

Two Polynomial Time Algorithms for the Metro-line Crossing Minimization Problem^{*}

Evmorfia Argyriou¹, Michael A. Bekos¹, Michael Kaufmann²,
and Antonios Symvonis¹

¹ School of Applied Mathematical & Physical Sciences,
National Technical University of Athens, Greece
{fargyriou,mikebekos,symvonis}@math.ntua.gr

² University of Tübingen, Institute for Informatics, Germany
mk@informatik.uni-tuebingen.de

Abstract. The *metro-line crossing minimization (MLCM)* problem was recently introduced as a response to the problem of drawing metro maps or public transportation networks, in general. According to this problem, we are given a planar, embedded graph $G = (V, E)$ and a set L of simple paths on G , called *lines*. The main task is to place the lines on G , so that the number of crossings among pairs of lines is minimized.

Our main contribution is two polynomial time algorithms. The first solves the general case of the MLCM problem, where the lines that traverse a particular vertex of G are allowed to use any side of it to either “enter” or “exit”, assuming that the endpoints of the lines are located at vertices of degree one. The second one solves more efficiently the restricted case, where only the left and the right side of each vertex can be used.

To the best of our knowledge, this is the first time where the general case of the MLCM problem is solved. Previous work was devoted to the restricted case of the MLCM problem under the additional assumption that the endpoints of the lines are either the topmost or the bottommost in their corresponding vertices, i.e., they are either on top or below the lines that pass through the vertex. Even for this case, we improve a known result of Asquith et al. from $O(|E|^{5/2}|L|^3)$ to $O(|V|(|E| + |L|))$.

1 Introduction

A metro map can be modeled as a tuple (G, L) , which consists of a connected graph $G = (V, E)$, referred to as the *underlying network*, and a set L of simple paths on G . 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. Then, the process of constructing a metro map consists of a sequence of steps. Initially, one has to draw the underlying network nicely. Then, the lines have to be properly added

^{*} This work has been funded by the project PENED-2003. PENED-2003 is co - funded by the European Social Fund (75%) and Greek National Resources (25%).

into the visualization and, finally, a labeling of the map has to be performed over the most important features.

In the graph drawing and computational geometry literature, the focus so far has been nearly exclusively on the first and the third step. Closely related to the first step are the works of Hong et al. [5], Merrick and Gudmundsson [6], Nöllenburg and Wolff [7] and Stott and Rodgers [8]. The map labeling problem has also attracted the interest of several researchers. An extensive bibliography on map labeling is maintained on-line by Strijk and Wolff [9]. Interestingly enough, the intermediate problem of adding the line set into the underlying network was recently introduced by Benkert et al. [3], followed by [2]. Since crossings within a visualization are often considered as the main source of confusion, the main goal is to draw the lines, so that they cross each other as few times as possible. This problem is referred to as the *metro-line crossing minimization problem (MLCM)*.

1.1 Problem Definition

The input of the metro-line crossing minimization problem consists of a connected, embedded, planar graph $G = (V, E)$ and a set $L = \{l_1, l_2 \dots l_k\}$ of simple paths on G , called *lines*. We will refer to G as the underlying network and to the nodes of G as *stations*. We also refer to the endpoints of each line as its *terminals*. In this paper, we study the case where all line terminals are located at stations of degree one, which are referred to as *terminal stations*. Stations of degree greater than one are referred to as *internal stations*. The stations are represented as particular shapes (usually as rectangles but in general as polygons). The sides of each station that each line may use to either “enter” or “exit” the station are also specified as part of the input. Motivated by the fact that a line cannot make a 180° turn within a station, we do not permit a line to use the same side of a station to both “enter” and “exit”.

The output of the MLCM problem should specify an ordering of the lines at each side of each station, so that the number of crossings is minimized.

Each line l_i consists of a sequence of edges $e_1 = (v_0, v_1), \dots, e_d = (v_{d-1}, v_d)$. Stations v_0 and v_d are the terminals of line l_i . Equivalently, we say that l_i *terminates* or *has terminals at* v_0 and v_d . By $|l_i|$ we denote the length of line l_i .

Each line that traverses a station 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* (see the dark-gray colored bullets on the boundary of each station in Fig. 1b). In general, we may permit tracks to all sides of each station, (see Fig. 1a). In the case where the stations are represented as rectangles, this model is referred to as the *4-side* model. In the general case where the stations are represented as polygons of at most k sides, this model is referred to as the *k-side* model. A more restricted model, referred to as the *2-side* model, is the one where i) the stations are represented as rectangles and ii) all lines that traverse a station may use only its left and right side (see Fig. 1b).

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




Fig. 1. The underlying network is the gray colored graph

(see Fig. 1b). This is to emphasize that the line terminates at that station. The variant of the MLCM problem that fulfills this restriction is referred to as the *metro-line crossing minimization problem with station ends (MLCM-SE)*. If additionally, the information whether a line terminates at a top or at a bottom station end in its terminal station is specified as part of the input, the corresponding problem is referred to as *metro-line crossing minimization problem with fixed station ends (MLCM-FixedSE)*.

A further refinement of the MLCM problem concerns the location of the crossings among pairs of lines. If the relative order of two lines changes between two consecutive stations, then the two lines must intersect between these stations (see Fig. 1b). We call this an *edge crossing*. As opposed to an edge crossing, a *station crossing* occurs inside a station. For aesthetic reasons, we want to avoid station crossings whenever this is possible (e.g. in the case of 4-side model this is not always feasible; see Fig. 1a).

1.2 Previous Work and Our Results

The first results on the MLCM problem were presented by Benkert et al. in [3], who devised a dynamic-programming algorithm that runs in $O(|L|^2)$ time for the restricted case where the crossings are minimized along a single edge of G . Bekos et al. [2] proved that the MLCM-SE problem is NP-complete even in the case where the underlying network is a path. They also proved that the MLCM-FixedSE problem can be solved in $O(|V| + \log d \sum_{i=1}^{|L|} |l_i|)$, in the case where the underlying network is a tree of degree d . Extending the work of Bekos et al., Asquith et al. [1] proved that the MLCM-FixedSE problem is also solvable in polynomial time in the case where the underlying network is an arbitrary planar graph. The time complexity of their algorithm was $O(|E|^{5/2} |L|^3)$. They also proposed an integer linear program which solves the MLCM-SE problem.

This paper is structured as follows: In Section 2, we present a polynomial time algorithm, which runs in $O((|E| + |L|^2)|E|)$ time for the MLCM problem under the k -side model, assuming that the line terminals are located at stations of degree one. To the best of our knowledge no results are currently known regarding this general model. In Section 3, we present a faster algorithm for the special case of 2-side restriction. The time complexity of the proposed algorithm is $O(|V||E| + \sum_{i=1}^{|L|} |l_i|)$. It can also be employed to solve the MLCM-FixedSE problem, which drastically improves the running time of the algorithm

of Asquith et al. [1] from $O(|E|^{5/2}|L|^3)$ to $O(|V||E| + |V||L|)$. We conclude in Section 4 with open problems and future work.

2 The MLCM Problem under the k -Side Model

To simplify the description of our algorithm and to make the accompanying figures simpler, we restrict our presentation to the MLCM problem under the 4-side model, i.e., we assume that each station is represented as a rectangle and we permit tracks to all four sides of each station. Our algorithm for the case of k -side model is identical, since it is based on recursion over the edges of the underlying network. Recall that all line terminals are located at stations of degree one, the lines can terminate at any track of their terminal stations, and, finally, the sides of each station that each line may use to either “enter” or “exit” are specified as part of the input. We further assume that an internal station always exists within the underlying network, otherwise the problem can be solved trivially.

The basic idea of our algorithm is to decompose the underlying network by removing an arbitrary edge out of the edges that connect two internal stations (and, consequently, appropriately partitioning the set of lines that traverse this edge), then recursively solve the subproblem and, finally, derive a solution of the initial problem by i) re-inserting the removed edge and ii) connecting the partitioned lines along the re-inserted edge.

2.1 Base of Recursion

The base of the recursion corresponds to the case of a graph G_B consisting of a “central station” u containing no terminals and a particular number of terminal stations, say $v_1, v_2 \dots v_d$, incident to u (see Fig. 2c). To cope with this case, we first group all lines that have exactly the same terminals into a single line, which is referred to as *bundle*. The notion of bundles corresponds to the fact that lines with same terminals are drawn in a uniform fashion, i.e., occupying consecutive tracks at their common stations. So, in an optimal solution once a bundle is drawn, it can be safely replaced by its corresponding lines without affecting the optimality of the solution. In Fig. 2c, lines belonging to the same bundle have been drawn with the same type of non-solid line. Note that single lines are also treated as bundles in order to maintain a uniform terminology (refer to the solid lines of Fig. 2c). Then, the number of bundles of each terminal station is bounded by the degree of the “central station” u .

In order to route the bundles along the edges of G_B , we will make use of the *Euler tour numbering* that was proposed by Bekos et al. [2]. Let v be a terminal station of G_B . The Euler tour numbering of the terminal stations v_1, v_2, \dots, v_d of G_B with respect to v is a function $\text{ETN}_v : \{v_1, v_2, \dots, v_d\} \rightarrow \{0, 1, \dots, d-1\}$. More precisely, given a terminal station v of G_B , we number all terminal stations of G_B according to the order of first appearance when moving clockwise along the external face of G_B starting from station v , which is assigned the zero value. Note that such a numbering is unique with respect to v and we refer to it as




Fig. 2. Illustration of the base of the recursion. The numbering of the lines is arbitrary.

the Euler tour numbering starting from station v or simply as ETN_v . Also, note that the computation of only one numbering is enough in order to compute the corresponding Euler tour numberings from any other terminal station of G_B , since $ETN_{v'}(w) = (ETN_v(w) - ETN_v(v')) \pmod d$.

Our approach is outlined as follows: We first sort in ascending order the bundles at each terminal station v based on the Euler tour numbering ETN_v of their destinations (see Fig. 2a). This implies the desired ordering of the bundles along the side of each terminal station that is incident to the “central station” u . We will denote by $BND(v)$ the ordered set of bundles of each terminal station v . Then, we pass these bundles from each terminal station to the “central station” u along their common edge without introducing any crossings (see Fig. 2b). This will also imply an ordering of the bundles at each side of the “central station” u . To complete the routing procedure, it remains to connect equal bundles in the interior of the “central station” u , which may imply crossings (see Fig. 2c). Note that only necessary station crossings are created, since the underlying network is planar and from the Euler tour numbering it follows that no edge crossings will eventually occur. So, the optimality of the solution follows trivially.

2.2 Description of the Recursive Algorithm

Having specified the base of the recursion, we now proceed to describe our recursive algorithm in detail. Let $e = (v, w)$ be an edge which connects two internal stations v and w of the underlying network. If no such edge exists, then the problem can be solved by employing the algorithm of the base of the recursion.

Let L_e be the set of lines that traverse e . Any line $l_{e,i} \in L_e$ originates from a terminal station, passes through a sequence of edges, then enters station v , traverses edge e , exits station w and, finally, passes through a second sequence of edges until it terminates at another terminal station. Let $p : E \times L \rightarrow \mathbb{N}$ be a function, such that $p(e, l)$ denotes the position of edge e along line l . Formally, $L_e = \{l_{e,1}, l_{e,2}, \dots, l_{e,|L_e|}\}$, where $l_{e,i}$ denotes the i -th line of L_e . Since each line of L_e consists of a sequence of edges, set L_e can be written in the form $\{l_{e,i} = e_{e,i}^1 e_{e,i}^2 \dots e_{e,i}^{k-1} e_{e,i}^{k+1} \dots e_{e,i}^{|l_{e,i}|}\}; k = p(e, l_{e,i}), i = 1, 2, \dots, |L_e|\}$. We proceed by removing edge e from the underlying network and by inserting two




Fig. 3. Illustration of the removal of an edge that connects two internal stations

new terminal stations t_e^v and t_e^w incident to the stations v and w , respectively (see the dark-gray colored stations of the right drawing of Fig. 3). Let $G^* = (V \cup \{t_e^v, t_e^w\}, (E - \{e\}) \cup \{(v, t_e^v), (t_e^w, w)\})$ be the new underlying network obtained in this manner.

Since the edge e has been removed from the underlying network, the lines of L_e cannot traverse it any more. So, we force them to terminate at t_e^v and t_e^w , as it is depicted in the right drawing of Fig. 3. This is done by splitting the set L_e into two new sets L_e^v and L_e^w (see Fig. 3), which are formally defined as follows:

- $L_e^v = \{e_{e,i}^1, e_{e,i}^2, \dots, e_{e,i}^{k-1}(v, t_e^v); k = p(e, l_{e,i}), i = 1, 2, \dots, |L_e|\}$
- $L_e^w = \{(t_e^w, w), e_{e,i}^{k+1}, \dots, e_{e,i}^{|L_e|}; k = p(e, l_{e,i}), i = 1, 2, \dots, |L_e|\}$

The new set of lines that it is obtained after the removal of the edge e is $L^* = (L - L_e) \cup (L_e^v \cup L_e^w)$. Observe that the removal of edge e from the underlying network may disconnect it. In the case where G^* is connected, we recursively solve the MLCM problem on (G^*, L^*) . Otherwise, since G^* was obtained from G by the removal of a single edge, it has exactly two connected components, say G_1^* and G_2^* . Let $L(G_i^*)$ denotes the lines of L^* induced by G_i^* . In this case, we recursively solve the MLCM problem on $(G_1^*, L(G_1^*))$ and $(G_2^*, L(G_2^*))$.

The recursion will lead to a solution of (G^*, L^*) . Part of the solution consists of two ordered sets of bundles $\text{BND}(t_e^v)$ and $\text{BND}(t_e^w)$ at each of the terminal stations t_e^v and t_e^w , respectively. Recall that, in the base of the recursion, all lines in a bundle have exactly the same terminals. In the recursive step, a bundle corresponds to a set of lines whose relative positions cannot be determined. In order to obtain a solution of (G, L) , we restore the removed edge e and remove the terminal stations t_e^v and t_e^w . The bundles $\text{BND}(t_e^v)$ and $\text{BND}(t_e^w)$ of t_e^v and t_e^w have also to be connected appropriately along the edge e . Note that the order of the bundles of t_e^v and t_e^w is equal to those of v and w , due to the base of the recursion. Thus, the removal of t_e^v and t_e^w will not produce unnecessary crossings.

We now proceed to describe the procedure of connecting the ordered bundle sets $\text{BND}(t_e^v)$ and $\text{BND}(t_e^w)$ along edge e . We say that a *bundle is of size k* iff it contains exactly k lines. We also say that two bundles are *equal* iff they contain the same set of lines, i.e., the parts of the lines that each bundle contains correspond to the same set of lines. First, we connect all equal bundles. Let $b \in \text{BND}(t_e^v)$ and $b' \in \text{BND}(t_e^w)$ be two equal bundles. The connection of b and b' will result into a new bundle which i) contains the lines of b (or equivalently of b') and ii) its terminals are the terminals of b and b' that do not participate in




Fig. 4. Splitting the largest bundle. Note that no equal bundles exist.

the connection. Note that a bundle is specified as a set of lines and a pair of stations, that correspond to its terminals. When the connection of b and b' is completed, we remove both b and b' from $BND(t_e^v)$ and $BND(t_e^w)$.

If both $BND(t_e^v)$ and $BND(t_e^w)$ are empty, all bundles are connected. In the case where they still contain bundles, we determine the largest in size bundle, say b_{max} , of $BND(t_e^v) \cup BND(t_e^w)$. W.l.o.g. we assume that $b_{max} \in BND(t_e^v)$ (see the left drawing of Fig. 4). Since b_{max} is the largest bundle among the bundles of $BND(t_e^v) \cup BND(t_e^w)$ and all equal bundles have been removed from both $BND(t_e^v)$ and $BND(t_e^w)$, b_{max} contains at least two lines that belong to different bundles of $BND(t_e^v)$. So, it can be split into smaller bundles, each of which contains a set of lines belonging to the same bundle in $BND(t_e^v)$ (see the right drawing of Fig. 4). Also, the order of the new bundles in $BND(t_e^v)$ should follow the order of their corresponding bundles in $BND(t_e^v)$ in order to avoid unnecessary crossings (refer to the order of the bundles within the dotted rectangle of Fig. 4). In particular, the information that a bundle was split should be propagated to all stations that this bundle traverses, i.e., splitting a bundle is not a local procedure that takes place along a single edge but it requires greater effort. Note that the crossings between lines of b_{max} and bundles in $BND(t_e^w)$ cannot be avoided. In addition, no crossings among lines of b_{max} occur.

We repeat these two steps (i.e. connection of equal bundles and splitting the largest bundle) until both $BND(t_e^v)$ and $BND(t_e^w)$ are empty. Since we always split the largest bundle into smaller ones, this guarantees that our algorithm regarding the connection of the bundles along the edge e will eventually terminate.

Theorem 1. *Given a graph $G = (V, E)$ and a set of lines L on G that terminate at stations of degree 1, the metro-line crossing minimization problem under the 4-side model can be solved in $O((|E| + |L|^2)|E|)$ time.*

Proof. The base of the recursion trivially takes $O(|V| + \sum_{i=1}^{|L|} |l_i|)$, or simply, $O(|V||L|)$ total time. The complexity of our algorithm is actually determined by the connection of the bundles along a particular edge, which is performed at most $O(|E|)$ times, since we always remove an edge that connects two internal stations. The previous steps of our algorithm (i.e., the construction of graph G^* and the necessary recursive calls) need a total of $O((|V| + |E|)|E| + |V||L|)$ time.

In order to connect equal bundles, we initially sort the lines of $BND(t_e^v)$ using counting sort [4] in $O(|L| + |L_e|)$ time, assuming that the lines are numbered from 1 to $|L|$, and we store them in an array, say B , such that the i -th numbered

line occupies the i -th position of B . Then, all equal bundles can be connected by performing a single pass over the lines of each bundle of $\text{BND}(t_e^v)$. Note that, given a line l that belongs to a particular bundle of $\text{BND}(t_e^v)$, say b , we can determine in constant time to which bundle of $\text{BND}(t_e^w)$ it belongs by employing array B . So, in a total of $O(|b|)$ time, we decide whether b is equal to one of the bundles of $\text{BND}(t_e^w)$, which yields into an $O(|L_e|)$ total time for all bundles of $\text{BND}(t_e^v)$. Thus, the connection of equal bundles can be accomplished in $O(|L| + |L_e|)$ time.

Having connected all equal bundles, the largest bundle is then determined in $O(|m_e|)$ time, where $m_e = \text{BND}(t_e^v) \cup \text{BND}(t_e^w)$. Using counting sort, we can split the largest bundle in $O(|L| + |L_e|)$ time. The propagation of the splitting of the largest bundle needs $O(|V||L_e|)$ time. The connection of the equal bundles and the splitting of the largest bundle will take place at most $O(|L_e|)$ times. Since $|m_e| \leq 2|L_e|$ and $|L_e| \leq |L|$, the total time needed for our algorithm is $O((|E| + |V||L|^2)|E| + |V||L|)$.

Note that the above straight-forward analysis can be improved by a factor of $|V|$. This is accomplished by propagating the splitting of each bundle only to its endpoints (i.e., not to all stations that each individual bundle traverses). This immediately implies that some stations of G may still contain bundles after the termination of the algorithm. So, we now need an extra post-processing step to fix this problem. We use the fact that the terminals of G do not contain bundles, since they are always at the endpoints of each bundle, when it is split. This suggests that we can split –up to lines– all bundles at stations incident to the terminal stations. We continue in the same manner until all bundles are eventually split. Note that this extra step needs a total of $O(|E||L|)$ time and consequently does not affect the total complexity, which is now reduced to $O((|V| + |E| + |L|^2)|E| + |V||L|)$. Since G is connected, $|E| \geq |V| - 1$ and therefore our algorithm needs $O((|E| + |L|^2)|E|)$ time, as desired. \square

Corollary 1. *Given a graph $G = (V, E)$ and a set of lines L on G that terminate at stations of degree 1, the metro-line crossing minimization problem under the k -side model can be solved in $O((|E| + |L|^2)|E|)$ time.*

3 The MLCM Problem under the 2-Side Model

In this section, we adopt the scenario of Section 2 under the 2-side model, i.e., we study the MLCM problem assuming that each station is represented as a rectangle and we permit tracks to the left and the right side of each station, i.e., one of the rectangle’s sides is devoted to “incoming” edges/lines while the other is devoted to “outgoing” edges/lines (see Fig. 5a). This assumption, combined with the fact that we do not permit a line to use the same side of a station to both “enter” and “exit”, implies that all lines should be x -monotone.

Since the lines are x -monotone, we refer to the leftmost (rightmost) terminal of each line as its *origin* (*destination*). We also say that a line uses the left side of a station to *enter* it and the right side to *exit* it. Furthermore, we refer to the edges incident to the left (right) side of each station u in the embedding of G as




Fig. 5. (a) Incoming/outgoing edges of u . (b) Construction of graph G' when G consists of a single internal station u . (c) An edge numbering of G .

incoming (*outgoing*) edges of station u (see Fig. 5a). For each station u of G , the embedding of G also specifies an order of both the incoming and outgoing edges of u . We denote these orders by $E_{in}(u)$ and $E_{out}(u)$, respectively (see Fig. 5a).

A key component of our algorithm is a numbering of the edges of G , i.e., a function $EN : E \rightarrow \{1, 2, \dots, |E|\}$. In order to obtain this numbering, we first construct a directed graph $G' = (V', E')$, as follows: For each edge $e \in E$ of G , we introduce a new vertex v_e in G' (refer to the black-colored bullets of Figures 5b and 5c). Therefore, $|V'| = |E|$. Also, for each pair of edges e_i and e_{i+1} of G that are consecutive in that order in $E_{in}(u)$ or $E_{out}(u)$, where $u \in V$ is an internal station of G , we introduce an edge $(v_{e_i}, v_{e_{i+1}})$ in G' (refer to the black-colored solid edges of Fig. 5b). Finally, we introduce an edge connecting the vertex of G' associated with the last edge of $E_{in}(u)$ to the vertex of G' associated with the first edge of $E_{out}(u)$ (refer to the black-colored dashed edge of Fig. 5b). Then, $|E'| = O(|E|)$. An illustration of the proposed construction is depicted in Fig. 5c. Note that all edges of G' are either directed “downward” or “left-to-right” w.r.t. an internal station. Thus there exist no cycles within the constructed graph (no “right-to-left” edges exist to form cycles). The desired numbering of the edges of G is then implied by performing a topological sorting on G' (see Fig. 5c).

Since each line is a sequence of edges, it can be expressed as a sequence of numbers based on the edge numbering $EN : E \rightarrow \{1, 2, \dots, |E|\}$. We refer to the sequence of numbers assigned to each line as its *numerical representation*. Note that the numerical representation of each line is sorted in ascending order.

Let l and l' be two lines that share a common path of the underlying network G . Let also $a_1 \dots a_k c_1 \dots c_m b_1 \dots b_n$ and $g_1 \dots g_q c_1 \dots c_m h_1 \dots h_r$ be their numerical representations, respectively, where the subsequence $c_1 c_2 \dots c_m$ corresponds to their common path. Then, l and l' inevitably cross iff $(a_k - g_q) \times (b_1 - h_1) < 0$ (see Fig. 6a). Note that their crossing can be placed along any edge of their common path. This is because we aim to avoid unnecessary station crossings.

Consider now two lines l and l' that share only a single internal station u of G . We assume that u is incident to –at least– four edges, say e_v^1, e_v^2, e_v^3 and e_v^4 , where e_v^1 and e_v^2 are incoming edges of u , whereas e_v^3 and e_v^4 outgoing. We further assume that l enters u using e_v^1 and exits u using e_v^4 . Similarly, l' enters u using e_v^2 and exits u using e_v^3 (see Fig. 6b). Then, l and l' form a station crossing which




Fig. 6. Crossings that cannot be avoided. Note that in Fig. 6a, $a_k < g_q < h_1 < b_1$, whereas in Fig. 6b, $EN(e_v^1) < EN(e_v^2) < EN(e_v^3) < EN(e_v^4)$.

cannot be avoided iff $(EN(e_v^1) - EN(e_v^2)) \times (EN(e_v^4) - EN(e_v^3)) < 0$. In this case, the crossing of l and l' can only be placed in the interior of station u .

Our intention is to construct a solution where only crossings that cannot be avoided are present. We will draw the lines of G incrementally by appropriately iterating over the stations of G and by extending the lines from previously iterated stations to the next station. Assuming that the edges of G are directed from left to right in the embedding of G , we first perform a topological sorting of the stations of G . Note that since all edges are directed from left to right, the graph does not contain cycles (no right to left edges exist to form cycles) and therefore a topological order exists. We consider the stations of G in their topological order. This ensures that whenever we consider the next station, its incoming lines have already been routed up to its left neighbors. Let u be the next station in the order. We distinguish the following cases:

Case (a) : $indegree(u) = 0$ (i.e. terminal station).

A station u with $indegree(u) = 0$ corresponds to a station which only contains the origins of some lines. In this case, we simply sort in ascending order these lines lexicographically with respect to their numerical representations. This implies the desired ordering of the lines along the right side of station u . It also ensures that these lines do not cross along their first common path.

Case (b) : $indegree(u) > 0$.

Let $e_u^1, e_u^2, \dots, e_u^k$ be the incoming edges of station u , where $k = indegree(u)$ and $e_u^i = (u_i, u)$, $i = 1, \dots, k$. W.l.o.g. we assume that $EN(e_u^i) < EN(e_u^j)$, $\forall i < j$. The lines that enter u from e_u^1 will occupy the topmost tracks of the left side of station u . Then, the lines that enter u from e_u^2 will occupy the next available tracks and so on. This ensures that the lines that enter u from different edges will not cross with each other, when entering u .

Let L_u^i be the lines that enter u from edge e_u^i , $i = 1, 2, \dots, k$, ordered according to the order of the lines along the right side of station u_i . In order to specify the order of all lines along the left side of station u , it remains to describe how the lines of L_u^i are ordered when entering u , for each $i = 1, 2, \dots, k$. We stably sort in ascending order the lines of L_u^i based on the numbering of the edges that they use when exit station u . Note that in order to perform this sorting we only consider the number following $EN(e_u^i)$ in the numerical representation of each line. Also, the stable sorting ensures that only unavoidable edge crossings will occur along e_u^i . To see this

consider two lines $l, l' \in L_u^i$ which use the same edge to exit station u . Since the sorting is stable their relative position will not change when they enter u , which implies that they will not cross along the edge e_u^i .

Up to this point, we have specified the order of the lines along the left side of station u , say L_{in}^u . In order to complete the description of this case it remains to specify the order, say L_{out}^u , of these lines along the right side of u . Again, the desired order L_{out}^u is implied by stably sorting the lines of L_{in}^u based on the numbering of the edges that they use when they exit station u . Note that also in this case the sorting of the lines is performed by considering only the EN-number of the edges used by the lines when exit station u . Again, the stable sorting ensures that only unavoidable station crossings will occur in the interior of station u .

Note that the stable sortings that are performed at each terminal station ensure that only unavoidable station and edge crossings eventually occur. Also, an unavoidable edge crossing between two lines is always placed along the last edge of their common path.

Theorem 2. *Given a graph $G = (V, E)$ and a set of lines L on G that terminate at stations of degree 1, the metro-line crossing minimization problem under the 2-side model can be solved in $O(|V||E| + \sum_{i=1}^{|L|} |l_i|)$ time.*

Proof. The topological sorting on G needs $O(|V| + |E|)$ time. The construction of graph G' and the computation of a topological sorting on it need $O(|E|)$ time, since both the number of nodes and the number of edges of G' are bounded by $|E|$. Having computed the EN-number of each individual edge of the underlying network, the numerical representations of all lines can be computed in $O(\sum_{i=1}^{|L|} |l_i|)$ time. Using radix sort [4], we can lexicographically sort all lines at each terminal station v of indegree zero in $O((|E| + |L_v|)|l_{max}^v|)$ total time, where L_v is the set of lines that originate at v and l_{max}^v is the longest line of L_v . Therefore, the sorting of all lines at stations of indegree zero needs a total of $O((|E| + |L|)|V|)$ time, since the length of the longest line of L is at most $|V|$. We can –stably– sort the lines of each set $L_u^i, i = 1, 2, \dots, k$ based on the numbering of the edges that they use when exit u , using counting sort. This can be done in $O(|E| + |L_u|)$ total time, where L_u denotes the set of lines that traverse station u . Recall that counting sort is stable. Similarly, can –stably– sort the lines of each set L_{in}^u based on the numbering of the edges they use when exit station u in $O(|E| + |L_u|)$ time. Summing over all internal stations, our algorithm needs $O(|V||E| + \sum_{i=1}^{|L|} |l_i|)$. □

As already stated, our algorithm can be employed to solve the MLCM-FixedSE problem. Our approach is as follows: For each station u of G , we introduce four new stations, say u_l^t, u_l^b, u_r^t and u_r^b , adjacent to u . Station u_l^t (u_l^b) is placed on top (below) and to the left of u in the embedding of G and contains all lines that originate at u 's top (bottom) station end. Similarly, station u_r^t (u_r^b) is placed on top (below) and to the right of u in the embedding of G and contains all lines that are destined for u 's top (bottom) station end. In the case where some of the

newly introduced stations contain no lines, we simply ignore their existence. So, instead of restricting each line to terminate at a top or at a bottom station end in its terminal stations, we equivalently assume that it terminates to one of the newly introduced stations. The following theorem summarizes this result.

Theorem 3. *Given a graph $G = (V, E)$ and a set of lines L on G , the metro-line crossing minimization problem with fixed station ends under the 2-side model can be solved in $O(|V||E| + \sum_{i=1}^{|L|} |l_i|)$ time.*

4 Conclusions

In this paper, we studied the MLCM problem under the k -side model for which we presented an $O((|E| + |L|^2)|E|)$ algorithm, and a more efficient algorithm for the special case of 2-side model. Possible extensions would be to study the problem where the lines are not simple, and/or the underlying network is not planar. Our first approach seems to work even for these cases, although the time complexity is harder to analyze and cannot be estimated so easily. The focus of our work was on the case where all line terminals are located at specific stations of the underlying network. Allowing the line terminals anywhere within the underlying network would hinder the use of the proposed algorithms in both models. Therefore, it would be of particular interest to study the computational complexity of this problem.

References

1. Asquith, M., Gudmundsson, J., Merrick, D.: An ILP for the metro-line crossing problem. In: 14th Computing: The Australian Theory Symposium, pp. 49–56 (2008)
2. Bekos, M.A., Kaufmann, M., Potika, K., Symvonis, A.: Line crossing minimization on metro maps. In: Hong, S.-H., Nishizeki, T., Quan, W. (eds.) GD 2007. LNCS, vol. 4875, pp. 231–242. Springer, Heidelberg (2008)
3. Benkert, M., Nöllenburg, M., Uno, T., Wolff, A.: Minimizing intra-edge crossings in wiring diagrams and public transport maps. In: Kaufmann, M., Wagner, D. (eds.) GD 2006. LNCS, vol. 4372, pp. 270–281. Springer, Heidelberg (2007)
4. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms, 2nd edn. MIT Press, Cambridge (2001)
5. Hong, S.H., Merrick, D., Nascimento, H.: The metro map layout problem. In: Churcher, N., Churcher, C. (eds.) Invis.au 2004. CRPIT, vol. 35, pp. 91–100 (2004)
6. Merrick, D., Gudmundsson, J.: Path simplification for metro map layout. In: Kaufmann, M., Wagner, D. (eds.) GD 2006. LNCS, vol. 4372, pp. 258–269. Springer, Heidelberg (2007)
7. Nöllenburg, M., Wolff, A.: A mixed-integer program for drawing high-quality metro maps. In: Healy, P., Nikolov, N.S. (eds.) GD 2005. LNCS, vol. 3843, pp. 321–333. Springer, Heidelberg (2006)
8. Stott, J.M., Rodgers, P.: Metro Map Layout Using Multicriteria Optimization. In: 8th International Conference on Information Visualisation (IV 2004), pp. 355–362. IEEE, Los Alamitos (2004)
9. Wolff, A., Strijk, T.: The Map-Labeling Bibliography, maintained since (1996), <http://i11www.ira.uka.de/map-labeling/bibliography>