

Two-Level Rectilinear Steiner Trees

Stephan Held*

Nicolas Kämmerling*

Abstract

Given a set P of terminals in the plane and a partition of P into k subsets P_1, \dots, P_k , a two-level rectilinear Steiner tree consists of a rectilinear Steiner tree T_i connecting the terminals in each set P_i ($i = 1, \dots, k$) and a top-level tree T_{top} connecting the trees T_1, \dots, T_k . The goal is to minimize the total length of all trees. This problem arises naturally in the design of low-power physical implementations of parity functions on a computer chip.

For bounded k we present a polynomial time approximation scheme (PTAS) that is based on Arora's PTAS for rectilinear Steiner trees after lifting each partition into an extra dimension.

For the general case we propose an algorithm that predetermines a connection point for each T_i and T_{top} ($i = 1, \dots, k$). Then, we apply any approximation algorithm for minimum rectilinear Steiner trees in the plane to compute each T_i and T_{top} independently.

This gives us a 2.37-factor approximation with a running time of $\mathcal{O}(|P| \log |P|)$ suitable for fast practical computations. The approximation factor reduces to 1.63 by applying Arora's approximation scheme in the plane.

1 Introduction

We consider the *two-level rectilinear Steiner tree problem* (R2STP) that arises from an application in VLSI design. Consider the computation of a parity function of k input bits using 2-input XOR-gates. Due to the symmetry, associativity, and commutativity of the XOR function, this can be realized by an arbitrary binary tree with k leaves, rooted at the output simply by inserting an XOR-gate at every internal vertex [11, 12]. Throughout this paper we consider the parity function as a placeholder for any fan-in function of the type $x_1 \circ x_2 \circ \dots \circ x_k$, where \circ is a symmetric, associative, and commutative 2-input operator, i. e. $\circ \in \{\oplus, \vee, \wedge\}$.

On a chip such a tree has to be embedded into the plane and all connections must be realized by rectilinear segments. If each input and the output are single points on the chip, a realization of minimum length and thus power consumption is given by a




Figure 1: On the left, we have two inputs p_1 and p_2 and a single output p_3 . The XOR-gate should be placed at the median of the three terminals. If the inputs have the side outputs p'_1 and p'_2 , the XOR-gate should be placed at p_3 , saving the horizontal length.

minimum length rectilinear Steiner tree. This is a tree connecting the inputs and the output by horizontal and vertical line segments using additional so-called Steiner vertices to achieve a shorter length than a minimum spanning tree. At each Steiner vertex of degree three an XOR-gate is placed. Higher degree vertices can be dissolved into degree three vertices sharing their position. Figure 1 shows an example of an embedded parity function on the left.

In practice input signals may be needed for other computations on the chip and thus delivered to other side outputs. Similarly, the result may have to be delivered to multiple output terminals. Thus, each input and its successors and the output terminals must be connected by separate Steiner trees as well. These trees are then connected by a top-level Steiner tree into which the XOR-gates will be inserted. Considering the additional terminals allows to construct a potentially shorter Steiner tree as shown in Figure 1 on the right. Algorithms ignoring the side outputs cannot guarantee an approximation factor better than two, as we will see in Section 2.

This motivates the definition of the *minimum two-level rectilinear Steiner tree problem*, where we are given a set $P \subset \mathbb{R}^2$ of n terminals and a partition of P into k subsets P_1, \dots, P_k .

A *two-level rectilinear Steiner tree* $T = (T_{top}, T_1, \dots, T_k)$ consists of a Steiner tree T_i for each $i \in \{1, \dots, k\}$ connecting the terminals in P_i and a (group) Steiner tree T_{top} connecting the embedded trees $\{T_1, \dots, T_k\}$. We call T_{top} the *top-level tree*. Note that all trees are allowed to cross. The objective is to minimize the total length of all trees

$$l(T) := \sum_{i=1}^k l(T_i) + l(T_{top}),$$

where $l(T') := \sum_{\{x,y\} \in E(T')} \|x - y\|_1$ is the ℓ_1 -length

*Research Institute for Discrete Mathematics, University of Bonn,
held@or.uni-bonn.de, kaemmerling@or.uni-bonn.de.

of a Steiner tree T' .

For each $i \in \{1, \dots, k\}$ the top-level tree and T_i intersect in at least one point. We can select one such point $q_i \in T_{top} \cap T_i$ and call it *connection point* for T_i and T_{top} . Then T_{top} is a Steiner tree for the terminals $\{q_1, \dots, q_k\}$ and each T_i is a Steiner tree for $P_i \cup \{q_i\}$.

Obviously, this problem is NP-hard as it contains the minimum rectilinear Steiner tree problem in two ways: if $k = 1$ or if $|P_i| = 1$ for $i \in \{1, \dots, k\}$.

Designing the top-level tree as a stand-alone problem is hard. If all subtrees T_i ($i \in \{1, \dots, k\}$) are fixed, T_{top} cannot be approximated to arbitrary quality, as the group Steiner tree problem for connected groups in the Euclidean plane cannot be approximated within a factor of $(2 - \epsilon)$ [9]. However we are in a more lucky situation as we can tradeoff the lengths of bottom-level and top-level trees.

To the best of our knowledge the two-level rectilinear Steiner tree problem has not been considered before despite its practical importance [11, 12]. It is loosely related to the hierarchical network design problems [1, 5, 6] or multi-level facility location problems [3, 4]. However, those problems are structurally different, typically considering problems in graphs, and do not apply to our case.

In [11], ordinary rectilinear Steiner trees were used to build power efficient fan-in trees, when each input and the output consists of a single terminal. In practice designers are also interested in the depth of the constructed circuit [12]. However, for finding good power versus depth tradeoffs a better understanding of short solutions is an essential prerequisite and the aim of our work.

1.1 Our Contribution

In Section 2 we show that the naïve approach of picking a random terminal from each partition as a connection point to the top-level tree and building the bottom-level trees and top-level as separate instances gives as a 2α -factor approximation, where α is the approximation factor of the used minimum Steiner tree algorithm.

Then in Section 3 we show how to lift our instance into an equivalent $(2+k)$ -dimensional rectilinear Steiner tree instance. If the number k of partitions is bounded by a constant, we obtain a PTAS by applying Arora's PTAS for rectilinear Steiner trees [2].

As our main result we improve the approximation guarantee for unbounded k from $(2 + \epsilon)$ to 1.63 in Section 4. Using spanning tree heuristics this approach turns also into a fast practical algorithm with running time $\mathcal{O}(n \log n)$ and approximation factor 2.37.

2 Simple Bottom-Up Construction

A simple bottom-up approach, which works for any metric space, is to compute a Steiner tree T_i for P_i



Figure 2: A tight example when choosing connection points as arbitrary points of P_i .

($i = 1, \dots, k$). In each T_i we fix a connection point $q_i \in P_i$ arbitrarily, compute a Steiner tree T_{top} for $\{q_1, \dots, q_k\}$, and return $T = (T_{top}, T_1, \dots, T_k)$.

Theorem 1 *The simple bottom-up approach is a 2α -factor approximation algorithm for the minimum two-level Steiner tree problem, if we use an α -factor approximation algorithm for the minimum Steiner tree problem as a subroutine.*

Proof. Let T be the two level Steiner tree computed by the simple bottom-up approach and let $T^* = (T_{top}^*, T_1^*, \dots, T_k^*)$ be a minimum two-level Steiner tree. Let be $q_i^* \in T_{top}^* \cap T_i^*$ the connection point of the optimum two-level Steiner tree. Since T_i^* is a Steiner tree on $\{q_i^*\} \cup P_i$, we have $dist(q_i^*, q_i) \leq l(T_i^*)$. Thus,

$$\begin{aligned} l(T) &\leq \alpha \cdot l(T_{top}) + \alpha \sum_{i=1}^k dist(q_i^*, q_i) + \sum_{i=1}^k \alpha \cdot l(T_i^*) \\ &\leq \alpha \cdot l(T_{top}^*) + 2\alpha \sum_{i=1}^k l(T_i^*) \leq 2\alpha \cdot l(T^*). \end{aligned}$$

□

Figure 2 shows that the factor $(2 + \epsilon)$ is sharp. For the instance $P_1 = \{(0, 0), (1, 0)\}$, $P_2 = \{(0, 0), (-1, 0)\}$ a minimum two-level rectilinear Steiner tree of length 2 is shown on the left with $l(T_{top}) = 0$. On the right, a bad choice of connection points and minimum Steiner trees T_{top}, T_1 , and T_2 yield a total length of 4.

3 PTAS for a bounded number of partitions

We can reduce the two-level rectilinear Steiner tree problem in the plane to an ordinary rectilinear Steiner tree problem in a higher dimensional space, where we can apply Arora's PTAS [2].

The idea of the PTAS is to lift every subset P_1, \dots, P_k to an additional dimension. We assume $k > 1$. Otherwise the two-level Steiner tree problem is an ordinary Steiner tree problem. Let $P_1, \dots, P_k \subset \mathbb{R}^2$ be the subsets of a two-level Steiner tree instance, we define a Steiner tree instance in \mathbb{R}^{2+k} . The set of terminals P' is comprised as follows.

For each original terminal $x \in P_i \subset \mathbb{R}^2$ ($i \in \{1, \dots, k\}$), we add a terminal $x' := (x, K \cdot e_i) \in \mathbb{R}^{2+k}$, where $e_i \in \mathbb{R}^k$ is the unit vector with value one at the i -th coordinate and K is a large constant, e. g. we could choose K as $l(B(P))$. Now for $x \in P_h$ and $y \in P_i$ the distance of their high dimensional copies $x', y' \in P'$ is $\|x' - y'\|_1 = \|x - y\|_1 + 2K \|e_h - e_i\|_1 = \|x - y\|_1 + 2K \delta_{h,i}$,




Figure 3: A flat Steiner Tree in a lifted instance.

where $\delta_{h,i}$ is one if $h = i$ and zero otherwise. An example of a lifted two-level Steiner tree is given in Figure 3.

A $(k + 2)$ -dimensional Steiner tree is called *flat* if all Steiner points have either the form $(x, 0) \in \mathbb{R}^{2+k}$ or $(x, K \cdot e_i) \in \mathbb{R}^{2+k}$, where $x \in \mathbb{R}^2$ and $e_i \in \mathbb{R}^k$ is a unit vector. The following Lemma has essentially been proven by Snyder [10], who shows that an optimum Steiner tree can be found in the d -dimensional Hanan grid [7].

Lemma 2 *A $(k + 2)$ -dimensional Steiner Tree T for P' of length $l(T)$ can be transformed in strongly polynomial time into a $(k + 2)$ -dimensional flat Steiner tree T' of length at most $l(T)$.*

The following lemma shows the equivalence between the original two-level rectilinear Steiner tree problem in the plane and the lifted regular rectilinear Steiner tree problem.

Lemma 3 *If $k > 1$, a two-level rectilinear Steiner tree T for P_1, \dots, P_k of length $l(T)$ can be transformed into a $(k + 2)$ -dimensional Steiner tree T' for P' of length at most $l(T) + kK$ and vice versa.*

Proof. We only sketch the proof, details can be found in [8]. We get T' by connecting the lifted components of T at its connection points by an edge of length K . Conversely, we flat T by applying Lemma 2 and replacing edges, s. t. for each $j \in \{3, \dots, k + 2\}$, T' contains at most one edge in direction j , without increasing the length of the tree. Projecting the edges onto the first two coordinates we obtain a feasible two-level Steiner tree of length at most $l(T') - kK$. \square

Theorem 4 *For bounded k there is a PTAS for the two-level rectilinear Steiner tree problem.*

Proof. Choose $K = l(B(P))$ and for $\epsilon > 0$ set $\epsilon' := \frac{1}{k+1}\epsilon$. Then compute an $(1 + \epsilon')$ -approximate $(k + 2)$ -dimensional Steiner tree T' for the lifted terminal set P' with Aroras PTAS [2] that has a polynomial running in bounded dimension. Then we apply Lemma 3 to obtain a two-level Steiner tree $T = (T_{top}, T_1, \dots, T_k)$ for P_1, \dots, P_k with length at most $l(T') - kK$. Let T'^* and T^* be optimum Steiner

trees for P' and P . Since $l(T^*) \geq K$ the length of T is

$$\begin{aligned} l(T) &\leq l(T') - kK \leq (1 + \epsilon')l(T'^*) - kK \\ &\leq (1 + \epsilon')l(T^*) + (1 + \epsilon')kK - kK \\ &\leq (1 + (k + 1)\epsilon')l(T^*) = (1 + \epsilon)l(T^*). \end{aligned}$$

□

4 Predetermined Connection Points

In all algorithms of this section we predetermine a connection point q_i for each set P_i ($i = 1, \dots, k$) and then call a Steiner tree approximation algorithm for $\{q_1, \dots, q_k\}$ to get T_{top} and $P_i \cup \{q_i\}$ to get T_i ($i = 1 \dots, k$). We use the fact that we consider rectilinear instances to obtain better approximation factors than in Section 2.

4.1 Bounding Box Center

A natural approach is to choose each connection point as the center of the bounding box $B(P_i)$.

Theorem 5 *Using bounding box centers as connection points, we get a 1.75α -factor approximation algorithm for the two-level rectilinear Steiner tree problem, when using an α -factor approximation algorithm for rectilinear Steiner trees as a subroutine.*

A detailed analysis can be found in [8], which also provides a simple example attaining this factor.

4.2 Adjusted Bounding Box Center

We can improve the approximation factor by a more careful choice of the connection point. For a set P_i ($i \in \{1, \dots, k\}$), we call the coordinate system with origin in the central point of its bounding box the *coordinate system of P_i* .

If a set P_i of terminals contains an element in each quadrant of its bounding box $B(P_i)$, we call the bounding box $B(P_i)$ *complete*. For subtrees with a complete bounding box we choose the connection point to the top-level tree as the central point of the bounding box as in Section 4.1 and compute a Steiner tree T'_i for $P_i \cup \{q_i\}$ as follows: We compute a Steiner tree T'_i for P_i . Thereby, we embed maximal paths in T'_i containing only Steiner vertices with degree two so that each such path has minimum distance to q_i while preserving its length. We then add an edge from q_i to a_i , where a_i is a point in T'_i minimizing the distance to q_i .

Let $T^* = \{T_{top}^*, T_1^*, \dots, T_k^*\}$ be an minimum two-level Steiner tree. Again we choose T_{top}^* under all minimum two-level Steiner trees as large as possible so that there is a connection point $q_i^* \in T_{top}^* \cap T_i^* \subseteq B(P_i)$.

We present the main ingredient for the analysis of the approximation ratio.




Figure 4: An example of the situation in the proof of Lemma 6. The green diamond is the l_1 -circle with radius h around q_i .

Lemma 6 Let $i \in \{1, \dots, k\}$ and T'_i, a_i be constructed as above. If $B(P_i)$ is complete, then $l(B(P_i)) + \|q_i - a_i\|_1 \leq l(T_i^*)$.

Proof. Define $h := \|q_i - a_i\|_1$. Since $B(P_i)$ is complete, T'_i intersects at least three of the four axes of the coordinate system to P_i . We assume w.l.o.g. that T'_i intersects the left, upper and right axis.

By the choice of T'_i and h there exist (see also Figure 4) for all $z \in [0, h]$

$$\begin{aligned} p &\in \{(x, y) \in P_i : x \leq -h, y \leq 0\}, \\ p' &\in \{(x, y) \in P_i : x \geq h, y \leq 0\}, \\ u &\in \{(x, y) \in P_i : x < 0, y \geq h\}, \\ v_z &\in V_z := \{(x, y) \in P_i : x \geq z, y \geq h - z\}. \end{aligned}$$

Let $v_h \in V_h$. If the unique T_i^* -paths from p to u and from p' to v_h intersect, then T_i^* connects the lines $\{(x, y) : x = 0\}$ and $\{(x, y) : x = h\}$ twice, and therefore $l(B(P_i)) + h \leq l(T_i^*)$.

Otherwise, we can choose a minimum $z \geq 0$ such that there is a $v \in V_z$ and the unique T_i^* -paths from p to u and from p' to v are disjoint. The lines $\{(x, y) : y = 0\}$ and $\{(x, y) : y = h - z\}$ are connected twice in T_i^* . Therefore we get $l(B(P_i)) + h - z \leq l(T_i^*)$.

If $z = 0$ we are done. Otherwise, if our statement is false there is an $0 < \epsilon \leq z$ such that $\epsilon = l(B(P_i)) + h - l(T_i^*)$ and a $w \in V_{z - \frac{\epsilon}{2}}$. Since the unique T_i^* -paths from p to u and from p' to w are not disjoint, T_i^* connects the lines $\{(x, y) : x = 0\}$ and $\{(x, y) : x = z - \frac{\epsilon}{2}\}$ twice. Therefore we get the contradiction

$$\begin{aligned} \epsilon &= l(B(P_i)) + h - l(T_i^*) \\ &\leq l(B(P_i)) + h - l(B(P_i)) - h + z - z + \frac{\epsilon}{2} = \frac{\epsilon}{2}. \end{aligned}$$

□

With Lemma 6 we tradeoff between the cost $\|q_i^* - q_i\|_1 \leq l(B(P))$ to connect q_i to T_{top}^* and the cost $\|q_i - a_i\|_1$ to connect q_i to T'_i . From this we could derive an approximation factor of $13/8$ if all partitions are complete. In general this is not the case and for incomplete bounding boxes we shift the connection points (and the coordinate system) towards the actual terminals as in Figure 5. A careful analysis using a similar version of Lemma 6 on the shifted coordinate system (details can be found in [8]) gives us the following result:




Figure 5: An example for an incomplete bounding box. The connection point q_i is shifted diagonally to the upper right until the red box hits a terminal.

Theorem 7 There is a 2.37-factor approximation algorithm with runtime $\mathcal{O}(n \log n)$ and an 1.63-factor approximation algorithm in poly-time for the two-level rectilinear Steiner tree problem.

References

- [1] E. Alvarez-Miranda, I. Ljubic, S. Raghavan, and P. Toth: Recoverable Robust Two-Level Network Design Problem. *INFORMS Journal on Computing*, to appear, 2014.
- [2] S. Arora: Polynomial time approximation schemes for Euclidian Traveling Salesman and other geometric problems. *Journal of the ACM*, 45(5): 753–782, 1998.
- [3] M. Baou and F. Barahona. A polyhedral study of a two level facility location model. *RAIRO - Operations Research*, 48: 153–165, 2014.
- [4] J. Byrka and B. Rybicki. Improved LP-Rounding Approximation Algorithm for k-level Uncapacitated Facility Location. *ICALP '12*, 157–169, 2012.
- [5] A. Balakrishnan, T. L. Magnanti, and P. Mirchandani: A dual-based algorithm for multi-level network design. *Management Science*, 40(5): 567–581, 1994.
- [6] J. R. Current, C. S. ReVelle, and J. L. Cohon: The hierarchical network design problem. *European Journal of Operational Research*, 27(1), 57–66, 1986.
- [7] M. Hanan: On Steiner's problem with rectilinear distance. *SIAM Journal on Applied Mathematics*, 14(2): 255–265, 1966.
- [8] S. Held and N. Kämmerling: Two-Level Rectilinear Steiner Trees. <http://arxiv.org/abs/1501.00933>, 2015.
- [9] S. Safra and O. Schwartz: On the complexity of approximating TSP with neighborhoods and related problems. *Computational Complexity*, 14(4): 281–307, 2006.
- [10] T.L. Snyder: On the exact location of Steiner points in general dimension. *SIAM Journal on Computing*, 21(1): 163–180, 1991.
- [11] H. Xiang, H. Ren, L. Trevillyan, L. Reddy, R. Puri, and M. Cho: Logical and physical restructuring of fan-in trees. *ISPD '10*, 67–74, 2010.
- [12] H. Xiang, L. Reddy, L. Trevillyan, and R. Puri: Depth Controlled Symmetric Function Fanin Tree Restructure. *ICCAD '13*, 585–591, 2013.