A rounding algorithm for approximating minimum Manhattan networks

Victor Chepoi, Karim Nouioua, Yann Vaxès *

Abstract. For a set T of n points (terminals) in the plane, a *Manhattan network* on T is a network $N(T) = (V, E)$ with the property that its edges are horizontal or vertical segments connecting points in $V \supseteq T$ and for every pair of terminals, the network $N(T)$ contains a shortest l_1-path between them. A *minimum Manhattan network* on T is a Manhattan network of minimum possible length. The problem of finding minimum Manhattan networks has been introduced by Gudmundsson, Levcopoulos, and Narasimhan (APPROX'99) and its complexity status is unknown. Several approximation algorithms (with factors 8,4, and 3) have been proposed; recently Kato, Imai, and Asano (ISAAC'02) have given a factor 2 approximation algorithm, however their correctness proof is incomplete. In this note, we propose a rounding 2-approximation algorithm based on a LP-formulation of the minimum Manhattan network problem.

1 Introduction

A *rectilinear path* P between two points p, q of the plane $\mathbb{R}^2$ is a path connecting p and q and consisting of only horizontal and vertical line segments. More generally, a *rectilinear network* $N = (V, E)$ consists of a finite set V of points of $\mathbb{R}^2$ (the vertices of N) and of a finite set of horizontal and vertical segments connecting pairs of points of V (the edges of N). The *length* $l(P)$ (or $l(N)$) of a rectilinear path P (or of a rectilinear network N) is the sum of lengths of its edges. The l_1-distance between two points $p = (p^x, p^y)$ and $q = (q^x, q^y)$ in the plane $\mathbb{R}^2$ is $d(p, q) := ||p - q||_1 = |p^x - q^x| + |p^y - q^y|$. An l_1-*path* between two points $p, q \in \mathbb{R}^2$ is a rectilinear path connecting p, q and having length $d(p, q)$.

Given a set $T = \{t_1, \ldots, t_n\}$ of n points (*terminals*) in the plane, a *Manhattan network* [3] on T is a rectilinear network $N(T) = (V, E)$ such that $T \subseteq V$ and for every pair of points in T, the network $N(T)$ contains an l_1-path between them. A *minimum Manhattan network* on T is a Manhattan network of minimum possible length and the Minimum Manhattan Network problem (*MMN problem*) is to find such a network.

The minimum Manhattan network problem has been introduced by Gudmundsson, Levcopoulos, and Narasimhan [3] and its complexity status is unknown. Gudmundsson et al. [3] proposed a factor 4 and a factor 8 approximation algorithms with different time complexity. They also conjectured that there exists a 2-approximation algorithm for this problem. Kato, Imai, and Asano [4] presented a factor 2 approximation algorithm, however, their correctness proof is incomplete. Following [4], Benkert, Shirabe, and Wolf [1] outlined a factor 3 approximation algorithm and presented a mixed-integer programming formulation of the MMN problem. Notice that all four mentioned algorithms are geometric and some of them employ results from computational geometry. Nouioua [6] presented another factor 3 approximation algorithm based on the primal-dual method from linear programming. In this note we present a rounding method applied to the optimal solution of the linear program described in [1, 6] and leading to a 2-approximation algorithm for the minimum Manhattan network problem. For approximation algorithms based on rounding techniques, see the book by Vazirani [10].

Gudmundsson et al. [3] introduced the minimum Manhattan networks in connection with the construction of sparse geometric spanners preserving the l_1-distances between the terminals. Such spanners have applications in VLSI circuit design, network design, distributed algorithms and other areas. Lam, Alexandersson, and Pachter [5] suggested to apply minimum Manhattan networks to design efficient search spaces for pair hidden Markov model (PHMM) alignment algorithms.

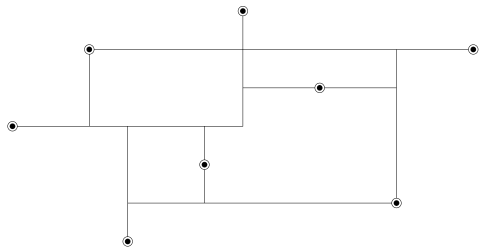

Figure 1: A minimum Manhattan network

2 Properties and LP-formulation

In this section, we present several properties of minimum Manhattan networks. First, we define some notations. Denote by $[p, q]$ the linear segment having p and q as endpoints. The set of all points of $\mathbb{R}^2$ lying on l_1-paths between p and q constitute the smallest axis-parallel rectangle $R(p, q)$ containing the points p, q. For two terminals $t_i, t_j \in T$, set $R_{ij} := R(t_i, t_j)$. (This rectangle is degenerated if t_i and t_j have the same x- or y-coordinate.) We say that R_{ij} is an *empty rectangle* if $R_{ij} \cap T = \{t_i, t_j\}$. The *complete grid* is obtained by drawing in the smallest axis-parallel rectangle containing the set T a horizontal segment and a vertical segment through every terminal. Using standard methods for

*Laboratoire d'Informatique Fondamentale de Marseille, Université de la Méditerranée, Faculté des Sciences de Luminy, F-13288 Marseille Cedex 9, France, {*chepoi,nouioua,vaxes*}*@lif.univ-mrs.fr*

establishing Hanan grid-type results [12], it can be shown that the complete grid contains at least one minimum Manhattan network [3].

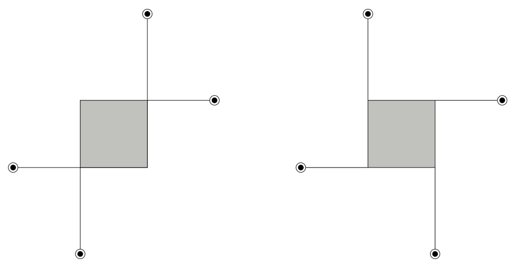

(a) (b)

Figure 2: Pareto envelope of four points

A point $p \in \mathbb{R}^2$ is said to be an *efficient* point of T [2, 11] if there does not exist any other point $q \in \mathbb{R}^2$ such that $d(q, t_i) \le d(p, t_i)$ for $t_i \in T$ and $d(q, t_j) < d(p, t_j)$ for at least one $t_j \in T$. Denote the set of all efficient points by $\mathcal{P}$, called the *Pareto envelope* of T. An optimal $O(n \log n)$ time algorithm to compute the Pareto envelope of n points in the l_1-plane is presented in [2] (for properties of $\mathcal{P}$ and an $O(n^2)$ time algorithm see also [11]). In particular, it is known that $\mathcal{P}$ is ortho-convex, i.e. the intersection of $\mathcal{P}$ with any vertical or horizontal line is convex, and that every two points of $\mathcal{P}$ can be joined in $\mathcal{P}$ by an l_1-path. $\mathcal{P}$, being ortho-convex, is a union of ortho-convex (possibly degenerated) rectilinear polygons (called *blocks*) glued together along vertices (they become cut points of $\mathcal{P}$); Fig. 2 presents two generic forms of the Pareto envelope of four points.

Lemma 2.1 *The Pareto envelope $\mathcal{P}$ contains at least one minimum Manhattan network on T.*

PROOF. Pick a minimum Manhattan network $N(T)$ having some vertices and edges outside $\mathcal{P}$. We can find two points $p, q \in N(T)$ on the boundary of $\mathcal{P}$ that are connected in $N(T)$ by an l_1-path L lying outside $\mathcal{P}$. Replace L by the l_1-path L' connecting p and q along the boundary of the Pareto envelope. Since all points of T are located inside or on the boundary of $\mathcal{P}$, the network $N'(T)$ obtained from $N(T)$ by replacing the path L by L' is Manhattan and its length is at most the length of $N(T)$. Applying the same procedure to $N'(T)$, after a finite number of steps we will obtain a minimum Manhattan network contained in $\mathcal{P}$. □

By this result, in order to solve the MMN problem on T it suffices to complete the set of terminals by adding to T the cut points of $\mathcal{P}$ and to solve a MMN problem on each block of $\mathcal{P}$ with respect to the new and old terminals located inside or on its boundary. Due to this decomposition of the MMN problem into smaller subproblems, further we can assume without loss of generality that $\mathcal{P}$ consists of a single block with at least 3 terminals; denote by $\partial\mathcal{P}$ the boundary of this ortho-convex rectilinear polygon. Then every convex vertex of $\mathcal{P}$ is a terminal. Since the sub-path of $\partial\mathcal{P}$ between two consecutive convex vertices of $\partial\mathcal{P}$ is the unique l_1-path connecting these vertices inside $\mathcal{P}$ and $\partial\mathcal{P}$ is covered by such l_1-paths, from Lemma 2.1 we conclude that the edges of $\partial\mathcal{P}$ belong to any minimum Manhattan network inside $\mathcal{P}$.

From Lemma 2.1 and the result of [3] mentioned above we conclude that the part $\Gamma = (V, E)$ of the complete grid contained in $\mathcal{P}$ hosts at least one minimum Manhattan network. Two edges of Γ are called *twins* if they are opposite edges of a rectangular face of the grid Γ. Two edges e, f of Γ are called *parallel* if there exists a sequence $e = e_1, e_2, \ldots, e_{m+1} = f$ of edges such that for $i = 1, \ldots, m$ the edges e_i, e_{i+1} are twins. By definition, any edge e is parallel to itself and all edges parallel to e have the same length. Notice also that exactly two edges parallel to a given edge e belong to $\partial\mathcal{P}$.

We continue with the notion of generating set introduced in [4] and used in approximation algorithms from [1, 6]. A *generating set* is a subset F of pairs of terminals (or, more compactly, of their indices) with the property that a rectilinear network containing l_1-paths for all pairs in F is a Manhattan network on T. For example, $F_\varnothing$ consisting of all pairs ij with R_{ij} empty is a generating set. In the next section, we will describe a sparse generating set contained in $F_\varnothing$.

To give an LP-formulation of the minimum Manhattan network problem, let F be an arbitrary generating set; for each pair $ij \in F$, let $\Gamma_{ij} := \Gamma \cap R(t_i, t_j)$ and set $\Gamma_{ij} = (V_{ij}, E_{ij})$. We formulate the MMN problem as a cut covering problem using an exponential number of constraints, which we further convert into an equivalent formulation that employs only a polynomial number of variables and constraints. In both formulations, l_e will denote the length of an edge e of the network $\Gamma = (V, E)$ and x_e will be a 0-1 decision variable associated with e. A subset of edges C of E_{ij} is called a (t_i, t_j)-*cut* if every l_1-path between t_i and t_j in Γ_{ij} meets C. Let $\mathcal{C}_{ij}$ denote the collection of all (t_i, t_j)-cuts and set $\mathcal{C} := \cup_{ij \in F} \mathcal{C}_{ij}$. Then the minimum Manhattan networks can be viewed as the optimal solutions of the following integer linear program (the dual of the relaxation of this program is a packing problem of the cuts from $\mathcal{C}$):

$$\begin{aligned} &\text{minimize} && \sum_{e \in E} l_e x_e && (1) \\ &\text{subject to} && \forall C \in \mathcal{C} : \sum_{e \in C} x_e \ge 1 \\ & && \forall e \in E : x_e \in \{0, 1\}. \end{aligned}$$

Indeed, every Manhattan network is a feasible solution of (1). Conversely, let $x_e, e \in E$, be a feasible solution for (1). Considering x_e's as capacities of the edges e of Γ, and applying the covering constraints and the Ford-Fulkerson's theorem to each network $\Gamma_{ij}, ij \in F$, oriented as described below, we conclude the existence in Γ_{ij} of an integer (t_i, t_j)-flow of value 1, i.e., of an l_1-path between t_i and t_j. As a consequence, we obtain a Manhattan network of the same cost. This observation leads to the second integer programming formulation for the MMN problem (but this time, having a polynomial size). For each pair $ij \in F$ and each edge

$e \in E_{ij}$ introduce a (flow) variable f_e^{ij}. Orient the edges of Γ_{ij} so that the oriented paths connecting t_i and t_j are exactly the l_1-paths between those terminals. For a vertex $v \in V_{ij} \setminus \{t_i, t_j\}$ denote by $\Gamma_{ij}^+(v)$ the oriented edges of Γ_{ij} entering v and by $\Gamma_{ij}^-(v)$ the oriented edges of Γ_{ij} out of v. We are lead to the following integer program:

$$\begin{aligned}
&\text{minimize} && \sum_{e\in E} l_e x_e && (2)\\
&\text{subject to} && \forall ij \in F, \forall v \in V_{ij} \setminus \{t_i, t_j\} : \\
& && \sum_{e\in\Gamma_{ij}^+(v)} f_e^{ij} = \sum_{e\in\Gamma_{ij}^-(v)} f_e^{ij} \\
& && \forall ij \in F : \sum_{e\in\Gamma_{ij}^-(t_i)} f_e^{ij} = 1 \\
& && \forall ij \in F, \forall e \in E_{ij} : 0 \le f_e^{ij} \le x_e \\
& && \forall e \in E : x_e \in \{0, 1\}.
\end{aligned}$$

Denote by $(1')$ and $(2')$ the LP-relaxation of (1) and (2) obtained by replacing the boolean constrains $x_e \in \{0, 1\}$ by the linear constraints $x_e \ge 0$. Since $(2')$ contains a polynomial number of variables and inequalities, it can be solved in strongly polynomial time using the algorithm of Tardos [8]. The x-part of any optimal solution of $(2')$ is an optimal solution of $(1')$. Notice also that there exist instances of the MMN problem for which the cost of an optimal (fractional) solution of $(1')$ or $(2')$ is smaller than the cost of an optimal (integer) solution of (1) or (2). Fig. 3 shows such an example ($x_e = 1$ for bolded edges and $x_e = \frac{1}{2}$ for dashed edges). Finally observe that in any feasible solution of $(1')$ and $(2')$ for any edge $e \in \partial\mathcal{P}$ holds $x_e = 1$.

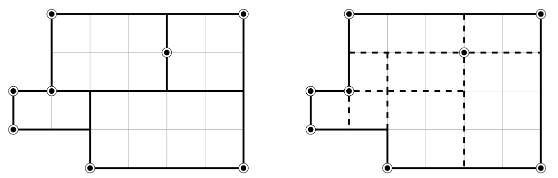

Integer optimum = 28 Fractional optimum = 27.5

Figure 3: Integrality gap

3 Strips and staircases

A degenerated empty rectangle R_{ij} is called a *degenerated vertical* or *horizontal strip*. A non-degenerated empty rectangle R_{ij} is called a *vertical strip* if the x-coordinates of t_i and t_j take consecutive values in the sorted list of x-coordinates of the terminals and the intersection of R_{ij} with degenerated vertical strips is either empty or one of the points t_i or t_j. Analogously, a non-degenerated empty rectangle R_{ij} is called a *horizontal strip* if the y-coordinates of t_i and t_j take consecutive values in the sorted list of y-coordinates of the terminals and the intersection of horizontal sides of R_{ij} with degenerated horizontal strips is either empty or one of the points t_i or t_j. The *sides* of a vertical (resp., horizontal) strip R_{ij} are the vertical (resp., horizontal) sides of R_{ij}. We say that the strips $R_{ii'}$ and $R_{jj'}$ (degenerated or not) form a *crossing configuration* if they intersect and the Pareto envelope of the points $t_i, t_{i'}, t_j, t_{j'}$ is of type (a); see Fig. 2. The importance of such configurations resides in the following property whose proof is straightforward:

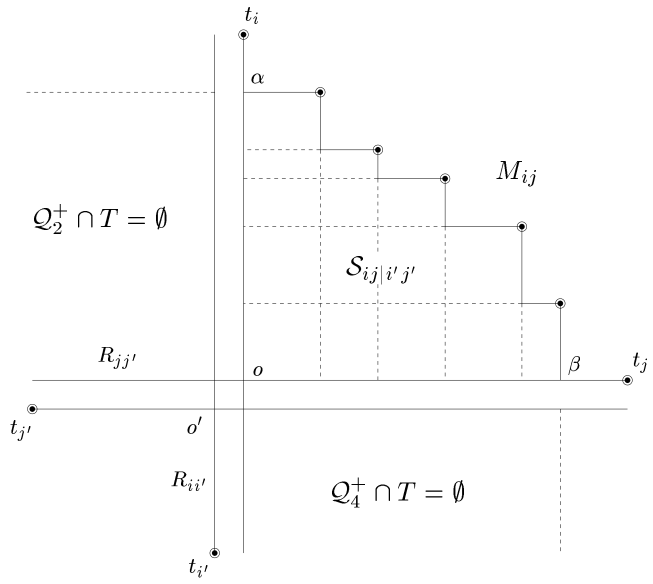

Figure 4: Staircase $\mathcal{S}_{ij|i'j'}$

Lemma 3.1 *If the strips $R_{ii'}$ and $R_{jj'}$ form a crossing configuration as in Fig. 4, then from the l_1-paths between t_i and $t_{i'}$ and between t_j and $t_{j'}$ one can derive an l_1-path connecting t_i and $t_{j'}$ and an l_1-path connecting $t_{i'}$ and t_j.*

For a crossing configuration $R_{ii'}, R_{jj'}$, denote by o and o' the cut points of the rectangular block of the Pareto envelope of $t_i, t_{i'}, t_j, t_{j'}$, and assume that the four tips of this envelope connect o with t_i, t_j and o' with $t_{i'}, t_{j'}$. Additionally, suppose without loss of generality, that t_i and t_j belong to the first quadrant $\mathcal{Q}_1$ with respect to the origin o (the remaining quadrants are labelled $\mathcal{Q}_2, \mathcal{Q}_3$, and $\mathcal{Q}_4$). Then $t_{i'}$ and $t_{j'}$ belong to the third quadrant with respect to the origin o'. Denote by T_{ij} the set of all terminals $t_k \in (T \setminus \{t_i, t_j\}) \cap \mathcal{Q}_1$ such that (i) $R(t_k, o) \cap T = \{t_k\}$ and (ii) the region $\{q \in \mathcal{Q}_2 : q^y \le t_k^y\} \cup \{q \in \mathcal{Q}_4 : q^x \le t_k^x\}$ does not contain any terminal of T. If T_{ij} is nonempty, then all its terminals belong to the rectangle R_{ij}, more precisely, they are all located on a common shortest rectilinear path between t_i and t_j. Denote by $\mathcal{S}_{ij|i'j'}$ the non-degenerated block of the Pareto envelope of the set $T_{ij} \cup \{o, t_i, t_j\}$ and call this rectilinear polygon a *staircase*; see Fig. 4 for an illustration. The point o is called the *origin* of this staircase. Analogously one can define the set $T_{i'j'}$ and the staircase $\mathcal{S}_{i'j'|ij}$ with origin o'. Two other types of staircases will be defined if t_i, t_j belong to the second quadrant and $t_{i'}, t_{j'}$ belong to the fourth quadrant. In order to simplify the presentation, further we will assume that after a suitable geometric transformation every staircase is located in the first quadrant. (Notice that our staircases are different from the staircase polygons occurring in the algorithms from [3].)

Let α be the leftmost highest point of the staircase $\mathcal{S}_{ij|i'j'}$ and let β be the rightmost lowest point of this staircase. Denote by M_{ij} the monotone boundary path of $\mathcal{S}_{ij|i'j'}$ between α and β and passing via the terminals of T_{ij}. By definition, $\mathcal{S}_{ij|i'j'} \cap T = T_{ij}$. By the choice of T_{ij}, there are no terminals of T located in the regions $\mathcal{Q}_2^+ := \{q \in \mathcal{Q}_2 : q^y \leq \alpha^y\}$ and $\mathcal{Q}_4^+ := \{q \in \mathcal{Q}_4 : q^x \leq \beta^x\}$. In particular, no strip traverses a staircase. From the definition of staircase immediately follows that two staircases either are disjoint or their intersection is a subset of terminals; in particular, every edge of the grid Γ belongs to at most one staircase.

Let F' be the set of all pairs ij such that R_{ij} is a strip. Let F'' be the set of all pairs $i'k$ such that there exists a staircase $\mathcal{S}_{ij|i'j'}$ such that t_k belongs to the set T_{ij}.

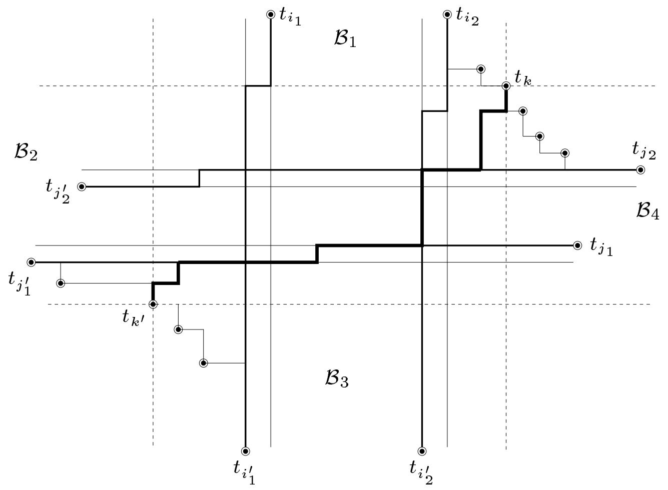

Figure 5: To the proof of Lemma 3.2

Lemma 3.2 *$F := F' \cup F'' \subseteq F_\varnothing$ is a generating set.*

Proof. Let N be a rectilinear network containing l_1-paths for all pairs in F. To prove that N is a Manhattan network on T, it suffices to establish that for an arbitrary pair $kk' \in F_\varnothing \setminus F$, the terminals t_k and $t_{k'}$ can be joined in N by an l_1-path. Assume without loss of generality that $t_{k'}^x \leq t_k^x$ and $t_{k'}^y \leq t_k^y$. The vertical and horizontal lines through the points t_k and $t_{k'}$ partition the plane into the rectangle $R_{kk'}$, four open quadrants and four closed unbounded half-bands labelled counterclockwise $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$, and $\mathcal{B}_4$. Since $t_k \in \mathcal{B}_1 \cap T$ and $t_{k'} \in \mathcal{B}_3 \cap T$, there should exist at least one vertical strip between a terminal from $\mathcal{B}_1$ and a terminal from $\mathcal{B}_3$. Denote by $R_{i_1 i'_1}$ the leftmost strip and by $R_{i_2 i'_2}$ the rightmost strip traversing the rectangle $R_{kk'}$. These two strips may coincide (and one or both of them may be degenerated), however they are both different from $R_{kk'}$ because $kk' \notin F$. Suppose without loss of generality that $t_{i_1}, t_{i_2} \in \mathcal{B}_1$ and $t_{i'_1}, t_{i'_2} \in \mathcal{B}_3$. Analogously define the lowest horizontal strip $R_{j_1 j'_1}$ and the highest horizontal strip $R_{j_2 j'_2}$ traversing $R_{kk'}$. Again these strips may coincide and/or may be degenerated but they must be different from $R_{kk'}$. Let $t_{j_1}, t_{j_2} \in \mathcal{B}_4$ and $t_{j'_1}, t_{j'_2} \in \mathcal{B}_2$; see Fig. 5. From the choice of the strips in question, we conclude that each of four combinations of horizontal and vertical strips constitute crossing configurations. Moreover, $\mathcal{S}_{i_2 j_2 | i'_2 j'_2}$ and $\mathcal{S}_{i'_1 j'_1 | i_1 j_1}$ must be staircases. The vertex t_k either belongs to $T_{i_2 j_2}$ or coincide with one of the vertices t_{i_2}, t_{j_2}. Analogously, $t_{k'} \in T_{i'_1 j'_1} \cup \{t_{i'_1}, t_{j'_1}\}$. By Lemma 3.1 there is an l_1-path connecting each of the terminals $t_{i_1}, t_{i_2}, t_{j_1}, t_{j_2}$ to each of the terminals $t_{i'_1}, t_{i'_2}, t_{j'_1}, t_{j'_2}$. Also there exist l_1-paths between t_k and $t_{i'_2}, t_{j'_2}$ and between $t_{k'}$ and t_{i_1}, t_{j_2}. Combining certain pieces of these l_1-paths we will produce an l_1-path connecting t_k and $t_{k'}$. $\square$

4 The rounding algorithm

Let $(\mathbf{x},\mathbf{f})=((x_e)_{e\in E}, (f_e^{ij})_{e\in E, ij\in F})$ be an optimal solution of the linear program (2′) (in general, this solution is not half-integral). The algorithm rounds up the solution $(\mathbf{x},\mathbf{f})$ in three phases. In **Phase 0**, we insert all edges of $\partial\mathcal{P}$ in the integer solution. In **Phase 1**, the rounding is performed inside every strip $R_{ii'}$, in order to ensure the existence of an l_1-path $P_{ii'}$ between the terminals t_i and $t_{i'}$. In **Phase 2**, an iterative rounding procedure is applied to each staircase.

Let $R_{ii'}$ be a strip. If $R_{ii'}$ is degenerated, then $[t_i, t_{i'}]$ is the unique l_1-path between t_i and $t_{i'}$, yielding $x_e = f_e^{ii'} = 1$ for any edge $e \in [t_i, t_{i'}]$. If $R_{ii'}$ is not degenerated, then any l_1-path in Γ between t_i and $t_{i'}$ has a simple form: it goes along the side of $R_{ii'}$ containing t_i, then it makes a turn by following an edge of Γ traversing $R_{ii'}$ (called further a *switch* of $R_{ii'}$), and continues its way on the side containing $t_{i'}$ until it reaches $t_{i'}$. Although, it may happen that several such l_1-paths have been used by the fractional flow $f^{ii'}$ between t_i and $t_{i'}$, the cut condition ensures that $x_e + x_{e'} \geq f_e^{ii'} + f_{e'}^{ii'} \geq 1$ for any pair e, e' of twins on opposite sides of the strip $R_{ii'}$, yielding $\max\{x_e, x_{e'}\} \geq \frac{1}{2}$.

Let p be the furthest from t_i vertex on the side of $R_{ii'}$ containing t_i such that $x_e \geq \frac{1}{2}$ for every edge e of the segment $[t_i, p]$. Let pp' be the edge of Γ incident to p that traverses the strip $R_{ii'}$. By the choice of p we have $x_e \geq \frac{1}{2}$ for all edges e of the segment $[p', t_{i'}]$.

Phase 1 (procedure `RoundStrip`). For each strip $R_{ii'}$, if $R_{ii'}$ is degenerated, then take in the integer solution all edges of $[t_i, t_{i'}]$, otherwise round up the edges of $[t_i, p]$ and $[p', t_{i'}]$ and take the edge pp' as a switch of $R_{ii'}$; in both cases, denote by $P_{ii'}$ the resulting l_1-path between t_i and $t_{i'}$.

Let $\mathcal{S}_{ii'|jj'}$ be a staircase. Denote by ϕ the closest to t_i common point of the l_1-paths $P_{ii'}$ and $P_{jj'}$ (this point is a corner of the rectangular face of Γ containing the vertices o and o'). Let $P_{ii'}^+$ and $P_{jj'}^+$ be the sub-paths of $P_{ii'}$ and $P_{jj'}$ comprised between ϕ and the terminals t_i and t_j, respectively. Now we slightly expand the staircase $\mathcal{S}_{ii'|jj'}$ by considering as $\mathcal{S}_{ii'|jj'}$ the region bounded by the paths $P_{ii'}^+, P_{jj'}^+$, and M_{ij} ($P_{ii'}^+$ and $P_{jj'}^+$ are not included in the staircase but M_{ij} and the terminals from the set T_{ij} are). Inside $\mathcal{S}_{ii'|jj'}$, any flow $f^{ki'}$ (or $f^{kj'}$), $k \in T_{ij}$, may be as fractional as possible: it may happen that several l_1-paths between t_k and $t_{i'}$ carry over flow $f^{ki'}$. Any such l_1-path intersects one of the paths $P_{ii'}^+$ or $P_{jj'}^+$, therefore the total $f^{ki'}$-flow arriving at $P^+_{ii'} \cup P^+_{jj'}$ is equal to 1. (This flow can be redirected to ϕ via the paths $P^+_{ii'}$ and $P^+_{jj'}$, and further, along the path $P_{ii'}$, to the terminal $t_{i'}$). Therefore it remains to decide how to round up the flow $f^{ki'}$ inside the expanded staircase $\mathcal{S}_{ii'|jj'}$. For this, notice that either the total $f^{ki'}$-flow carried over by the l_1-paths that arrive at $P^+_{ii'}$ is at least $\frac{1}{2}$ or the total $f^{ki'}$-flow on the l_1-paths that arrive at $P^+_{jj'}$ is at least $\frac{1}{2}$.

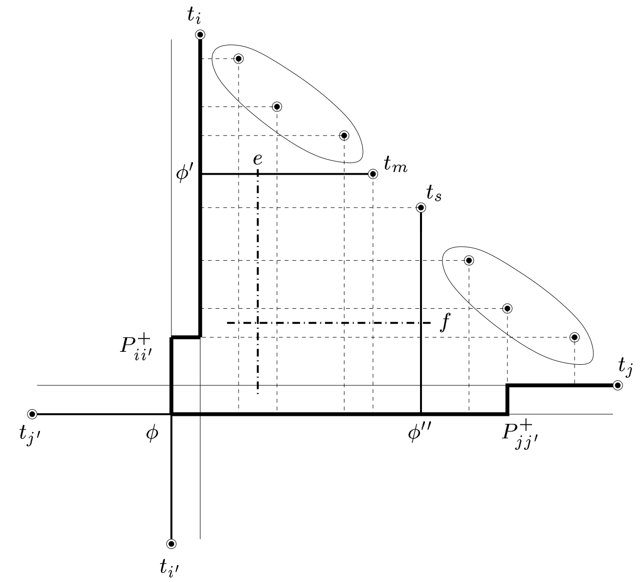

Figure 6: Procedure `RoundStaircase`

Phase 2 (procedure `RoundStaircase`). For a staircase $\mathcal{S}_{ii'|jj'}$ defined by the l_1-paths $P^+_{ii'}$ and $P^+_{jj'}$ and the monotone path M_{ij}, find the lowest terminal $t_m \in T_{ij}$ such that the $f^{mi'}$-flow on l_1-paths between t_m and $t_{i'}$ that arrive first at $P^+_{ii'}$ is $\geq \frac{1}{2}$ (we may suppose without loss of generality that this terminal exists). Let t_s be the terminal of T_{ij} immediately below t_m (this terminal may not exist). By the choice of t_m, the $f^{si'}$-flow on paths which arrive at $P_{jj'}$ is $\geq \frac{1}{2}$. Denote by ϕ' the intersection of the horizontal line passing via the terminal t_m with the path $P^+_{ii'}$. Analogously, let ϕ'' denote the intersection of the vertical line passing via t_s with the path $P^+_{jj'}$. Round up all edges of the horizontal segment $[t_m, \phi']$ and all edges of the vertical segment $[t_s, \phi'']$. If T_{ij} contains terminals located above the horizontal line (t_m, ϕ'), then recursively call `RoundStaircase` to the expanded staircase defined by $[t_m, \phi']$, the sub-path of $P^+_{ii'}$ comprised between ϕ' and t_i, and the sub-path M_{im} of the monotone path M_{ij} between t_m and α. Analogously, if T_{ij} contains terminals located to the right of the vertical line (t_s, ϕ''), then recursively call `RoundStaircase` to the expanded staircase defined by $[t_s, \phi'']$, the sub-path of $P^+_{jj'}$ comprised between ϕ'' and t_j, and the sub-path M_{sj} of the monotone path M_{ij} between t_s and β; see Fig. 6 for an illustration.

Let E_0 denote the edges of Γ which belong to the boundary of the Pareto envelope of T. Let E_1 be the set of all edges picked by the procedure `RoundStrip` and which do not belong to E_0, and let E_2 be the set of all edges picked by the recursive procedure `RoundStaircase` and which do not belong to $E_0 \cup E_1$. Denote by $N^* = (V^*, E_0 \cup E_1 \cup E_2)$ the resulting rectilinear network. From Lemma 3.2 and the rounding procedures presented above we infer that N^* is a Manhattan network. Let $\mathbf{x}^*$ be the integer solution of (1) associated with N^*, i.e., $x^*_e = 1$ if $e \in E_0 \cup E_1 \cup E_2$ and $x^*_e = 0$ otherwise.

5 Analysis

In this section, we will show that the length of the Manhattan network N^* is at most twice the cost of the optimal fractional solution of ($1'$), i.e., that

$$\mathrm{cost}(\mathbf{x}^*) = \sum_{e \in E} l_e x^*_e \leq 2 \sum_{e \in E} l_e x_e = \mathrm{cost}(\mathbf{x}). \qquad (3)$$

Recall that $x_e = x^*_e = 1$ holds for every edge $e \in E_0$. To establish the inequality (3), to every edge $e \in E_1 \cup E_2$ we will assign a set E_e of parallel to e edges such that (i) $\sum_{e' \in E_e} x_{e'} \geq \frac{1}{2}$ and (ii) $E_e \cap E_f = \emptyset$ for any two edges $e, f \in E_1 \cup E_2$.

First pick an edge $e \in E_1$, say $e \in P_{ii'}$ for a strip $R_{ii'}$. If e belongs to a side of this strip, then $x_e \geq \frac{1}{2}$, and in this case we can set $E_e := \{e\}$. Now, if e is the switch of $R_{ii'}$, then E_e consists of anyone of the two edges of $\partial\mathcal{P}$ parallel to e. From the definition of strips one conclude that no other switch can be parallel to these edges of $\partial\mathcal{P}$. Therefore each pair of parallel edges of ∂P may appear in at most one set E_e for a switch e.

Finally suppose that $e \in E_2$, say e belongs to the expanded staircase $\mathcal{S}_{ii'|jj'}$. If e belongs to the segment $[t_m, \phi']$, then E_e consists of e and all parallel to e edges of $\mathcal{S}_{ii'|jj'}$ located below e; see Fig. 6. Since every l_1-path between t_m and $t_{i'}$ intersecting the path $P^+_{ii'}$ contains an edge of E_e, we infer that the value of the $f^{mi'}$-flow traversing the set E_e is at least $\frac{1}{2}$, therefore $\sum_{e' \in E_e} x_{e'} \geq \frac{1}{2}$, thus establishing (i). Analogously, if f is an edge of the vertical segment $[t_s, \phi'']$, then E_f consists of f and all parallel to f edges of $\mathcal{S}_{ii'|jj'}$ located to the left of f. Obviously, $E_e \cap E_f = \emptyset$. Since E_e and E_f belong to the region of $\mathcal{S}_{ii'|jj'}$ delimited by the segments $[t_m, \phi']$ and $[t_s, \phi'']$ and the recursive calls of the procedure `RoundStaircase` concern the staircases disjoint from this region, we deduce that E_e and E_f are disjoint from the sets $E_{e'}$ for all edges e' picked by the recursive calls of `RoundStaircase` to the staircase $\mathcal{S}_{ii'|jj'}$. Every edge of Γ belongs to at most one staircase, therefore $E_e \cap E_f = \emptyset$ if the edges $e, f \in E$ belong to different staircases. Finally, since there are no terminals of T located below or to the left of the staircase $\mathcal{S}_{ii'|jj'}$, no strip traverses this staircase (a strip intersecting $\mathcal{S}_{ii'|jj'}$ either coincides with $R_{ii'}$ and $R_{jj'}$, or intersects the staircase along segments of the boundary path M_{ij}). Therefore, no edge from E_1 can be assigned to a set E_e for some $e \in E_2 \cap \mathcal{S}_{ii'|jj'}$, thus establishing (ii) and the desired inequality (3). Now, we are in position to formulate the main result of this note:

Theorem 5.1 *The rounding algorithm described in Section 4 achieves an approximation guarantee of 2 for the minimum Manhattan network problem.*

Remark. Given a staircase $\mathcal{S}_{ii'|jj'}$ and the paths $P^+_{ii'}$ and $P^+_{jj'}$, the problem of constructing a minimum rectilinear network such that every terminal of T_{ij} can be connected by an l_1-path to $P^+_{ii'} \cup P^+_{jj'}$ can be solved in polynomial time using dynamic programming (for example, by adapting the algorithm from [7] for the Rectilinear Steiner Arborescence problem on staircases). However, we do not know how to analyze this solution via linear programming. Furthermore, we do not have examples of staircases having an integrality gap in $(1')$.

6 Conclusions and perspectives

In this paper, we presented a simple rounding algorithm for the minimum Manhattan network problem and we established that the length of the Manhattan network returned by this algorithm is at most twice the cost of the optimal fractional solution of the MMN problem. Nevertheless, experiences show that the ratio between the costs of the solution returned by our algorithm and the optimal solution of the linear programs $(1')$ and $(2')$ is much better than 2. We do not know the worst integrality gap of (1) (the worst gap obtained by computer experiences is about 1.087). Say, is this gap smaller or equal than 1.5? Does there exist a gap in the case when the terminals are the origin and the corners of a staircase?

The minimum Manhattan network problem can be compared with the (NP-complete) Rectilinear Steiner Arborescence problem (*RSA problem*) [7]. In this problem, given n terminals (lying in the first quadrant), one search for a minimum rectilinear network comprising an l_1-path between the origin of coordinates and each of the n terminals (clearly, such an optimal network will be a tree). The LP-formulation for the MMN problem can be viewed as a generalization of the LP-formulation of the RSA problem given in [9]. The paper [7] presents an instance of the RSA problem having an integrality gap. To our knowledge, the worst integrality gap for this problem is also not known.

Consider now the following common generalization of the MMN and RSA problems which we call the *F-restricted MMN problem*: given a set of n terminals and a collection F of pairs of terminals, find a minimum rectilinear network $N_F(T)$, such that for every pair $t_i t_j \in F$, the network $N_F(T)$ contains an l_1-path between t_i and t_j. If (T, F) is a complete graph, then we obtain the MMN problem and if (T, F) is a star, then we obtain the RSA problem. We can show that there exists a minimum F-restricted Manhattan network contained in the sub-grid of Γ generated by all empty rectangles. Using this grid, one can view (1) and (2) as integer programming formulations for the F-restricted MMN problem.

Notice that the rounding algorithm presented in our note (as well as all other approximation algorithms for the MMN or RSA problems) cannot be extended in a direct way to get an approximation algorithm for the F-restricted MMN problem. Developing such an algorithm seems to be an interesting question. A simple example shows that the integrality gap in this case is at least 1.5: consider the four corners of a unit square as the set T of terminals, and let F consists of the two pairs of opposite corners of this square. Then $x_e = \frac{1}{2}$ for every side e of the square is an optimal solution of $(1')$ having cost 2, while an optimal integer solution uses three edges of the square and has cost 3.

References

[1] M. Benkert, T. Shirabe, and A. Wolff, The minimum Manhattan network problem - approximations and exact solutions, *20th EWCG,* March 25-26, Seville (Spain), 2004.

[2] G. Chalmet, L. Francis, and A. Kolen, Finding efficient solutions for rectilinear distance location problems efficiently, *European J. Operations Research,* **6** (1981), 117–124.

[3] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan, Approximating a minimum Manhattan network, *Nordic J. Computing,* **8** (2001), 219–232 and *Proc. APPROX'99,* 1999, pp. 28–37.

[4] R. Kato, K. Imai, and T. Asano, An improved algorithm for the minimum Manhattan network problem, *ISAAC'02,* Lecture Notes Computer Science, **2518**, 2002, pp. 344-356.

[5] F. Lam, M. Alexanderson, and L. Pachter, Picking alignements from (Steiner) trees, *J. Computational Biology,* **10** (2003), 509–520.

[6] K. Nouioua, Une approche primale-duale pour le problème du réseau de Manhattan minimal, *RAIRO Operations Research* (submitted).

[7] S.K. Rao, P. Sadayappan, F.K. Hwang, and P.W. Shor, The rectilinear Steiner arborescence problem, *Algorithmica,* **7** (1992), 277–288.

[8] E. Tardos, A strongly polynomial algorithm to solve combinatorial linear programs, *Operations Research,* **34** (1986), 250–256.

[9] V. Trubin, Subclass of the Steiner problem on a plane with rectilinear metric, *Cybernetics,* **21** (1985), 320–322.

[10] V.V. Vazirani, *Approximation Algorithms*, Springer-Verlag, Berlin, 2001.

[11] R.E. Wendell, A.P. Hurter, and T.J. Lowe, Efficient points in location theory, *AIEE Transactions* **9** (1973), 314-321.

[12] M. Zachariasen, A catalog of Hanan grid problems, *Networks* **38** (2001), 76–83.