

Consistent Labeling of Rotating Maps

Andreas Gemsa*

Martin Nöllenburg*

Ignaz Rutter*

Abstract

Dynamic maps that allow continuous map rotations encounter new issues unseen in static map labeling before. We study the following dynamic map labeling problem: Given a set of points P in the plane with non-overlapping axis-aligned rectangular labels attached to them, the goal is to find a *consistent* labeling of P under rotation that maximizes the number of visible labels for all possible rotation angles. A labeling is called consistent if a single *active* interval of angles is selected for each label such that no two labels intersect at any rotation angle and no point in P is ever occluded by a label.

We 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 and present an efficient polynomial-time approximation scheme.

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—ideally in a continuous fashion. 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, 10], 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 appli-

cations such as car or pedestrian navigation often rotate the map view dynamically to be always forward facing. Still, the labels should 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, 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.



Figure 1: Input map with five points (a) and two rotated views with some occluded labels (b),(c).

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, 7], 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 whether all points can be labeled in some of the simpler models (see the discussion by Klau and Mutzel [6]). Approximation results [1, 7], heuristics [12], and exact approaches [6] 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. [11] compute a conflict-free labeling for any fixed scale and

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

map region from preprocessed conflict information. Mote [9] presents a 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, 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. [10] recently studied a dynamic version of the alternative boundary labeling model and presented an algorithm that supports continuous zooming and panning. None of the existing dynamic map labeling approaches supports map rotation.

2 Model

Let M be a labeled input map, i.e., a set of points $P = \{p_1, \dots, p_n\}$ in the plane together with a set of pairwise disjoint and axis-aligned rectangular labels $L = \{\ell_1, \dots, \ell_n\}$, 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* in p_i . As M is rotated, each label ℓ_i in L remains horizontally aligned and anchored in p_i . Thus, label intersections form and disappear during rotation of M . We take the following alternative perspective. Rather than rotating the points, say clockwise, and keeping the labels fixed 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.

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]: having a fixed anchor point labels do not jump and the labeling is independent on the rotation history, but labels may still *pop* during a full rotation from 0 to 2π , i.e., appear and disappear more than once. To avoid this, each label must 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

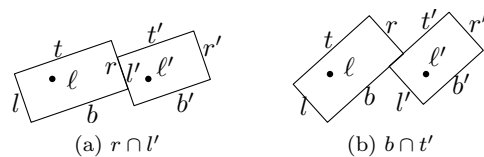


Figure 2: Two labels ℓ and ℓ' and two of their eight possible boundary intersection events. Anchor points are marked as black dots.

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**.

3 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 slightly 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 two of at most eight boundary intersection events. Each of those eight 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

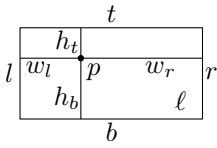


Figure 3: Parameters of label ℓ anchored at p .

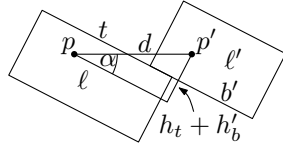


Figure 4: Conflict event for $t \cap b'$.

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' .

Lemma 2 *Let ℓ and ℓ' be two labels anchored at points p and p' . Assume without loss of generality that p and p' lie on a horizontal line and p is to the left of p' . Then there is an even number of at most eight conflict events; they depend on the distance d of p and p' .*

Proof. 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 by simple trigonometric reasoning the two rotation angles at which the conflict events occur are $2\pi - \arcsin((h_t + h'_b)/d)$ (see Figure 4) and $\pi + \arcsin((h_t + h'_b)/d)$. Obviously, we need $d \geq h_t + h'_b$. Furthermore, for the two intersection induced by $t \cap b'$ to be non-empty we need $d^2 \leq (w_r + w'_l)^2 + (h_t + h'_b)^2$ for the case depicted in Figure 4 and $d^2 \leq (w_l + w'_r)^2 + (h_t + h'_b)^2$ for the other one. Similar reasoning yields the three remaining pairs of conflict events.

It remains to show that for any distance d between p and p' the number of actual conflict events is even. Considering the distance constraints for all eight conflict events, we observe that when growing the distance d from 0 we initially have no conflict events and in fact $C(\ell, \ell') = [0, 2\pi)$, i.e., the two labels intersect at any rotation. Every time a lower bound on d is reached, we add two conflict events to a set \mathcal{C} of conflict events. Similarly, every time an upper bound on d is reached, we remove two events from \mathcal{C} and hence the cardinality of \mathcal{C} is always even. \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. Obviously, every hard conflict is also a soft conflict but not vice versa. Similar to (soft) conflict events we can define hard conflict

events and their union as the set of *label events*. If we set the label ℓ' in Lemma 2 to have height and width 0, we obtain a result analogous to Lemma 2 for the number of hard conflict events.

Next 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.

Lemma 3 *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 some $\varepsilon > 0$ and obtain a better labeling.

We define a new labeling ϕ' that coincides with ϕ except for the labels in L_l and L_r . We set the left border of the active ranges of all labels in L_l and the right border of the active ranges of all labels to 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 values but ϕ' has fewer active range borders that are non-label events—a contradiction. \square

4 Complexity

By a reduction from planar 3-SAT [8] it can be shown that finding an optimal solution for MAXTOTAL is NP-hard even if all labels are unit squares and their anchor points are their lower-left corners. Due to space constraints we omit the proof.

Theorem 4 *MAXTOTAL is NP-complete.*

5 Approximation Scheme

Initially we assume that labels are congruent unit-height rectangles with constant width $w > 1$ and that the anchor points are the lower-left corners of the labels and afterwards explain how to adapt it to more general labeling models.

We present an efficient polynomial-time approximation scheme (EPTAS) that relies on four geometric key properties; they hold for the restricted case of congruent rectangles as well as for the general model of Section 2, where anchor points are arbitrary points within each 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 rectangle is proportional to its area, ii) any label has $O(1)$ conflicts with other labels, 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 be proved by a simple packing argument. Property (iii) follows from property (ii) and Lemma 1. Property (iv) follows immediately from Lemma 3.

Our EPTAS uses the line stabbing technique of Hochbaum and Maass [5]. 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. For any integer k we can remove every k -th row and every k -th column of 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 no two labels whose anchor points lie in different cells of the same subset can have a conflict. In total we obtain k^2 such collections of meta cells. We say that a grid cell c covers a label ℓ if the anchor point of ℓ lies inside c .

Determining an optimal solution for the labels covered by a grid cell c works as follows. Calculate for the set of labels $L_c \subseteq L$ covered by c the set of events E_c . Due to Lemma 3 we know that there exists an optimal solution with only regular active ranges. 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 given $\varepsilon \in (0, 1)$ we set $k = \lceil 1/\varepsilon \rceil$. Let c be a meta cell for the given k and let L_c be the set of labels covered by c . By a packing argument, we can prove that the number of labels covered by c (and thus also the number of events of labels in L_c) is in $O(1/\varepsilon^2)$.

Since we need to test all possible active ranges for all labels in L_c we require $O(2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)})$ time to determine an optimal solution for the meta cell c .

The number of cells we need to consider is linear and we use a range query data structure to locate all points within a meta cell. So for one collection of meta cells the time complexity for finding an optimal solution is $O(n2^{O(1/\varepsilon^2 \log 1/\varepsilon^2)} + n \log n)$. There are k^2 such collections and the optimal solution for at least one of them is a $(1 - \varepsilon)$ -approximation.

This yields the result for the simple case of congruent rectangles. In fact we have only used properties (i)–(iv), and there is nothing special about congruent

rectangles anchored at their lower-left corners. Hence we can extend 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 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. All properties (i)–(iv) still hold. 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 .

Theorem 5 *There exists an efficient polynomial-time approximation scheme that computes a $(1 - \varepsilon)$ -approximation of MAXTOTAL. Its time complexity 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 Transactions on Visualization and Computer 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 Ann. ACM Sympos. Comput. Geom. (SoCG'91)*, pp. 281–288, 1991.
- [5] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM* 32(1):130–136, 1985.
- [6] G. W. Klau and P. Mutzel. Optimal labeling of point features in rectangular labeling models. *Mathematical Programming (Series B)* pp. 435–458, 2003.
- [7] M. van Kreveld, T. Strijk, and A. Wolff. Point labeling with sliding labels. *Comput. Geom. Theory Appl.* 13:21–47, 1999.
- [8] D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.* 11(2):329–343, 1982.
- [9] K. D. Mote. Fast point-feature label placement for dynamic visualizations. *Information Visualization* 6(4):249–260, 2007.
- [10] M. Nöllenburg, V. Polishchuk, and M. Sysikaski. Dynamic one-sided boundary labeling. *Proc. 18th ACM GIS*, pp. 310–319. ACM Press, November 2010.
- [11] I. Petzold, G. Gröger, and L. Plümer. Fast screen map labeling—data-structures and algorithms. *Proc. 23rd Int'l Cartogr. Conf (ICC'03)*, pp. 288–298, 2003.
- [12] F. Wagner, A. Wolff, V. Kapoor, and T. Strijk. Three rules suffice for good label placement. *Algorithmica* 30(2):334–349, 2001.