

Exact Bipartite Crossing Minimization under Tree Constraints^{*}

Frank Baumann¹, Christoph Buchheim¹, and Frauke Liers²

¹ Technische Universität Dortmund, Fakultät für Mathematik, Vogelpothsweg 87, 44227 Dortmund, Germany

² Universität zu Köln, Institut für Informatik, Pohligstraße 1, 50969 Köln, Germany

Abstract. A *tanglegram* consists of a pair of (not necessarily binary) trees. Additional edges, called *tangles*, may connect the leaves of the first with those of the second tree. The task is to draw a tanglegram with a minimum number of tangle crossings while making sure that the trees are drawn crossing-free. This problem has relevant applications in computational biology, e.g., for the comparison of phylogenetic trees. Most existing approaches are only applicable for binary trees. In this work, we show that the problem can be formulated as a quadratic linear ordering problem (QLO) with side constraints. Buchheim et al. (INFORMS J. Computing, to appear) showed that, appropriately reformulated, the QLO polytope is a face of some cut polytope. It turns out that the additional side constraints do not destroy this property. Therefore, any polyhedral approach to max-cut can be used in our context. We present experimental results for drawing random and real-world tanglegrams defined on both binary and general trees. We evaluate linear as well as semidefinite programming techniques. By extensive experiments, we show that our approach is very efficient in practice.

Key words: tanglegram, graph drawing, computational biology, crossing minimization, quadratic programming, maximum cut problem

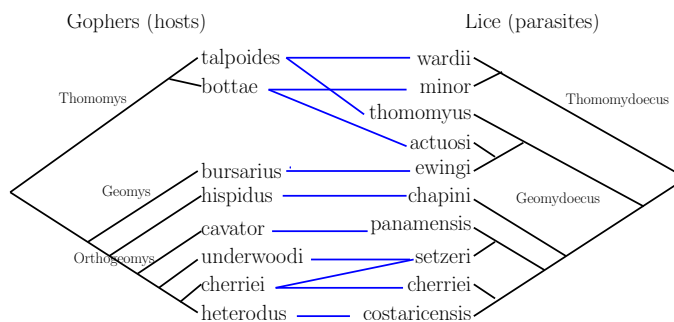
1 Introduction

A *tanglegram* [11] consists of a pair of trees T_1, T_2 and a correspondences between the leaf sets L_1 and L_2 of T_1 and T_2 , respectively. The correspondence is represented by edges between leaves in L_1 and L_2 called *tangles*. When visualizing a tanglegram, it is natural to ask for a drawing in which no edge crossings occur within either of the trees, while the number of tangle crossings is minimized. We require that the leaves in L_1 and L_2 are drawn on two parallel lines, while the trees are drawn outside the strip bounded by these lines.

^{*} Financial support from the German Science Foundation (DFG) is acknowledged under contracts Bu 2313/1-1 and Li 1675/1-1.

The task of drawing tanglegrams arises in several relevant applications, e.g., in computational biology for the comparison of phylogenetic trees [11]. A phylogenetic tree represents a hypothesis of the evolutionary history of a set of species. These species are drawn as the leaves of the tree, their ancestors as inner nodes. Different reconstruction methods may lead to a set of different candidate trees; a tanglegram layout then allows to compare a pair of such trees visually.

As another application, consider a phylogenetic tree of some set of species that serve as hosts for a certain set of parasites. The hypothesis that the evolution of hosts and their parasites is strongly correlated can be tested by analyzing a tanglegram layout. A tangle then specifies which host is affected by which parasite. Whereas in the first application the number of tangles incident to a leaf is always one, in the latter it can be higher, as shown below in a tanglegram from Hafner et al. [6] (here the hypothesis seems to be true).



Tanglegrams also occur in hierarchical clusterings, which can be visualized by so-called *dendrograms*. Dendrograms consist of trees where the elements to be clustered are identified with the leaves. Internal nodes determine clusters that contain the elements or sub-clusters. A tanglegram layout helps comparing the results of different clustering methods. Moreover, tanglegrams occur when analyzing software projects in which a tree represents package, class and method hierarchies. Hierarchy changes are analyzed over time, or automatically generated decompositions are compared with human-made ones. This application yields tanglegrams on trees that are not binary in general [10].

In the next section, we review related work. In Section 3, we introduce an exact model for tanglegrams that can be applied to pairs of general (not necessarily binary) trees with arbitrary tangle density. To this end,

we show that the task is to optimize a quadratic function over the linear ordering polytope intersected with further hyperplanes. We show that the corresponding polytope is isomorphic to a face of a cut polytope. We compare and evaluate different solution methods based on both linear and semidefinite approaches in Section 4. We show results for random as well as real-world instances. The results prove that our approach is very efficient in practice.

2 Related Work

Most of the literature is concerned with the case of binary trees and leaves that are in one-to-one correspondence. Whereas several of the presented methods could easily be generalized to arbitrary tangle layers, an extension to non-binary trees is usually not possible. When allowing general trees, one extreme case would be a star, where all leaves are adjacent to the root node. If both T_1 and T_2 are stars, there are no constraints on the orders of leaves on either shore, so that the problem specializes to the bipartite crossing minimization problem [9, 2].

We are not aware of any implementation of an exact method for drawing tanglegrams with non-binary trees. Fernau et al. [5] showed the NP-hardness of tanglegram layout, even in the case of binary trees. They also presented a fixed-parameter algorithm for binary tanglegrams. Recently, an improved fixed-parameter algorithm was presented by Böcker et al. [1] which can solve large binary instances quickly in practice, provided that the number of crossings is not too large. Finally, while in the recent paper by Venkatachalam et al. [13] the focus is on binary instances, a fixed-parameter algorithm for general tanglegram instances is presented. According to our knowledge, this is the only algorithm that could deal with non-binary trees; however, no implementation or running times are provided making it impossible to evaluate its practical performance.

Besides analyzing the performance and quality of several heuristics in a computational study for binary tanglegrams with one-to-one tangles, Nöllenburg et al. [10] also implemented a branch-and-bound algorithm and an exact integer-programming (IP) based approach for this case.

As we will compare our approach with the exact IP-approach of Nöllenburg et al. [10], we describe it in more detail in the following. A feasible but not necessarily optimal tanglegram layout is given as an input. For each inner node, a binary variable x_i is introduced. In the case of complete binary trees with n leaves each, this gives rise to $2(n - 1)$ variables. If $x_i = 1$, the subtree rooted in node i is flipped with respect to the

input drawing, otherwise it remains unchanged. As by definition there are no crossings within the trees, the number of crossings can be determined by counting the number of tangle crossings. Let (a, c) and (b, d) be tangles with $a, b \in L_1$ and $c, d \in L_2$. Let i be the lowest common ancestor of a, b in T_1 and j that of c, d in T_2 . If the tangles cross each other in the input drawing, then a crossing occurs in the output drawing if and only if either both subtrees below i and j are flipped or both remain unchanged. This can be expressed as $x_i x_j = 1$ or $(1 - x_i)(1 - x_j) = 1$. Similarly, if the edges do not cross each other in the input drawing, then there is a crossing in the output drawing if and only if either $(1 - x_i)x_j = 1$ or $x_i(1 - x_j) = 1$.

Thus minimizing the number of tangle crossings reduces to minimizing the sum of the given products. The latter is an instance of the unconstrained quadratic binary optimization problem, which is well-known to be equivalent to a maximum cut problem in some associated graph with an additional node [3]. In an undirected graph $G = (V, E)$, the cut $\delta(W)$ induced by a set $W \subseteq V$ is defined as the set of edges (u, v) such that $u \in W$ and $v \notin W$. If edge weights are given, the weight of a cut is the total weight of edges in the cut. Now the maximum cut problem asks for a cut of maximal weight or cardinality.

While Nöllenburg et al. used this model only for instances with one-to-one tangles, they briefly note that it could be extended to leaves of higher degree as well. However, their model cannot be generalized to instances with non-binary trees in a straightforward way. In many applications, the trees are not necessarily binary. In the next section we will present an exact model for tanglegrams that neither restricts the degree of inner nodes in the trees nor the number of tangles incident to a leaf.

3 An Exact Model for General Tanglegrams

The problem of drawing tanglegrams is closely related to bipartite crossing minimization. As argued above, the latter problem can be considered a special case of the former. Therefore, we first review approaches for drawing bipartite graphs.

3.1 Bipartite Crossing Minimization

Let $G = (V_1 \cup V_2, E)$ be a bipartite graph. The task is to draw G with straight line edges. The nodes in V_1 and V_2 have to be placed on two parallel lines H_1 and H_2 such that the number of edge crossings is minimal. Both heuristic and exact methods [9] exist for this problem.

Assume for a moment that the nodes on the first layer H_1 are fixed, and only the nodes on layer H_2 are permuted. For each pair of nodes on H_2 , we introduce a variable x_{uv} such that $x_{uv} = 1$ if u is drawn to the left of v and $x_{uv} = 0$ otherwise. For edges (i, k) and (j, l) with $i, j \in H_1$ and $k, l \in H_2$, such that i is left of j , a crossing exists if and only if l is left of k . We thus have to punish x_{lk} in the objective function. The task of minimizing the number of crossings is now equivalent to determining a minimum linear ordering on the nodes of H_2 . Exploiting $x_{uv} = 1 - x_{vu}$, we can eliminate half of the variables and only keep those with $u < v$. Note that bipartite crossing minimization with one fixed layer is already NP-hard [4].

If the nodes on both layers are allowed to permute, the number of crossings depends on the order of the nodes on each layer. Therefore, the problem can be modeled as a quadratic optimization problem over linear ordering variables. We write the quadratic linear ordering problem (QLO) in its general form as

$$(QLO) \quad \begin{array}{ll} \min & \sum_{(i,j,k,l) \in I} c_{ijkl} x_{ij} x_{kl} \\ \text{s.t.} & x \in P_{LO} \\ & x_{ij} \in \{0, 1\} \text{ for all } (i, j) \in J \end{array}$$

where P_{LO} is the linear ordering polytope, i.e. the convex hull of the incidence vectors of all linear orderings. The index set I consists of all quadruples (i, j, k, l) such that $x_{ij} x_{kl}$ occurs as a product in the objective function, while J is the set of all pairs (i, j) for which a linear ordering variable x_{ij} is needed. For the bipartite crossing minimization case, I and J are given as

$$\begin{aligned} I &= \{(i, j, k, l) \mid i, j \in H_1, i < j, \text{ and } k, l \in H_2, k < l\} \\ J &= \{(i, j) \mid i, j \in H_1 \text{ or } i, j \in H_2, i < j\} \end{aligned}$$

In order to linearize the objective function, we introduce a new binary variable y_{ijkl} for each $(i, j, k, l) \in I$, modeling the product $x_{ij} x_{kl}$. (Note that $y_{ijkl} = y_{kl ij}$.) Applying the standard linearization, the corresponding linearized quadratic linear ordering problem (LQLO) can be written as

$$(LQLO) \quad \begin{array}{ll} \min & \sum_{(i,j,k,l) \in I} c_{ijkl} y_{ijkl} \\ \text{s.t.} & x \in P_{LO} \\ & x_{ij} \in \{0, 1\} \quad \text{for all } (i, j) \in J \\ & y_{ijkl} \leq x_{ij}, x_{kl} \quad \text{for all } (i, j, k, l) \in I \\ & y_{ijkl} \geq x_{ij} + x_{kl} - 1 \quad \text{for all } (i, j, k, l) \in I \\ & y_{ijkl} \in \{0, 1\} \quad \text{for all } (i, j, k, l) \in I. \end{array}$$

Buchheim et al. [2] introduced the above model for bipartite crossing minimization. Additionally, a quadratic reformulation of the constraints defining P_{LO} was given: it was shown that a 0/1 vector (x, y) satisfying $y_{ijkl} = x_{ij}x_{kl}$ is contained in (LQLO) if and only if

$$x_{ik} - y_{ijik} - y_{ikjk} + y_{ijjk} = 0 \text{ for all } (i, j, k, l) \in I. \quad (1)$$

Note that (LQLO) is a quadratic binary optimization problem where the feasible solutions need to satisfy further side constraints, namely those restricting the set of feasible solutions to linear orderings. As unconstrained binary quadratic optimization is equivalent to the maximum cut problem [3], the task is to intersect a cut polytope with a set of hyperplanes.

In general, the convex hull of the corresponding feasible incidence vectors has a structure that is very different from that of a cut polytope. In the above context, however, Buchheim et al. [2] showed that the hyperplanes (1) cut out faces of the cut polytope.

3.2 Modeling Tanglegrams

Crossing minimization in tanglegrams can be seen as a generalization of bipartite crossing minimization. The set of feasible orderings is implicitly restricted by the given tree structures. Starting from the model discussed above, we formalize these restrictions as follows: let us consider a triple of leaves a, b, c in one of the trees, say T_1 . In case all pairwise lowest common ancestors coincide, all relative orderings between a, b , and c are feasible. However, if the lowest common ancestor of, say, a and b is on a lower level than that of, say, a and c (in this case, the former is a descendant of the latter), then c must not be placed between a and b , as an intra-tree crossing would be induced; see Figure 1.




Fig. 1. Leaf c is not allowed to lie between a and b .




Fig. 2. Variables x_{ac} and x_{bd} can be identified.

Therefore, we derive a betweenness restriction for every triple of leaves such that two of the leaf pairs have different lowest common ancestors.

Each such betweenness restriction of the form ‘ c cannot be placed between a and b ’ can be written in linear ordering variables as $x_{ac}x_{cb} = 0$ and $x_{bc}x_{ca} = 0$. In the linearized model (LQLO), the latter amounts to requiring

$$y_{acb} = 0 \text{ and } y_{cab} = 0 . \quad (2)$$

For binary trees with n leaves each, all triples of leaves have two different lowest common ancestors, so in this case the number of additional equations is $2\binom{n}{3}$.

In summary, we now obtain a quadratic linear ordering problem (QLO) on a smaller number of variables, with additional constraints of the form (2), where

$$\begin{aligned} J &= \{(i, j) \mid i, j \text{ are leaves of the same tree, } i < j\} \\ I &= \{(i, j, k, l) \mid (i, j), (k, l) \in J \text{ belong to different trees}\} . \end{aligned}$$

For complete binary trees with n leaves each, the total number of linear ordering variables is $2\binom{n}{2}$. The same number of variables is necessary in the corresponding bipartite crossing minimization model [2].

As mentioned above, the polytope corresponding to the linearized problem (LQLO) is isomorphic to a face of a cut polytope [2]. Since all y -variables are binary, constraints of the form (2) are always face-inducing for (LQLO). In summary, we derive the following result:

Theorem 1. *The problem of drawing tanglegrams with a minimum number of edge crossings can be solved by optimizing over a face of a suitable cut polytope.*

3.3 Binary Case

In the binary case, the model introduced in the preceding sections is closely related to the model presented in [10]. To see this, first observe that the two equations (2) can be written as

$$x_{ac} = x_{bc} . \quad (3)$$

This replacement does not affect the set of feasible solutions, even the corresponding LP-relaxations of (LQLO) are equivalent. Note however that introducing the y -variables allows to strengthen the model, see Theorem 1.

When using the linear equations (3) instead of the quadratic equations (2), we end up with a set of equivalence classes of linear ordering variables such that all pairwise orderings corresponding to variables in the

same class can only be flipped simultaneously. Two variables x_{ac} and x_{bd} belong to the same class if and only if there is a node r such that a, b and c, d are descendants of different children of r ; see Figure 2. In the binary case, a class of linear ordering variables thus corresponds to the decision of flipping the children of node r or not, which is modeled explicitly by a single variable in the model of Nöllenburg et al. [10].

In the general case, however, where node r has k children, there are $k!$ different orderings. As these cannot be modeled by a single binary variable, the model of Nöllenburg et al. [10] cannot be applied here.

4 Computational Results

We implemented the model detailed in Section 3.2. Instead of adding equations (3) explicitly, we used one variable for each equivalence class of linear ordering variables, thereby significantly reducing the number of variables. For evaluating the IP-based methods, we used CPLEX 11.2 [8]. The naive approach is to solve the linearized model (LQLO) using a standard integer programming solver. A more advanced approach is to solve the quadratic reformulation (1), using separation of cutting planes for max-cut, both in the context of integer and semidefinite programming. For the SDP approaches, we used the bundle method by Rendl et al. [12]. For comparison, we also implemented the IP model by Nöllenburg et al. [10] that only works for binary tanglegrams. For the tested binary instances, the running times for solving the latter are very comparable to those for our model. This can be expected since our model generalizes [10].

We generated random instances on general binary, ternary and quad trees. I.e., the degree of each internal node is at most 2, 3 or 4, respectively. Each tree has n leaves, either having one-to-one tangles or a certain tangle density $d\%$. Instances are generated following the description in [10], with obvious extensions to the more general cases considered here. Finally, we solved real-world binary tanglegram instances from [10] arising in applications in biology and general real-world instances from visualizing software hierarchies [7].

Average results are always computed over 5 randomly generated instances. For each instance, we imposed an upper limit of 10h of CPU time. Instances that could not be solved within this limit count with 10h in the averages. Runs were performed on Intel Xeon machines with 2.33GHz.

In Table 1, we present the average cpu time in seconds for real-world binary instances. Table 2 shows results for random ternary and quad trees, respectively. Figure 4 visualizes the results from Table 2 for $n = 128$.

n	# inst	SDP	IP	QP	IP+cyc	QP+cyc
0–49	2988	<1	<1	<1	<1	14
50–99	717	3	<1	2	<1	428
100–149	102	31	<1	7	<1	1100
150–199	42	125	1	32	1	1282
200–249	18	437	1	66	2	2704
250–299	18	4786	2	2529	7	10602
300–349	9	6483	2	80	9	8862
400–449	3	20508	12	12067	69	14931

Table 1. Average cpu time in seconds for real-world one-to-one binary trees having n leaves each [10]. Instances are grouped by their number of leaves. The second column shows the number of instances in each group. The nine largest instances with up to 540 leaves could not all be computed within the time and memory constraints and are omitted.

Running times for real-world general tanglegram instances are presented in Table 3.

n	d	SDP	IP	QP	IP+cyc	QP+cyc	SDP	IP	QP	IP+cyc	QP+cyc
64	1	3	1	<1	<1	1	<1	<1	<1	<1	<1
	5	15	12	4	67	6	2	3	1	1	<1
	10	18	27	19	1301	17	3	8	2	5	1
	15	78	54	23	3893	37	4	16	6	57	3
	20	41	54	48	12508	66	4	19	7	65	3
128	1	165	16	32	78	63	32	19	3	2	4
	5	362	448	280	22253	448	80	59	35	1045	51
	10	436	2179	791	36000	2746	73	1406	119	10147	495
	15	4049	3247	1551	36000	5583	77	844	352	31582	1184
	20	8293	2326	1408	36000	15426	1420	1560	1908	36000	10111

Table 2. Average cpu time in seconds for random general ternary (left) and quad trees (right) having n leaves each, density $d\%$.

The first column *SDP* shows results obtained by semidefinite optimization, whereas the remaining columns refer to IP-based approaches for solving the model from Section 3. *IP* refers to solving the standard linearization using CPLEX default, *QP* its quadratic reformulation (1). In the options *IP+cyc* and *QP+cyc*, cycle inequalities for max-cut are additionally separated, and all CPLEX cuts are switched off.

Clearly, for real-world binary trees with one-to-one tangles, the SDP approach usually needs considerably more time than the IP-based methods. Furthermore, memory requirements strongly increase with system size and so the largest instances could not be solved. On average, the

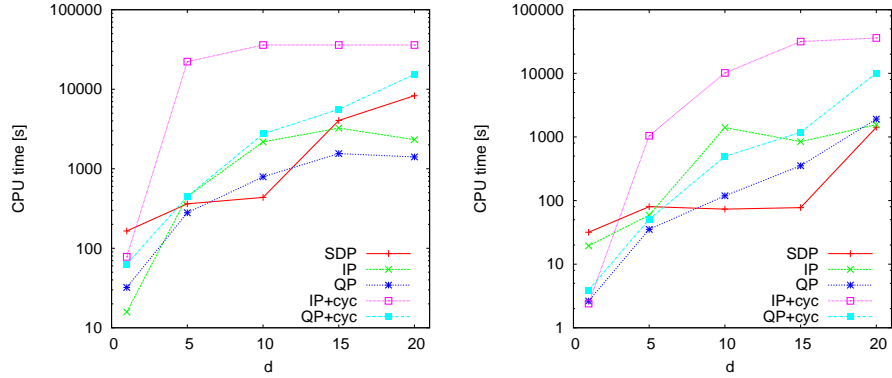


Fig. 3. Plot showing results for $n = 128$ of Table 2 for ternary (left) and quad trees (right).

instance	δ	SDP	IP	QP	IP+cyc	QP+cyc
philips_orig_4a	14	27618	4725	73	4321	32692
philips_orig_4b	14	26093	4205	74	2756	23443
philips_orig_4c	14	36000	2666	122	4030	36000
philips_orig_4d	14	27891	3208	314	4178	28902
philips_orig_4e	14	36000	2789	90	5665	36000
philips_4a_4b	9	5238	1395	4	420	5
philips_4a_4c	9	4769	1858	3	637	4
philips_4a_4d	9	2924	1467	4	494	3
philips_4a_4e	9	2575	827	2	338	3
philips_4b_4c	8	6526	1965	8	510	7
philips_4b_4d	8	4872	2070	5	577	5
philips_4b_4e	8	2127	649	4	124	36
philips_4c_4d	7	4611	804	3	217	5
philips_4c_4e	7	6403	1074	3	396	5
philips_4d_4e	7	3738	891	4	811	11

Table 3. CPU time in seconds for real-world general tanglegram instances [7]. δ specifies the maximum node degree.

fastest approaches for solving the largest instances are the pure standard linearization IP and the quadratic reformulation QP .

In fact, we can optimize tanglegrams with more than 500 leaves in each tree. This is the range of sizes arising in real-world applications. The real-world instances can be solved particularly fast. Interestingly, cycle separation for max-cut usually does not pay off for binary one-to-one tanglegrams: the running time increases, although the dual bounds are usually considerably better when cycle separation is included. Often, an

optimum solution can be determined in the root node. However, although the root bound is weak in the standard linearization, after few branching steps the optimum LP solution is often feasible and the program can stop. We found similar characteristics for random binary tanglegram instances.

The picture changes when varying the density of the tangles: for big enough tangle density the SDP approach usually outperforms the IP approaches. Memory requirements, however, usually prohibit solving instances with more than 500 leaf nodes and tangle density of 1%. On the IP side, reformulation is often preferable. Indeed, the best performance is observed when the problems are quadratically reformulated. For $n = 512$ and 1% tangle density, the average solution time is 2120.42 seconds. These instances cannot be solved within the given time limits when using only the standard linearization, with or without separation of cycle inequalities.

The instances for ternary and quad trees are computationally slightly more difficult. This is mainly due to the fact that the number of betweenness restrictions decreases in comparison to binary trees. Here again, the SDP approach performs well for denser instances however memory requirements strongly increase with system size. For larger instances, best performance is often found for the quadratic reformulation.

Comparing IP with $IP+cyc$ and QP with $QP+cyc$ for not necessarily binary trees, it turns out that the performance of separating cycle inequalities improves for ternary and quad trees. The special case of a star, where the degree is maximal, is equivalent to the quadratic linear ordering problem, for which we know that separation of cycle inequalities improves over IP [2].

The real-world general instances have 131 one-to-one tangles and between 371 and 414 nodes. The average degree of an internal node ranged between 2.71 and 3.38. We show the results in Table 3. Here, solving the reformulation usually yields best performance. Note that these non-binary instances could not be solved before by any other exact method.

5 Acknowledgments

We are grateful to Franz Rendl, Giovanni Rinaldi, and Angelika Wiegele for the permission to use their bundle program for the SDP-based approach, as well as to Martin Nöllenburg and Danny Holten for sending us real-world tanglegram instances. We thank Henning Fernau for drawing our attention to tanglegrams.

References

- [1] S. Böcker, F. Hüffner, A. Truss, and M. Wahlström. A faster fixed-parameter approach to drawing binary tanglegrams. In *Proc. of International Workshop on Parameterized and Exact Computation (IWPEC 2009)*, 2009. To appear.
- [2] C. Buchheim, A. Wiegele, and L. Zheng. Exact algorithms for the quadratic linear ordering problem. *INFORMS J. on Computing*. To appear.
- [3] C. De Simone. The cut polytope and the boolean quadric polytope. *Discrete Mathematics*, 79:71–75, 1989.
- [4] P. Eades and N. C. Wormald. Edge crossings in drawing bipartite graphs. *Algorithmica*, 11:379–403, 1994.
- [5] H. Fernau, M. Kaufmann, and M. Poths. Comparing trees via crossing minimization. *J. of Computer and System Sciences*, In Press, 2009.
- [6] M. S. Hafner, P. D. Sudman, F. X. Villablanca, T. A. Spradling, J. W. Demastes, and S. A. Nadler. Disparate rates of molecular evolution in cospeciating hosts and parasites. *Science*, 265:1087–1090, 1994.
- [7] D. Holten. personal communication, 2009.
- [8] ILOG, Inc. ILOG CPLEX 11.2, 2007. www.ilog.com/products/cplex.
- [9] M. Jünger and P. Mutzel. 2-layer straightline crossing minimization: performance of exact and heuristic algorithms. *J. Graph Algorithms Appl.*, 1:1–25, 1997.
- [10] M. Nöllenburg, M. Völker, A. Wolff, and D. Holten. Drawing binary tanglegrams: An experimental evaluation. In *Proc. of the Workshop on Algorithm Engineering and Experiments, ALENEX 2009*, pages 106–119. SIAM, 2009.
- [11] R. D. M. Page. *Tangled Trees: Phylogeny, Cospeciation, and Coevolution*. University of Chicago Press, 2002.
- [12] F. Rendl, G. Rinaldi, and A. Wiegele. A branch and bound algorithm for max-cut based on combining semidefinite and polyhedral relaxations. In M. Fischetti and D. P. Williamson, editors, *IPCO 2007*, volume 4513 of *Lecture Notes in Computer Science*, pages 295–309. Springer, 2007.
- [13] B. Venkatchalam, J. Apple, K. St. John, and D. Gusfield. Untangling tanglegrams: Comparing trees by their drawings. In I. Mandoiu, G. Narasimhan, and Y. Zhang, editors, *Bioinformatics Research and Applications, 5th International Symposium, ISBRA 2009, Fort Lauderdale, FL, USA, May 13-16, 2009, Proceedings*, volume 5542 of *Lecture Notes in Computer Science*, pages 88–99. Springer, 2009.