan Langerman and Pat Morin for discussions about the off-the-grid version of GEODESIC MATCHING.

References

1. S. Cabello. Planar embeddability of the vertices of a graph using a fixed point set is NP-hard. *J. Graph Algorithms Appl.*, 10(2):353–363, 2006.
2. E. Demaine. Simple polygonizations. <http://erikdemaine.org/polygonization/>, 2007. Accessed May 30, 2009.
3. M. R. Garey and D. S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.
4. X. Goaoc, J. Kratochvíl, Y. Okamoto, C.-S. Shin, A. Spillner, and A. Wolff. Untangling a planar graph. *Discrete Comput. Geom.*, 2009. Appeared online at <http://dx.doi.org/10.1007/s00454-008-9130-6>.
5. F. Hurtado. Personal communication, 2006.
6. B. Katz, M. Krug, I. Rutter, and A. Wolff. Manhattan-geodesic point-set embeddability and polygonization. Technical Report 2009-XX, Universität Karlsruhe, 2009. Appeared online at <http://digbib.ubka.uni-karlsruhe.de/volltexte/1000000000>.
7. M. Kaufmann and R. Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms Appl.*, 6(1):115–129, 2002.
8. Y. Liu, P. Marchioro, R. Petreschi, and B. Simeone. Theoretical results on at most 1-bend embeddability of graphs. *Acta Math. Appl. Sinica (English Ser.)*, 8(2):188–192, 1992.
9. J. O’Rourke. Uniqueness of orthogonal connect-the-dots. In G. Toussaint, editor, *Computational Morphology*, pages 97–104. North-Holland, 1988.
10. J. Pach and R. Wenger. Embedding planar graphs at fixed vertex locations. *Graph. Combinator.*, 17(4):717–728, 2001.
11. R. Raghavan, J. Cohoon, and S. Sahni. Single bend wiring. *J. Algorithms*, 7(2):232–257, 1986.
12. D. Rappaport. On the complexity of computing orthogonal polygons from a set of points. Technical Report SOCS-86.9, McGill University, Montréal, 1986.
13. F. Rendl and G. Woeginger. Reconstructing sets of orthogonal line segments in the plane. *Discrete Math.*, 119(1-3):167–174, 1993.
14. W. Schnyder. Embedding planar graphs on the grid. In *Proc. 1st ACM-SIAM Symp. on Discrete Algorithms (SODA’90)*, pages 138–148, 1990.
15. R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput.*, 16(3):421–444, 1987.