

Line Crossing Minimization on Metro Maps

Michael A. Bekos

School of Applied Mathematics & Physical Sciences
National Technical University of Athens, Greece
mikebekos@math.ntua.gr

Michael Kaufmann

Institute for Informatics, University of Tübingen, Germany
mk@informatik.uni-tuebingen.de

Katerina Potika

School of Electrical & Computer Engineering
National Technical University of Athens, Greece
epotik@cs.ntua.gr

Antonios Symvonis

School of Applied Mathematics & Physical Sciences
National Technical University of Athens, Greece
symvonis@math.ntua.gr

Abstract

We consider the problem of drawing a set of simple paths along the edges of an embedded underlying graph $G = (V, E)$, so that the total number of crossings among pairs of paths is minimized. This problem arises when drawing metro maps, where the embedding of G depicts the structure of the underlying network, the nodes of G correspond to train stations, an edge connecting two nodes implies that there exists a railway track connecting them, whereas the paths illustrate the lines connecting terminal stations. We call this the *metro-line crossing minimization problem (MLCM)*. As a first step towards solving MLCM in arbitrary graphs, we study path and tree networks. We examine several variations of the problem for which we develop algorithms for obtaining optimal solutions.

Article Type	Communicated by	Submitted	Revised
...	...	...	...

1 Motivation

We consider a relatively new problem that arises when drawing metro maps or, in general, public transportation networks. In such drawings, we are given an undirected embedded graph $G = (V, E)$, which depicts the structure of the underlying network. In the case of metro maps, the nodes of G correspond to train stations whereas an edge connecting two nodes implies that there exists a railway track which connects them. The problem we consider is motivated by the fact that an edge within the underlying network may be used by several metro lines. Since crossings are often considered as the main source of confusion in a visualization, we want to draw the lines along the edges of G so that they cross each other as few times as possible.

In the graph drawing literature, the focus has been so far exclusively on drawing the underlying graph nicely and not on how to embed the bus or the metro lines along the underlying network. The latter problem was recently introduced by Benkert et. al in [4]. Following their approach, we assume that the underlying network has already received an embedding. The problem of determining a solution of the general metro-line routing problem, in which the graph drawing and line routing are solved simultaneously would be of particular interest as a second step in the process of automated metro map drawing.

2 Problem Definition

We are given an undirected embedded graph $G = (V, E)$. We will refer to G as the *underlying network*. We are also given a set $\mathcal{L} = \{l_1, l_2, \dots, l_k\}$ of simple paths of G (in the following, referred to as *lines*). Each line l_i consists of a sequence of edges $e_1 = (v_0, v_1)$, $e_2 = (v_1, v_2), \dots, e_d = (v_{d-1}, v_d)$ of G . The nodes v_0 and v_d are referred to as the *terminals* of line l_i . Equivalently, we say that the line *terminates at* or *has terminals at* nodes v_0 and v_d . By $|l_i|$ we denote the length of line l_i . The main task is to draw the lines along the edges of G , so that the number of crossings among pairs of lines is minimized. We call this the *metro-line crossing minimization problem (MLCM)*. Formally, the MLCM problem is defined as a tuple $(G, \mathcal{L})$, where G is the underlying network and $\mathcal{L}$ is the set of lines.

One can define several variations of the MLCM problem based on the type of the underlying network, the location of the crossings and/or the location of the terminals (refer to Figure 1). In general, the underlying network is an undirected graph. In this paper, as a first step towards the study of the MLCM problem in arbitrary graphs, we focus on path and tree networks.

For aesthetic reasons, we insist that the crossings between lines that traverse a node of the underlying network should not be hidden under the area occupied by that node. This implies that the relative order of the lines should not change within the nodes and therefore, all possible crossings have to take place along the edges of the underlying network.

In our approach, we assume that the nodes which correspond to stations are




Figure 1: A map illustrating Munich’s S-Bahn and U-Bahn networks obtained from <http://www.mvv-muenchen.de>.

drawn as rectangles, which is a quite usual convention in metro maps. Each line that traverses a node u has to touch two of the sides of u at some points (one when it “enters” u and one when it “exits” u). These points are referred to as *tracks*. In general, we may permit tracks to all four sides of a node, i.e. a line that traverses a node may use any side of it to either “enter” or “exit”. This model is referred to as *4-side model* (see Figure 2a). A more restricted model, referred to as *2-side model*, is the one where all lines that traverse a node may use only its left and right side (see Figure 2b). In the latter case, we only allow tracks at the left and right side of the node.




Figure 2: Illustration of different models defined based on the location of the tracks.

Note that a solution for the MLCM problem should first specify the number of tracks that enter each side of each station and, for each track, the line of $\mathcal{L}$ that uses it.




Figure 3: Station ends and middle tracks.

A further refinement of the MLCM problem concerns the location of the line terminals at the nodes. A particularly interesting case - that arises under the 2-side model - is the one where the lines that terminate at a station occupy its topmost and bottommost tracks, in the following referred to as *top* and *bottom station ends*, respectively. The remaining tracks on the left and right side of the station are referred to as *middle tracks* and are occupied by the lines that pass through the station. Figure 3 illustrates the notions of top and bottom station ends and middle tracks on the left and right side of a station (solid lines correspond to lines that terminate, whereas the dashed lines correspond to lines that go through the station). Based on the above, we introduce the following two variants of the MLCM problem:

- (a) *The MLCM problem with terminals at station ends* (MLCM-SE), where we ask for a drawing of the lines along the edges of G so that (i) all lines terminate at station ends and (ii) the number of crossings among pairs of lines is minimized.
- (b) *The MLCM problem with terminals at fixed station ends* (MLCM-FixedSE), where all lines terminate at station ends and the information whether a line terminates at a top or at a bottom station end in its terminal stations is specified as part of the input. We ask for a drawing of the lines along the edges of G so that the number of crossings among pairs of lines is minimized.

2.1 Related Literature

The problem of drawing a graph with a minimum number of crossings has been extensively studied in the graph drawing literature. For a quick survey refer to [2] and [10]. However, in the problems we study in this paper we assume that the underlying graph has already received an embedding and we seek to draw the lines along the graph's edges, so that the number of crossings among pairs of lines is minimized.

This problem was recently introduced by Benkert et. al in [4]. In their work, they proposed a dynamic-programming based algorithm that runs in $O(n^2)$ time

for the *one-edge layout problem*, which is defined as follows: Given a graph $G = (V, E)$ and an edge $e = (u, v) \in E$, let L_e be the set of lines that traverse e . L_e is divided into three subsets L_u , L_v and L_{uv} . Set L_u (L_v) consists of the lines that traverse u (v) and terminate at v (u). Set L_{uv} consists of the lines that traverse both u and v and do not terminate either at u or at v . The lines for which u is an intermediate station, i.e., $L_{uv} \cup L_u$, enter u in a predefined order S_u . Analogously, the lines for which v is an intermediate station, i.e., $L_{uv} \cup L_v$, enter v in a predefined order S_v . The number of pairs of crossing lines is then determined by inserting the lines of L_u into the order S_v and by inserting the lines of L_v into the order S_u . The task is to determine appropriate insertion orders so that the number of pairs of crossing lines is minimized. However, in their work Benkert et. al [4] do not address the case of larger graphs and they leave as an open problem the case where the lines that terminate at a station occupy its station ends.

For the latter problem, Asquith et al. [1] proposed an integer linear program, which always determines an optimal solution regardless of the type of the underlying network. They mention that their approach can be generalized to support the case where the set of lines consists of subgraphs of the underlying network of maximum degree 3.

A closely related problem to the one we consider is the problem of drawing a metro map nicely, widely known as *metro map layout problem*. Hong et al. [8] implemented five methods for drawing metro maps using modifications of spring-based graph drawing algorithms. Stott and Rodgers [13] have approached the problem by using a hill climbing multi-criteria optimization technique. The quality of a layout is a weighted sum over five metrics that were defined for evaluating the niceness of the resulting drawing. Nöllenburg and Wolff [12] specified the niceness of a metro map by listing a number of hard and soft constraints and they proposed a mixed-integer program which always determines a drawing that fulfills all hard constraints (if such exists) and optimizes a weighted sum of costs corresponding to the soft constraints.

Relevant to the problem we study is also the problem of computing a confluent drawing of a given non-planar graph (see [6], [9]). Confluent drawings capture the connection properties of train tracks as follows: Edge crossings are allowed, however they are hidden in the drawing. This is accomplished by allowing group of edges to merge together and drawn as “tracks”, so that edge crossings are turned into overlapping paths, like merging train tracks into a single track.

2.2 Outline

In Section 3, we consider the MLCM problem on a path. We show that the MLCM-SE problem is NP-Hard and we also present a polynomial time algorithm for the MLCM-FixedSE problem. In Section 4, we consider the MLCM problem on a tree and we present polynomial time algorithms for two variations of it. We conclude in Section 5 with open problems and future work. A preliminary version of this paper appeared in Graph Drawing 2007 [3].

3 The Metro-Line Crossing Minimization Problem on a Path

We first consider the case where the underlying network G is a path and its nodes are restricted to lie on a horizontal line. We adopt the 2-side model where each line uses the left side of a node to “enter” it and the right one to “exit” it. Then, assuming that there exist no restrictions on the location of the line terminals at the nodes, it is easy to see that there exist solutions without any crossing among lines. In fact, using a simple reduction from the interval graph coloring problem [7], we can determine a solution, which also minimizes the number of tracks used at each individual node. So, in the remaining of this section, we further assume that the lines that terminate at a station occupy its top and bottom station ends. In particular, we consider the MLCM-SE problem on a path. Since the order of the stations is fixed as part of the input of the problem, the only remaining choice is whether each line terminates at the top or at the bottom station end in its terminal stations. In the following we show that under this assumption, the problem of determining a solution so that the total number of crossings among pairs of lines is minimized is NP-Hard. Our proof is based on a reduction from the *fixed linear crossing number problem* [11].

Definition 1 *Given a simple graph $G = (V, E)$, a linear embedding of G is a special type of embedding in which the nodes of V are placed on a horizontal line L and the edges are drawn as semicircles either above or below L .*

Definition 2 *A node ordering (or a node permutation) of a graph G is a bijection $\delta : V \rightarrow \{1, 2, \dots, n\}$, where $n = |V|$. For each pair of nodes u and v , with $\delta(u) < \delta(v)$ we shortly write $u < v$.*

Masuda et al. [11] proved that it is NP-Hard to determine a linear embedding of a given graph with minimum number of crossings, even if the ordering of the nodes on L is fixed. The latter problem is referred to as *fixed linear crossing number problem*.

Theorem 1 *The MLCM-SE problem on a path is NP-Hard.*

Proof: Let I be an instance of the fixed linear crossing number problem, consisting of a graph $G = (V, E)$ and a horizontal input line L , where $V = \{u_1, u_2, \dots, u_n\}$ and $E = \{e_1, e_2, \dots, e_m\}$. Without loss of generality, we assume that $u_1 < u_2 < \dots < u_n$. We construct an instance I' of the MLCM-SE problem on a path as follows: The underlying network $G' = (V', E')$ is a path consisting of $n + 2$ nodes and $n + 1$ edges, where $V' = V \cup \{u_0, u_{n+1}\}$ and $E' = \{(u_{i-1}, u_i); 1 \leq i \leq n + 1\}$. The set of lines $\mathcal{L}$ is partitioned into two sets $\mathcal{L}^A$ and $\mathcal{L}^B$:

- $\mathcal{L}^A$ consists of a sufficiently large number of lines (e.g. $2nm^2$ lines) connecting u_0 with u_{n+1} .

- $\mathcal{L}^B$ contains m lines $l_1, l_2, \dots, l_m$ one for each edge of G . Line l_i which corresponds to edge e_i of G , has terminals at the end points of e_i .

Figure 4 illustrates the construction. First observe that all lines of $\mathcal{L}^A$ can be routed “in parallel” without any crossing among them (see Figure 4b). Also observe that in an optimal solution none of the lines $l_1, l_2, \dots, l_m$ crosses the lines of $\mathcal{L}^A$, since that would contribute a very large number of crossings. Thus, in an optimal solution each line of $\mathcal{L}^B$ has both of its terminals either at top or at bottom station ends. So, in a sense, we exclude the case where a line $l_i \in \mathcal{L}^B$ has one of its terminals at a top station end, whereas the second one at a bottom station end.




Figure 4: (a) A linear embedding, (b) An instance of the MLCM-SE problem.

Assume now that there exists an optimal linear embedding of I , where the number of crossings is k . We will construct a solution for I' with k crossings among pairs of lines. We first route the lines of $\mathcal{L}^A$ without introducing any crossing among them. The remaining lines $l_1, l_2, \dots, l_m$ will be routed either above or below the lines of $\mathcal{L}^A$ (refer to Figure 4b). We choose to route line l_i above (below) the lines of the set $\mathcal{L}^A$, if in the embedding of I edge e_i is placed above (below) the input line L . Additionally, if in the embedding of I , two edges e_i and e_j are placed above (below) L so that they do not cross each other and the semicircle of e_i contains the one of e_j (refer to the edges (u_1, u_6) and (u_2, u_5) in Figure 4), then we route line l_i below (above) line l_j . This implies that these lines will not cross each other. If two edges cross each other in the embedding of I (refer to edges (u_1, u_4) and (u_3, u_6) in Figure 4), then their corresponding lines will also cross in I' . This implies an one-to-one correspondence between the crossings among the edges of I and the crossings among pairs of lines in I' , as desired.

Consider now the case where we have determined an optimal solution of I' with k pairs of crossing lines. As already mentioned, lines $l_1, l_2, \dots, l_m$ do not cross the lines of the set $\mathcal{L}^A$, since a solution including such a crossing is not optimal. Therefore, each line of $\mathcal{L}^B$ either lies entirely above the lines of $\mathcal{L}^A$ or below them. In the case where a line $l_i \in \mathcal{L}^B$ lies above them, we draw edge e_i of graph G above the horizontal line L , otherwise below that. Again, it is easy to see that there exists an one-to-one correspondence between the crossings among the edges of I and the crossings among pairs of lines in I' . Therefore, instance I has a linear embedding with k crossings among its edges. $\square$

3.1 The Metro-Line Crossing Minimization Problem with fixed positioned terminals

Theorem 1 implies that, unless $P = NP$, we can not efficiently determine an optimal solution of MLCM-SE problem on a path. The main reason for this is that the information whether each line terminates at the top or at the bottom station end in its terminal stations is not known in advance. In the following, we assume that this information is part of the input, which is a reasonable assumption, since terminals may represent physical locations within a station. In particular, we show that the MLCM-FixedSE problem on a path can be solved in polynomial time.

To simplify the description of our algorithm, we assume that each node u_i of the path G is adjacent to two nodes u_i^t and u_i^b , each of which will be the terminal of the lines that terminated at the top and bottom station ends of node u_i , respectively¹. In the drawing of G , u_i^t is placed directly on top of u_i (we refer to it as the *top leg* of u_i), whereas u_i^b directly below it (*bottom leg* of u_i). So, instead of restricting each line to have the terminals at a top or at a bottom station end in its terminal stations, we will equivalently consider that the line terminals are located at two leg nodes. We refer to this special type of graph which is implied by the addition of the leg nodes as *caterpillar with at most two legs per node* (see an example in Figure 5).




Figure 5: A caterpillar with at most 2 legs per node.

A caterpillar with at most two legs per node consists of two sets of nodes. The first set, denoted by V_b , contains n nodes $u_1, u_2, \dots, u_n$ (referred to as *backbone nodes*), which form a path. In the embedding of G , these nodes are collinear and more precisely they are located on a horizontal line so that $u_1 < u_2 < \dots < u_n$. The second set of nodes, denoted by V_l , contains n' nodes $v_1, v_2, \dots, v_{n'}$ of degree 1 (referred to as *leg nodes* or simply as *legs*) each of which is connected to one backbone node. In the embedding of G , we assume that for each backbone u one of its legs is placed directly on top of it, whereas the second one directly below it. Since each backbone node is adjacent to at most two legs, $n' \leq 2n$.

If v is a leg node, we will refer to its neighbor backbone node as $bn(v)$. Edges that connect backbone nodes are called *backbone edges*. Edges that connect backbone nodes with legs are called *leg edges*.

¹In the degenerated case, where there exists no lines terminating either at the top or bottom terminal tracks of node u_i , we assume that either u_i^t or u_i^b does not exist, respectively.

Definition 3 Let $l \in \mathcal{L}$ be a line that connects two terminals v and v' . If v is located to the left of v' in the embedding of the underlying network, i.e. $v < v'$, then we consider v to be the origin of line l , whereas v' to be its destination. We also denote by $\mathcal{L}_i^t$ ($\mathcal{L}_i^b$) the lines that have as origin the top (bottom) leg node adjacent to backbone node u_i .

Definition 4 Let l and l' be a pair of lines that have the same origin w and destination nodes v and v' , respectively. We say that l precedes l' , if when we start moving from w along the external face of G in counterclockwise direction we meet v before v' . The notion of precedence defines an order $\preceq$ among the lines that have the same origin, namely $l \preceq l'$, if and only if l precedes l' .

Lemma 1 The number of tracks in the left and right side of each backbone node that are needed in order to route all lines in $\mathcal{L}$ can be computed in $O(n + |\mathcal{L}|)$ time.

Proof: The number of tracks in the right side of the leftmost backbone node u_1 is $|\mathcal{L}_1^t| + |\mathcal{L}_1^b|$. Due to the fact that no lines have as terminal a backbone node, the same number of tracks are needed in the left side of node u_2 . We index the needed tracks from top to bottom (refer to Figure 6a). We compute the number of tracks in the left side of any backbone node u_i as the number of lines originating at nodes $< u_i$ and destined for nodes $\geq u_i$. Similarly, we compute the number of tracks in the right side of any backbone node u_i as the number of lines originating at nodes $\leq u_i$ and destined for nodes $> u_i$.

By performing a left to right pass over the set of backbone nodes, we can compute the number of tracks in the left and right side of each backbone node in $O(n + |\mathcal{L}|)$ total time, assuming that each backbone node u keeps appropriate references to the number of lines originating at and destined for u . $\square$

The basic steps of our algorithm that solves the MLCM-FixedSE problem are outlined in Algorithm-1. The lines of $\mathcal{L}$ are drawn incrementally by performing a left to right pass over the set of backbone nodes and by extending them from station to station with small horizontal or diagonal line segments. Therefore, each line $l \in \mathcal{L}$ is drawn as a polygonal line.

We now proceed in detail to explain the routing in Step C of Algorithm-1. Let $\mathcal{L}_v$ be the set of the lines that either originate at or are destined for leg node v . In each leg edge, that connects leg node v to $bn(v)$, we use $|\mathcal{L}_v|$ tracks indexed from right to left (refer to Figure 6a). These tracks will be used in order to route the lines that either originate at or are destined for leg node v .

In each backbone node u_i , we have to route the newly “introduced” lines, i.e. the ones that originate either at the top or at bottom leg of u_i . This procedure is illustrated in Figure 6a. We first consider the top leg node u_i^t of u_i . We sort the set $\mathcal{L}_i^t$ of the lines that originate at u_i^t in increasing order $\preceq$ of their destinations and store them in ascending order, in $Sort(\mathcal{L}_i^t)$. Based on this sorting we route the j -th line l in $Sort(\mathcal{L}_i^t)$ through the j -th rightmost track at the top of u_i . l is then routed to the j -th top track in the right side of u_i . We proceed by considering the bottom leg node u_i^b of u_i . Again, we sort the set $\mathcal{L}_i^b$ of the lines

Algorithm 1: MLCM-FIXEDSE ON PATHS

Input : A path $G = (V, E)$ and a set of lines $\mathcal{L} = \{l_1, \dots, l_k\}$.

Output: A routing of the lines along the edges of G , so that they cross each other as few times as possible.

Step A: Based on G build caterpillar $G_c = (V_b \cup V_l, E_c)$, where:
 $V_b = \{u_1, \dots, u_n\}$ and $V_l = \{v_1, \dots, v_{n'}\}$, as described above.

Step B: For each backbone node, determine the number of tracks that are used by the lines along its left and right sides;

Step C: Route the lines as follows:

foreach backbone node u_i , where $i = 1$ **to** $n - 1$ **do**

1. “Introduce” the new lines that originate at the top and bottom legs of u_i to node u_i ;
2. Route the lines from the left side of u_i to the right side of u_i ;
3. Route the lines from the right side of u_i to the left side of u_{i+1} ;

Step D: Remove all leg nodes.

that originate at u_i^b in decreasing order $\preceq$ of their destinations and store them in descending order, in $\text{Sort}(\mathcal{L}_i^b)$. Based on the sorting, we route the j -th line l in $\text{Sort}(\mathcal{L}_i^b)$ through the j -th rightmost track at the the bottom of u_i and then to the j -th bottom track in the right side of u_i . We then route the lines that go from the tracks of the left side to the tracks of the right side of u_i , by preserving their relative order.

The next step is to route the lines from the right side of u_i to the left side of u_{i+1} . This is done by performing three passes over the set of tracks of the right side of u_i .

In the first pass, we consider the tracks of the right side of u_i from top to bottom and we check whether the line l that occupies the j -th track is destined for the leg node u_{i+1}^t . In this case, we route l to the topmost available track of the right side of u_{i+1} and then to the leftmost available track in the leg edge which connects u_{i+1} with u_{i+1}^t (see the dotted lines of Figure 6b).

In the second pass, we consider the remaining tracks of the right side of u_i from bottom to top and we check whether the line l that occupies the j -th track is destined for the leg node u_{i+1}^b . In this case, we route l to the bottommost available track of the right side of u_{i+1} and then to the leftmost available track in the leg edge which connects u_{i+1} with u_{i+1}^b (see the dash dotted lines of Figure 6b).

The remaining tracks of the right side of u_i are obviously occupied by the lines that are not destined for either u_{i+1}^t or for u_{i+1}^b . We consider these tracks from top to bottom and we route the line l that occupies the j -th track to the




Figure 6: (a) Introduction of the new lines to a station, (b) Routing lines along a backbone edge.

topmost available track of the left side of u_{i+1} (see the dashed lines of Figure 6b). The construction of Algorithm-1 guarantees the following two properties:

Property of common destinations: Lines that are destined for the same top (bottom) leg node u_i^t (u_i^b) do not cross each other along the backbone edge which connects u_{i-1} with u_i .

Property of parallel routing: Two lines that both traverse a backbone node u_i (i.e. none of them are destined either for u_i^t or for u_i^b) do not cross each other along the backbone edge which connects u_{i-1} with u_i .

By combining the property of *common destinations* and the property of *parallel routing*, we easily obtain the following lemma.

Lemma 2 *In a solution produced by Algorithm-1, the following statements hold:*

- (i) *Two lines l and l' cross each other at most once.*
- (ii) *Two lines l and l' with the same origin do not cross each other.*
- (iii) *Two lines l and l' with the same destination do not cross each other.*
- (iv) *Let l and l' be two lines that cross each other, destined for leg nodes v and v' , respectively, where v is to the left of v' in the embedding of G . Then, l and l' will cross along the backbone edge which connects u_{k-1} and u_k , where $u_k = \text{bn}(v)$.*

□

Consider an instance $I = (G, \mathcal{L})$ of the problem we study consisting of a caterpillar graph G with at most 2 legs per node and a set of lines $\mathcal{L}$ terminating at leg nodes only. To prove the optimality of the solution produced by Algorithm-1, we transform I to a circle graph $(R_G, \mathcal{C})$ [7], consisting of a set of nodes R_G and a set of chords $\mathcal{C}$, so that:

- Each leg node in G corresponds to a node of R_G .
- The order of the leg nodes in G when moving clockwise in the external face of G and the order of the nodes of R_G when moving in clockwise direction along its boundary is identical.
- Each line in $\mathcal{L}$ which connects two leg nodes corresponds to a chord in $\mathcal{C}$ connecting the corresponding nodes.

Let $R_G = \{w_1, w_2, \dots, w_{n'}\}$, where n' is the number of leg nodes of G . W.l.o.g. we further assume that the nodes of R_G are indexed so that w_{i+1} immediately follows w_i in the clockwise direction for $i = 1, 2, \dots, n' - 1$. Figure 7 illustrates such a transformation. We denote by f the function which given a leg node in G returns the corresponding node in R_G .

The transformation described above implies a lower bound on the number of crossings among the lines in an optimal solution. More precisely, the number of chords crossing in circle R_G is a lower bound for the number of crossings in an optimal solution of I . To see this, consider two crossing chords $ch_i = (w_{s_i}, w_{d_i})$ and $ch_j = (w_{s_j}, w_{d_j})$ in the circle graph, where $s_i < d_i$ and $s_j < d_j$. One can see that ch_i and ch_j cross, if and only if

$$s_i < s_j < d_i < d_j \quad \text{or} \quad s_j < s_i < d_j < d_i.$$




Figure 7: An instance $I = (G, \mathcal{L})$ is transformed to a circle graph.

Translating these inequalities back to our problem, we observe that there exists crossings which can not be avoided. More precisely, consider the lines l_i and l_j , which fulfill the following: (i) l_i originates at $f^{-1}(w_{s_i})$ and is destined for $f^{-1}(w_{d_i})$ and (ii) l_j originates at $f^{-1}(w_{s_j})$ and is destined for $f^{-1}(w_{d_j})$. It is easy to see that l_i and l_j form a “cross” within the underlying network and they inevitably intersect (see the lines that correspond to the crossing chords of Figure 7).

Lemma 3 *The number of crossings produced by Algorithm-1 for instance I is the same as the number of chords crossing in circle R_G .*

Proof: Assume to the contrary that there exists two lines l and l' which, when routed by Algorithm-1, cross each other in G and their associated chords do not

cross in R_G . By Lemma 2.ii and Lemma 2.iii, l and l' do not share a terminal at a common leg node, which also holds in the circle and thus in R_G (i.e. chords that share a common end point do not cross each other).

Without loss of generality, we assume that l (l') is destined for leg node v (v'), where v is to the left of v' in the embedding of G . Algorithm-1 ensures that l and l' will cross along the backbone edge which connects u_{k-1} and u_k , where $u_k = bn(v)$ and that l' is not destined for v , i.e. $v \neq v'$. The contradiction is implied by the fact that such a pair of lines corresponds to two crossing chords in R_G . $\square$

Lemma 4 *Assume a caterpillar with at most two legs per node and let the lines terminate at leg nodes only. Then, the problem of drawing a set of lines $\mathcal{L}$ in an n -node caterpillar with two legs per node so that the number of crossings among pairs of lines is minimized can be solved in $O(n + \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.*

Proof: By Lemma 1, Step B of Algorithm-1 needs $O(n + |\mathcal{L}|)$ time. Using bucket sort [5], we can sort all lines at each individual leg node v in $O(|\mathcal{L}_v|)$ time, where $\mathcal{L}_v$ is the set of lines that originate at v . Therefore, Step C.1 needs a total of $O(|\mathcal{L}|)$ time. Steps C.2 and C.3 need $O(\sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time in total. $\square$

Theorem 2 *An instance $(P, \mathcal{L})$ of the MLCM-FixedSE problem on an n -node path P can be solved in $O(n + \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.*

4 The Metro-Line Crossing Minimization Problem on a Tree

In this Section, we consider the MLCM problem on a tree $T = (V, E)$, where $V = \{u_1, \dots, u_n\}$ and $E = \{e_1, \dots, e_{n-1}\}$. In the embedding of T , we assume that the neighbors of each node u of T are located either to the left or to the right of u (refer to Figure 8). In particular, we consider a “left-to-right tree structured network” to represent the underlying network. In such a network, we do not allow lines which make “right-to-right” or “left-to-left” turns, which implies that all lines should be x -monotone. For instance, in Figure 8, we would not permit a line connecting the two leftmost leaves. This assumption is motivated by the fact that a train can not make a 180° turn within a station. We seek to route all lines along the edges of T , so that the total number of crossings among pairs of lines is minimum.

We adopt the 2-side model, where each line uses the left side of a node to “enter” it and the right side to “exit” it. We refer to the edges that are adjacent to the left (right) side of node u in the embedding of T as *incoming (outgoing) edges* of u (see Figure 9). Since we assume that the lines are x -monotone, the notions of the origin and the destination of a line, as defined in Section 3.1, also apply in the case of line crossing minimization on “left-to-right tree structured network”.




Figure 8: A “left-to-right tree structured network”. Nodes containing diamonds or cycles correspond to terminal stations. Diamonds correspond to potential origins. Cycles correspond to potential destinations.

We consider the case where all line terminals are located only at nodes of degree 1 (see Figure 8) and the lines can terminate at any track of their terminal stations². Since the number of lines that “enter” an internal node is equal to the number of lines that “exit” it, we simply have to specify either the order of the lines that enter the node or the corresponding order when they leave it. Recall that we do not permit crossings inside the nodes.




Figure 9: Incoming and outgoing edges of a node u .

Assuming that the edges of T are directed from left to right in the embedding of T , we first perform a topological sorting over the nodes of T . As in the preceding section, we route the lines along the edges of T incrementally. We consider the nodes of T in their topological order. This ensures that whenever we consider the next node u all of its incoming lines have already been routed up to its left neighbor nodes.

A key component of our algorithm is a numbering of the nodes of the underlying network T based on the node u , that we are currently processing, e.g. a function $\text{ETN}_u : V \rightarrow \{0, 1, \dots, |V| - 1\}$. More precisely, given a node

²Recall that, in the case of a path network, this problem was quite easy due to the structure of the path.




Figure 10: A sample Euler-tour-numbering of a “left-to-right tree structured network” starting from the leftmost node

u of T , we number all nodes of T according to the order of first appearance when moving clockwise along the external face of T starting from node u . Note that such a numbering is unique with respect to the node u and we refer to it as the *Euler tour numbering* starting from node u or simply as ETN_u . Also, note that the computation of only one numbering is enough (say a numbering from the node with the smallest topological order number) in order to compute the corresponding Euler tour numberings from any other node of T , since $ETN_v(w) = (ETN_u(w) - ETN_u(v)) \bmod |V|$. Figure 10 illustrates such a numbering.

The basic steps of our algorithm are outlined in Algorithm-2. As already mentioned, we consider the nodes of T in their topological order. Each time we consider a new node u , we distinguish the following cases:

Case 1: $\text{indegree}(u) = 0$

If node u is of indegree zero (i.e. u is a leaf containing the origins of some lines), we simply sort the lines that originate from u based on the Euler tour numbering ETN_u (starting from u) of their destinations in ascending order. They are assigned tracks based on their sorted order from top to bottom.

Case 2: $\text{indegree}(u) = 1$

We simply pass the lines from the left neighbor node of u to u without introducing any crossing (i.e. by keeping the order of the lines unchanged).

Case 3: $\text{indegree}(u) > 1$

In the case where node u is of indegree greater than one, we have to “merge” its incoming lines and thus, we may introduce crossings. We “stably merge”³ the incoming lines based on the Euler tour numbering

³Given two sorted lists A and B of n and m elements, respectively, a *stable merging* of A and B is a sequence S of $n + m$ elements, consisting of the elements of A and B , such that i) S is sorted and ii) the internal ordering of A and B is maintained in the resulting sequence S , i.e. elements with the same value appear in sequence S in the same order in which they appear in sequences A and B before the merging [5, pp 170], [14].

Algorithm 2: MLCM ON TREES

Input : A set of lines $\mathcal{L} = \{l_1, \dots, l_k\}$ and an embedded tree $T = \{V, E\}$.

Output: A routing of the lines along the edges of T , so that they cross each other as few times as possible.

Require: In the embedding of T , we assume that the neighbors of each node u of T are located either to the left or to the right of u . The terminals are located only at leaf nodes.

{**Step A:** Topological Sorting}

Perform a topological sorting over the nodes of T , assuming that the edges of T are directed from left to right in the embedding of T ;

{**Step B:** Compute an Euler Tour Numbering}

Compute an Euler tour numbering of the nodes of T starting from the node with the smallest topological order number;

{**Step C:** Line Routing}

foreach node u in topological order **do**

if ($\text{inDegree}(u) == 0$) **then**

 Sort the lines that originate at u based on the Euler tour numbering ETN_u (starting from u) of their destination nodes in ascending order and draw their tracks in this order;

else

if ($\text{inDegree}(u) == 1$) **then**

 Pass the lines from the left neighbor node of u to u without introducing any crossing;

else

 { $\text{inDegree}(u) > 1$ }

 “Stably merge” the incoming lines of u based on the Euler tour numbering ETN_u (starting from u) of their destination nodes in ascending order;

 Connect the lines to their corresponding tracks in the left side of u , introducing the necessary crossings;

(obtained from u) of their destinations so that: i) Lines coming along the same edge do not change order, and ii) if two lines with the same destination come along different edges, the one coming from the topmost edge is considered to be smaller.

Figure 11 illustrates a sample routing produced by Algorithm-2. We use different types of lines to denote lines that originate at a common leaf node. The construction of Algorithm-2 supports the following Lemma:

Lemma 5 *In a solution produced by Algorithm-2 the following statements hold:*

- (i) *Two lines l and l' cross each other at most once.*
- (ii) *Two lines l and l' with the same origin/destination do not cross each other.*




Figure 11: A sample routing obtained from Algorithm-2.

- (iii) Let l and l' be two lines that cross each other. Then, l and l' will cross just before entering their leftmost common node.

Proof: The first property is implied by the fact that the underlying network is a tree. The second property is due to the sorting of the lines taking place at nodes of indegree zero and the stable merging taking place at nodes of indegree greater than one. Finally, the latter property holds because of the numbering of the nodes. To see this consider two lines l and l' that cross each other and let u be their leftmost common node. Without loss of generality, we further assume that l enters u from edge e_l , l' enters u from edge $e_{l'}$ and e_l is located to the top of $e_{l'}$. Since l and l' cross each other, $\text{ETN}_u(l) < \text{ETN}_u(l')$ ⁴. Therefore, the merging step of the Algorithm-2 will imply a crossing among these lines before entering u . $\square$

By using Lemma 5 and by following a similar approach as in the proof of Theorem 4, we can show that Algorithm-2 produces an optimal solution, in terms of line crossings.

Theorem 3 Assuming that each line terminates at leaf nodes, an instance $(T, \mathcal{L})$ of the MLCM problem on a “left-to-right” n -node tree T of maximum indegree d can be solved in $O(n + \log d \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.

Proof: Steps A and B of Algorithm-2 need $O(n)$ total time. Using bucket sort [5], we can sort all lines at each individual leaf node v in $O(|\mathcal{L}_v|)$ time, where $\mathcal{L}_v$ is the set of lines that originate at v . Therefore, the sorting of the lines in Step C (taking place at nodes of indegree zero) needs a total of $O(|\mathcal{L}|)$ time. The stable merging of the incoming lines of a node can be accomplished by successively merging (in a bottom up fashion, and in an order similar to that employed by a traditional non-recursive implementation of merge-sort) pairs of arrays built by concatenating the destinations of the lines that enter the node

⁴We denote by $\text{ETN}_u(l)$ the euler tour numbering (obtained from u) of the destination of l .

from top to bottom. If the merging of two arrays is performed by employing the “counting sort” stable sorting method [5] and since each line will participate in at most $\log d$ mergings, the merging of all incoming lines will be completed in $O(\log d \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time, where d is the maximum indegree of T . Thus, the algorithm terminates after $O(n + \log d \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time. $\square$

Corollary 1 *Assuming that each line terminates at leaf nodes, an instance $(T, \mathcal{L})$ of the MLCM problem on a “left-to-right” n -node tree T with bounded degree can be solved in $O(n + \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.*

4.1 The MLCM-SE and MLCM-FixedSE problems on a Tree

Since a path can be viewed as a degenerated case of a tree, Theorem 1 implies that MLCM-SE problem on a tree is NP-Hard. However, for the MLCM-FixedSE problem we can obtain a polynomial time algorithm adopting a similar approach as the one of Section 3.1. For each node u of T we introduce at most four new nodes u_L^t, u_L^b, u_R^t and u_R^b adjacent to u . Node u_L^t (u_L^b) is placed on top (below) and to the left of u in the embedding of T and contains all lines that originate at u ’s top (bottom) station end. Similarly, node u_R^t (u_R^b) is placed on top (below) and to the right of u in the embedding of T and contains all lines that are destined for u ’s top (bottom) station end. In the case where any of the u_L^t, u_L^b, u_R^t or u_R^b does not contain any lines we ignore its existence. So, instead of restricting each line to terminate at a top or at a bottom station end in its terminal stations, we equivalently consider that it terminates to some of the newly introduced nodes. Note that the underlying network remains a tree after the introduction of the new nodes, so our algorithm can be applied in this case, too. The following Theorem summarizes our result.

Theorem 4 *An instance $(T, \mathcal{L})$ of the MLCM-FixedSE problem on a “left-to-right” n -node tree T of maximum indegree d can be solved in $O(n + \log d \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.*

Corollary 2 *An instance $(T, \mathcal{L})$ of the MLCM-FixedSE problem on a “left-to-right” n -node tree T with bounded degree can be solved in $O(n + \sum_{i=1}^{|\mathcal{L}|} |l_i|)$ time.*

5 Conclusions

Clearly, our work is a first step towards solving the MLCM problem and its variants in arbitrary graphs. Extending the work of Benkert et al. [4] we studied path and tree networks. However, we did not consider the case where the underlying network is an arbitrary graph. Additionally, for the case where the underlying network is a tree we only considered the case, where the terminals are located at nodes of degree 1. No results are known regarding the case where we permit terminals at internal nodes of the tree.

Another line of research would be to develop approximation algorithms for the MLCM-SE problem on paths and trees. The problem of determining a solution of the general metro-line routing problem, in which the graph drawing and line routing are solved simultaneously is also of particular interest as a second step in the process of automated metro map drawing.

References

- [1] M. Asquith, J. Gudmundsson, and D. Merrick. An ILP for the line ordering problem. Technical Report PA006288, National ICT Australia, 2007.
- [2] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall, 1999.
- [3] M. A. Bekos, M. Kaufmann, K. Potika, and A. Symvonis. Line crossing minimization on metro maps. In S.-H. Hong and T. Nishizeki, editors, *Proc. 15th International Symposium on Graph Drawing (GD2007)*, LNCS 4875, pages 231–242, 2007.
- [4] M. Benkert, M. Nöllenburg, T. Uno, and A. Wolff. Minimizing intra-edge crossings in wiring diagrams and public transport maps. In M. Kaufmann and D. Wagner, editors, *Proc. 14th Int. Symposium on Graph Drawing (GD’06)*, volume 4372 of *Lecture Notes in Computer Science*, pages 270–281. Springer-Verlag, 2007.
- [5] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. The MIT Press, 2001.
- [6] M. Dickerson, D. Eppstein, M. T. Goodrich, and J. Y. Meng. Confluent drawings: Visualizing non-planar diagrams in a planar way. *Journal of Graph Algorithms Applications*, 9(1):31–52, 2005.
- [7] M. C. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, 1980.
- [8] S.-H. Hong, D. Merrick, and H. A. D. d. Nascimento. The metro map layout problem. In N. Churcher and C. Churcher, editors, *Australasian Symposium on Information Visualisation, (invis.au’04)*, volume 35 of *CRPIT*, pages 91–100. ACS, 2004.
- [9] P. Hui, M. J. Pelsmajer, M. Schaefer, and D. Stefankovic. Train tracks and confluent drawings. *Algorithmica*, 47(15):465–479, 2007.
- [10] M. Kaufmann and D. Wagner, editors. *Drawing Graphs: Methods and Models*, volume 2025 of *Lecture Notes in Computer Science*. Springer-Verlag, 2001.

- [11] S. Masuda, K. Nakajima, T. Kashiwabara, and T. Fujisawa. Crossing minimization in linear embeddings of graphs. *IEEE Trans. Comput.*, 39(1):124–127, 1990.
- [12] M. Nöllenburg and A. Wolff. A mixed-integer program for drawing high-quality metro maps. In P. Healy and N. S. Nikolov, editors, *Proc. 13th Int. Symposium on Graph Drawing (GD'05)*, volume 3843 of *Lecture Notes in Computer Science*, pages 321–333. Springer-Verlag, 2006.
- [13] J. M. Stott and P. Rodgers. Metro map layout using multicriteria optimization. In *Proc. 8th International Conference on Information Visualisation*, pages 355–362. IEEE Computer Society, 2004.
- [14] A. Symvonis. Optimal stable merging. *The Computer Journal*, 38(8):681–690, 1995.