

An Efficient Fixed Parameter Tractable Algorithm for 1-Sided Crossing Minimization*

Vida Dujmović and Sue Whitesides

School of Computer Science, McGill University, 3480 University St.,
Montreal, Québec, Canada H3A 2A7,
{vida, sue}@cs.mcgill.ca

Abstract. We give an $\mathcal{O}(\phi^k \cdot n^2)$ algorithm for the 1-SIDED CROSSING MINIMIZATION problem, thus showing that the problem is Fixed Parameter Tractable. The constant ϕ in the running time is the golden ratio $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$. The constant k is the parameter of the problem: the number of allowed edge crossings.

1 Introduction

A common method for drawing directed acyclic graphs is to produce *layered drawings* or *hierarchical drawings* as introduced by Tomii *et al.* [17], Carpano [1], and Sugiyama *et al.* [16]. In these drawings, the vertices are arranged on two or more “layers”, i.e. on parallel horizontal lines, and edges are drawn straight between vertices on adjacent layers. Edges between vertices on the same layer are not permitted, and no point between layers may lie on more than two edges. Layouts of this kind have applications, for example, in visualization, in DNA mapping, and in row-based VLSI layout.

The readability of layered drawings is believed to depend crucially on the number of edge crossings. Once vertices have been assigned to layers, this number is determined by the orderings of the vertices within the layers. Unfortunately, the problem of choosing vertex orderings that minimize the number of edge crossings in layered drawings is in fact an NP-complete problem [12] even if there are only two layers [11]. The problem of choosing vertex orderings that minimize the number of edges whose removal leaves the graph planar is also NP-complete, even for two layers [10].

Most techniques for producing layered drawings first assign vertices to layers (sometimes this is determined by the context), and then do a *layer-by-layer sweep*. A permutation π_1 for the vertices in the top layer L_1 is chosen and fixed. Then for each succeeding layer L_i , a permutation π_i is sought that keeps to a minimum the number of edge crossings among the edges between L_{i-1} and L_i .

A key step in this method is to minimize crossings between two adjacent layers when the ordering on one layer is fixed. This problem is called 1-SIDED CROSSING MINIMIZATION. Unfortunately, the 1-SIDED CROSSING MINIMIZATION problem is also NP-complete [11]. The problem is NP-complete even for

* Research supported by FCAR and NSERC.

graphs with only degree-1 vertices in the fixed layer and vertices of degree at most 4 in the other layer [14], i.e. for a forest of 4-stars.

The 1-SIDED CROSSING MINIMIZATION problem is the focus of this paper. Many heuristics have been proposed (e.g. [19,16,9,11,6,2,18]). Jünger and Mutzel [13] gave an exact integer linear programming algorithm for the 1-SIDED CROSSING MINIMIZATION problem. They also surveyed heuristics and made performance comparisons with optimal solutions generated by their methods. They reported that the *iterated barycentre* method of Sugiyama *et al.* [16] performs best in practice. However, from a theoretical point of view the *median heuristic* of Eades and Wormald [11] is a linear 3-approximation algorithm, whereas the barycentre heuristic is a $\Theta(\sqrt{n})$ -approximation algorithm.

When only a small number, k , of edge crossings is acceptable, then an algorithm for 1-SIDED CROSSING MINIMIZATION whose running time is exponential in k but polynomial in the size of the graph may be useful. The theory of parameterized complexity [5] addresses complexity issues of this nature, in which a problem is specified in terms of one or more parameters. Such a problem with input size n and parameter size k is *fixed parameter tractable*, or in the class *FPT*, if there is an algorithm to solve the problem in $f(k) \cdot n^\alpha$ time, where α is a constant independent of k and n , and f is an arbitrary function dependent only on parameter k . A problem in FPT is thus solvable in polynomial time for any fixed value of k .

The more general h -LAYER CROSSING MINIMIZATION problem as well as the related h -LAYER PLANARIZATION problem have been studied from the fixed parameter tractability point of view in [8]. Here h represents the number of layers and planarization means to remove some number k of edges so that the remaining graph can be drawn without crossings (see [10]). It has been proven in [8], using bounded pathwidth techniques, that both these general problems (which include 1-SIDED CROSSING MINIMIZATION) are in the class FPT. Unfortunately, the pathwidth-based approach is only of theoretical interest, since the running time of the algorithms is $O(2^{32(h+2k)^3} n)$. In [7], other FPT techniques are used to derive an $O(k \cdot 6^k + |G|)$ time algorithm for 2-LAYER PLANARIZATION of a graph G , and an $O(3^k \cdot |G|)$ time algorithm for 1-LAYER PLANARIZATION.

In this paper, we give an algorithm for 1-SIDED CROSSING MINIMIZATION that runs in $\mathcal{O}(\phi^k |L_2|^2 + |L_1| |L_2|)$ time, where L_2 is the set of vertices on the free layer where the permutation π_2 of vertices is to be chosen, L_1 is the set of vertices on the fixed layer where the permutation π_1 of vertices is fixed, and the constant ϕ is the golden ratio. The algorithm is based on the FTP technique called *Bounded Search Tree*. This technique relies on exploring some search space, and then proving that its size depends only upon the parameter k . The search space thus becomes *constant size* and the algorithm is then polynomial time for each fixed k .

The remainder of this paper is organized as follows. After definitions and preliminary results in Section 2, we study properties of optimal drawings in Section 3. Our algorithm for the 1-SIDED CROSSING MINIMIZATION problem is then given in Section 4.

2 Problem Statement, Notation, and Some Facts

A graph $G = (V, E)$ with vertex set V and edge set $E \subseteq V \times V$ is called *bipartite* if there is a partition of V into two disjoint non-empty sets L_1 and L_2 such that $V = L_1 \cup L_2$ and $E \subseteq L_1 \times L_2$. The number of vertices and edges of G are respectively denoted by $n = |V|$ and $m = |E|$. Let d_v denote the degree of a vertex v . We assume $d_v \geq 1$.

In a *2-layer drawing* of a graph $G = (L_1, L_2; E)$, the vertices in L_1 and L_2 are positioned on two distinct parallel lines (layers), and the edges are drawn straight. Since edges are not allowed within a layer, G must be bipartite. Let L_1 denote the top, *fixed* layer, whose vertex ordering π_1 is fixed. Let L_2 denote the bottom, *free* layer, whose vertices are free to be permuted.

We study the following problem:

Problem: 1-SIDED CROSSING MINIMIZATION

Instance: a bipartite graph $G = (L_1, L_2; E)$, an integer k , and a fixed ordering π_1 for the vertex set L_1 on the top layer.

Question: Is there a 2-layer drawing of G that respects π_1 and that has at most k crossings?




Fig. 1. A 2-layer drawing. The ordering π_1 of L_1 is *fixed*. Vertices of L_2 are *free*.

From now on, we assume that input graphs are bipartite, with minimum degree at least 1, and that an ordering π_1 has been specified for the top layer. We do not consider multiple edges, although these are easy to handle.

Let $\langle G, \pi_1, k \rangle$ denote an instance of the 1-SIDED CROSSING MINIMIZATION problem, and let (G, π_1, π_2) denote a combinatorial representation of a 2-layer drawing of G , with π_1 and π_2 giving the permutations for the vertices on layers L_1 and L_2 , respectively. Let the number of crossings in the drawing (G, π_1, π_2) be denoted by $cr(G, \pi_1, \pi_2)$, and let the minimum possible number of crossings subject to the vertices of L_1 being ordered by π_1 be denoted by $cr(G, \pi_1, \pi_{opt})$. Note that $cr(G, \pi_1, \pi_{opt}) = \min_{\pi_2} \{ cr(G, \pi_1, \pi_2) \}$, where π_2 ranges over all permutations for L_2 . Let $v < w$ denote an ordered pair of vertices on the same layer, and let v, w denote an unordered pair of vertices. Sometimes it is convenient to denote unordered pairs of vertices by (v, w) , which is also used to denote an edge. The meaning will be clear from the context. Throughout the paper, the term “pair” refers to a pair of *distinct* objects.

Fact 1. [4] Two edges (v, v') and (w, w') , where $v, w \in L_2$ and $v', w' \in L_1$, cross in a 2-layer drawing if and only if $v < w$ and $w' < v'$, or $w < v$ and $v' < w'$.

Fact 2. [4] For vertices v and w in the free layer L_2 , the number of crossings of the edges incident to v with the edges incident to w is completely determined by the relative ordering of v and w .

Definition 1. Consider a problem instance $\langle G, \pi_1, k \rangle$, and let v and w be vertices in L_2 . The crossing number c_{vw} is the number of crossings that edges incident with v make with edges incident with w in drawings having $v < w$; the crossing number c_{wv} is for $w < v$.

Fact 3. [4] The total number of crossings in a 2-layer drawing (G, π_1, π_2) is:

$$cr(G, \pi_1, \pi_2) = \sum_{\forall v < w \in \pi_2} c_{vw}, \quad (1)$$

where the summation is over all ordered pairs $v < w$ of elements of π_2 ; furthermore,

$$\sum_{v, w \in L_2} \min(c_{vw}, c_{wv}) \leq cr(G, \pi_1, \pi_{opt}) \leq \sum_{v, w \in L_2} \max(c_{vw}, c_{wv}), \quad (2)$$

where the summations are over all unordered pairs v, w of vertices of L_2 .

3 Properties of Optimal Drawings

In this section, we establish a property of optimal drawings (G, π_1, π_{opt}) (see Lemma 1) that will be fundamental for our algorithm in the next section.

For each vertex v in L_2 , let l_v denote the leftmost neighbor of v in L_1 , and let r_v denote the rightmost neighbor of v in L_1 . Note that if $v \in L_2$ has degree 1, then $l_v = r_v$.

Now consider two vertices v and w in L_2 . We say that v and w are a *suited pair* if $r_v \leq l_w$, or if $r_w \leq l_v$; otherwise we call the pair *unsuited*. For example, in Fig. 2 v, u is a suited pair and so is pair w, u ; pair v, w is not suited. If v and w each have degree 1 and have the same neighbor in L_1 (i.e. $l_v = r_v = l_w = r_w$), we say that v and w are a *trivial suited pair*. The following fact, which is an immediate consequence of Fact 1 and the definition of unsuited pair, indicates the importance of the notion of suitable pairs.

Fact 4. A pair of vertices $v, w \in L_2$ is *unsuited* if and only if $c_{vw} \geq 1$ and $c_{wv} \geq 1$.

On the other hand, the edges of a suited pair v, w do not cross if v and w appear in their *natural ordering* in π_2 , i.e., if $v < w$ when $r_v \leq l_w$, and $w < v$ when $r_w \leq l_v$. If v, w is a trivial suited pair, then $c_{vw} = 0$ and $c_{wv} = 0$, and we say that both $v < w$ and $w < v$ are natural orderings for the pair v, w .

The notion of natural ordering leads to a useful fact for our algorithm.




Fig. 2. Pair v, w is unsuited, while v, u and w, u are suited pairs.

Fact 5. Suppose v, w is a suited pair for π_1 with natural ordering $v < w$. Then for any π_2 , the drawing (G, π_1, π_2) satisfies: (i) $c_{vw} = 0$; (ii) if $r_v \neq l_w$, then $c_{wv} = d_v \cdot d_w$; (iii) if $r_v = l_w$, then $c_{wv} = (d_v \cdot d_w) - 1$; and finally, (iv) unless v and w are a trivial suited pair, $c_{wv} > 0$.

Note that natural ordering is only defined for pairs of suited vertices. For general pairs, we say that $v < w$ is the *preferred ordering* for a pair v, w if $c_{vw} < c_{wv}$. Thus the natural ordering is the preferred ordering for non-trivial suited pairs. The following Lemma is the basis for our algorithm.

Lemma 1. For fixed π_1 , let $\Gamma_{opt} = (G, \pi_1, \pi_{opt})$ be a drawing with the minimum possible number of crossings. Then all suited pairs appear in π_{opt} in their natural ordering.

To prove Lemma 1 the following straight forward lemma will be useful.

Lemma 2. For $1 \leq i \leq |L_2|$, let v_i denote the vertex in the i^{th} position in π_2 of some drawing (G, π_1, π_2) with π_1 fixed. Moving any vertex $v_i \in L_2$ from its starting position i across the t consecutive vertices $v_{i+1}, v_{i+2}, \dots, v_{i+t}$ to the right creates a new drawing (G, π_1, π'_2) with:

$$cr(G, \pi_1, \pi'_2) = cr(G, \pi_1, \pi_2) + \sum_{j=1}^t (c_{v_{i+j}v_i} - c_{v_iv_{i+j}}). \quad (3)$$

Similarly, if v_i is moved to the left over t consecutive vertices, then the above summation is from $j = -1$ to $j = -t$ and the sign in front of the summation is “−”.

Proof. Assume, without loss of generality, that vertex v_i moves to the right across t consecutive vertices in (G, π_1, π_2) . This creates a new drawing (G, π_1, π'_2) . The only pairs of vertices in (G, π_1, π_2) whose relative ordering changes in (G, π_1, π'_2) are the pairs v_i, v_j for $i+1 \leq j \leq i+t$. Hence by Fact 2, these are the only pairs whose crossing number might change. In particular, the crossing number for a pair v_i, v_j for j in the range $[i+1, i+t]$ changes from $c_{v_iv_j}$ to $c_{v_{i+j}v_i}$. Substituting these changes into Equation 1 of Fact 3 gives Equation 3 above. \square

Now we give the proof of Lemma 1.

Proof of Lemma 1: The proof is by contradiction. Assume that in $\Gamma_{opt} = (G, \pi_1, \pi_{opt})$ there is a suited pair v, w whose ordering in π_{opt} is not its natural

ordering. Note that v and w are not degree-1 vertices with a common neighbor, as both orderings would be natural in that case. Assume that $v < w$ is the (unique) natural ordering of v, w .

By Fact 5 the crossing number $c_{vw} = 0$ and the crossing number $c_{wv} > 0$.

For a contradiction, we now prove that either v or w can be moved in Γ_{opt} such that the resulting drawing $\Gamma_{new} = (G, \pi_1, \pi_{new})$ satisfies $cr(\Gamma_{new}) < cr(\Gamma_{opt})$.

Let i and j denote the positions of w and v , respectively, in π_{opt} . Here $i < j$ since v and w appear in the order $w < v$ in π_{opt} .

If $|j - i| = 1$, we can interchange v and w without affecting any other pair of vertices in Γ_{opt} . Equation 3 in Lemma 2 gives the number of crossings in the resulting drawing Γ_{new} :

$$cr(\Gamma_{new}) = cr(\Gamma_{opt}) - c_{wv} + c_{vw} = cr(\Gamma_{opt}) - c_{wv} + 0.$$

Since $c_{wv} > 0$, we have $cr(\Gamma_{new}) < cr(\Gamma_{opt})$, which contradicts the optimality of Γ_{opt} .

If $|j - i| > 1$, let $u_{i+1}, u_{i+2}, \dots, u_{j-1}$ denote the vertices between w and v in π_{opt} , listed in order of appearance in π_{opt} . Regard these vertices as a frozen block U inside which no changes are made. See Fig. 3.




Fig. 3. Proof of Lemma 1

According to Lemma 2, moving v or w from one side of block U to the other may only affect the crossing number contributions of pairs of the form u, w and u, v for $u \in U$. Let c_{Up} denote the number of crossings that the edges incident to vertices in U have with the edges incident to a vertex p to the right of U : $c_{Up} = \sum_{u \in U} c_{up}$. Similarly, let c_{pU} denote this number of crossings when p lies to the left of U .

Since Γ_{opt} is optimal, we claim we have the strict inequality

$$c_{Uv} < c_{vU}. \quad (4)$$

Otherwise, we could move v to the left side of U and then interchange v with w to obtain a drawing with the following total number of crossings: $cr(\Gamma_{new}) = \Gamma_{opt} - c_{Uv} + c_{vU} - c_{wv} + 0$. Since $c_{wv} > 0$, if $c_{Uv} \geq c_{vU}$, then $cr(\Gamma_{new}) < cr(\Gamma_{opt})$, a contradiction.

observation: To conclude the case $|j - i| > 1$, and hence the proof of the lemma, it suffices to show that $c_{Uv} < c_{vU}$ implies $c_{wU} \geq c_{Uw}$, for this means

we can move w to the right side of U without increasing the total number of crossings in the resulting drawing and then interchange w and v to produce a drawing with fewer crossings than Γ_{opt} . This gives a contradiction, and so proves that the assumption that π_{opt} contains a suited pair not ordered by its natural ordering cannot hold.

To establish the desired inequality $c_{wU} \geq c_{Uw}$, we first derive some intermediate inequalities for c_{Uv} , c_{vU} , c_{wU} , and c_{Uw} in terms of sizes of the following sets: E_R = the set of edges in Γ_{opt} with one endpoint in U and the other endpoint strictly greater than r_v in the ordering π_1 ; E_L = the set of edges with one endpoint in U and the other endpoint strictly less than r_v in the ordering π_1 ; N_v = the neighbors of v ; and N_w = the neighbors of w .

By the definition of E_R , all the vertices in N_v occur in π_1 strictly before the L_1 endpoint of each edge in E_R . By the definition of E_L and by the fact that v, w is a suited pair with the natural ordering $v < w$, the vertices in N_w occur in π_1 strictly after the L_1 endpoints of the edges in E_L .

Fact 2 implies the following inequalities for crossing numbers:

$c_{Uv} \geq d_v \cdot |E_R|$: The edges incident to v and the edges in E_R all pairwise intersect, creating $d_v \cdot |E_R|$ crossings. Since E_R is a subset of the edges incident to U , $c_{Uv} \geq d_v \cdot |E_R|$.

$c_{vU} \leq d_v \cdot |E_L|$: This holds because no edge incident to v crosses any edge incident to U that is not in E_L .

$c_{wU} \geq d_w \cdot |E_L|$: The edges incident to w and the edges in E_L all pairwise intersect, so $c_{wU} \geq d_w \cdot |E_L|$.

$c_{Uw} \leq d_w \cdot |E_R|$: This holds because no edge incident to w intersects any edge incident to U that is not in E_R .

Recall inequality (4), that $c_{Uv} < c_{vU}$. Since $c_{Uv} \geq d_v \cdot |E_R|$ and $c_{vU} \leq d_v \cdot |E_L|$, we have $d_v \cdot |E_R| \leq c_{Uv} < c_{vU} \leq d_v \cdot |E_L|$, which implies that $|E_R| < |E_L|$. This, and the fact that $c_{wU} \geq d_w \cdot |E_L|$, and the fact that $c_{Uw} \leq d_w \cdot |E_R|$ together imply that $c_{wU} > c_{Uw}$. By the *observation* above, this completes the proof. \square

4 An Efficient FPT Algorithm

4.1 The Bounded Search Tree Approach for the Algorithm

One of the basic methods for developing FPT algorithms is the method of *bounded search trees* (BST) (see Downey and Fellows [5, Chapter 3.1]). In this method one builds a search tree which is exhaustively traversed for a solution. The critical observation for many parameterized problems is that, while the computation done at each node of the tree may depend on the problem size, the size of the tree itself depends only on the parameters.

In this section we present an FPT algorithm for the 1-SIDED CROSSING MINIMIZATION problem based on the bounded search tree approach. The key observations for building a BST for this problem lie in Lemma 1 and Fact 4. Here is an overview of our algorithm.

Lemma 1 allows us, at the start, to fix the relative ordering of each non-trivial suited pair of vertices in L_2 according to its unique natural ordering. The remaining unordered pairs of vertices in L_2 are either trivial suited pairs, or unsuited pairs which will each, by Fact 4, create a crossing no matter which relative ordering is chosen. We build a bounded search tree (i.e., a BST) based on the unsuited pairs. (It turns out that the trivial pairs neighbours can be dealt with later in the algorithm.) The input to every node of the BST is a budget B giving the remaining number of allowed edge crossings, and a relation D containing all pairs of L_2 ordered thus far. We will formally define relation D shortly. At each node of the BST some unordered pair (v, w) is chosen (i.e. a pair not in D) such that $c_{vw} \neq c_{wv}$. Then the node branches to two recursive subproblems. In one branch, the ordering of (v, w) is fixed to $v < w$ and the budget B is reduced by c_{vw} . In the other branch the ordering of (v, w) is fixed to $w < v$ and B is reduced by c_{wv} . Since we only work with unsuited pairs in building the tree, we know that $c_{vw} \geq 1$ and $c_{wv} \geq 1$. Therefore, since the initial budget $B = k$, the height of the BST is at most k .

As a matter of fact the situation is better than that, for two reasons. Firstly, since $c_{vw} \neq c_{wv}$, then either c_{vw} or c_{wv} is at least 2, so one of the two branches of the BST node reduces B by at least 2. Secondly, since $<$ is a transitive relation, fixing an ordering of the pair (v, w) at a node of the BST may in fact impose an ordering of another as yet unordered pair (p, q) in the relation D at that node. Hence B can be reduced not only by c_{vw} but also by either c_{pq} or c_{qp} , depending on which relative ordering is imposed on (p, q) .

4.2 The Algorithm

The following definitions will be useful for the description of the algorithm.

Let D be a directed acyclic graph (digraph) that represents a binary relation “ $<$ ” on the set of vertices L_2 . In particular, the vertices of L_2 are represented by nodes of a digraph D and an ordered pair of vertices $v < w$ is represented by a directed edge from v to w (denoted henceforth by vw) in D . The digraph D is stored as an $|L_2| \times |L_2|$ matrix. We use D to denote both the set of pairs in the current binary relation “ $<$ ” and the associated acyclic directed graph that represents these pairs as directed edges. The algorithm labels nodes in the BST with digraphs. The digraph associated with the root will be transitively closed, acyclic, and directed. As the algorithm progresses, it computes a digraph label for each child node it generates in the BST by choosing a directed edge to add to the digraph of the parent and then taking the transitive closure of this.

The following algorithm solves the 1-SIDED CROSSING MINIMIZATION problem.

Algorithm: 1-Sided Crossing Minimization

Input: $\langle G, \pi_1, k \rangle$

Output: π_{opt} if $\langle G, \pi_1, k \rangle$ is a YES instance, else NO

Step 0. Computing crossing numbers: Compute the crossing numbers c_{vw} and c_{wv} for all pairs of vertices in L_2 , stopping the computation of a particular crossing number as soon as it is known to exceed k .

Step 1. Checking for extreme values: Compare k with the upper and lower bound as per Fact 3.

if $k < \sum_{(v,w)} \min(c_{vw}, c_{wv})$ then output NO and HALT;

if $k \geq \sum_{(v,w)} \max(c_{vw}, c_{wv})$ then output an arbitrary π_2 and HALT.

Step 2. Initialization: Precompute the following information required by the bounded search tree.

$C = \{(v, w) | c_{vw} = c_{wv}\};$

$D_0 =$ a directed acyclic graph (V, E) , where $V = L_2$, and the directed edges $vw \in E$ correspond to the naturally ordered pairs (v, w) that satisfy $c_{vw} = 0$ and $c_{wv} \neq 0$. (It is easy to check that D_0 is transitively closed.);

$B_0 =$ initial budget $= k - \sum_{vw \in D_0} c_{vw} - \sum_{(v,w) \in C} c_{vw}$. Note that $\sum_{vw \in D_0} c_{vw} = 0$. Also note that we reduce the budget k by the eventual cost of the pairs in C even though these pairs do not appear in D_0 .

Step 3. Building the search tree: A node of the search tree has at most two children. Each node has a label (D, B) . The label D of a node represents a “possible” partial solution, i.e. a partial ordering of vertices of L_2 . The label B represents the remaining budget for crossings. For instance the label of the root is: $D = D_0$ and $B = B_0$.

We now build the search tree as follows. Label the **root** of the tree with (D, B) where $D = D_0$ and $B = B_0$. In general, for a non-leaf node labeled (D, B) , choose a pair (v, w) such that D contains no edge joining v and w and such that $c_{vw} \neq c_{wv}$. A pair (v, w) is thus an unordered pair not in C . In any ordering π_2 , the pair (v, w) is ordered as either vw or wv , so we create at most two children (D_1, B_1) and (D_2, B_2) of the non-leaf node (D, B) corresponding to these two possibilities. No child is created if its budget would be negative. Thus a node labeled (D, B) is a **leaf** if and only if, either it does not have an unordered pair $(v, w) \notin C$; or there remains any pair (v, w) for which both B_1 and B_2 are negative.

For a non-leaf node (D, B) , we label one of its two children by (D_1, B_1) where:

$D_1 =$ transitive closure of $D \cup vw$, and $B_1 = B - c_{vw} - \sum_{pq} c_{pq}$.

Here $D \cup vw$ represents the addition of directed edge vw to D . The summation in B_1 is over the directed edges that are added to $D \cup vw$ by the transitive closure and that have $c_{pq} \neq c_{qp}$. That is, the sum is over pq s.t. $pq \in D_1$ and $pq \notin D \cup vw$ and $c_{pq} \neq c_{qp}$.

Similarly, we label the other child of node (D, B) with (D_2, B_2) , where:

$D_2 =$ transitive closure of $D \cup vw$, and $B_2 = B - c_{vw} - \sum_{pq} c_{pq}$.

If a leaf is created whose label D has the property that

$\forall (v, w)$ if $vw \notin D$ and $wv \notin D$ then $(v, w) \in C$,

then **output $\pi_2 =$ topological sort of D** . Also, update the minimum number of crossings found so far to $k - B$, where B is the budget of the leaf, and update the best ordering so far to π_2 . We call such a node a *solution leaf*.

If after exploring the entire tree, no solution leaf is found, **output NO** and **HALT**; otherwise, output the best ordering found, which is π_{opt} , and **HALT**.

Remarks. In Step 0, we stop computing c_{vw} as soon as it becomes $k + 1$, even though c_{vw} may be bigger than that. This is because a child with $v < w$ would have a negative budget. Hence it suffices to know that $c_{vw} \geq k + 1$.

Step 3 of the algorithm effectively creates and explores the search tree simultaneously. This can be done by depth-first search, or by breath-first search. The depth-first way requires less space and is thus the preferred choice.

Also notice that, when creating a child by choosing, say, to order v and w as vw , we reduce the budget for the child by an amount computed not only for the ordered pair vw , but also for the pairs that are newly ordered by the transitive closure of $D \cup vw$. However, we only do this for newly ordered pairs whose two crossing numbers are not the same. Those whose crossing numbers are the same have already been accounted for in B_0 .

Theorem 6. *Given a bipartite graph $G = (L_1, L_2; E)$, a fixed permutation π_1 of L_1 , and an integer k , algorithm 1-Sided Crossing Minimization(G, π_1, k) determines in $O(\phi^k \cdot |L_2|^2 + |L_1||L_2|)$ time if $cr(G, \pi_1, \pi_{opt}) \leq k$ and if yes produces a 2-layer drawing (G, π_1, π_2) with the optimum number of crossings. The constant ϕ in the running time is the golden ratio $\phi = \frac{1+\sqrt{5}}{2} \approx 1.618$.*

Proof. Step 3 of the algorithm creates and explores the search tree simultaneously. For every node (D, B) of the search tree we maintain the following two invariants: (i) D is a transitively closed, directed, acyclic graph; (ii) The budget B at node (D, B) is $B = k - \sum_{vw \in C} c_{vw} - \sum_{vw \in D \& vw \notin C} c_{vw}$. This is true for the root node (D_0, B_0) . Suppose this is true for a node labeled (D, B) . At this node, the algorithm chooses an unordered pair (v, w) with $c_{vw} \neq c_{wv}$. This pair is used to create up to two child nodes. We claim that both $D \cup vw$ and $D \cup wv$ are acyclic. Suppose, on the contrary, that $D \cup vw$ contains a directed cycle. Then D must contain a directed path from w to v . Since D is transitively closed, it contains edge wv , contradicting the fact that (v, w) is unordered in D . Similarly for $D \cup wv$. Since the transitive closure of a directed acyclic graph is again acyclic, the graph labels D_1, D_2 for any child nodes created at a node labeled D are again transitive and acyclic. Thus all the graph labels in the search tree are directed, acyclic, and transitively closed. The fact that labels B_1 and B_2 agree with formula (ii), follows directly from the formulas used to compute these two labels from the parent label B in Step 3 of the algorithm.

As the tree is built, either a solution leaf is found, or the tree is completely explored without finding such a leaf. A solution leaf (D, B) has, by definition, a non-negative budget B . By the invariant (ii), the cost of all the crossings arising from the ordered pairs in D has been taken into account. By the definition of a solution leaf, all pairs (v, w) not ordered by directed edges in D are in C and satisfy $c_{vw} = c_{wv}$; hence the total cost directly attributable to them has already been deducted from the initial budget k . Hence any topological sort of D produces a total ordering consistent with D and having total cost $k - B$, where

$0 \leq B \leq k$. By the invariant (i), the label D of every node of the search tree is an acyclic graph and it necessarily has a topological sort. Based on this argument, a solution leaf (D, B) encodes an ordering π_2 such that $cr(G, \pi_1, \pi_2) \leq k$.

It is not difficult to verify that the solution leaves of the BST implicitly store all the orderings π_2 for which $cr(G, \pi_1, \pi_2) \leq k$ and in which all the suited pairs are ordered by their natural ordering. Lemma 1 implies that in order to decide if $\langle G, \pi_1, k \rangle$ is a YES or NO instance it is enough to consider only such orderings π_2 . Therefore, if there is an ordering π_2 such that $cr(G, \pi_1, \pi_2) \leq k$ the algorithm finds one. In fact, since the algorithm updates the best solution found so far, when it terminates it outputs an optimal ordering π_{opt} .

We now discuss the running time of the algorithm. Only for an unordered unsuited pair v, w that has $c_{vw} \neq c_{wv}$ are child nodes created, of which there are at most two. Therefore, in one child node the budget is reduced by at least 1 and in the other by at least 2. A node with $B = 0$ must be a leaf node, because any child of such a node would have a negative budget. Therefore, no further branching is allowed. At a node for which budget $B = 1$, at most one child can have a budget B_1 that is non-negative, and in this case, $B_1 = 0$ and the child must be a leaf. Thus a recurrence relation that generates an upper bound for the number of nodes in this search tree is:

$$s_B = s_{B-1} + s_{B-2} + 1 \text{ for } k \geq 2; s_0 = 1, s_1 = 2.$$

It can be verified by induction that for $B \geq 0$, $s_B = F_{B+2} + F_{B+1} - 1$ where F_B is the B -th Fibonacci number. From the bound on Fibonacci numbers and given that $B_0 \leq k$, it follows that $s_{B_0} < \frac{\phi^{k+2}}{\sqrt{5}} + \frac{\phi^{k+1}}{\sqrt{5}} - 1 < 1.9 \cdot \phi^k$. Thus the search tree has $\mathcal{O}(\phi^k)$ nodes.

The time taken at each node of the search tree is dominated by updating a transitive closure of its label D after insertion of one ordered pair $v < w$ (or $w < v$). Updating the transitive closure after one insertion can be done in $\mathcal{O}(|L_2|^2)$ time [3]. These updates are needed to generate the labels for the children, of which there are at most two. Thus the time taken in the third step of the 1-Sided Crossing Minimization algorithm is $\mathcal{O}(\phi^k \cdot |L_2|^2)$.

It can be shown that the time taken in steps 0 - 2 of the algorithm is no more than $\mathcal{O}(k \cdot |L_2|^2 + |L_1||L_2|)$. Thus the total running time of the algorithm is $\mathcal{O}(\phi^k \cdot |L_2|^2 + |L_1||L_2|)$. \square

5 Conclusion

We have proved that the 1-SIDED CROSSING MINIMIZATION problem is fixed parameter tractable by presenting a very easy-to-implement FPT algorithm for its solution. Moreover, the algorithm finds a drawing with the smallest possible number of crossings in the case that this number does not exceed k . The exponential part of the running time of the algorithm is 1.618^k . In many instances the base of this exponent will be even smaller. The reason is that a pair of vertices v, w will often have both crossing numbers c_{vw} and c_{wv} bigger than 1, or at least one of them bigger than 2. Each time a node of the BST branches on such a pair of vertices, the resulting search tree will be even smaller.

An interesting investigation for future research would be to compare experimentally the performance of the two known methods for optimal 1-SIDED CROSSING MINIMIZATION, namely, integer linear programming [13] versus our FPT algorithm.

References

1. M. J. Carpano. Automatic display of hierarchized graphs for computer aided decision analysis. *IEEE Trans. Syst. Man Cybern.*, SMC-10(11):705–715, 1980.
2. C. Catarci. The assignment heuristic for crossing reduction. *IEEE Trans. on Systems, Man, and Cybernetics*, 25(3), 1995.
3. T. Corman, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press, 1990.
4. G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.
5. R. G. Downey and M. R. Fellows. *Parametrized complexity*. Springer, 1999.
6. S. Dresbach. A new heuristic layout algorithm for directed acyclic graphs. In U. Derigs, A. Bachem, and A. B. A. Drexler, editors, *Operations Research Proceedings (1994)*, pages 121–126. Springer, 1995.
7. V. Dujmović, M. Fellows, M. Hallett, M. Kitching, G. Liotta, C. McCartin, N. Nishimura, P. Ragde, F. Rosemand, M. Suderman, S. Whitesides, and D. R. Wood. A fixed-parameter approach to two-layer planarization. In Mutzel [15], pages 1–15.
8. V. Dujmović, M. Fellows, M. Hallett, M. Kitching, G. Liotta, C. McCartin, N. Nishimura, P. Ragde, F. Rosemand, M. Suderman, S. Whitesides, and D. R. Wood. On the parameterized complexity of layered graph drawing. In *Proc. 9th European Symposium on Algorithms (ESA 2001)*, volume 2161 of *Lecture Notes in Comput. Sci.*, pages 488–499. Springer, 2001.
9. P. Eades and D. Kelly. Heuristics for drawing 2-layered networks. *Ars Combin.*, 21(A):89–98, 1986.
10. P. Eades and S. Whitesides. Drawing graphs in two layers. *Theoret. Comput. Sci.*, 131(2):361–374, 1994.
11. P. Eades and N. C. Wormald. Edge crossings in drawings of bipartite graphs. *Algorithmica*, 11(4):379–403, 1994.
12. M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM J. Algebraic Discrete Methods*, 4(3):312–316, 1983.
13. M. Jünger and P. Mutzel. 2-layer straightline crossing minimization: performance of exact and heuristic algorithms. *J. Graph Algorithms Appl.*, 1(1):1–25, 1997.
14. X. Muñoz, U. Unger, and I. Vrto. One sided crossing minimization is np-hard for sparse graphs. In Mutzel [15]. to appear.
15. P. Mutzel, editor. *Proc. Graph Drawing: 9th International Symposium (GD'01)*, volume 2265 of *Lecture Notes in Comput. Sci.* Springer, 2001.
16. K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Trans. Systems Man Cybernet.*, 11(2):109–125, 1981.
17. N. Tomii, Y. Kambayashi, and S. Yajima. On planarization algorithms of 2-level graphs. *Papers of tech. group on elect. comp., IECEJ*, EC77-38:1–12, 1977.
18. V. Valls, R. Marti, and P. Lino. A branch and bound algorithm for minimizing the number of crossing arcs in bipartite graphs. *J. Operational Res.*, 90:303–319, 1996.
19. J. N. Warfield. Crossing theory and hierarchy mapping. *IEEE Trans. Systems Man Cybernet.*, SMC-7(7):505–523, 1977.