

Consistent Labeling of Rotating Maps

Andreas Gemsa, Martin Nöllenburg, and Ignaz Rutter

Institute of Theoretical Informatics, Karlsruhe Institute of Technology (KIT), Germany

Abstract. Dynamic maps that allow continuous map rotations, e.g., on mobile devices, encounter new issues unseen in static map labeling before. We study the following dynamic map labeling problem: The input is a static, labeled map, i.e., a set P of points in the plane with attached non-overlapping horizontal rectangular labels. The goal is to find a *consistent* labeling of P under rotation that maximizes the number of visible labels for all rotation angles such that the labels remain horizontal while the map is rotated. A labeling is consistent if a single *active* interval of angles is selected for each label such that labels neither intersect each other nor occlude points in P at any rotation angle.

We first introduce a general model for labeling rotating maps and derive basic geometric properties of consistent solutions. We show NP-completeness of the active interval maximization problem even for unit-square labels. We then present a constant-factor approximation for this problem based on line stabbing, and refine it further into an EPTAS. Finally, we extend the EPTAS to the more general setting of rectangular labels of bounded size and aspect ratio.

1 Introduction

Dynamic maps, in which the user can navigate continuously through space, are becoming increasingly important in scientific and commercial GIS applications as well as in personal mapping applications. In particular GPS-equipped mobile devices offer various new possibilities for interactive, location-aware maps. A common principle in dynamic maps is that users can pan, rotate, and zoom the map view. Despite the popularity of several commercial and free applications, relatively little attention has been paid to provably good labeling algorithms for dynamic maps.

Been et al. [2] identified a set of consistency desiderata for dynamic map labeling. Labels should neither “jump” (suddenly change position or size) nor “pop” (appear and disappear more than once) during monotonous map navigation; moreover, the labeling should be a function of the selected map viewport and not depend on the user’s navigation history. Previous work on the topic has focused solely on supporting zooming and/or panning of the map [2, 3, 12], whereas consistent labeling under map rotations has not been considered prior to this paper.

Most maps come with a natural orientation (usually the northern direction facing upward), but applications such as car or pedestrian navigation often rotate the map view dynamically to be always forward facing [6]. Still, the labels must remain horizontally aligned for best readability regardless of the actual rotation angle of the map. A basic requirement in static and dynamic label placement is that labels are pairwise disjoint, i.e., in general not all labels can be placed simultaneously. For labeling point features,



Fig. 1: Input map with five points (a) and three rotated views with some partially occluded labels (b)–(d).

it is further required that each label, usually modeled as a rectangle, touches the labeled point on its boundary. It is often not allowed that labels occlude the input point of another label. Figure 1 shows an example of a map that is rotated and labeled. The objective in map labeling is usually to place as many labels as possible. Translating this into the context of rotating maps means that, integrated over one full rotation from 0 to 2π , we want to maximize the number of visible labels. The consistency requirements of Been et al. [2] can immediately be applied for rotating maps.

Our Results. Initially, we define a model for rotating maps and show some basic properties of the different types of conflicts that may arise during rotation. Next, we prove that consistently labeling rotating maps is NP-complete, for the maximization of the total number of visible labels in one full rotation and NP-hard for the maximization of the visibility range of the least visible label. Finally, we present a new $1/4$ -approximation algorithm and an efficient polynomial-time approximation scheme (EPTAS) for unit-height rectangles. A PTAS is called *efficient* if its running time is $O(f(\epsilon) \cdot \text{poly}(n))$. Both algorithms can be extended to the case of rectangular labels with the property that the ratio of the smallest and largest width, the ratio of the smallest and largest height, as well as the aspect ratio of every label is bounded by a constant, even if we allow the anchor point of each label to be an arbitrary point of the label. This applies to most practical scenarios where labels typically consist of few and relatively short lines of text.

Related Work. Most previous algorithmic research efforts on automated label placement cover *static* labeling models for point, line, or area features. For static point labeling, fixed-position models and slider models have been introduced [4, 9], in which the label, represented by its bounding box, needs to touch the labeled point along its boundary. The label number maximization problem is NP-hard even for the simplest labeling models, whereas there are efficient algorithms for the decision problem that asks whether all points can be labeled in some of the simpler models (see, e.g., the discussion by Klau and Mutzel [8]). Approximation results [1, 9], heuristics [14], and exact approaches [8] are known for many variants of the static label number maximization problem.

In recent years, *dynamic* map labeling has emerged as a new research topic that gives rise to many unsolved algorithmic problems. Petzold et al. [13] used a preprocessing step to generate a reactive conflict graph that represents possible label overlaps for maps of all scales. For any fixed scale and map region, their method computes a

conflict-free labeling using heuristics. Mote [11] presents another fast heuristic method for dynamic conflict resolution in label placement that does not require preprocessing. The consistency desiderata of Been et al. [2] for dynamic labeling (no popping and jumping effects when panning and zooming), however, are not satisfied by either of the methods. Been et al. [3] showed NP-hardness of the label number maximization problem in the consistent labeling model and presented several approximation algorithms for the problem. Nöllenburg et al. [12] recently studied a dynamic version of the alternative boundary labeling model, in which labels are placed at the sides of the map and connected to their points by leaders. They presented an algorithm to precompute a data structure that represents an optimal one-sided labeling for all possible scales and thus allows continuous zooming and panning. None of the existing dynamic map labeling approaches supports map rotation.

2 Model

In this section we describe a general model for rotating maps with axis-aligned rectangular labels. Let M be a labeled input map, i.e., a set $P = \{p_1, \dots, p_n\}$ of points in the plane together with a set $L = \{\ell_1, \dots, \ell_n\}$ of pairwise disjoint, closed, and axis-aligned rectangular labels, where each point p_i is a point on the boundary $\partial\ell_i$ of its label ℓ_i . We say ℓ_i is *anchored* at p_i . As M is rotated, each label ℓ_i in L remains horizontally aligned and anchored at p_i . Thus, label intersections form and disappear during rotation of M . We take the following alternative perspective on the rotation of M . Rather than rotating the points, say clockwise, and keeping labels horizontal we may instead rotate each label around its anchor point counterclockwise and keep the set of points fixed. It is easy to see that both rotations are equivalent and yield exactly the same results.

A *rotation* of L is defined by a rotation angle $\alpha \in [0, 2\pi)$; a *rotation labeling* of M is a function $\phi: L \times [0, 2\pi) \rightarrow \{0, 1\}$ such that $\phi(\ell, \alpha) = 1$ if label ℓ is visible or *active* in the rotation of L by α , and $\phi(\ell, \alpha) = 0$ otherwise. We call a labeling ϕ *valid* if, for any rotation α , the set of labels $L(\alpha) = \{\ell \in L \mid \phi(\ell, \alpha) = 1\}$ consists of pairwise disjoint labels and no label in $L(\alpha)$ contains any point in P (other than its anchor point). We note that a valid labeling is not yet consistent in terms of the definition of Been et al. [2, 3]: given fixed anchor points, labels clearly do not jump and the labeling is independent of the rotation history, but labels may still *pop* during a full rotation from 0 to 2π , i.e., appear and disappear more than once. In order to avoid popping effects, each label may be active only in a single contiguous range of $[0, 2\pi)$, where ranges are circular ranges modulo 2π so that they may span the input rotation $\alpha = 0$. A valid labeling ϕ , in which for every label ℓ the set $A_\phi(\ell) = \{\alpha \in [0, 2\pi) \mid \phi(\ell, \alpha) = 1\}$ is a contiguous range modulo 2π , is called a *consistent* labeling. For a consistent labeling ϕ the set $A_\phi(\ell)$ is called the *active range* of ℓ . The *length* $|A_\phi(\ell)|$ of an active range $A_\phi(\ell)$ is defined as the length of the circular arc $\{(\cos \alpha, \sin \alpha) \mid \alpha \in A_\phi(\ell)\}$ on the unit circle.

The objective in static map labeling is usually to find a maximum subset of pairwise disjoint labels, i.e., to label as many points as possible. Generalizing this objective to rotating maps means that integrated over all rotations $\alpha \in [0, 2\pi)$ we want to display as many labels as possible. This corresponds to maximizing the sum $\sum_{\ell \in L} |A_\phi(\ell)|$ over all consistent labelings ϕ of M ; we call this optimization problem **MAXTOTAL**. An




Fig. 2: Two labels ℓ and ℓ' and three of their eight possible boundary intersection events. Anchor points are marked as black dots.

alternative objective is to maximize over all consistent labelings ϕ the minimum length $\min_{\ell} |A_{\phi}(\ell)|$ of all active ranges; this problem is called MAXMIN.

3 Properties of consistent labelings

In this section we show basic properties of consistent labelings. If two labels ℓ and ℓ' intersect in a rotation of α they have a (regular) *conflict* at α , i.e., in a consistent labeling at most one of them can be active at α . The set $C(\ell, \ell') = \{\alpha \in [0, 2\pi) \mid \ell \text{ and } \ell' \text{ are in conflict at } \alpha\}$ is called the *conflict set* of ℓ and ℓ' .

We show the following lemma in a more general model, in which the anchor point p of a label ℓ can be *any* point within ℓ and not necessarily a point on the boundary $\partial\ell$.

Lemma 1. *For any two labels ℓ and ℓ' with anchor points $p \in \ell$ and $p' \in \ell'$ the set $C(\ell, \ell')$ consists of at most four disjoint contiguous conflict ranges.*

Proof. The first observation is that due to the simultaneous rotation of all initially axis-parallel labels in L , ℓ and ℓ' remain “parallel” at any rotation angle α . Rotation is a continuous movement and hence any maximal contiguous conflict range in $C(\ell, \ell')$ must be a closed “interval” $[\alpha, \beta]$, where $0 \leq \alpha, \beta < 2\pi$. Here we explicitly allow $\alpha > \beta$ by defining, in that case, $[\alpha, \beta] = [\alpha, 2\pi) \cup [0, \beta]$. At a rotation of α (resp. β) the two labels ℓ and ℓ' intersect only on their boundary. Let l, r, t, b be the left, right, top, and bottom sides of ℓ and let l', r', t', b' be the left, right, top, and bottom sides of ℓ' (defined at a rotation of 0). Since ℓ and ℓ' are parallel, the only possible cases, in which they intersect on their boundary but not in their interior are $t \cap b'$, $b \cap t'$, $l \cap r'$, and $r \cap l'$. Each of those four cases may appear twice, once for each pair of opposite corners contained in the intersection. Figure 2 illustrates three of these eight boundary intersection events. Each of the conflicts defines a unique rotation angle and obviously at most four disjoint conflict ranges can be defined with these eight rotation angles as their endpoints. \square

In the following we look more closely at the conditions under which the boundary intersection events (also called *conflict events*) occur and at the rotation angles defining them. Let h_t and h_b be the distances from p to t and b , respectively. Similarly, let w_l and w_b be the distances from p to l and r , respectively (see Figure 3). By $h'_t, h'_b, w'_l,$ and w'_r we denote the corresponding values for label ℓ' . Finally, let d be the distance of the two anchor points p and p' . To improve readability of the following lemmas we define two functions $f_d(x) = \arcsin(x/d)$ and $g_d(x) = \arccos(x/d)$.




Fig. 3: Parameters of label ℓ anchored at p .




Fig. 4: Boundary intersection events for $t \cap b'$.

Lemma 2. *Let ℓ and ℓ' be two labels anchored at points p and p' . Then the conflict events in $C(\ell, \ell')$ are a subset of $\mathcal{C} = \{2\pi - f_d(h_t + h'_b), \pi + f_d(h_t + h'_b), f_d(h_b + h'_t), \pi - f_d(h_b + h'_t), 2\pi - g_d(w_r + w'_l), g_d(w_r + w'_l), \pi - g_d(w_l + w'_r), \pi + g_d(w_l + w'_r)\}$.*

Proof. Assume without loss of generality that p and p' lie on a horizontal line. First we show that the possible conflict events are precisely the rotation angles in \mathcal{C} . We start considering the intersection of the two sides t and b' . If there is a rotation angle under which t and b' intersect then we have the situation depicted in Figure 4 and by simple trigonometric reasoning the two rotation angles at which the conflict events occur are $2\pi - \arcsin((h_t + h'_b)/d)$ and $\pi + \arcsin((h_t + h'_b)/d)$. Obviously, we need $d \geq h_t + h'_b$. Furthermore, for the intersection in Figure 4a to be non-empty, we need $d^2 \leq (w_r + w'_l)^2 + (h_t + h'_b)^2$; similarly, for the intersection in Figure 4b, we need $d^2 \leq (w_l + w'_r)^2 + (h_t + h'_b)^2$.

From an analogous argument we obtain that the rotation angles under which b and t' intersect are $\arcsin((h_b + h'_t)/d)$ and $\pi - \arcsin((h_b + h'_t)/d)$. Clearly, we need $d \geq h_b + h'_t$. Furthermore, we need $d^2 \leq (w_r + w'_l)^2 + (h_b + h'_t)^2$ for the first intersection and $d^2 \leq (w_l + w'_r)^2 + (h_b + h'_t)^2$ for the second intersection to be non-empty under the above rotations.

The next case is the intersection of the two sides r and l' . Here the two rotation angles at which the conflict events occur are $2\pi - \arccos((w_r + w'_l)/d)$ and $\arccos((w_r + w'_l)/d)$. For the first conflict event we need $d^2 \leq (w_r + w'_l)^2 + (h_t + h'_b)^2$, and for the second we need $d^2 \leq (w_r + w'_l)^2 + (h_b + h'_t)^2$. For each of the intersections to be non-empty we additionally require that $d \geq w_r + w'_l$.

Similar reasoning for the final conflict events of $l \cap r'$ yields the rotation angles $\pi - \arccos((w_l + w'_r)/d)$ and $\pi + \arccos((w_l + w'_r)/d)$. The additional constraints are $d \geq w_l + w'_r$ for both events and $d^2 \leq (w_l + w'_r)^2 + (h_b + h'_t)^2$ for the first intersection and $d^2 \leq (w_l + w'_r)^2 + (h_t + h'_b)^2$. Thus, \mathcal{C} contains all possible conflict events. \square

One of the requirements for a valid labeling is that no label may contain a point in P other than its anchor point. For each label ℓ this gives rise to a special class of conflict ranges, called *hard* conflict ranges, in which ℓ may never be active. The rotation angles at which hard conflicts start or end are called *hard* conflict events. Every angle that is a (hard) conflict event is called a *label event*. Obviously, every hard conflict is also a regular conflict. Regular conflicts that are not hard conflicts are also called *soft* conflicts. We note that by definition regular conflicts are symmetric, i.e., $C(\ell, \ell') = C(\ell', \ell)$, whereas hard conflicts are not symmetric. The next lemma characterizes the hard conflict ranges.

Lemma 3. For a label ℓ anchored at point p and a point $q \neq p$ in P , the hard conflict events of ℓ and q are a subset of $\mathcal{H} = \{2\pi - f_d(h_t), \pi + f_d(h_t), f_d(h_b), \pi - f_d(h_b), 2\pi - g_d(w_r), g_d(w_r), \pi - g_d(w_l), \pi + g_d(w_l)\}$.

Proof. We define a label of width and height 0 for q , i.e., we set $h'_t = h'_b = w'_l = w'_r = 0$. Then the result follows immediately from Lemma 2. \square

A simple way to visualize conflict ranges and hard conflict ranges is to mark, for each label ℓ anchored at p and each of its (hard) conflict ranges, the circular arcs on the circle centered at p and enclosing ℓ . Figure 5 shows an example.

In the following we show that the MAXTOTAL problem can be discretized in the sense that there exists an optimal solution whose active ranges are defined as intervals whose borders are label events. An active range *border* of a label ℓ is an angle α that is characterized by the property that the labeling ϕ is not constant in any ε -neighborhood of α . We call an active range where both borders are label events a *regular active range*.




Fig. 5: Conflict ranges of two labels ℓ and ℓ' marked in bold on the enclosing circles.

Lemma 4. Given a labeled map M there is an optimal rotation labeling of M consisting of only regular active ranges.

Proof. Let ϕ be an optimal labeling with a minimum number of active range borders that are no label events. Assume that there is at least one active range border β that is no label event. Let α and γ be the two adjacent active range borders of β , i.e., $\alpha < \beta < \gamma$, where α and γ are active range borders, but not necessarily label events. Then let L_l be the set of labels whose active ranges have left border β and let L_r be the set of labels whose active ranges have right border β . For ϕ to be optimal L_l and L_r must have the same cardinality since otherwise we could increase the active ranges of the larger set and decrease the active ranges of the smaller set by an $\varepsilon > 0$ and obtain a better labeling.

So define a new labeling ϕ' that is equal to ϕ except for the labels in L_l and L_r : define the left border of the active ranges of all labels in L_l and the right border of the active ranges of all labels in L_r as γ instead of β . Since $|L_l| = |L_r|$ we shrink and grow an equal number of active ranges by the same amount. Thus the two labelings ϕ and ϕ' have the same objective value $\sum_{\ell \in L} |A_\phi(\ell)| = \sum_{\ell \in L} |A_{\phi'}(\ell)|$. Because ϕ' uses as active range borders one non-label event less than ϕ this number was not minimum in ϕ —a contradiction. As a consequence ϕ has only label events as active range borders. \square

4 Complexity

In this section we show that finding an optimal solution for MAXTOTAL (and also MAXMIN) is NP-hard even if all labels are unit squares and their anchor points are their lower-left corners. We present a gadget proof reducing from the NP-complete

problem planar 3-SAT [10]. Proofs of the lemmas in this section are found in the full version of the paper [5]. Before constructing the gadgets, we show a special property of unit-square labels.

Lemma 5. *If two unit-square labels ℓ and ℓ' whose anchor points are their lower-left corners have a conflict at a rotation angle α , then they have conflicts at all angles $\alpha + i \cdot \pi/2$ for $i \in \mathbb{Z}$.*

For every label ℓ we define the *outer circle* of ℓ as the circle of radius $\sqrt{2}$ centered at the anchor point of ℓ . Since the top-right corner of ℓ traces the outer circle we will use the locus of that corner to visualize active ranges or conflict ranges on the outer circle. Note that due to the fact that at the initial rotation of 0 the diagonal from the anchor point to the top-right corner of ℓ forms an angle of $\pi/4$ all marked ranges are actually offset by $\pi/4$.

4.1 Basic Building Blocks

Chain. A *chain* consists of at least four labels anchored at collinear points that are evenly spaced with distance $\sqrt{2}$. Hence, each point is placed on the outer circles of its neighbors. We call the first and last two labels of a chain *terminals* and the remaining part *inner chain*, see Figure 6a. We denote an assignment of active ranges to the labels as the *state* of the chain. The important observation is that in any optimal solution of MAXTOTAL an inner chain has only two different states, whereas terminals have multiple optimal states that are all equivalent for our purposes; see Figure 6a. In particular, in an optimal solution each label of an inner chain has an active range of length π and active ranges alternate between adjacent labels. We will use the two states of chains as a way to encode truth values in our reduction.

Lemma 6. *In any optimal solution, any label of an inner chain has an active range of length π . The active ranges of consecutive labels alternate between $(0, \pi)$ and $(\pi, 2\pi)$.*

Inverter. The second basic building block is an *inverter*. It consists of five collinear labels that are evenly spaced with distance $3/4 \cdot \sqrt{2}$ as depicted in Figure 6c. This means that the five labels together take up the same space as four labels in a usual inner chain. Similar to Lemma 6 the active ranges in an optimal solution also alternate. By replacing four labels of an inner chain with an inverter we can alter the parity of an inner chain.

Turn. The third building block is a *turn* that consists of four labels, see Figure 6b. The anchor points p_a and p_b are at distance $\sqrt{2}$ and the pairwise distances between p_b, p_c , and p_d are also $\sqrt{2}$ such that the whole structure is symmetric with respect to the line through p_a and p_b . The central point p_b is called *turn point*, and the two points p_c and p_d are called *outgoing points*. Due to the hard conflicts created by the four points we observe that the outer circle of p_b is divided into two ranges of length $5\pi/6$ and one range of length $\pi/3$. The outer circles of the outgoing points are divided into ranges of length $\pi, 2\pi/3$, and $\pi/3$. The outer circle of p_a is divided into two ranges of length π . The outgoing points serve as connectors to terminals, inner chains, or further turns. Note, by coupling multiple turns we can divert an inner chain by any multiple of 30° .

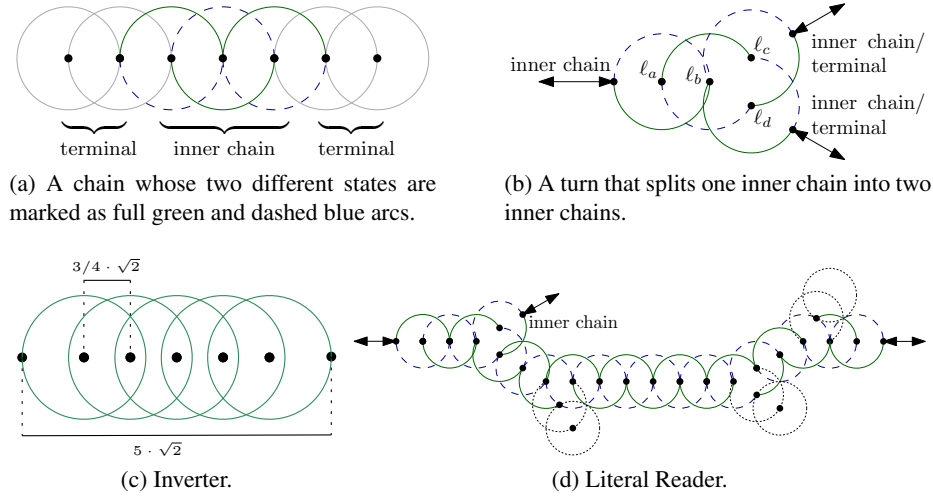


Fig. 6: Basic Building Blocks.

Lemma 7. *A turn has only two optimal states and allows to split an inner chain into two equivalent parts in an optimal solution.*

4.2 Gadgets of the Reduction

Variable Gadget. The variable gadget consists of an alternating sequence of two building blocks: horizontal chains and *literal readers*. A literal reader is a structure that allows us to split the truth value of a variable into one part running towards a clause and the part that continues the variable gadget, see Figure 6d. The literal reader consists of four turns, the first of which connects to a literal pipe and the other three are dummy turns needed to lead the variable gadget back to our grid. Note that some of the distances between anchor points in the literal reader need to be slightly less than $\sqrt{2}$ in order to reach a grid point at the end of the structure.

In order to encode truth values we define the state in which the first label of the first horizontal chain has active range $(0, \pi)$ as *true* and the state with active range $(\pi, 2\pi)$ as *false*.

Clause Gadget. The clause gadget consists of one inner and three outer labels, where the anchor points of the outer labels split the outer circle of the inner label into three equal parts of length $2\pi/3$, see Figure 7. Each outer label further connects to an incoming literal pipe and a terminal. These two connector labels are placed so that the outer circle of the outer label is split into two ranges of length $3\pi/4$ and one range of length $\pi/2$.

The general idea behind the clause gadget is as follows. The inner label obviously cannot have an active range larger than $2\pi/3$. Each outer label is placed in such a way that if it carries the value *false* it has a soft conflict with the inner label in one of the three possible active ranges of length $2\pi/3$. Hence, if all three labels transmit the value




Fig. 7: Clause gadget with one inner and three outer labels.




Fig. 8: Sketch of the gadget placement for the reduction.

false then every possible active range of the inner label of length $2\pi/3$ is affected by a soft conflict. Consequently, its active range can be at most $\pi/2$.

On the other hand, if at least one of the pipes transmits *true*, the inner label can be assigned an active range of maximum length $2\pi/3$.

Lemma 8. *There must be a label in a clause or one of the incoming pipes with an active range of length at most $\pi/2$ if and only if all three literals of that clause evaluate to false.*

Pipes. Pipes propagate truth values of variable gadgets to clause gadgets. We use three different types of pipes, which we call *left arm*, *middle arm*, and *right arm*, depending on where the pipe attaches to the clause.

One end of each pipe attaches to a variable at the open outgoing label of a literal reader. Initially, the pipe leaves the variable gadget at an angle of 30° . By using sequences of turns, we can route the pipes at any angle that is an integer multiple of 30° . Thus we can make sure that for a clause above the variables the left arm enters the clause gadget at an angle of 150° , the middle arm at an angle of 270° , and the right arm at an angle of 30° with respect to the positive x -axis. For clauses below the variables the pipes are mirrored.

In order to transmit the correct truth value into the clause we first need to place the literal reader such that the turn point of the first turn corresponds to an even position in the variable chain. Next, for a positive literal we need a pipe of even length, whereas for a negative literal the pipe must have odd length. Note that we can always achieve the correct parity by making use of the inverter gadgets.

Gadget Placement. We place all variable gadgets on the same y -coordinate such that each anchor point of variable labels (except for literal readers) lies on integer x - and y -coordinates with respect to a grid of width and height $\sqrt{2}$. Clause gadgets and pipes lie below and above the variables and form three-legged “combs”. The overall structure of the gadget arrangement is sketched in Figure 8.

Theorem 1. *MAXTOTAL is NP-complete.*

Proof. For a given planar 3-SAT formula φ we construct the MAXTOTAL instance as described above. For this instance we can compute the maximum possible sum K of

active ranges assuming that each clause is satisfiable. By Lemma 8 every unsatisfied clause forces one label to have an active range of only $\pi/2$. Thus we know that φ is satisfiable if and only if the MAXTOTAL instance has a total active range sum of at least K . Constructing and placing the gadgets can be done in polynomial time and space.

Due to Lemma 4 we can discretize the MAXTOTAL problem. Thus we can construct an oracle that guesses an active range assignment, which we can verify in polynomial time. So MAXTOTAL is in \mathcal{NP} . \square

We note that the same construction as for the NP-hardness of MAXTOTAL can also be applied to prove NP-hardness of MAXMIN. The maximally achievable minimum length of an active range for a satisfiable formula is $2\pi/3$, whereas for an unsatisfiable formula it is $\pi/2$ due to Lemma 8. This observation also yields that MAXMIN cannot be efficiently approximated within a factor of $3/4$.

Corollary 1. *MAXMIN is NP-hard and it has no efficient approximation algorithm with an approximation factor larger than $3/4$ unless $\mathcal{P} = \mathcal{NP}$.*

5 Approximation Algorithms

In the previous section we have established that MAXTOTAL is NP-complete. Unless $\mathcal{P} = \mathcal{NP}$ we cannot hope for an efficient exact algorithm to solve the problem. In the following we devise a $1/4$ -approximation algorithm for MAXTOTAL and refine it to an EPTAS. For both algorithms we initially assume that labels are congruent unit-height rectangles with constant width $w \geq 1$ and that the anchor points are the lower-left corners of the labels. Let d be the length of the label's diagonal, i.e., $d = \sqrt{w^2 + 1}$.

Before we describe the algorithms we state four important properties that apply even to the more general labeling model, where anchor points are arbitrary points within the label or on its boundary, and where the ratio of the smallest and largest width and height, as well as the aspect ratio are bounded by constants: i) the number of anchor points contained in a square is proportional to its area, ii) the number of conflicts a label can have with other labels is bounded by a constant, iii) any two conflicting labels produce only $O(1)$ conflict regions, and finally, iv) there is an optimal MAXTOTAL solution where the borders of all active ranges are events.

Properties (i) and (ii) can easily be proved with a simple packing argument (see full version of the paper [5] for details). Property (iii) follows from property (ii) and Lemma 1. Property (iv) follows immediately from Lemma 4.

5.1 A $1/4$ -approximation for MAXTOTAL

The basis for our algorithm is the *line stabbing* or *shifting* technique by Hochbaum and Maass [7], which has been applied before to *static* labeling problems for (non-rotating) unit-height labels [1, 9]. Consider a grid G where each grid cell is a square with side length $2d$. We can address every grid cell by its row and column index. Now we can partition G into four subsets by deleting every other row and every other column with either even or odd parity. Within each of these subsets we have the property that any two grid cells have a distance of at least $2d$. Thus no two labels whose anchor points lie

in different cells of the same subset can have a conflict. We say that a grid cell c *covers* a label ℓ if the anchor point of ℓ lies inside c . By property (i) only $O(1)$ labels are covered by a single grid cell. Combining this with property (ii) we see that the number of conflicts of the labels covered by a single grid cell is constant. This implies that the number of events in that cell (cf. Lemma 4) is also constant.

The four different subsets of grid cells divide a MAXTOTAL instance into four subinstances, each of which decomposes into independent grid cells. If we solve all subsets optimally, at least one of the solutions is a $1/4$ -approximation for the initial instance due to the pigeon-hole principle.

Determining an optimal solution for the labels covered by a grid cell c works as follows. We compute, for the set of labels $L_c \subseteq L$ covered by c , the set E_c of label events. Due to Lemma 4 we know that there exists an optimal solution where all borders of active ranges are label events. Thus, to compute an optimal active range assignment for the labels in L_c we need to test all possible combinations of active ranges for all labels $\ell \in L_c$. For a single cell this requires only constant time.

We can precompute the non-empty grid cells by simple arithmetic operations on the coordinates of the anchor points and store those cells in a binary search tree. Since we have n anchor points there are at most n non-empty grid cells in the tree, and each of the cells holds a list of the covered anchor points. Building this data structure takes $O(n \log n)$ time and then optimally solving the active range assignment problem in the non-empty cells takes $O(n)$ time.

Theorem 2. *There exists an $O(n \log n)$ -time algorithm that yields a $1/4$ -approximation of MAXTOTAL for congruent unit-height rectangles with their lower-left corners as anchor points.*

5.2 An Efficient Polynomial-Time Approximation Scheme for MAXTOTAL

We extend the technique for the $1/4$ -approximation to achieve a $(1 - \varepsilon)$ -approximation. Let again G be a grid whose grid cells are squares of side length $2d$. For any integer k we can remove every k -th row and every k -th column of the grid cells, starting at two offsets i and j ($0 \leq i, j \leq k - 1$). This yields collections of meta cells of side length $(k - 1) \cdot 2d$ that are pairwise separated by a distance of at least $2d$ and thus independent. In total, we obtain k^2 such collections of meta cells.

For a given $\varepsilon \in (0, 1)$ we set $k = \lceil 2/\varepsilon \rceil$. Let c be a meta cell for the given k and let again L_c be the set of labels covered by c , and E_c the set of label events for L_c . Then, by properties (i) and (ii), both $|L_c|$ and $|E_c|$ are $O(1/\varepsilon^2)$. Since we need to test all possible active ranges for all labels in L_c , it takes $O(2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)})$ time to determine an optimal solution for the meta cell c .

For a given collection of disjoint meta cells we determine (as in Section 5.1) all $O(n)$ non-empty meta cells and store them in a binary search tree such that each cell holds a list of its covered anchor points. This requires again $O(n \log n)$ time. So for one collection of meta cells the time complexity for finding an optimal solution is $O(n 2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)} + n \log n)$. There are k^2 such collections and, by the pigeon hole principle, the optimal solution for at least one of them is a $(1 - \varepsilon)$ -approximation of the original instance. This yields the following theorem.

Theorem 3. *There exists an EPTAS that computes a $(1 - \varepsilon)$ -approximation of MAX-TOTAL for congruent unit-height rectangles with their lower-left corners as anchor points. Its time complexity is $O((n2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)} + n \log n)/\varepsilon^2)$.*

We note that this EPTAS basically relies on properties (i)–(iv) and that there is nothing special about congruent rectangles anchored at their lower-left corners. Hence we can generalize the algorithm to the more general labeling model, in which the ratio of the label heights, the ratio of the label widths, and the aspect ratios of all labels are bounded by constants. Furthermore, the anchor points are not required to be label corners; rather they can be any point on the boundary or in the interior of the labels. Finally, we can even ignore the distinction between hard and soft conflicts, i.e., allow that anchor points of non-active labels are occluded. Properties (i)–(iv) still hold in this general model. The only change in the EPTAS is to set the width and height of the grid cells to twice the maximum diameter of all labels in L .

Corollary 2. *There exists an EPTAS that computes a $(1 - \varepsilon)$ -approximation of MAX-TOTAL in the general labeling model with rectangular labels of bounded height ratio, width ratio, and aspect ratio, whose anchor points are arbitrary points in the respective labels. The time complexity of the EPTAS is $O((n2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)} + n \log n)/\varepsilon^2)$.*

References

1. P. K. Agarwal, M. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Comput. Geom. Theory Appl.* 11:209–218, 1998.
2. K. Been, E. Daiches, and C. Yap. Dynamic map labeling. *IEEE Trans. Visual. and Comput. Graphics* 12(5):773–780, 2006.
3. K. Been, M. Nöllenburg, S.-H. Poon, and A. Wolff. Optimizing active ranges for consistent dynamic map labeling. *Comput. Geom. Theory Appl.* 43(3):312–328, 2010.
4. M. Formann and F. Wagner. A packing problem with applications to lettering of maps. *Proc. 7th Annu. ACM Sympos. Comput. Geom. (SoCG'91)*, pp. 281–288, 1991.
5. A. Gemsa, M. Nöllenburg, I. Rutter. Consistent Labeling of Rotating Maps. *ArXiv e-prints, arXiv:1104.5634*, April 2011.
6. E. Gervais, D. Nussbaum, and J.-R. Sack. DynaMap: a context aware dynamic map application. *Proc. GISPlanet, Estoril, Lisbon, Portugal*, 2005.
7. D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM* 32(1):130–136, January 1985.
8. G. W. Klau and P. Mutzel. Optimal labeling of point features in rectangular labeling models. *Math. Programming (Series B)* pp. 435–458, 2003.
9. M. van Kreveld, T. Strijk, and A. Wolff. Point labeling with sliding labels. *Comput. Geom. Theory Appl.* 13:21–47, 1999.
10. D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.* 11(2):329–343, 1982.
11. K. D. Mote. Fast point-feature label placement for dynamic visualizations. *Inform. Visual.* 6(4):249–260, 2007.
12. M. Nöllenburg, V. Polishchuk, and M. Sysikaski. Dynamic one-sided boundary labeling. *Proc. 18th ACM SIGSPATIAL Int'l Conf. Adv. Geo. Inform. Syst.*, pp. 310–319. ACM, 2010.
13. I. Petzold, G. Gröger, and L. Plümer. Fast screen map labeling—data-structures and algorithms. *Proc. 23rd Int'l. Cartographic Conf. (ICC'03)*, pp. 288–298, 2003.
14. F. Wagner, A. Wolff, V. Kapoor, and T. Strijk. Three rules suffice for good label placement. *Algorithmica* 30(2):334–349, 2001.