

Trimming of Graphs, with Application to Point Labeling*

Thomas Erlebach[†] Torben Hagerup[‡] Klaus Jansen[§]

Moritz Minzlaff[¶] Alexander Wolff^{||}

September 21, 2008

Abstract

For $t > 0$ and $g \geq 0$, a vertex-weighted graph of total weight W is (t, g) -trimmable if it contains a vertex-induced subgraph of total weight at least $(1 - 1/t)W$ and with no simple path of more than g edges. A family of graphs is *trimmable* if for every constant $t > 0$, there is a constant $g \geq 0$ such that every vertex-weighted graph in the family is (t, g) -trimmable. We show that every family of graphs of bounded domino treewidth is trimmable. This implies that every family of graphs of bounded degree is trimmable if the graphs in the family have bounded treewidth or are planar. We also show that every family of directed graphs of bounded layer bandwidth (a less restrictive condition than bounded directed bandwidth) is trimmable. As an application of these results, we derive polynomial-time approximation schemes for various forms of the problem of labeling a subset of given weighted point features with nonoverlapping sliding axes-parallel rectangular labels so as to maximize the total weight of the labeled features, provided that the ratios of label heights or the ratios of label lengths are bounded by a constant. This settles one of the last major open questions in the theory of map labeling.

*Preliminary versions of this work appeared in Proc. 22nd European Workshop on Computational Geometry (EWCG 2006), pp. 137–140, and in Proc. 25th International Symposium on Theoretical Aspects of Computer Science (STACS 2008), pp. 265–276.

[†]Department of Computer Science, University of Leicester, Leicester LE1 7RH, England. t.erlebach@mcs.le.ac.uk

[‡]Institut für Informatik, Universität Augsburg, 86135 Augsburg, Germany. hagerup@informatik.uni-augsburg.de

[§]Institut für Informatik und Praktische Mathematik, Universität Kiel, 24098 Kiel, Germany. kj@informatik.uni-kiel.de

[¶]Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, Germany. minzlaff@math.tu-berlin.de. The author was supported by the Berlin Mathematical School, which is funded by the German Research Foundation (DFG) as a graduate school in the framework of the “Excellence Initiative”.

^{||}Faculteit Wiskunde en Informatica, Technische Universiteit Eindhoven, The Netherlands. <http://www.win.tue.nl/~awolff>

Keywords: Trimming weighted graphs, domino treewidth, planar graphs, layer bandwidth, point-feature label placement, map labeling, sliding labels, polynomial-time approximation schemes

1 Introduction

In this paper we first show that certain families of vertex-weighted graphs have the property that a vertex subset of small weight suffices to hit all long simple paths. This finding allows us to address an application in map labeling, namely the problem of labeling a subset of given weighted point features with axes-parallel rectangular labels. In the following we discuss these two research directions in turn.

1.1 Graph Trimming

We study the following problem: Given a graph in which each vertex has a nonnegative weight, delete vertices of small total weight such that the remaining graph does not contain any long simple paths. Whereas there is an extensive literature on separators, which can be viewed as serving to destroy all large connected components, we are not aware of previous work on vertex sets that destroy all long simple paths. Let us make our notions precise. The *length* of a path π is the number of edges on π .

Definition 1.1. For $t > 0$ and $g \geq 0$, a (t, g) -*trimming* of a vertex-weighted graph $G = (V, E)$ of total weight W is a set $U \subseteq V$ of weight at most W/t such that every simple path in G of length more than g contains a vertex in U . If G has a (t, g) -trimming, we also say that G is (t, g) -*trimmable*. A family of graphs is *trimmable* if, for every constant $t > 0$, there is a constant $g \geq 0$ (that depends only on t) such that every vertex-weighted graph in the family is (t, g) -trimmable.

Definition 1.1 applies to directed and undirected graphs. Of course, trimming undirected graphs is the harder task in the sense that a (t, g) -trimming of the undirected version of a directed graph G , for arbitrary $t > 0$ and $g \geq 0$, is also a (t, g) -trimming of G . In order to demonstrate the trimmability of a family of graphs, it suffices to verify that the condition of Definition 1.1 holds for all integers t larger than an arbitrary constant.

Not every family of graphs is trimmable, even in the *unweighted* case where all vertices are taken to have weight 1. For example, for $n, t \geq 2$, if we delete a $(1/t)$ -fraction of the vertices in an unweighted n -clique K_n , the remaining graph still has a simple path of length $n(1 - 1/t) - 1$. This expression is not bounded by a function of t alone, so the family of complete graphs is not trimmable.

With a little effort, one can show the family of trees to be trimmable. One popular generalization of trees is based on the definition below. Given a graph $G = (V, E)$ and a

set $U \subseteq V$, we denote by $G[U]$ the subgraph of G induced by U . The *union* of graphs $G_i = (V_i, E_i)$, for $i = 1, \dots, m$, is the graph $\bigcup_{i=1}^m G_i = (\bigcup_{i=1}^m V_i, \bigcup_{i=1}^m E_i)$.

Definition 1.2. A *tree decomposition* of an undirected graph $G = (V, E)$ is a pair (T, B) , where $T = (X, E_T)$ is an undirected tree and $B : X \rightarrow 2^V$ maps each node x of T to a subset of V , called the *bag* of x , such that

- (1) $\bigcup_{x \in X} G[B(x)] = G$, and
- (2) for all $x, y, z \in X$, if y is on the path from x to z in T , then $B(x) \cap B(z) \subseteq B(y)$.

The *width* of the tree decomposition (T, B) is $\max_{x \in X} |B(x)| - 1$, and the *treewidth* of G is the smallest width of any tree decomposition of G .

The notions related to treewidth were introduced by Robertson and Seymour [RS86]. We refer to condition (2) of Definition 1.2 as the *connectedness property*. The family of graphs of treewidth at most 1 coincides with the family of forests. By analogy with many other generalizations from the family of trees to families of graphs of bounded treewidth, it seems natural to ask whether every family of graphs of bounded treewidth is trimmable. At present we cannot answer this question; we need a concept stronger than bounded treewidth alone.

Definition 1.3. The *elongation* of a tree decomposition (T, B) is the maximum length of a simple path in T between two nodes with intersecting bags. For every $s \geq 0$, the *s-elongation treewidth* of an undirected graph G is the smallest width of a tree decomposition of G with elongation at most s .

Ding and Oporowski [DO95] use the term “diameter” to denote what we call elongation; our different terminology is motivated by a desire to avoid any possible confusion with the diameter of the tree T . Since every graph has a trivial tree decomposition of elongation 0, the s -elongation treewidth of every graph is well-defined for every $s \geq 0$. The 1-elongation treewidth is the *domino treewidth* studied, e.g., by Bodlaender [Bod99]. While every family of bounded domino treewidth trivially has bounded s -elongation treewidth for every $s \geq 1$, the converse is not true. For example, for $n \geq 2$, the n -vertex star graph has 2-elongation treewidth 1, but domino treewidth $\lfloor n/2 \rfloor$.

Our main result about graph trimming, proved in Section 2, is that for all fixed $s \geq 0$, every family of graphs of bounded s -elongation treewidth is trimmable. Ding and Oporowski [DO95] showed that the domino treewidth of a graph can be bounded by a function of its usual treewidth and its maximum degree. It follows that every family of graphs of bounded treewidth and bounded degree is also trimmable. We derive from this that all families of planar graphs of bounded degree are trimmable. We also consider the following variation of directed bandwidth.

Definition 1.4. The *layer bandwidth* of a directed acyclic graph $G = (V, E)$ is the smallest integer D for which there exists an integer-valued mapping f defined on V such that $1 \leq f(v) - f(u) \leq D$ for all $(u, v) \in E$.

If the mapping f is additionally required to be a bijection from V to $\{1, 2, \dots, |V|\}$, this definition yields the known concept of *directed bandwidth* [GGJK78]. Of course, the layer bandwidth of a directed acyclic graph is at most its directed bandwidth. We are not aware of previous studies of layer bandwidth. Our second result about graph trimming, also proved in Section 2, is that every family of directed acyclic graphs of bounded layer bandwidth is trimmable. These results have applications described in the following subsection.

One may phrase the definition of a (t, g) -trimming of an unweighted n -vertex graph $G = (V, E)$ in the language of ε -nets [HW87]. A *range space* is a pair (X, R) , where X is a set and R is a set of subsets of X . In our case, we would take $X = V$ and let R be the set of vertex sets of simple paths in G . In the context of a range space $S = (X, R)$, a subset $A \subseteq X$ is *shattered* by R if every subset of A is of the form $A \cap r$ for some $r \in R$. If some integer d bounds the cardinality of every subset of X shattered by R , the smallest such d is called the (*Vapnik-Chervonenkis*) *dimension* of S ; otherwise the dimension of S is infinite. For $0 \leq \varepsilon \leq 1$, an ε -net of X is a subset of X that contains at least one element of every $r \in R$ with $|r| > \varepsilon|X|$. With the choice of (X, R) indicated above and for $\varepsilon = (g + 1)/n$, an ε -net of X of size at most n/t is precisely a (t, g) -trimming of G . For range spaces of finite dimension d , ε -nets of size roughly d/ε are known to exist. Results of this kind do not appear useful in our context, however, because the relevant range spaces have dimension $\Omega(n)$ for even very simple graphs, e.g., graphs of domino treewidth 2.

1.2 Label Placement

Our main motivation for investigating trimmable graph families arose in the context of labeling maps with sliding labels. Generally speaking, map labeling is the problem of placing a set of labels, each in the vicinity of the object that it labels, while satisfying certain conditions. For an overview, see the map-labeling bibliography of Wolff and Strijk [WS08]. A fundamental requirement in map labeling is that labels are not allowed to overlap. As a consequence, it may not be possible to label all objects in a map, and the goal is to make an optimal selection according to some criterion. We consider the labeling of point features such as towns or mountain tops, each of which is located at a point in the plane called a *site*. The label of such a feature can usually be approximated without much loss by an open axes-parallel rectangular shape that must be placed in the plane without rotation so that its boundary touches the site of the feature. One distinguishes between *fixed-position models* and *slider models*. In fixed-position models, each label has a predetermined finite set of *anchor points* on its boundary (e.g., the four corners), and the label must be placed so that one of its anchor points coincides with the site of the feature to be labeled. In slider models, the anchor points form *anchor segments* on the boundary of the label (e.g., its bottom edge).

Van Kreveld et al. [vKSW99] introduced a taxonomy of fixed-position and slider models, which was later extended by Poon et al. [PSS⁺03]. We use the slider models 1SH, 2SH, 1SV, 2SV and 4S of Poon et al., which define the anchor segments of a label to be its bottom edge, its top and bottom edges, its left edge, its left and right edges, and its entire boundary,

respectively. An illustration of the three slider models 1SH, 2SH and 4S is given in the top row of Fig. 1, adapted from [PSS⁺03]. We assume that each feature comes equipped with a nonnegative weight, which may be used to express priorities among the features. If features represent villages, towns and cities on a map, priorities may correspond to the number of inhabitants, for example. Our objective is to label features with nonoverlapping labels so as to maximize the sum of the weights of those features that actually receive a label. This objective function favors the labeling of features of large weight (e.g., large cities) over those of smaller weight. We refer to the specific map-labeling problems described in this paragraph as *weighted 1SH-labeling*, etc.

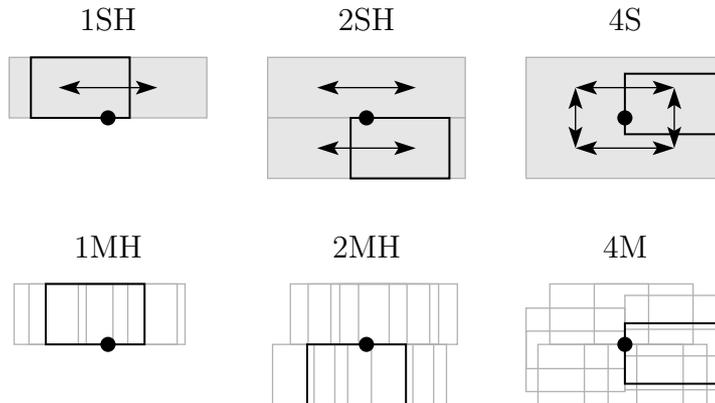


Fig. 1: Slider models (top row) and fixed-position models (bottom row). Possible positions of the label boundary are indicated in gray.

We define the *height ratio* of an instance of a map-labeling problem as the ratio of the maximum height of a label to the minimum such height. If the height ratio is bounded by a constant in a class of instances, the class is of *bounded height ratio*. If all labels are of the same height, we use the term *unit-height*. Instances of bounded height ratio and, in particular, unit-height instances are of great practical importance because they model the common case in which each label contains a single or a few lines of text of a common character height. We apply the qualifier “unit-height” to map-labeling problems to indicate that the input is restricted to be a unit-height instance. The *length ratio* of an instance of a map-labeling problem is the ratio of the maximum length of a label to the minimum such length, and similarly to above we can consider classes of instances with *bounded length ratio*. Finally, we say that an instance has *height or length ratio* ρ if its height ratio or its length ratio is ρ , and a class of instances has *bounded height or length ratio* if there is a fixed $\rho \geq 1$ such that every instance in the class has height or length ratio at most ρ .

For $c \leq 1$, a *c-approximation algorithm* for a maximization problem is an algorithm that always outputs a solution whose value under the objective function is at least c times the optimal value. An algorithm that takes an additional parameter $\varepsilon > 0$ and, for each fixed ε , is a polynomial-time $(1 - \varepsilon)$ -approximation algorithm is called a *polynomial-time approximation scheme (PTAS)*. If the running time depends polynomially on ε as well, the algorithm is a *fully polynomial-time approximation scheme (FPTAS)*.

Poon et al. [PSS⁺03] show weighted unit-height 1SH-labeling to be NP-hard, even if all sites lie on a horizontal line and the weight of each feature equals the length of its label. For the one-dimensional case, in which all n sites lie on a horizontal line, they give an FPTAS, which yields an $O(n^2/\varepsilon)$ -time $(1/2 - \varepsilon)$ -approximation algorithm for the two-dimensional unit-height case for arbitrary $\varepsilon > 0$. Poon et al. also describe a PTAS for unit-square labels. They raise the question of whether a PTAS exists for rectangular labels of arbitrary lengths and unit height. This is known to be the case for fixed-position models [AvKS98] and for sliding labels of unit weight [vKSW99]. The corresponding $(1 - \varepsilon)$ -approximation algorithms run in $n^{O(1/\varepsilon)}$ and in $n^{O(1/\varepsilon^2)}$ time, respectively, for arbitrary $\varepsilon > 0$. The question of whether the combination of both sliding labels and arbitrary weights allows a PTAS in the unit-height case has been one of the last major open problems in theoretical map labeling.

In Section 3 we settle the open question of Poon et al. and, in fact, a slightly more general question by presenting, for every fixed $\rho \geq 1$, a PTAS for the weighted 1SH-labeling problem for instances of height ratio at most ρ . There are no restrictions on label lengths and weights. Our approach is to discretize a given instance I of the weighted 1SH-labeling problem, i.e., to turn it into a fixed-position instance I' , after which we can apply a generalization of a known fixed-position algorithm to I' . The main difficulty lies in finding a suitable set of discrete label positions for each site. “Suitable” means that the weight of an optimal labeling of I' must be close enough to the weight of an optimal labeling of I . Dependencies between labels can be modeled via a graph, and long paths in this graph translate into large sets of anchor points that cannot be left out of consideration. Here our results from Section 2 come into play. We prove that the family of dependency graphs, if suitably defined, is trimmable, and we show how this may be used to bound the number of anchor points by a polynomial. This yields the PTAS. Then we show how to obtain a PTAS for weighted 1SH-labeling also on classes of instances with bounded length ratio, and for weighted 2SH-labeling, 1SV-labeling, 2SV-labeling and 4S-labeling on classes of instances with bounded height or length ratio.

In this paper, our objective is to maximize the sum of the weights of those features that receive a label. Let us call this objective *label-weight maximization*. In the literature, a different objective has also been considered. In *label-size maximization* one insists that *all* features receive a label, and the objective is to maximize a factor by which each label is scaled before it is attached to its feature. Label-size maximization has also been combined with *multi-label map labeling*, where each feature may receive several labels. Approximation algorithms have been given for labeling points with maximum-size congruent squares or disks, two per site [QWXZ00, Jia06]. Labeling points with maximum-size squares, three per site, can be solved exactly in polynomial time [DQVZ03]. In this paper we combine, for the first time, multi-label map labeling with label-weight maximization. Our labeling models and approximation schemes are flexible enough to allow the user to specify several features with sites at the same position, each with its own label and weight.

We use \mathbb{Z} and \mathbb{N} to denote the set of integers and the set of positive integers, respectively. By \mathbb{R} , $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ we denote the sets of real numbers, of positive real numbers and of nonnegative real numbers, respectively, and \mathbb{R}^2 is the two-dimensional Euclidean plane.

2 Trimming of Graphs

In this section we show that two generalizations of trees are trimmable. First, we prove that for every constant $s \geq 0$, every family of graphs of bounded s -elongation treewidth is trimmable. This implies that every family of graphs of bounded degree is trimmable if the graphs in the family have bounded treewidth or are planar. Subsequently we show that for every constant $D \geq 1$, the family of directed acyclic graphs of layer bandwidth at most D is trimmable.

Theorem 2.1. *Let $k, s \geq 0$ and suppose that a vertex-weighted undirected graph G has a tree decomposition of width k and elongation s . Take $a = k + 1$ if $s \geq 2$ and $a = \lceil k/2 \rceil$ if $s \leq 1$. Then, for every integer $t \geq 2$, G has a (t, g) -trimming, where*

$$g = \begin{cases} (2(s+1)t - 3)(k+1) - 1 & \text{if } a \leq 1; \\ (a^{(s+1)t-2}(a+1) - 2)(k+1)/(a-1) - 1 & \text{otherwise.} \end{cases}$$

Therefore, for every constant s , every family of graphs of bounded s -elongation treewidth is trimmable.

Proof. Let (T, B) be a tree decomposition of G of width k and elongation s , root T at an arbitrary node and let U be the set of vertices in bags of nodes whose depth d in T satisfies $d \bmod (s+1)t = i$, with the integer i chosen to minimize the weight of U . We show that U is a (t, g) -trimming of G .

Let $G = (V, E)$ and denote the total weight of the vertices in V by W . Since each vertex in V occurs in bags of nodes on at most $s+1$ levels in T , the sum, over all levels, of the weight of the vertices occurring in bags of nodes on the level under consideration is at most $(s+1)W$. Therefore, by the choice of i , the weight of U is at most $(s+1)W/((s+1)t) = W/t$, as desired.

Let $\pi = (v_0, \dots, v_m)$ be a simple path in G of length $m \geq 1$ and, for $i = 1, \dots, m$, choose a node x_i in T whose bag contains both v_{i-1} and v_i . For $i = 1, \dots, m-1$, we call the unique path in T from x_i to x_{i+1} the *stroke* of v_i . By the connectedness property of T , every bag of a node on the stroke of a vertex v contains v . Concatenating the strokes of v_1, \dots, v_{m-1} in this order, we obtain a walk π' in T (that, informally, can be viewed as induced by π). The walk π' may visit a node x in T several times. Every edge on π' that has x as an endpoint, however, must lie on the stroke of a vertex in $B(x)$, and two such edges can lie on the stroke of the same vertex only if they are consecutive on π' . It follows that x occurs at most $|B(x)| \leq k+1$ times on π' . If $s \leq 1$, we can strengthen this statement as follows: Every stroke is of length at most 1, so every visit to x by π' “uses” either the strokes of at least two vertices in $B(x)$, rather than one, or—at the ends of π' —a stroke and a vertex in $B(x)$ that has no stroke. It follows that if $s \leq 1$, the number of occurrences of x on π' is bounded by $\lfloor (k+1)/2 \rfloor = \lceil k/2 \rceil$. Since T is a tree, if π' leaves x over an edge e , its next return to x , if any, must also happen over e . Therefore the nodes on π' span a subtree T' of T in which no node has more than $a+1$ neighbors, where a is defined in the statement of the theorem. In other words, no node in T' has more than a children, except

that the root may have $a + 1$ children. The number of nodes at depth d in such a tree is bounded by $(a + 1)a^{d-1}$, for all $d \geq 0$, and therefore the number of nodes at depth at most d is bounded by $2d + 1$ if $a = 1$ and by $1 + (a + 1)(a^d - 1)/(a - 1) = ((a + 1)a^d - 2)/(a - 1)$ if $a \geq 2$.

Suppose that π contains no vertex in U . Then, by the choice of U , the depth of T' is at most $(s + 1)t - 2$, and the number of nodes in T' is at most $2(s + 1)t - 3$ if $a = 1$ and at most $(a^{(s+1)t-2}(a + 1) - 2)/(a - 1)$ if $a \geq 2$. Since each bag contains at most $k + 1$ vertices, it follows that $m + 1 \leq (2(s + 1)t - 3)(k + 1)$ if $a = 1$ and that $m + 1 \leq (a^{(s+1)t-2}(a + 1) - 2)(k + 1)/(a - 1)$ if $a \geq 2$. \square

Corollary 2.2. *For all integers $k \geq 0$, $d \geq 1$ and $t \geq 2$, every vertex-weighted undirected graph of treewidth k with maximum degree d has a $(t, \lceil K/2 \rceil^{2t})$ -trimming, where $K = (9k + 7)d(d + 1) - 1$. Therefore every family of graphs with bounded degree and bounded treewidth is trimmable.*

Proof. According to [Bod99, Theorem 3.1], every such graph has a domino tree decomposition of width at most K . Except in the trivial case $k = 0$, we have $K \geq 31$. By Theorem 2.1, used with $s = 1$, the graph has a (t, g) -trimming, where

$$g = \frac{(\lceil K/2 \rceil^{2t-2}(\lceil K/2 \rceil + 1) - 2)(K + 1)}{\lceil K/2 \rceil - 1} - 1 \leq \lceil K/2 \rceil^{2t}.$$

\square

We can extend this result to planar graphs of bounded degree.

Corollary 2.3. *For all integers $d, t \geq 1$, every vertex-weighted undirected planar graph of maximum degree d has a $(t, \lceil K/2 \rceil^{4t})$ -trimming, where $K = (54t - 29)d(d + 1) - 1$. Therefore every family of planar graphs of bounded degree is trimmable.*

Proof. Let $G = (V, E)$ be a planar graph with maximum degree d and denote the total weight of the vertices in V by W . We first follow the approach of Baker [Bak94] to obtain a $(2t - 1)$ -outerplanar subgraph of G by deleting vertices of total weight at most $W/(2t)$. Process an arbitrary planar embedding of G by repeatedly deleting the vertices on the boundary of the outer face until no vertex remains. The vertices deleted in one iteration of this process form a *layer*. Number the layers R_1, R_2, \dots in the order of their deletion. For $j = 0, \dots, 2t - 1$, let V_j be the set of vertices in layers R_i with $i \bmod (2t) = j$, choose j such that the total weight of V_j is at most $W/(2t)$ and consider the subgraph H_j of G induced by $V \setminus V_j$.

H_j is $(2t - 1)$ -outerplanar and thus has treewidth at most $6t - 4$ [Bod98, Theorem 83]. By Corollary 2.2, H_j has a $(2t, \lceil K/2 \rceil^{4t})$ -trimming U . The set $V_j \cup U$ has weight at most $W/(2t) + W/(2t) = W/t$ and therefore is a $(t, \lceil K/2 \rceil^{4t})$ -trimming of G . \square

Finally, we consider directed graphs of bounded layer bandwidth.

Lemma 2.4. *Let $G = (V, E)$ be a vertex-weighted directed acyclic graph of layer bandwidth D . Then, for every integer $t \geq 2$, G has a (t, g) -trimming, where $g = D(t - 1) - 1$.*

Proof. Let $f : V \rightarrow \mathbb{Z}$ be a mapping such that $1 \leq f(v) - f(u) \leq D$ for all $(u, v) \in E$. For each $i \in \mathbb{Z}$, we call $f^{-1}(i) = \{v \in V \mid f(v) = i\}$ the i th *layer* of G and define $V(i) = \bigcup_{i \leq j < i+D} f^{-1}(j)$ as the union of the D consecutive layers of G starting with the i th layer. For $k = 0, \dots, t-1$, let $U_k = \bigcup_{i \in \mathbb{Z}} V((k+it)D)$. The set U_k consists of groups of D consecutive layers, with any two consecutive groups separated by a *band* of $(t-1)D$ layers that are not in U_k . By the properties of f , the layer numbers of the vertices on a path in G form a strictly increasing sequence with gaps of at most D . For $k = 0, \dots, t-1$, therefore, a path in $G[V \setminus U_k]$ must lie entirely within one band and be of length at most $D(t-1) - 1 = g$. Moreover, the t sets U_0, \dots, U_{t-1} are disjoint, so one of them of minimum weight is a (t, g) -trimming of G . \square

3 Labeling Weighted Point Features with Sliding Labels

In this section we define the labeling problems of principal relevance to us formally and show that there are, for every fixed $\rho \geq 1$, polynomial-time approximation schemes for weighted 1SH-labeling, 2SH-labeling, 1SV-labeling, 2SV-labeling and 4S-labeling on instances of height or length ratio at most ρ .

3.1 Problem Definitions

Instances of all of the labeling problems corresponding to slider models can be formalized in the uniform way set out in the following definition.

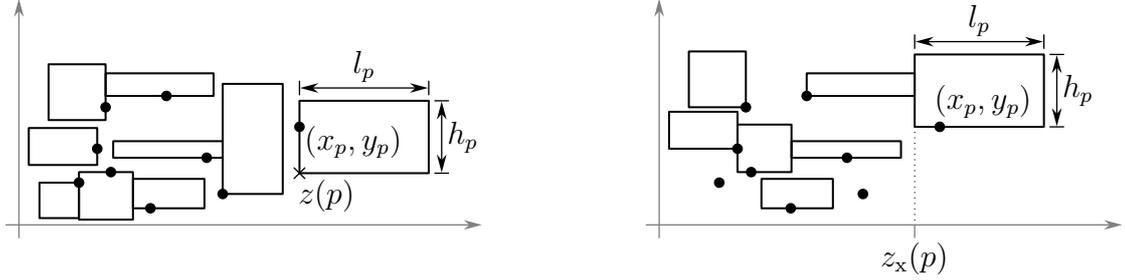
Definition 3.1. A *sliding-label instance* is a tuple $I = (F, x, y, l, h, w)$, where F is a finite set and $x, y : F \rightarrow \mathbb{R}$, $l, h : F \rightarrow \mathbb{R}_{>0}$ and $w : F \rightarrow \mathbb{R}_{\geq 0}$ are functions defined on F . The *size* of I is $|F|$.

For a sliding-label instance $I = (F, x, y, l, h, w)$ and a $p \in F$, we write x_p, y_p, l_p, h_p and w_p for $x(p), y(p), l(p), h(p)$ and $w(p)$, respectively.

In Definition 3.1, the set F represents the set of (point) features to be labeled. For an instance of size n , we can typically take F to be the set $\{1, 2, \dots, n\}$. For all $p \in F$, (x_p, y_p) is the site of the feature p , l_p and h_p are the length and the height of the label of p , respectively, and w_p is the weight of p . The definition allows different features to have identical sites; this can be useful if different features to be labeled are located at the same point in the plane. For each set $Q \subseteq F$, we call $w(Q) = \sum_{p \in Q} w_p$ the *weight* of Q .

We next define the most general problem, weighted 4S-labeling, and then derive the other labeling problems from it.

Definition 3.2. A *4S-labeling* of a sliding-label instance $I = (F, x, y, l, h, w)$ is a pair $L = (Q, z)$, where $Q \subseteq F$ and $z : Q \rightarrow \mathbb{R}^2$ is a function that maps each feature $p \in Q$ to a point $z(p)$ in such a way that, if we let $R(p)$ denote the open axes-parallel rectangle with bottom left corner $z(p)$, width l_p and height h_p , then for all $p, q \in Q$ with $p \neq q$,



(a) A 4S-labeling of a sliding-label instance I .

(b) A 1SH-labeling of I (optimal for unit weights).

Fig. 2: Two labelings of a sliding-label instance.

Type of labeling	Additional constraints
4S	–
2SH	$z_y(p) \in \{y_p - h_p, y_p\}$
1SH	$z_y(p) = y_p$
2SV	$z_x(p) \in \{x_p - l_p, x_p\}$
1SV	$z_x(p) = x_p$

Table 1: Additional constraints on $z(p)$ for all $p \in Q$ that a 4S-labeling (Q, z) must satisfy in order to be a 2SH-labeling, a 1SH-labeling, etc.

the rectangles $R(p)$ and $R(q)$ are disjoint, and for all $p \in Q$, the site (x_p, y_p) lies on the boundary of $R(p)$. The *weight* of L is the weight of Q .

Informally, Q is the set of features that receive a label, and the label of each $p \in Q$ is placed with $z(p)$ at its bottom left corner; see Fig. 2(a). When considering a 4S-labeling (Q, z) , we let z_x and z_y be the functions that map each $p \in Q$ to the x - and y -coordinate of $z(p)$, respectively, so that $z(p) = (z_x(p), z_y(p))$ for each $p \in Q$.

The (*weighted*) 4S-labeling problem is the optimization problem of, given a sliding-label instance I , computing a 4S-labeling of I of largest possible weight. The corresponding definitions for (*weighted*) 2SH-labeling, 1SH-labeling, 2SV-labeling and 1SV-labeling are similar, the only difference being additional constraints on z as listed in Table 1.

When considering a sliding-label instance $I = (F, x, y, l, h, w)$ in the context of the 1SH-labeling problem, we say that two features $p, q \in F$ *y-overlap* if $y_p \leq y_q < y_p + h_p$ or $y_q \leq y_p < y_q + h_q$, i.e., if their labels, when placed with (x_p, y_p) and (x_q, y_q) on their respective bottom edges, have overlapping projections on the y -axis. In a 1SH-labeling $L = (Q, z)$ of I , the second component of z is determined by I and therefore redundant, for which reason we may also specify L through the pair (Q, z_x) and call $z_x(p)$ the *position* of the label of p for each $p \in Q$; see Fig. 2(b). For a given sliding-label instance $I =$

(F, x, y, l, h, w) , a pair (Q, z_x) with $Q \subseteq F$ and $z_x : Q \rightarrow \mathbb{R}$ is a 1SH-labeling of I if and only if $x_p - l_p \leq z_x(p) \leq x_p$ for all $p \in Q$ and for all y -overlapping features $p, q \in Q$ with $p \neq q$, either $z_x(p) + l_p \leq z_x(q)$ or $z_x(q) + l_q \leq z_x(p)$.

Interchanging the roles of the x - and y -dimensions or, equivalently, *mirroring* the Euclidean plane in the line through the origin of slope 1, one can translate 1SH- and 2SH-labeling to 1SV- and 2SV-labeling, respectively, or vice versa. For example, to compute a 2SV-labeling of a sliding-label instance (F, x, y, l, h, w) , compute a 2SH-labeling of (F, y, x, h, l, w) and interchange its x - and y -components.

Our proofs operate not only with slider models, but also with the fixed-position models 1MH, 2MH, 1MV, 2MV and 4M, which allow the set of anchor points of a label to be an arbitrary finite subset of its bottom edge, of its bottom and top edges, of its left edge, of its left and right edges, and of its entire boundary, respectively. Some of these models are illustrated in the bottom row of Fig. 1. Formally, we define a *fixed-position instance* as a pair (I, \mathcal{M}) , where $I = (F, x, y, l, h, w)$ is a sliding-label instance and \mathcal{M} is a function that maps each feature in F to a finite subset of \mathbb{R}^2 . The *size* of (I, \mathcal{M}) is defined as $|F| + \sum_{p \in F} |\mathcal{M}(p)|$. A *4M-labeling* of (I, \mathcal{M}) is a 4S-labeling (Q, z) of I that is *consistent* with \mathcal{M} , i.e., that satisfies $z(p) \in \mathcal{M}(p)$ for all $p \in Q$. The *(weighted) 4M-labeling problem* is the optimization problem of, given a fixed-position instance (I, \mathcal{M}) , computing a 4M-labeling of (I, \mathcal{M}) of largest possible weight. A 2MH-labeling of a fixed-position instance (I, \mathcal{M}) is a 2SH-labeling of I that is consistent with \mathcal{M} , and labelings for the other fixed-position models are defined analogously. The mirroring transformation discussed above applies to fixed-position labeling problems as well in an obvious way.

Similarly to our simplifying convention that omits the y -coordinates of 1SH-labelings, we may also, when dealing with 1MH-labelings, specify a fixed-position instance as a pair (I, \mathcal{M}_x) , where $I = (F, x, y, l, h, w)$ is a sliding-label instance and \mathcal{M}_x maps each $p \in F$ to a finite subset of \mathbb{R} that represents the possible x -coordinates of the left edge of the label of p . In this case, if \mathcal{M}_x maps all $p \in F$ to the same set M , we may write (I, \mathcal{M}_x) as (I, M) .

The principal technical contribution of this section is a reduction of weighted 1SH-labeling to weighted 1MH-labeling. Once this reduction has been established, corresponding reductions from 2SH- to 2MH-labeling, from 1SV- to 1MV-labeling, from 2SV- to 2MV-labeling, and from 4S- to 4M-labeling follow with little additional effort. Under the assumption that the reductions are applied to instances of bounded height or length ratio, they work in polynomial time, and the resulting fixed-position instances can be solved using an adaptation of the PTAS of Agarwal et al. [AvKS98], so that we obtain a PTAS for each of the slider models. We first present our results for 1SH-labeling and then discuss the extensions to the other slider models.

3.2 Normalization, Dependency Graphs, and Trimming

This subsection introduces the notions and preliminary results that form the backbone of our main reduction of 1SH-labeling to 1MH-labeling. It begins with a less formal overview that introduces and motivates the necessary complications one by one.

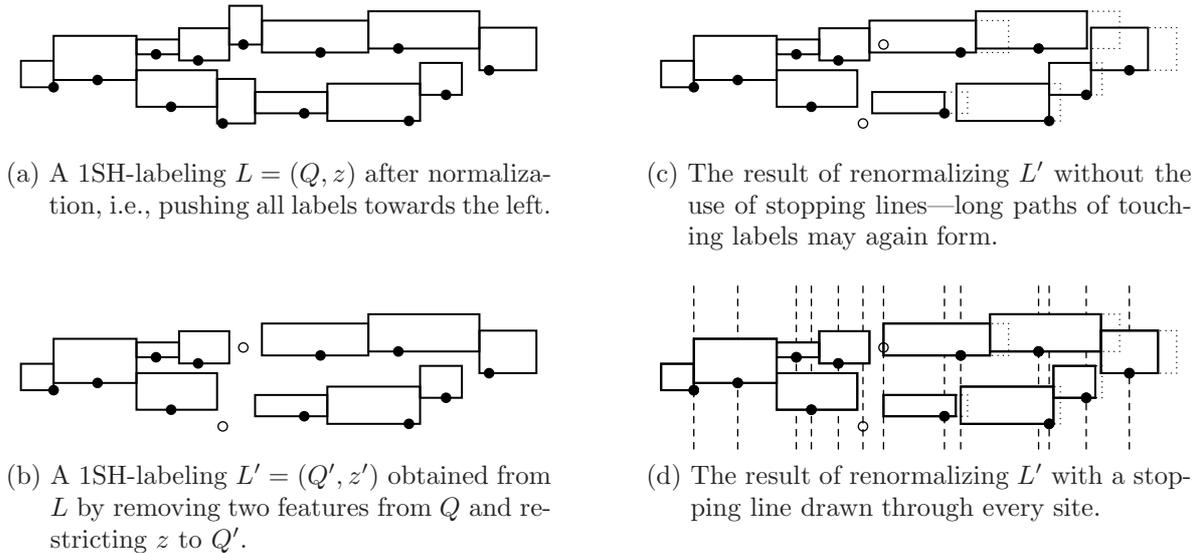


Fig. 3: The process of normalizing (a), trimming (b) and renormalizing without (c) and with (d) stopping lines.

Let a sliding-label instance $I = (F, x, y, l, h, w)$ of size n and a constant $\varepsilon > 0$ be given. Our goal is to describe a polynomial-time computation of a fixed-position instance (I, \mathcal{M}) that is almost as good as I in the sense that the weight of an optimal 1MH-labeling of (I, \mathcal{M}) is at least $1 - \varepsilon$ times that of an optimal 1SH-labeling of I . The fixed-position instance will in fact be of the form (I, M) , where $M \subseteq \mathbb{R}$. It therefore suffices to show that a suitable set M exists and can be computed sufficiently fast.

In a 1SH-labeling (Q, z_x) of I , a priori, $z_x(p)$ could assume any value in the continuum between $x_p - l_p$ and x_p for every $p \in Q$. A *normalization* procedure to be described next shows that nothing is lost by restricting attention to a finite set of candidate values. The normalization is introduced for the sake of argument only and is not actually carried out as part of the reduction.

The normalization can be applied to an arbitrary 1SH-labeling (Q, z_x) of I and results in a *normalized* labeling. The basic idea is to process the labels of the features in Q in the order from left to right, pushing each label as far to the left as it can go without bumping into another label or being separated from its site. Fig. 3(a) shows a possible outcome of this procedure. In a normalized labeling (Q, z'_x) , the position $z'_x(q)$ of the label of a feature $q \in Q$ is either $x_q - l_q$ (no other label blocked the movement of the label of q) or $z'_x(p) + l_p$ for some $p \in Q$ (whose label stopped the movement of that of q and therefore is to the left of it and was processed before it). In the latter case, we introduce the edge (p, q) with length l_p in an auxiliary graph G on the vertex set Q .

For every $q \in Q$, $z'_x(q)$ can be read off any maximal path π in G that ends in q . Denote the length of π , i.e., the sum of the lengths of its edges, by $l(\pi)$. Then, if π starts at p , we simply have $z'_x(q) = x_p - l_p + l(\pi)$. The auxiliary graph G depends on the original labeling

(Q, z_x) to which the normalization was applied. Even so, the expression just found for $z'_x(q)$ depends only on the sequence of the vertices on π , for which there are clearly no more than n^n choices. It follows that for every 1SH-labeling of I , in particular, for one of maximum weight, there is a 1MH-labeling of (I, M) of the same weight for an easily computable set M with $|M| \leq n^n$.

The set M found so far, though finite, is much too large for our intended use, which requires a set of size polynomial in n . If no path in G contains more than $g \geq 0$ edges, the number of such paths is bounded by n^{g+1} , and we obtain a valid set M of the same size. When g is a constant, the size of M is polynomial, as desired. However, paths in G may contain many more than a constant number of edges.

Accepting a small deviation from optimality, as allowed by the constant ε , we may try to bring the notion of graph trimming studied in Section 2 into play. Removing a vertex $p \in Q$ from G corresponds to excluding it from Q and losing its weight w_p in the solution—more intuitively, we will speak of *dropping* the label of p . With $t = \lceil 1/\varepsilon \rceil$, we can afford to remove vertices whose weight is $1/t$ of the total weight from G , and we would like this to destroy all paths in G with more than g edges for some constant g . If G belongs to a trimmable family of graphs, this is always possible. The resulting situation may be as shown in Fig. 3(b).

Apart from the question of whether the auxiliary graph G belongs to a trimmable family, the approach outlined in the previous paragraph meets with the following difficulty: After the trimming of G , i.e., after the dropping of some labels, the labeling defined by the remaining labels must be renormalized. If this is not done, of course, the labels have the positions that they had before the trimming, and the trimming buys us nothing. The renormalization, on the other hand, may create new long paths in the auxiliary graph of the resulting labeling, as shown in Fig. 3(c) for our running example, thus defeating the original purpose of the trimming. Informally, the problem stems from the fact that other labels may close the gap left by a dropped label. In order to counter this, we introduce vertical *stopping lines* and redefine the process of normalization to never push the left edge of a label past a stopping line (see Fig. 3(d)).

The exact choice of stopping lines is largely a technical matter that cannot be well motivated at this point. Each feature $p \in F$ gives rise to exactly three stopping lines, one passing through the site of p and the other two to its left and right at a distance of l_p . Two labels that are (disjoint from and) separated by a stopping line before a normalization can never influence each other in the normalization, so we redefine the auxiliary graph G to not have any such edges.

Even with stopping lines, it can happen that an edge (p, q) that is not present in the original auxiliary graph appears in the auxiliary graph of the labeling obtained by dropping some labels and renormalizing. The creation of new edges is undesirable because it may lead to new long paths. We therefore redefine the auxiliary graph one last time by including all such potential edges from the outset and call the resulting graph the *dependency graph* of (Q, z_x) . It turns out that the dependency graph G is planar and—if I is of bounded height ratio—of bounded degree, which implies that it is trimmable, as needed above. If I is of bounded length ratio, we show that G is trimmable by virtue of having bounded

layer bandwidth.

We now make these ideas precise and begin with a formal definition of dependency graphs.

Definition 3.3. Given a sliding-label instance $I = (F, x, y, l, h, w)$, the *dependency graph* of a 1SH-labeling (Q, z_x) of I is a directed graph on the vertex set Q that, for all $p, q \in Q$, contains the edge (p, q) exactly if $x_p < x_q$, p and q y -overlap, and there is no $\bar{x} \in S_I = \bigcup_{r \in F} \{x_r - l_r, x_r, x_r + l_r\}$ with $z_x(p) + l_p \leq \bar{x} \leq z_x(q)$.

The set S_I corresponds to the set of (vertical) stopping lines through the points $(\bar{x}, 0)$ for $\bar{x} \in S_I$. With this correspondence in mind, we may also refer to S_I as the set of stopping lines of I .

If the label of a feature $q \in Q$, moving left, may hit that of another feature $p \in Q$ without crossing a stopping line, then certainly $x_p < x_q$ holds, p and q y -overlap, and there is no stopping line whose x -coordinate lies between $z_x(p) + l_p$ and $z_x(q)$, inclusive. Conversely, if an edge (p, q) is present in the dependency graph of (Q, z_x) according to Definition 3.3, the label of q will indeed hit the label of p if all labels that y -overlap q and are located between the labels of p and q are dropped.

Lemma 3.4. *Let (Q, z_x) be a 1SH-labeling of a sliding-label instance of height ratio ρ . The dependency graph of (Q, z_x) is planar, and its in-degrees and out-degrees are bounded by $\lceil \rho + 1 \rceil$.*

Proof. To demonstrate the planarity of a graph $G = (V, E)$, it clearly suffices to map each vertex $u \in V$ to an open rectangle $R(u)$ in \mathbb{R}^2 and each edge in E to an open line segment in \mathbb{R}^2 in such a way that all of these rectangles and line segments are pairwise disjoint and that each edge $(u, v) \in E$ is mapped to a line segment with an endpoint on the boundary of each of $R(u)$ and $R(v)$. The reason is that arbitrary points on the boundary of an open rectangle R can be connected to an arbitrary point z in R (its center, say) with closed line segments that intersect only in z .

In the case of a dependency graph $G = (Q, E)$, such a mapping is immediate: For each feature $p \in Q$, take $R(p)$ to be the area occupied by the label of p , shrunk slightly horizontally to allow for labels that touch, and map each edge (p, q) to a part of a horizontal line ℓ that intersects both $R(p)$ and $R(q)$, namely the open line segment on ℓ between $R(p)$ and $R(q)$ (see Fig. 4). That this line segment intersects no $R(r)$ with $r \in Q$ follows from the fact that the stopping line through (x_r, y_r) would prevent (p, q) from being an edge of G , a contradiction.

If the vertices in a set $P \subseteq Q$ have a common out-neighbor or a common in-neighbor r in G , some vertical line intersects $R(p)$ for all $p \in P$. Otherwise $R(p)$ and $R(q)$ could be separated by a vertical line for some $p, q \in P$, and the stopping line through the site of one of the features p and q would prevent the other feature from being a neighbor of r in G . If $|P| \geq 3$, it is now easy to see that the height of $R(r)$ exceeds the total height of the $|P| - 2$ “middle” rectangles in $\{R(p) \mid p \in P\}$ (see Fig. 5), so that $|P| - 2 < \rho$ and therefore $|P| \leq \lceil \rho + 1 \rceil$. \square

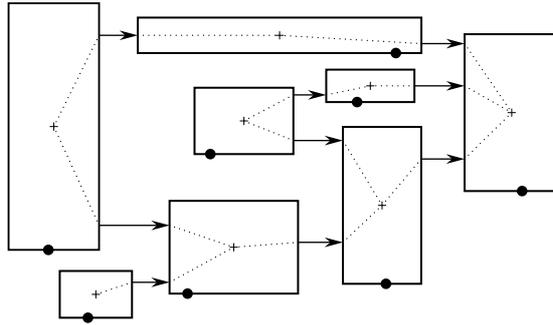


Fig. 4: A plane drawing of the dependency graph of a 1SH-labeling.

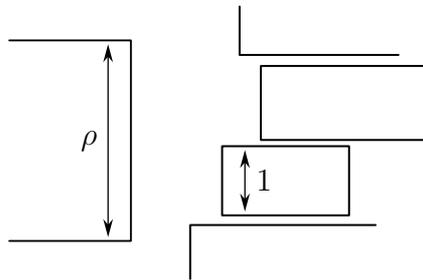


Fig. 5: A site with label height $\rho = 2.9$ can have at most $\lceil \rho + 1 \rceil = 4$ out-neighbors with label heights at least 1.

Corollary 3.5. *Let G be the dependency graph of a 1SH-labeling of a sliding-label instance with height ratio ρ . Then for every integer $t \geq 1$, G is (t, g) -trimmable, where $g = (\rho t)^{O(t)}$.*

Proof. According to Lemma 3.4, G is planar and of maximum degree at most $2(\rho + 2)$. Applying Corollary 2.3 now yields a (t, g) -trimming with $g = (\rho^2 t)^{O(t)} = (\rho t)^{O(t)}$. \square

Next we consider classes of sliding-label instances with bounded length ratio.

Lemma 3.6. *Let G be the dependency graph of a 1SH-labeling of a sliding-label instance with length ratio ρ . Then for every integer $t \geq 2$, G is (t, g) -trimmable, where $g = \lceil 2\rho \rceil (t - 1) - 1$.*

Proof. Let $G = (Q, E)$ be the dependency graph of a 1SH-labeling (Q, z) of a sliding-label instance $I = (F, x, y, l, h, w)$ with length ratio ρ . Without loss of generality, we assume label lengths to lie between 1 and ρ , inclusive. Consider the function $f : Q \rightarrow \mathbb{Z}$ with $f(p) = \lfloor z_x(p) \rfloor$ for all $p \in Q$. Let $(p, q) \in E$. Since p and q y -overlap, we have $z_x(q) - z_x(p) \geq l_p \geq 1$. Moreover, the stopping line through $(x_q - l_q, 0)$ forces $z_x(p) + l_p > x_q - l_q \geq z_x(q) - l_q$, implying $z_x(q) - z_x(p) < l_p + l_q \leq 2\rho$ and $\lfloor z_x(q) \rfloor - \lfloor z_x(p) \rfloor \leq \lceil 2\rho \rceil$. We can conclude that $1 \leq f(q) - f(p) \leq \lceil 2\rho \rceil$ for all $(p, q) \in E$. This shows that G has layer bandwidth at most $\lceil 2\rho \rceil$, and the claim now follows from Lemma 2.4. \square

3.3 Reduction to a Fixed-Position Model

After developing the necessary prerequisites in the previous subsection, in this subsection we complete the description of the reduction from 1SH-labeling to 1MH-labeling.

Recall that the basic intention of the dependency graph G was that the position of each label after normalization (which can be omitted), trimming and renormalization should be given essentially by the length of a path in G . This correspondence was invalidated by the introduction of stopping lines, but can be approximately re-established by adding an additional vertex O and, for every stopping line ℓ , passing through $(\bar{x}, 0)$, say, and for every feature $p \in Q$, an edge from O to p of length \bar{x} . The idea behind this new edge is simply that if the label of p moves to ℓ and stops there, it will be at a position of \bar{x} .

Suppose that, after the removal of the vertices corresponding to dropped labels, each label of a remaining feature p is moved to a position that is the largest length, no larger than the original position of the label, of a path from O to p in what remains of the graph. This procedure is closely related to the normalization discussed in the previous subsection. In actual fact, it may move some labels a shorter distance to the left than the normalization would. Nonetheless, it will be easy to establish the pairwise disjointness of the labels in their resulting positions, and the process immediately suggests a suitable set M of candidate label positions. The fact that the left edge of a label crosses no stopping line in S_I as it moves left—a property that we will need in the proof of Theorem 3.14—can be expressed by saying that the movement leaves invariant the rank in S_I of the position of the label.

If $I = (F, x, y, l, h, w)$ is a sliding-label instance and $I' = (F, x, y', l, h, w)$ differs from I only in the y -coordinates of the sites of an arbitrary (possibly empty) subset of the features, we call I' a y -modification of I , with x -modifications defined analogously. Of course, every sliding-label instance is an x -modification and a y -modification of itself. By considering 1SH-labelings of all y -modifications of a given sliding-label instance in the following lemma, we gain the additional flexibility that allows us to apply the lemma also in the context of 2SH-labeling, as needed in the proof of Lemma 3.11.

Lemma 3.7. *Let a sliding-label instance $I = (F, x, y, l, h, w)$ of size n and a $t \in \mathbb{N}$ be given such that the dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable for some computable $g = g(t) \geq 0$. Then, in $O(n^{g+1})$ time, we can compute a set $M \subseteq \mathbb{R}$ with $|M| \leq 3n^{g+1}$ that does not depend on the y -coordinates of the sites of I and satisfies that, for every y -modification I' of I , the fixed-position instance (I', M) has the following property: For every 1SH-labeling (Q, z_x) of I' , there is a 1MH-labeling (Q', z'_x) of (I', M) with $Q' \subseteq Q$ of weight at least $(1 - 1/t)w(Q)$ such that for all $p \in Q'$, $z'_x(p) \leq z_x(p)$ and $z'_x(p)$ and $z_x(p)$ have the same rank in S_I .*

Proof. Let $I' = (F, x, y', l, h, w)$ be a y -modification of I and let G be the dependency graph of a 1SH-labeling (Q, z_x) of I' . We give each edge (p, q) of G the length l_p . Let U be a (t, g) -trimming of G and take $Q' = Q \setminus U$. Moreover, let \overline{G} be the multigraph obtained from G by adding a new vertex O and, for each $\bar{x} \in S_I$ and each $p \in Q$, an edge from O to p of length \bar{x} .

For all $p \in Q'$, let a p -path be a path in $\overline{G}[\{O\} \cup Q']$ from O to p and define the *length* of a p -path as the sum of the lengths of its edges. For all $p \in Q'$, let $z'_x(p)$ be the largest length of a p -path that does not exceed $z_x(p)$ —this is well-defined since $z_x(p) \geq x_p - l_p$, while there is an edge, and hence a path, in \overline{G} from O to p of length $x_p - l_p$. We will show that (Q', z'_x) is a 1SH-labeling of I' . First, for each $p \in Q'$, the relation $x_p - l_p \leq z'_x(p) \leq z_x(p) \leq x_p$ was essentially argued above. Second, we must show that the labels of the sites in Q' , if placed as indicated by z'_x , do not overlap.

Let p and q be y -overlapping features in Q' and assume, without loss of generality, that $z_x(p) < z_x(q)$ and therefore $z_x(p) + l_p \leq z_x(q)$. If G contains the edge (p, q) , then, since $z'_x(p)$ is the length of a p -path, $z'_x(p) + l_p$ is the length of a q -path and, by definition of z'_x , we have $z'_x(q) \geq z'_x(p) + l_p$. If G does not contain the edge (p, q) , there is an $\bar{x} \in S_I$ with $z_x(p) + l_p \leq \bar{x} \leq z_x(q)$. Again by definition of z'_x and since \overline{G} contains an edge from O to q of length \bar{x} , we then have $z'_x(q) \geq \bar{x} \geq z_x(p) + l_p \geq z'_x(p) + l_p$. In either case, the labels of p and q , placed according to z'_x , do not overlap.

We have $w(Q') \geq (1 - 1/t)w(Q)$, and for each $p \in Q'$, $z'_x(p)$ is the length of a p -path. The length of every p -path belongs to the set M of all sums of an element of S_I and at most g elements of $\{l_p \mid p \in F\}$. M is of size at most $|S_I|n^g \leq 3n^{g+1}$, does not depend on the y -coordinates of the sites of the features in F , and can be computed in $O(n^{g+1})$ time. Let $p \in Q'$. Since for each $\bar{x} \in S_I$ there is a p -path of length \bar{x} , it is easy to see that stepping from $z_x(p)$ to $z'_x(p)$ does not descend strictly below any $\bar{x} \in S_I$, i.e., $z'_x(p)$ has the same rank in S_I as $z_x(p)$. \square

3.4 A Polynomial-Time Approximation Scheme for 4M-Labeling

We need to show how to solve the instances of weighted 1MH-labeling obtained using Lemma 3.7. Agarwal et al. [AvKS98] have given a PTAS that finds a near-maximum independent set in the intersection graph of any given set of closed axes-parallel unit-height rectangles. It is easy to see that their PTAS for maximum independent set at the same time is a PTAS for maximizing the number of features labeled with unit-height *closed* rectangular labels in a fixed-position model. The reason is simply that, by definition, any two label candidates of the same feature must touch the site of the feature. If label candidates are closed, one label candidate automatically excludes the other one from the solution. Unfortunately, this is not the case if labels are open, as we assume throughout; e.g., in the 1MH-model the leftmost and the rightmost label candidate of a site may *not* intersect, so an algorithm for maximum independent set that treats the labels as open rectangles would not automatically yield a feasible solution for the fixed-position labeling problem. Treating the labels as closed rectangles does not work either, because in our models we allow labels of different sites to touch. Fortunately, we can adapt the PTAS of Agarwal et al. to the fixed-position models arising in our setting and even extend it to problems of bounded height or length ratio. In fact, the adapted PTAS can deal with the most general fixed-position problem, that is, 4M-labeling.

Lemma 3.8. *Given a fixed-position instance (I, \mathcal{M}) of size n and with height or length*

ratio ρ and an $\varepsilon \in \mathbb{R}$ with $0 < \varepsilon \leq 1$, a 4M-labeling of (I, \mathcal{M}) of weight at least $1 - \varepsilon$ times the weight of an optimal 4M-labeling of (I, \mathcal{M}) can be computed in $n^{O(\rho/\varepsilon)}$ time. For every fixed $\rho \geq 1$, the weighted 4M-labeling problem for instances of height or length ratio at most ρ therefore admits a PTAS.

Proof. Let (I, \mathcal{M}) with $I = (F, x, y, l, h, w)$ be a fixed-position instance of size n . We can assume without loss of generality that F is a set of integers and, in the light of the mirroring transformation discussed in Section 3.1, that the height ratio of I is at most ρ . Define an *indexed rectangle* to be an open rectangle R that is associated with an integer index $i(R)$. Each placement of the label of a feature $p \in F$ as a rectangle R corresponds to an indexed rectangle R with $i(R) = p$. Two indexed rectangles R_1 and R_2 *intersect* if $R_1 \cap R_2 \neq \emptyset$ or $i(R_1) = i(R_2)$. Computing a 4M-labeling of (I, \mathcal{M}) of weight at least $1 - \varepsilon$ times the optimal weight is equivalent to computing an independent set of weight at least $1 - \varepsilon$ times the optimal weight in the intersection graph of a set \mathcal{R} of at most n weighted indexed rectangles whose height ratios are bounded by ρ and that have the following property: For each integer i , the indexed rectangles with index i share a common point on their boundary. We show how to solve the latter problem. Referring implicitly to the intersection graph, we will say that a subset \mathcal{S} of \mathcal{R} is *independent* if its elements are pairwise nonintersecting, and we denote its weight by $w(\mathcal{S})$.

Assume, without loss of generality, that the height h_R of every rectangle R in \mathcal{R} satisfies $1/\rho \leq h_R \leq 1$ and that no horizontal edge of a rectangle in \mathcal{R} has an integer y -coordinate. For every integer j , we call the horizontal line through $(0, j)$ the *stabbing line* of index j . We apply the shifting technique [Bak94, HM85]. Let $k = \lceil 1/\varepsilon \rceil$ and, for $j = 0, \dots, k-1$, denote by \mathcal{R}_j the set of indexed rectangles in \mathcal{R} that do not intersect any stabbing line whose index modulo k is j . Our algorithm computes a maximum-weight independent subset of each of $\mathcal{R}_0, \dots, \mathcal{R}_{k-1}$ and outputs a set of maximum weight among the k sets obtained.

Let \mathcal{R}^* be a maximum-weight independent subset of \mathcal{R} . Every indexed rectangle in \mathcal{R} is missing from at most one of the sets $\mathcal{R}_0, \dots, \mathcal{R}_{k-1}$, so the sets $\mathcal{R}^* \setminus \mathcal{R}_0, \dots, \mathcal{R}^* \setminus \mathcal{R}_{k-1}$ are disjoint subsets of \mathcal{R}^* . As a consequence, $w(\mathcal{R}^* \setminus \mathcal{R}_b) \leq (1/k)w(\mathcal{R}^*)$ and therefore $w(\mathcal{R}^* \cap \mathcal{R}_b) \geq (1 - 1/k)w(\mathcal{R}^*)$ for some $b \in \{0, \dots, k-1\}$. Since $\mathcal{R}^* \cap \mathcal{R}_b$ is independent and $\varepsilon \leq 1/k$, our algorithm indeed outputs an independent subset of \mathcal{R} of weight at least $(1 - \varepsilon)w(\mathcal{R}^*)$.

It remains to show how to compute a maximum-weight independent subset of \mathcal{R}_b efficiently for a fixed $b \in \{0, \dots, k-1\}$. Since all indexed rectangles intersecting a stabbing line whose index modulo k is b have been removed, \mathcal{R}_b decomposes into instances, each of which is completely contained between two stabbing lines at distance k . Because indexed rectangles from different instances do not intersect, an overall maximum-weight independent set can be obtained as the union of a maximum-weight independent set of each instance.

To compute a maximum-weight independent subset of a nonempty set $\mathcal{S} \subseteq \mathcal{R}$ of indexed rectangles that are all contained in a horizontal slab of height k , we apply a dynamic-programming approach. Intuitively, one can imagine moving a vertical swepline from left to right while considering all possible independent sets of indexed rectangles that intersect

the sweepline in its current position. Formally, we translate the problem into a longest-path problem in an acyclic auxiliary graph H .

Let $x_1 < \dots < x_m$ be the distinct x -coordinates of the left edges of indexed rectangles in \mathcal{S} . For $j = 1, \dots, m$, let \mathcal{S}_j be the set of indexed rectangles in \mathcal{S} whose left edge lies to the left of or on the vertical line ℓ through $(x_j, 0)$ and whose right edge lies strictly to the right of ℓ . Moreover, let \mathcal{I}_j be the family of all independent subsets of \mathcal{S}_j . Since the indexed rectangles in \mathcal{R} have height at least $1/\rho$ and two indexed rectangles in \mathcal{S}_j can be disjoint only if their projections on the y -axis are disjoint, the cardinality of every set in \mathcal{I}_j is bounded by $k\rho$.

The directed auxiliary graph H is defined as follows. For $j = 1, \dots, m$ and for all $A \in \mathcal{I}_j$, H contains a vertex $v_{j,A}$. In addition, H contains a start vertex s and a goal vertex t . The edges of H are the following:

- For each $A \in \mathcal{I}_1$, H contains the edge $(s, v_{1,A})$ with weight $w(A)$.
- For $j = 1, \dots, m-1$, H contains every edge of the form $(v_{j,A}, v_{j+1,B})$, where $A \in \mathcal{I}_j$, $B \in \mathcal{I}_{j+1}$, $A \cup B$ is independent and A and B are *consistent* in the sense that every indexed rectangle in $\mathcal{S}_j \cap \mathcal{S}_{j+1}$ is contained either in both A and B or in none of them. The weight of $(v_{j,A}, v_{j+1,B})$ is $w(B \setminus A)$.
- For each $A \in \mathcal{I}_m$, H contains the edge $(v_{m,A}, t)$ with weight 0.

For each $R \in \mathcal{S}$, $J(R) = \{j \mid 1 \leq j \leq m \text{ and } R \in \mathcal{S}_j\}$ is a nonempty set of consecutive integers. Given an s - t path $\pi = (s, v_{1,A_1}, \dots, v_{m,A_m}, t)$ in H , let $\mathcal{A}(\pi) = \bigcup_{j=1}^m A_j$. Because of the consistency requirement in the definition of H , if some $R \in \mathcal{S}$ belongs to $\mathcal{A}(\pi)$, it belongs to A_j for each $j \in J(R)$. Now observe that if R and R' are intersecting indexed rectangles in \mathcal{S} , then there are $j \in J(R)$ and $j' \in J(R')$ such that $|j - j'| \leq 1$ (with $|j - j'| = 1$ needed only in case R and R' are disjoint but have the same index). If R and R' both belong to $\mathcal{A}(\pi)$, we must have $R \in A_j$ and $R' \in A_{j'}$, which contradicts the independence of $A_j \cup A_{j'}$. Therefore $\mathcal{A}(\pi)$ is independent for every s - t path π in H .

On the other hand, if A is an independent set in \mathcal{S} , then $\pi = (s, v_{1,A \cap \mathcal{S}_1}, \dots, v_{m,A \cap \mathcal{S}_m}, t)$ is easily seen to be an s - t path in H with $\mathcal{A}(\pi) = A$ whose total edge weight is $w(A)$. Therefore finding a maximum-weight s - t path π in H and determining $\mathcal{A}(\pi)$ computes a maximum-weight independent set in \mathcal{S} .

With $N = |\mathcal{S}|$, H has $O(m \cdot N^{k\rho})$ vertices and $O(m \cdot N^{2k\rho})$ edges. Since $m \leq N$, H can be constructed in $N^{O(k\rho)}$ time. By processing the vertices in H in topological order, we can find a maximum-weight s - t path in H within the same time bound. Therefore the algorithm computes a maximum-weight independent set in \mathcal{S} in $N^{O(k\rho)}$ time and a maximum-weight independent set in \mathcal{R}_b in $n^{O(k\rho)}$ time, for $b = 0, \dots, k-1$. The overall running time is $n^{O(k\rho)} = n^{O(\rho/\varepsilon)}$. \square

3.5 Polynomial-Time Approximation Schemes for 1SH-Labeling

Clearly, the PTAS for 4M-labeling of Lemma 3.8 is also a PTAS for the more restricted 1MH-labeling problem. Therefore we now have all the ingredients that we need to obtain

a PTAS for weighted 1SH-labeling. We first treat instances of bounded height ratio and then instances of bounded length ratio.

Theorem 3.9. *Given a sliding-label instance I of size n and with height ratio ρ and an $\varepsilon \in \mathbb{R}$ with $0 < \varepsilon \leq 1$, a 1SH-labeling of I of weight at least $1 - \varepsilon$ times the weight of an optimal 1SH-labeling of I can be computed in $n^{(\rho t)^{O(t)}}$ time, where $t = \lceil 2/\varepsilon \rceil$. For every fixed $\rho \geq 1$, the weighted 1SH-labeling problem for instances of height ratio at most ρ therefore admits a PTAS.*

Proof. Let W^* be the weight of an optimal 1SH-labeling of I . By Corollary 3.5, the dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable, where $g = (\rho t)^{O(t)}$. By Lemma 3.7, we can compute a set $M \subseteq \mathbb{R}$ with $|M| \leq 3n^{g+1}$ such that the fixed-position instance (I, M) has a 1MH-labeling of weight at least $(1 - 1/t)W^*$. Applying the PTAS of Lemma 3.8 to (I, M) , we obtain a 1MH-labeling of (I, M) , and therefore a 1SH-labeling of I , of weight at least $(1 - 1/t)^2 W^* \geq (1 - 2/t)W^* \geq (1 - \varepsilon)W^*$ in time $(n^{g+2})^{O(\rho t)} = n^{(\rho t)^{O(t)}}$, which dominates the time needed by the first step. \square

Theorem 3.10. *Given a sliding-label instance I of size n and with length ratio ρ and an $\varepsilon \in \mathbb{R}$ with $0 < \varepsilon \leq 1$, a 1SH-labeling of I of weight at least $1 - \varepsilon$ times the weight of an optimal 1SH-labeling of I can be computed in $n^{O(\rho^2 t^2)}$ time, where $t = \lceil 2/\varepsilon \rceil$. For every constant $\rho \geq 1$, the weighted 1SH-labeling problem for instances of length ratio at most ρ therefore admits a PTAS.*

Proof. Let W^* be the weight of an optimal 1SH-labeling of I . By Lemma 3.6, the dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable, where $g = \lceil 2\rho \rceil(t - 1) - 1$. By Lemma 3.7, we can compute a set $M \subseteq \mathbb{R}$ with $|M| \leq 3n^{g+1}$ such that the fixed-position instance (I, M) has a 1MH-labeling of weight at least $(1 - 1/t)W^*$. Applying the PTAS of Lemma 3.8 to (I, M) , we obtain a 1MH-labeling of (I, M) , and therefore a 1SH-labeling of I , of weight at least $(1 - 1/t)^2 W^* \geq (1 - 2/t)W^* \geq (1 - \varepsilon)W^*$ in time $(n^{g+2})^{O(\rho t)} = n^{O(\rho^2 t^2)}$, which dominates the time needed by the first step. \square

3.6 Extension to Other Slider Models

Our results for 1SH-labeling can be extended with little additional effort to the other slider models—2SH, 1SV, 2SV, and 4S. First, we adapt Lemma 3.7 to obtain a reduction from 2SH-labeling to 2MH-labeling.

Lemma 3.11. *Let a sliding-label instance $I = (F, x, y, l, h, w)$ of size n and a $t \in \mathbb{N}$ be given such that the dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable for some computable $g = g(t) \geq 0$. Then, in $O(n^{g+2})$ time, we can compute a function $\mathcal{M} : F \rightarrow \mathbb{R}^2$ with $|\mathcal{M}(p)| \leq 6n^{g+1}$ for all $p \in F$ and with the following property: For every 2SH-labeling (Q, z) of I , there is a 2MH-labeling (Q', z') of (I, \mathcal{M}) with $Q' \subseteq Q$ of weight at least $(1 - 1/t)w(Q)$ such that for all $p \in Q'$, $z'_y(p) = z_y(p)$, $z'_x(p) \leq z_x(p)$ and $z'_x(p)$ and $z_x(p)$ have the same rank in S_I .*

Proof. Observe that every 2SH-labeling (Q, z) of I is a 1SH-labeling of the y -modification $I_{Q,z} = (F, x, y', l, h, w)$ of I , where for all $p \in F$, $y'_p = y_p - h_p$ if $p \in Q$ and $z_y(p) = y_p - h_p$ (i.e., p is labeled with a rectangle that has the site of p on its top edge), and $y'_p = y_p$ otherwise. Let M be the set of cardinality at most $3n^{g+1}$ that can be computed in time $O(n^{g+1})$ by applying the algorithm of Lemma 3.7 to I . Now let $\mathcal{M}(p) = \{(\bar{x}, y_p) \mid \bar{x} \in M\} \cup \{(\bar{x}, y_p - h_p) \mid \bar{x} \in M\}$ for all $p \in F$. Clearly, $|\mathcal{M}(p)| = 2|M| \leq 6n^{g+1}$ for all $p \in F$. Consider now an arbitrary 2SH-labeling (Q, z) of I . As argued above, (Q, z) is a 1SH-labeling of the y -modification $I_{Q,z}$ of I . By Lemma 3.7, there is a 1MH-labeling (Q', z') of $(I_{Q,z}, M)$ with $Q' \subseteq Q$ of weight at least $(1 - 1/t)w(Q)$ such that for all $p \in Q'$, $z'_x(p) \leq z_x(p)$ and $z'_y(p)$ and $z_x(p)$ have the same rank in S_I . That 1MH-labeling (Q', z') is a 2MH-labeling of (I, \mathcal{M}) with the required properties. \square

Theorem 3.12. *For every fixed $\rho \geq 1$, there is a PTAS for weighted 2SH-labeling on instances of height or length ratio at most ρ .*

Proof. Let a sliding-label instance $I = (F, x, y, l, h, w)$ of size n and with height or length ratio at most ρ and an ε with $0 < \varepsilon \leq 1$ be given and take $t = \lceil 2/\varepsilon \rceil$.

The dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable, where $g = (\rho t)^{O(t)}$ if the height ratio of I is at most ρ (Corollary 3.5) and $g = \lceil 2\rho \rceil(t-1) - 1$ if the length ratio of I is at most ρ (Lemma 3.6). In either case, apply Lemma 3.11 to I to obtain a fixed-position instance (I, \mathcal{M}) with $|\mathcal{M}(p)| \leq 6n^{g+1}$ for all $p \in F$.

Then apply the PTAS of Lemma 3.8 to (I, \mathcal{M}) to obtain a 2MH-labeling (\hat{Q}, \hat{z}) of (I, \mathcal{M}) of weight at least $1 - 1/t$ times the weight of an optimal 2MH-labeling of (I, \mathcal{M}) . Output (\hat{Q}, \hat{z}) as a 2SH-labeling of I .

Let (Q^*, z^*) be an optimal 2SH-labeling of I . By Lemma 3.11, there is a 2MH-labeling (Q', z') of (I, \mathcal{M}) of weight at least $(1 - 1/t)w(Q^*)$. As the weight of (\hat{Q}, \hat{z}) is at least $1 - 1/t$ times the weight of an optimal 2MH-labeling of (I, \mathcal{M}) , we have that $w(\hat{Q}) \geq (1 - 1/t)w(Q') \geq (1 - 1/t)^2w(Q^*) \geq (1 - \varepsilon)w(Q^*)$.

The running time of the algorithm is dominated by the application of the PTAS of Lemma 3.8 to (I, \mathcal{M}) and amounts to $(n^{g+2})^{O(\rho t)}$, which is $n^{(\rho t)^{O(t)}}$ if the height ratio of I is at most ρ and $n^{O(\rho^2 t^2)}$ if the length ratio of I is at most ρ . \square

In conjunction with the mirroring transformation, Theorems 3.9, 3.10 and 3.12 immediately imply the following result concerning vertically sliding labels.

Corollary 3.13. *For every fixed $\rho \geq 1$, there are polynomial-time approximation schemes for weighted 1SV-labeling and for weighted 2SV-labeling on instances of height or length ratio at most ρ .*

A further generalization is to consider the most general slider model, 4S, in which a label may have its site anywhere on its boundary. Informally, we deal with this case as follows. Given a sliding-label instance I with height or length ratio at most ρ , we apply to I two reductions, namely the one from 2SH-labeling to 2MH-labeling and the one from 2SV-labeling to 2MV-labeling, and form for each feature the union of the two sets of anchor points obtained for it. Then we run the PTAS of Lemma 3.8 for 4M-labeling with the

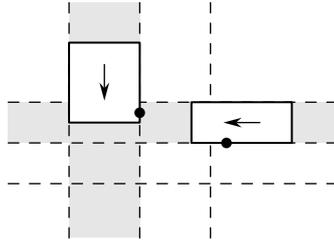


Fig. 6: Two labels overlap if and only if each intersects the “lane” (shown in gray) of the other. As no label enters the lane of another label during the normalization, the pairwise disjointness among the labels is preserved.

corresponding combined fixed-position instance as its input. To analyze the algorithm, we consider an optimal 4S-labeling and partition its labeled features into a 2SH-labeling and a 2SV-labeling depending on whether the site of a feature lies on a horizontal or on a vertical edge of its label. The two corresponding dependency graphs can be trimmed separately, and because our choice of stopping lines prevents each horizontally moving label from entering the “lane” of a vertically moving label, and vice versa (see Fig. 6), the union of the two resulting renormalized labelings is a 4M-labeling of the combined fixed-position instance and therefore a 4S-labeling of the original instance I .

In the proof we need to consider dependency graphs also of 1SV-labelings. Intuitively, the dependency graph of a 1SV-labeling of a sliding-label instance $I = (F, x, y, l, h, w)$ is analogous to that of a 1SH-labeling, but models vertically instead of horizontally sliding labels and horizontal stopping lines with y -coordinates in $\bigcup_{p \in F} \{y_p - h_p, y_p, y_p + h_p\}$. Formally, we can define the dependency graph of a 1SV-labeling (Q, z) of a sliding-label instance $I = (F, x, y, l, h, w)$ to be the dependency graph of the 1SH-labeling (Q, \bar{z}) of the mirror image $\bar{I} = (F, y, x, h, l, w)$ of I , where $\bar{z}(p) = (z_y(p), z_x(p))$ is the mirror image of $z(p)$ for all $p \in Q$.

Theorem 3.14. *For every fixed $\rho \geq 1$, there is a PTAS for weighted 4S-labeling on instances of height or length ratio at most ρ .*

Proof. Let a sliding-label instance $I = (F, x, y, l, h, w)$ of size n and an ε with $0 < \varepsilon \leq 1$ be given and take $t = \lceil 2/\varepsilon \rceil$. Without loss of generality (if necessary, apply the mirroring transformation), we assume the height ratio of I to be at most ρ .

By Corollary 3.5, the dependency graph of every 1SH-labeling of a y -modification of I is (t, g) -trimmable for some $g = (\rho t)^{O(t)}$. Applying the algorithm of Lemma 3.11 to I , construct a fixed-position instance (I, \mathcal{M}_h) of size $O(n^{g+2})$. By Lemma 3.6, applied to the mirror image $\bar{I} = (F, y, x, h, l, w)$ of I , which has length ratio bounded by ρ , the dependency graph of every 1SV-labeling of an x -modification of I is (t, \bar{g}) -trimmable for $\bar{g} = \lceil 2\rho \rceil (t - 1) - 1$. Applying the algorithm of a “mirrored” analogue of Lemma 3.11 for 2SV-labeling to I , construct a fixed-position instance (I, \mathcal{M}_v) of size $O(n^{\bar{g}+2})$.

Now create a fixed-position instance (I, \mathcal{M}) by letting $\mathcal{M}(p) = \mathcal{M}_h(p) \cup \mathcal{M}_v(p)$ for all $p \in F$. Note that $|\mathcal{M}(p)| \leq 6n^{g+1} + 6n^{\bar{g}+1}$ for all $p \in F$. Applying the PTAS of Lemma 3.8

to (I, \mathcal{M}) , obtain a 4M-labeling of (I, \mathcal{M}) of weight at least $1 - 1/t$ times the weight of an optimal 4M-labeling of (I, \mathcal{M}) and output it as the desired 4S-labeling of I .

To analyze the approximation ratio achieved, partition an optimal 4S-labeling (Q^*, z^*) of I into a 2SH-labeling (Q_h^*, z_h^*) and a 2SV-labeling (Q_v^*, z_v^*) by defining $Q_h^* = \{p \in Q^* \mid z_y^*(p) \in \{y_p, y_p - h_p\}\}$ and $Q_v^* = Q^* \setminus Q_h^*$ and letting z_h^* and z_v^* be the restrictions of z^* to Q_h^* and Q_v^* , respectively. By Lemma 3.11, there is a 2MH-labeling (Q_h, z') of (I, \mathcal{M}_h) with $Q_h \subseteq Q_h^*$ of weight at least $(1 - 1/t)w(Q_h^*)$ such that for all $p \in Q_h$, $z'_y(p) = z_y^*(p)$, $z'_x(p) \leq z_x^*(p)$ and $z'_x(p)$ and $z_x^*(p)$ have the same rank in $S_I = \bigcup_{p \in F} \{x_p - l_p, x_p, x_p + l_p\}$. By the “mirrored” analogue of Lemma 3.11 for 2SV-labeling, there is a 2MV-labeling (Q_v, z'') of (I, \mathcal{M}_v) with $Q_v \subseteq Q_v^*$ of weight at least $(1 - 1/t)w(Q_v^*)$ such that for all $p \in Q_v$, $z''_x(p) = z_x^*(p)$, $z''_y(p) \leq z_y^*(p)$ and $z''_y(p)$ and $z_y^*(p)$ have the same rank in $S_{\bar{I}} = \bigcup_{p \in F} \{y_p - h_p, y_p, y_p + h_p\}$.

Consider the pair (\hat{Q}, \hat{z}) with $\hat{Q} = Q_h \cup Q_v$, $\hat{z}(p) = z'(p)$ for $p \in Q_h$, and $\hat{z}(p) = z''(p)$ for $p \in Q_v$. The weight of (\hat{Q}, \hat{z}) is at least $(1 - 1/t)w(Q_h^*) + (1 - 1/t)w(Q_v^*) = (1 - 1/t)w(Q^*)$. We claim that (\hat{Q}, \hat{z}) is a 4M-labeling of (I, \mathcal{M}) . Assume for a contradiction that the label of a feature $p \in Q_h$ overlaps the label of a feature $q \in Q_v$. The labels of p and q are disjoint in the labeling (Q^*, z^*) . In stepping from (Q^*, z^*) to (\hat{Q}, \hat{z}) , the label of each $p \in Q_h$ reaches its position in (\hat{Q}, \hat{z}) by sliding left without crossing a stopping line in S_I , and the label of each $q \in Q_v$ reaches its position in (\hat{Q}, \hat{z}) by sliding down without crossing a stopping line in $S_{\bar{I}}$. By the definition of S_I and $S_{\bar{I}}$, it is impossible for the two labels to end in overlapping positions, so (\hat{Q}, \hat{z}) is indeed a 4M-labeling of (I, \mathcal{M}) .

Applying the PTAS of Lemma 3.8 to (I, \mathcal{M}) therefore gives a 4M-labeling of (I, \mathcal{M}) of weight at least $(1 - 1/t)w(\hat{Q}) \geq (1 - 1/t)^2 w(Q^*) \geq (1 - \varepsilon)w(Q^*)$.

The running time of the algorithm is dominated by the time needed for applying the PTAS of Lemma 3.8 to (I, \mathcal{M}) . As (I, \mathcal{M}) is of size $O(n^{g+2} + n^{\bar{g}+2})$, the running time amounts to $(n^{(\rho t)^{O(t)}} + n^{O(\rho t)})^{O(\rho t)} = n^{(\rho t)^{O(t)}}$. \square

4 Open Problems

Corollary 2.2 states that a family of graphs is trimmable if it is of bounded treewidth *and* bounded degree. We cannot exclude, however, that the bounded-degree condition is superfluous. In other words, is there a function $g : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that for all $k, t \in \mathbb{N}$, every weighted undirected graph of treewidth k has a $(t, g(k, t))$ -trimming? The answer is yes in the unweighted case, that is, if all weights are the same. If the answer were generally yes, it would follow by the argument in the proof of Corollary 2.3 that the family of planar graphs is also trimmable. This would then give a general polynomial-time reduction from weighted 1SH-labeling to weighted 1MH-labeling (albeit not, by itself, a PTAS for weighted 1SH-labeling), and similarly for the other slider models. More generally, the question of which families of graphs are trimmable deserves further study.

Acknowledgments

The foundations of this work were laid in many fruitful discussions at the research institutes in Dagstuhl and Oberwolfach. We thank Hans Bodlaender for pointing us to the concept of domino treewidth. The connection between trimming and ε -nets was suggested to us by Raimund Seidel.

References

- [AvKS98] Pankaj K. Agarwal, Marc van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. *Comput. Geom. Theory Appl.*, 11:209–218, 1998.
- [Bak94] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41:153–180, 1994.
- [Bod98] Hans L. Bodlaender. A partial k -arboretum of graphs with bounded treewidth. *Theoret. Comput. Sci.*, 209(1–2):1–45, 1998.
- [Bod99] Hans L. Bodlaender. A note on domino treewidth. *Discrete Math. Theor. Comput. Sci.*, 3(4):141–150, 1999.
- [DO95] Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *J. Graph Theory*, 20:481–499, 1995.
- [DQVZ03] Rob Duncan, Jianbo Qian, Antoine Vigneron, and Binhai Zhu. Polynomial time algorithms for three-label point labeling. *Theoret. Comput. Sci.*, 296(1):75–87, 2003.
- [GGJK78] M. R. Garey, R. L. Graham, D. S. Johnson, and D. E. Knuth. Complexity results for bandwidth minimization. *SIAM J. Appl. Math.*, 34(3):477–495, 1978.
- [HM85] Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32:130–136, 1985.
- [HW87] David Haussler and Emo Welzl. ε -nets and simplex range queries. *Discrete Comput. Geom.*, 2:127–151, 1987.
- [Jia06] Minghui Jiang. A new approximation algorithm for labeling points with circle pairs. *Inform. Process. Lett.*, 99(4):125–129, 2006.
- [PSS⁺03] Sheung-Hung Poon, Chan-Su Shin, Tycho Strijk, Takeaki Uno, and Alexander Wolff. Labeling points with weights. *Algorithmica*, 38(2):341–362, 2003.

- [QWXZ00] Zhongping Qin, Alexander Wolff, Yinfeng Xu, and Binhai Zhu. New algorithms for two-label point labeling. In Mike Paterson, editor, *Proc. 8th Annual European Symposium on Algorithms (ESA '00)*, volume 1879 of *Lecture Notes Comput. Sci.*, pages 368–379. Springer-Verlag, 2000.
- [RS86] Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
- [vKSW99] Marc van Kreveld, Tycho Strijk, and Alexander Wolff. Point labeling with sliding labels. *Comput. Geom. Theory Appl.*, 13:21–47, 1999.
- [WS08] Alexander Wolff and Tycho Strijk. The Map-Labeling Bibliography. <http://i11www.ira.uka.de/map-labeling/bibliography>, 1996–2008.