Universität Karlsruhe - Fakultät für Informatik - Bibliothek - Postfach 6980 - 76128 Karlsruhe

Manhattan-Geodesic Point-Set Embeddability and Polygonization

Bastian Katz, Marcus Krug, Ignaz Rutter, Alexander Wolff

Interner Bericht 2009-17

Universität Karlsruhe (TH)
Forschungsuniversität · gegründet 1825

ISSN 1432-7864

Fakultät für Informatik

Manhattan-Geodesic Point-Set Embeddability and Polygonization

Bastian Katz* Marcus Krug* Ignaz Rutter* Alexander Wolff†

Technical Report 2009-17
Faculty of Informatics, Karlsruhe Institute of Technology (KIT)

Abstract

The *Manhattan-geodesic* drawing convention for graphs 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. In this work, we consider the problem whether a given planar graph has a Manhattan-geodesic drawing such that the vertices are embedded onto a given set of points. We show that this problem, which we call Geodesic Point-Set Embeddability, is $\mathcal{NP}$-hard in general. On the positive side, we give a simple characterization of the *yes*-instances for the special case that the graph is a cycle. We call this problem *geodesic polygonization*. Our characterization can easily be checked in linear time, and we give an $O(n \log n)$ time algorithm to compute a polygonalization if one exists. The results are obtained for drawings on a grid, but can be extended to the setting without this restriction.

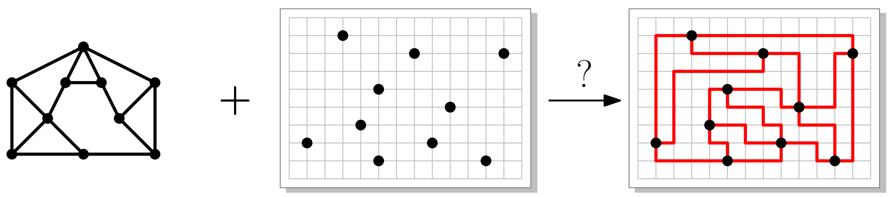

1 Introduction

One of the most popular conventions for drawing planar graphs 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. In the *Manhattan-geodesic* drawing convention, edges additionally have to be drawn as *monotone* chains. Such chains are called *Manhattan paths*. Manhattan paths are geodesics with respect to the Manhattan metric. 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. It also combines the concept of orthogonal drawings with the straight-line drawing convention, where edges must be embedded as Euclidean geodesics. In this work, we consider the setting where we are given not just a graph, but also a

*Faculty of Informatics, Karlsruhe Institute of Technology (KIT), {`firstname.lastname`}`@kit.edu`
†Institut für Informatik, Universität Würzburg, `www1.informatik.uni-wuerzburg.de/en/staff`

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. This report complements our work on Manhattan-geodesic drawings [KKRW10]. There, we also considered (a) the more restrictive case where one also gets a mapping between the vertices and the points and (b) the less restrictive case where one does not get a point set at all.

Related Work. Kaufmann and Wiese [KW02] 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 [Cab06] showed that zero bends are hard, too. In other words, 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 [Rap86] 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 [O'R88] 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 the reader to Demaine's survey [Dem07] 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 [RW93]. They showed that given a set P 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 [Hur06] gave a simple $O(n \log n)$-time algorithm for the matching problem under the geodesic drawing convention. The idea is to alternatingly go up and down the occupied grid columns.

In this paper, we investigate the problem GEODESIC PSE of deciding 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. Note that the variant where a bijection is given is $\mathcal{NP}$-hard as well [KKRW10]. First, we prove that in general, this problem is hard, using a two-step reduction from HAMILTONIAN CYCLE, see Section 2. Second, we show that there is a simple characterization of the *yes*-instances for the case that the graph is a cycle, see Section 3. This is the main result of this report. The proof of the correctness of our characterization is constructive and yields a simple $O(n \log n)$-time algorithm for this special case of GEODESIC PSE, which we call GEODESIC POLYGONIZATION.

2 Complexity of Geodesic Point-Set Embeddability

In this section, we show that GEODESIC PSE is $\mathcal{NP}$-hard by reduction from the problem HAMILTONIAN CYCLE COMPLETION (HCC), which 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. We first show that HCC is hard.

Lemma 1. HAMILTONIAN CYCLE COMPLETION *is $\mathcal{NP}$-hard.*

Proof. We reduce from the $\mathcal{NP}$-hard problem HAMILTONIAN CYCLE (HC), where the task is to decide whether a given planar cubic graph is Hamiltonian [GJ79]. Given an instance

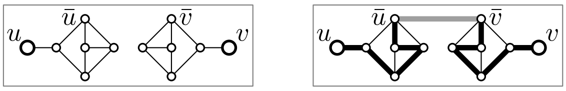

(a) (b)

Fig. 1: Gadget that replaces the edge uv of G in G_{uv} (when reducing from HC to HCC)

$G = (V, E)$ of HC, we construct, for each $uv \in E$, an instance G_{uv} of HCC. The graph G_{uv} is a copy of G where we replace uv by the gadget depicted in Fig. 1a. We claim that an edge uv lies on some Hamiltonian cycle in G if and only if G_{uv} is a *yes*-instance of HCC.

We first assume that G_{uv} is a *yes*-instance of HCC. Then there is a pair $\{a, b\}$ of vertices such that $G_{uv}+ab$ is Hamiltonian, see Fig. 1b. The vertices a and b must lie in our gadget, one on each side (albeit not necessarily $\{a, b\} = \{\bar{u}, \bar{v}\}$); otherwise u or v would remain separators. It is obvious how to transform a Hamiltonian cycle in $G_{uv}+ab$ into a Hamiltonian cycle in G.

Conversely, assume G contains a Hamiltonian cycle H that uses some edge uv. We observe two things. First, if we add the edge $\bar{u}\bar{v}$ to G_{uv}, then the concatenation of $\bar{u}\bar{v}$, the bold black edges in the gadget, and $H - uv$ forms a Hamiltonian cycle $\bar{H}$ in $G_{uv}+\bar{u}\bar{v}$. Second, the planar embedding that $G_{uv}+\bar{u}\bar{v}$ inherits from G and from the embedding of the gadget as depicted in Fig. 1b makes sure that $\bar{u}$ and $\bar{v}$ are incident to two faces on each side of $\bar{H}$.

Thus, we could apply a hypothetical algorithm for HCC to G_{uv} for each edge uv of G. As soon as the algorithm finds a vertex pair $\{a, b\}$ such that $G_{uv} + ab$ is Hamiltonian, it's easy to construct the corresponding Hamiltonian cycle in G. If, on the other hand, the algorithm decides for each edge uv of G that G_{uv} is a *no*-instance, we can conclude that G is not Hamiltonian. This yields the $\mathcal{NP}$-hardness of HCC. □

Now we are ready to show the hardness of Geodesic PSE.

Theorem 1. 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. Let $k = n/2 + 1$. Given three non-negative integers k_0, k_1, k_2, let $P_0 = \{(-j, 0) \mid j = 0, \ldots, k_0 - 1\}$, $P_1 = \{(j, nj) \mid j = 1, \ldots, k_1\}$, $P_2 = \{(j, -nj) \mid j = 1, \ldots, k_2\}$, and $P(k_0, k_1, k_2) = P_0 \cup P_1 \cup P_2$, see Fig. 2a. 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 uv is incident to the outer face. An example of a plane graph G' is depicted in Fig. 2b; 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. 2c. 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

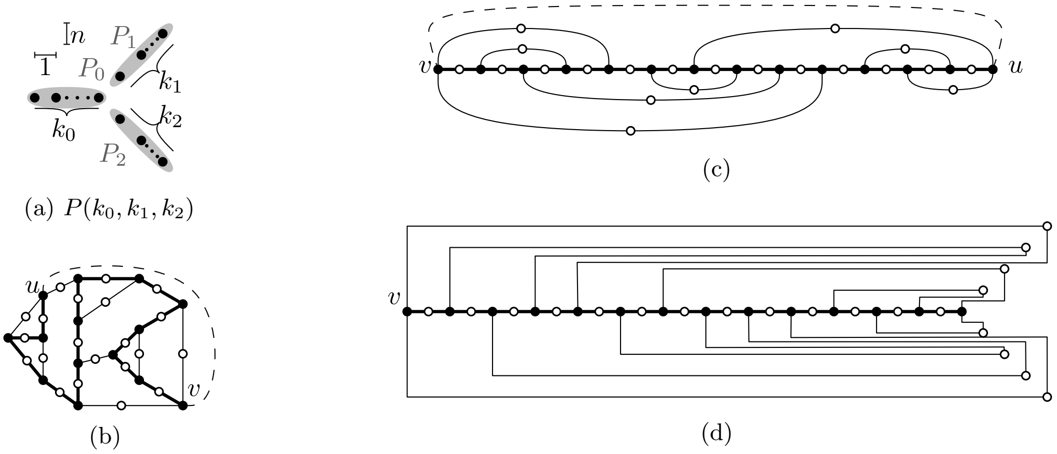

Fig. 2: Reduction of HCC to Geodesic PSE.

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. 2d. 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. □

Note that the proof of Theorem 1 implies that the result extends to the problem where the (Manhattan-) geodesics are not restricted to the grid.

3 A Special Case: Geodesic Polygonization

In this section, we present our main result—a simple characterization of the *yes*-instances of Geodesic Polygonization. The proof is constructive; it yields an efficient algorithm that, for a given set of grid points, computes a geodesic polygonization or proves that such a polygonization does not exist.

To this end, we partition the grid points in a given axis-parallel rectangle B into two groups as follows. We say that a grid point p in B is *even (with respect to B)* if its rectilinear distance to the lower left corner of B is even. Otherwise, we say that p is *odd (with respect to B)*, see Fig. 3a.

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.

Theorem 2. *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 numbers 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 with respect to $\mathcal{B}(P)$.*

We can test, in $O(n)$ time, whether a given set of n points has a geodesic polygonization, and if so, compute one within $O(n \log n)$ time.

Before we prove this theorem, let us quickly consider the case that we are not restricted to the grid. If P is a non-degenerate set of points in the plane, we can use the grid $\Gamma(P)$ induced by P. If P fulfills the requirements of Theorem 2 with respect to $\Gamma(P)$, we have a very natural polygonization of P. Otherwise—if $\Gamma(P)$ has an odd number of both rows and columns, and P contains all even points with respect to $\mathcal{B}(P)$—it is sufficient to introduce *one* additional column between any two existing columns of $\Gamma(P)$ to meet the requirements of Theorem 2. Hence, we obtain the following corollary of Theorem 2.

Corollary 1. *Every non-degenerate set of points in the plane has a geodesic polygonization off the grid. Such a polygonization can be computed in $O(n \log n)$ time.*

Proof of Theorem 2. Given the above characterization is correct, testing for the existence of a polygonalization can clearly be done in linear time: Determine the bounding box, count even points and return *yes* if either one dimension of the bounding box is even or if not all even points of the bounding box are occupied. Hence, the focus will be on the proof of the characterization. It is constructive and can be extended to an efficient algorithm. It is easy

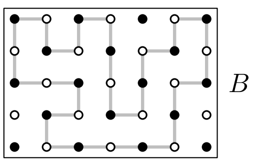

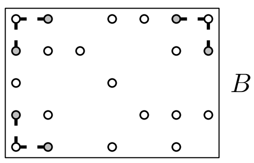

Fig. 3: (a) A polygon hits even grid points (black disks) and odd grid points (circles) alternatingly. (b) If a point set contains a corner of its bounding box, we can assume that it also contains the (marked) points at distance 1 from that corner.

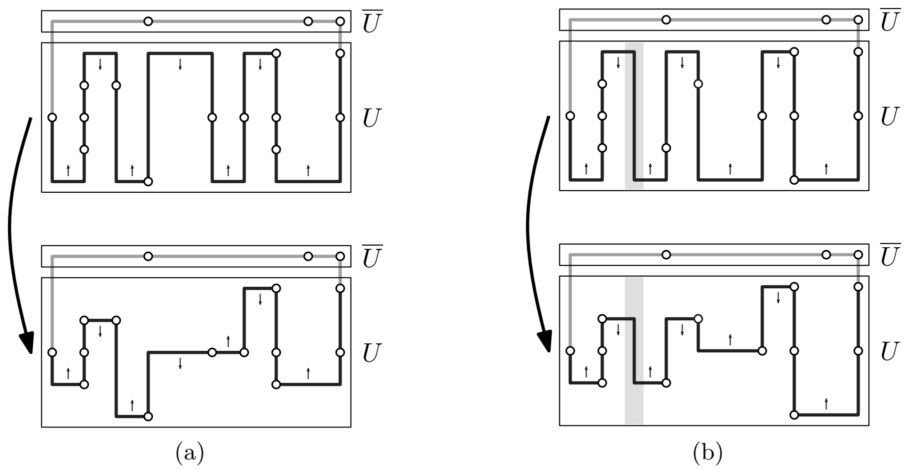

Fig. 4: Polygonalization according to case 1: U contains (a) an even number of occupied columns or (b) an odd number of occupied columns plus (at least) one empty column, marked gray.

to see that each of the cases considered in the proof can be solved by a simple algorithm with running time $\mathcal{O}(n \log n)$.

Unless stated otherwise, even and odd always refers to $\mathcal{B}(P)$. We first show that P does not have a polygonization if h and w are both odd *and* P contains all even points: Observe that any polygonization of P must contain an equal number of even and odd grid points on its boundary (see Fig. 3a). If h and w are both odd, the number of even points in $\mathcal{B}(P)$ exceeds the number of odd points in $\mathcal{B}(P)$ by one. Hence, P does not have a geodesic polygonization in this case.

In the remainder of the proof, we show that we can construct a polygonization in the other cases. Note that the fact that P has a polygonization is invariant under rotation by multiples of 90 degrees and reflection at vertical or horizontal lines.

The key idea of the proof is to partition (some rotation or reflection of) P into two sets U and $\overline{U}$ such that $\overline{U}$ contains all the points on the topmost occupied row and $U = P \setminus \overline{U}$. A *nice path* for U is a path that (a) connects all points in U by geodesics, (b) ends in the topmost point of U in the leftmost column and in the topmost point of U in the rightmost column, and (c) does not occupy the grid points above the two endpoints. To motivate (c), see the thick gray path in Fig. 4a. Note that, by definition, $\overline{U}$ is not empty, and by construction of the nice path, it is easy to connect the endpoints of the nice path in U by a path that contains all points in $\overline{U}$ such that the concatenation of the two paths yields the desired geodesic polygonization of P. Thus, the problem of finding a geodesic polygonization of P reduces to the problem of finding a nice path in U.

Without loss of generality, we can also assume the following. If a corner of $\mathcal{B}(P)$ lies in P, then both points in $\mathcal{B}(P)$ at distance 1 from the corner also lie in P. This follows from the fact that any polygonization containing the corner of $\mathcal{B}(P)$ must contain these two points as well (see Fig. 3b). This observation ensures that any partition of P into $\overline{U}$ and U as described above has the property that the leftmost column and the rightmost column of U are the same as those of P, that is, any nice path can be extended to a polygonization. We now consider two cases.

Case 1: We can partition P (or some rotation or reflection of P) into $\overline{U}$ and U as described above such that either (i) the number of occupied columns in U is even or (ii) it is odd and there is at least one unoccupied column in U.

First, assume that the number of occupied columns is even. Then we can sweep the points in U from left to right and alternatingly from top to bottom and vice versa, starting at the topmost point of the leftmost column in downward direction. Imagine that all bends of our tour lie on the boundary of $\mathcal{B}(U)$. Then some pairs of points of U that are consecutive in our tour may be connected by U-shaped pieces of the tour, which are not geodesics. This, however, can easily be fixed by shortening each U-shape in the sense that its horizontal part is moved away from the boundary of $\mathcal{B}(U)$ until it hits at least one of the two endpoints of the U-shape. The result is either an L-shaped or simply a horizontal connection, and hence a geodesic. This process is depicted in Fig. 4a.

Since the number of columns is even, the endpoints are exactly the topmost points on the leftmost and rightmost column, respectively, and the unoccupied points above the endpoints of the path are not used.

Next, assume that the number of occupied columns is odd and there is an unoccupied column (somewhere between, but not necessarily adjacent to two occupied columns). In this case we use the same approach with the only difference that we also alternate the vertical sweeping at exactly one of the unoccupied columns. Note that this does not necessarily mean that there is a bend in the unoccupied column. As illustrated in Fig. 4b, the last point before the unoccupied row is linked to the first point after the unoccupied row by two horizontal and one (possibly degenerated) vertical straight-line segment that uses the unoccupied column.

Case 2: We cannot partition P as described in case 1. Then the numbers of occupied columns and rows of P are both odd, and every partition (of a rotation or reflection) of P into $\overline{U}$ and U has the property that every row and every column in U has at least one occupied point. We know that $\mathcal{B}(P) \setminus P$ contains an even point, that is, there is an unoccupied even point in $\mathcal{B}(P)$.

Before we proceed, we introduce the following notation, see Fig. 5. Let X be a non-degenerate set of points on the grid, and let q be a grid point in $\mathcal{B}(X)$. Then we define the *height of* q *with respect to* X to be $\max_{p \in X}\{y(p)\} - y(q)$, where $y(r)$ denotes the y-coordinate of a point $r \in \mathbb{R}^2$. Similarly, we define the *depth of* q *with respect to* X to be $y(q) - \min_{p \in X}\{y(p)\}$. Finally, we define the *stretch of* X to be $\max_{p \in X}\{y(p)\} - \min_{p \in X}\{y(p)\} + 1$, that is, the stretch of X is the number of rows spanned by X.

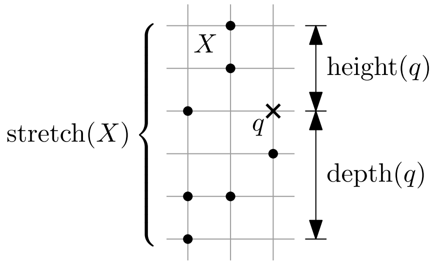

Fig. 5: Stretch of a set X of grid points; height and depth of a grid point q w.r.t. $\mathcal{B}(X)$

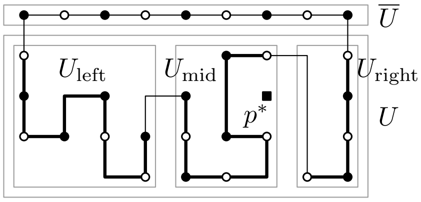

Fig. 6: Partition of P into U and $\overline{U}$, and partition of U into U_{left}, U_{mid}, U_{right}

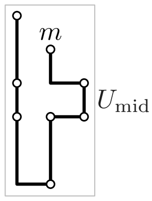

Fig. 7: Relaxed definition of a nice path for U_{mid}

We now claim the following. In the above situation, we can find a partition of some reflection or rotation of P into two sets $\overline{U}$ and U as well as a partition of U into three sets U_{left}, U_{mid}, and U_{right} as depicted in Fig. 6 such that the following three requirements are fulfilled.

(R1) The set $\overline{U}$ contains the points in the topmost row of P, and $U = P \setminus \overline{U}$.

(R2) The set U_{mid} consists of the points in three consecutive columns of U; the corresponding subsets of U_{mid} are $U^1_{\text{mid}}, U^2_{\text{mid}}$, and U^3_{mid} (from left to right). Let $U^{23}_{\text{mid}} = U^2_{\text{mid}} \cup U^3_{\text{mid}}$, and let ℓ^1 denote the lowest point in U^1_{mid}. Then, at least one of the following three statements holds:

- (U1) The stretch of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is odd.
- (U2) The middle column of U_{mid} contains an unoccupied point p^* with even height and odd depth with respect to $U^{23}_{\text{mid}} \cup \{\ell^1\}$.
- (U3) The rightmost column of U_{mid} contains an unoccupied point p^* with odd height and even depth with respect to $U^{23}_{\text{mid}} \cup \{\ell^1\}$.

(R3) The sets U_{left} and U_{right} are the (possibly empty) sets of points of U to the left and to the right of U_{mid}, respectively.

We call a partition of P fulfilling (R1)–(R3) an *odd* partition.

We postpone the proof of the claim for now. Given an odd partition, we first show how to find a nice path for U by finding and linking nice paths for U_{left}, U_{mid}, and U_{right}. We slightly modify the definition of a nice path for U_{mid} by allowing the path to end in a point m of the middle column of U_{mid} if all points in U_{mid} either lie to the left or below m and if the path does not go through the point to the right of m, see Fig. 7.

Note that we can find nice paths for U_{left} and U_{right} as in case 1 since both sets consist of an *even* number of occupied columns. In other words, we need only consider U_{mid}. The first part of the nice path for U_{mid} consists of a vertical straight-line segment containing all points on the leftmost column of U_{mid}. Now we follow the case distinction concerning the shape of U^{23}_{mid} in the above definition of an odd partition.

Case 2.U1: The stretch of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is odd.

We sweep the second and third column of U_{mid} row by row from bottom to top, starting at the bottommost row of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ and alternating the walking direction between right and left in each (not necessarily occupied) row. This yields an ordering of the points in U^{23}_{mid}. We link consecutive points in this ordering by geodesics (see Fig. 8a). The geodesics can

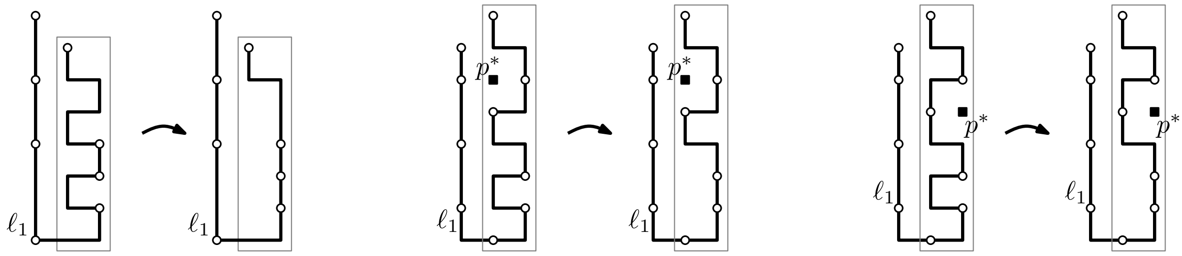

(a) (U1) (b) (U2) (c) (U3)

Fig. 8: Nice paths for U_{mid} according to subcases (U1)–(U3). Each figure shows the sweep starting in the bottommost row of $U^{23}_{\text{mid}} \cup \{\ell_1\}$ (left) and the resulting geodesic path (right).

be drawn such that they are interior-disjoint since, if two consecutive points aren't already connected by a geodesic, they are either in the same column or there must be at least one empty row between them. Let m be the last point on the resulting path. Since the stretch of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is odd, the path reaches m coming from below or from the left. Hence, our path is a nice path for U^{23}_{mid}. It can be connected to ℓ_1 since the sweep goes left-to-right through the bottommost row of $U^{23}_{\text{mid}} \cup \{\ell^1\}$.

Case 2.U2: The middle column of U_{mid} contains an unoccupied point p^* with even height and odd depth with respect to $U^{23}_{\text{mid}} \cup \{\ell^1\}$.

In this case we compute the nice path as illustrated in Fig. 8b. We sweep U^{23}_{mid} from bottom to top starting at the lowest (not necessarily occupied) row in $U^{23}_{\text{mid}} \cup \{\ell^1\}$ from left to right. We alternate the walking direction between right and left in each (not necessarily occupied) row *skipping the row that contains* p^*. Since the depth of p^* is odd, the walking direction is left-to-right in the row below p^*. Hence, our sweep leaves out p^*. We link points that were swept consecutively by geodesics as in the previous subcase. Again, the resulting path is nice.

Case 2.U3: The rightmost column of U_{mid} contains an unoccupied point p^* with odd height and even depth with respect to $U^{23}_{\text{mid}} \cup \{\ell^1\}$.

This case is very similar to the previous subcase. We skip the row containing p^* in the sweep. Since the depth of p^* is even, the walking direction in the row below p^* is right-to-left. Hence, we leave out a point in the right column, which is exactly p^*. An example for this case is depicted in Fig. 8c.

In all three cases the two geodesic paths for U^1_{mid} and U^{23}_{mid} can be combined to a nice path for U_{mid} since the path for U^{23}_{mid} (a) starts at the leftmost point on the lowest row of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ and (b) stops at the topmost point of U^{23}_{mid} without going through the point to the right of the topmost point. Hence, we have found a nice path for U_{mid}.

We now turn to the claim whose proof we postponed before.

Claim: Suppose that case 1 does not apply, that is, the numbers of occupied columns and rows of P are both odd and the partition of any rotation or reflection of P into $\overline{U}$ and U has the property that every row and every column in U has at least one occupied point and there is an even point in $\mathcal{B}(P) \setminus P$.

In this setting, we show how to find an odd partition of P, and in cases (U2) and (U3) a corresponding point p^* as stated in the claim. Note that any rotation and any reflection φ of the point set maps even points to even points and odd points to odd points with respect to the new positions, respectively. In order to prove the claim, we consider two cases:

Case (i): There is an unoccupied even point p^0 on the boundary of $\mathcal{B}(P)$.

We assume without loss of generality that p^0 is in the bottom row r_0 of $\mathcal{B}(P)$. Since r_0 is non-empty, there must be an occupied point on r_0, say, to the left of p^0. Let p^* be the leftmost unoccupied even point of r_0 such that there is an occupied point to the left of p^*. Since p^* is an even point in r_0, there is an even number of occupied columns in U to the left of p^*. By the choice of p^* this number of columns is at least 1 and hence there are at least two occupied columns to the left of p^*.

Let $\overline{U}$ be the set of points in the top row of P. Let $U = P \setminus \overline{U}$, and let U_{mid} be the subset of U in the column of p^* and the two columns to its left. Since p^* is an even point on the lowest row, there must be an even number of columns to the left and right of U_{mid}, respectively, and we fix U_{left} and U_{right} accordingly.

Since p^* is on the lowest row the depth of p^* with respect to $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is zero and, hence, even. If the stretch of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is even, then the height of p^* with respect to U^{23}_{mid} is odd and its depth is even. Since p^* is in the rightmost column with respect to U_{mid}, this yields an odd partition according to case (U3). If the stretch of $U^{23}_{\text{mid}} \cup \{\ell^1\}$ is odd, this yields an odd partition according to case (U1).

Case (ii): All even points on the boundary of $\mathcal{B}(P)$ are occupied.

Again, we let $\overline{U}$ be the top row of $\mathcal{B}(P)$ and let $U = P \setminus \overline{U}$. Let p^0 be a leftmost unoccupied even point in $\mathcal{B}(U)$. Such a point must exist, since all even points on the boundary are in P and $\overline{U}$ only contains points on the boundary. This also implies that there is at least one occupied column to the left of p^0. Note that by assumption all even points to the left of p^0 are occupied. We distinguish two cases:

First, suppose that there is an even number of columns to the left of p^0. Let U^{23}_{mid} consist of the points in the column of p^0 and the two columns to its left and choose U_{left} and U_{right} accordingly. Let p^* be the lowest unoccupied even point in the column of p^0. Since p^* is an even point with an even number of columns to the left and since all even points below p^* are occupied, p^* has an even depth with respect to U^{23}_{mid}. If the height of p^* is even, then the stretch of U_{mid} is odd, hence, this yields an odd partition according to case (U1). Otherwise, this yields an odd partition according to case (U3).

Next, assume that there is an odd number of columns to the left of p^0. In this case we mirror the instance on a vertical line, that is, p^0 is in the rightmost column containing an unoccupied even point and there is an odd number of columns to the right of p^0 whose even points are all occupied. We let U_{mid} consist of all the points of U in the columns of p^0 and its two neighboring columns and choose U_{left} and U_{right} accordingly. Let p^* be the lowest unoccupied even point on the column of p^0. Since p^* is an even point with an odd number of columns to the left and since all even points on the column to its right are occupied, p^* has an odd depth with respect to U_{mid}. If the height of p^* is odd as well, then the stretch of U_{mid} is odd, and therefore, this yields an odd partition according to case (U1). Otherwise the height of p^* is even, and this yields an odd partition according to case (U2). $\square$

References

[Cab06] Sergio 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.

[Dem07] Erik Demaine. Simple polygonizations. http://erikdemaine.org/polygonization/, 2007. Accessed May 30, 2009.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness.* W. H. Freeman and Company, 1979.

[Hur06] Ferran Hurtado. Personal communication, 2006.

[KKRW10] Bastian Katz, Marcus Krug, Ignaz Rutter, and Alexander Wolff. Manhattan-geodesic embedding of planar graphs. In David Eppstein and Emden Gansner, editors, *Proc. 17th Internat. Sympos. Graph Drawing (GD'09)*, *Lect. Notes Comput. Sci.* Springer-Verlag, 2010. To appear.

[KW02] Michael Kaufmann and Roland Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms Appl.*, 6(1):115–129, 2002.

[O'R88] Joseph O'Rourke. Uniqueness of orthogonal connect-the-dots. In G.T. Toussaint, editor, *Computational Morphology*, pages 97–104. North-Holland, 1988.

[Rap86] David Rappaport. On the complexity of computing orthogonal polygons from a set of points. Technical Report SOCS-86.9, McGill University, Montréal, 1986.

[RW93] Franz Rendl and Gerhard Woeginger. Reconstructing sets of orthogonal line segments in the plane. *Discrete Math.*, 119(1-3):167–174, 1993.