Scientific Annals of Computer Science vol. 31 (2), 2021, pp. 293–313
doi: 10.7561/SACS.2021.2.293

Maximizing the Number of Visible Labels on a Rotating Map

Ali Gholami Rudi[1]

Abstract

For a map that can be rotated, we consider the following problem. There are a number of feature points on the map, each having a geometric object as a label. The goal is to find the largest subset of these labels such that when the map is rotated and the labels remain vertical, no two labels in the subset intersect. We show that, even if the labels are vertical bars of zero width, this problem remains NP-hard, and present a polynomial approximation scheme for solving it. We also introduce a new variant of the problem for vertical labels of zero width, in which any label that does not appear in the output must be coalesced with a label that does. Coalescing a subset of the labels means to choose a representative among them and set its label height to the sum of the individual label heights.

Keywords: Computational geometry, geometric independent set, map labelling, rotating maps, polynomial-time approximation scheme.

1 Introduction

Map labelling is a classical optimization problem in cartography and graph drawing [9], whose goal is to place as many non-intersecting labels on a map as possible. For static maps, the problem of placing labels on a map can be stated as an instance of the geometric independent set problem: given a set

[1]Department of Electrical and Computer Engineering, Babol Noshirvani University of Technology, Babol, Mazandaran, Iran, Email: gholamirudi@nit.ac.ir

of geometric objects, the goal is to find its largest non-intersecting subset. In the weighted version, each object also has a weight and the goal is to find a non-intersecting subset of maximum possible weight.

This geometric problem can be described using intersection graphs by mapping each object to a vertex and adding an edge between vertices corresponding to intersecting objects. This converts the geometric problem to the classical maximum independent set for graphs, which is NP-hard and difficult to approximate even within a factor of $n^{1-\epsilon}$, where n is the number of vertices and ϵ is any non-zero positive constant [15]. Although the geometric version remains NP-hard even for unit disks [10], it is easier to approximate, and several polynomial-time approximation schemes (PTAS) have been presented for this problem [14, 1, 8, 5, 6]. Note that a PTAS finds a $(1 - \epsilon)$-approximate solution in time $O(n^{f(\epsilon)})$, for any $\epsilon > 0$ and some function f independent of n.

Maps may be dynamic, and allow zooming, panning, or rotation, as recent technology has made prevalent. Most work on labelling dynamic maps consider zooming and panning operations [2, 3], but only few results have been published for labelling rotating maps. Gemsa et al. [12] were the first to study this problem. They assumed the model presented by Been et al. [2] for zoomable maps, to define the consistency of a rotating map. For the kR-model, in which each label may disappear at most k times during rotation, they showed that labelling rotating maps is NP-hard, even for unit-height labels, when the goal is to maximise the total duration in which labels are visible without intersecting other labels. For unit-height labels, they also presented a 1/4-approximation algorithm and a PTAS, both under the assumption that the number of anchor points contained in any rectangle is bounded by a constant multiplied by its area, each label may intersect a constant number of other labels, and the aspect ratio of the labels is bounded. Note that the first two assumptions may not hold in real world maps. Subsequently, they extended their results by presenting heuristic algorithms, and an integer linear programming (ILP)-based solution for labelling rotating maps under the same assumptions [13]. They also experimentally evaluated these algorithms. The size of their ILP model of the problem was improved by Cano et al. [4].

Yokosuka and Imai [18] solved the problem of maximising the size of labels for rotating maps. Although this problem is NP-hard for static maps, they presented an exact $O(n \log n)$-time algorithm for the case where anchor points can be inside the labels, and also for the case when labels are of unit height and points can be on the boundary of labels. Gemsa et al. [11] also studied a trajectory-based labelling problem, when the trajectory of the viewport of the map is specified as an input.

In this paper, we also study the problem of labelling rotating maps. The geometric statement of our problem is as follows. The input is a set of points in the plane. To each of these points a vertical segment is assigned. The goal is to place the maximum possible number of these segments such that: i) each segment contains its corresponding point (the point is the anchor of the segment), ii) when the plane is rotated around a fixed centre point, each segment is rotated in the reverse direction around its anchor point to remain vertical, iii) during the rotation of the plane, no two segments intersect. Note that our model is different from the one assumed by Gemsa et al. [12], in which labels may disappear and reappear again.

Our contributions are as follows.

- We prove that this problem is NP-hard and present a PTAS for this problem; our PTAS is based on the one presented by Chan [5] for geometric maximum independent set.
- We extend our results to the general case, where the labels can be arbitrary objects. We make no assumptions about the distribution of the labels: a label may intersect any number of other labels, and the number of feature points in any rectangle may not be proportional to its area.
- We also discuss two new variants of the problem, in which labels that do not appear in the output must be coalesced with labels that do. We prove them to be NP-hard and model their solution as an integer linear programme.

This paper is organised as follows. We first prove that our problem is NP-hard in Section 2. Then, in Section 3, we present a PTAS for this problem, and then extend it to arbitrary labels. In Section 4 we discuss new variants of the problem, in which labels that do not appear in the output must be coalesced with labels that do. Finally, we conclude this paper in Section 5.

2 Notation and Preliminary Results

Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of points, and let ℓ_q be the length of the vertical segment corresponding to point q. A labelling ϕ for P assigns a vertical segment to some of the points in P. If a segment is assigned to point p_i in labelling ϕ, we say that p_i is included in ϕ, or equivalently, $p_i \in \phi$. The notation $\phi(p_i)$ denotes the segment assigned to p_i, and $|\phi|$ indicates the *size* of ϕ, that is, the number of points in P to which ϕ assigns segments. Note that the length of the segment assigned to p_i is ℓ_{p_i}.

If segment s is assigned to p_i in ϕ, then p_i must lie on s. The point of s that is identified with p_i is the *anchor point* of s; alternatively, we say that s is anchored at p_i. When the plane is rotated, s is rotated in the reverse direction around p_i to remain vertical. In the 1-position (1P) model, the anchor point of a segment must be its bottom end point. In the 2-position (2P) model, either the top or the bottom end points of a segment can be its anchor point. In the fixed-position (FP) model, the anchor point of each segment is fixed (but different segments may be anchored at different positions). In the slider model, any point on the segment can be its anchor point. Similar models have been introduced for labelling (static) points with axis-aligned rectangles, both in the unweighted and weighted cases [17, 7].

A labelling is *proper* if its assigned segments do not intersect during the rotation of the plane. In the Maximum Rotating Independent Set (MRIS) problem for vertical segments, the goal is to find the largest proper labelling. Instead of rotating the plane and keeping visible segments vertical, we can equivalently fix the plane and rotate all visible labels in the reverse direction. This is what we do in the rest of this paper.

We now show that MRIS for vertical segments is NP-hard in the 1P model by a reduction from the Geometric Maximum Independent Set (GMIS) problem for unit disks, which is NP-hard [10].

Theorem 1 *MRIS for a set of segments in 1P model is NP-hard.*

Proof: We reduce any instance of GMIS for unit disks to an instance of MRIS for segments. Let D be a set of n unit disks on the plane, and let P be the set of the centres of these disks. Also, let $\ell_{p_i} = 2$ for $1 \leq i \leq n$.

We first show that from every non-intersecting subset of disks in D we can obtain a proper labelling of P with the same size. Let D' be a non-intersecting subset of D, and let P' be their centres. Since the disks in D' are non-intersecting, the distance between any pair of points in P' is at least 2. Let ϕ be the labelling of size $|D'|$ that assigns a segment of length ℓ_{p_i}, anchored at its bottom end point, to each point p_i of P'. These segments cannot intersect during rotation: the segments are always parallel, and since the distance of their anchors is at least 2, they do not intersect; this is demonstrated in Figure 1. This implies that ϕ is proper.

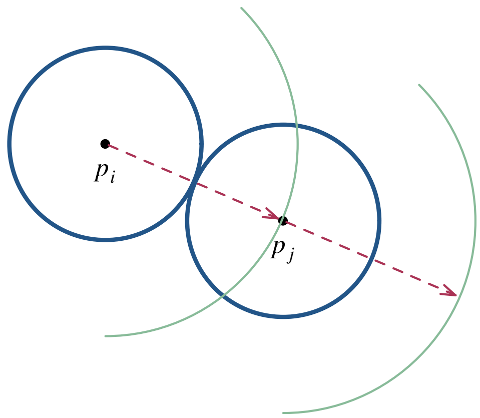

Figure 1: Two non-intersecting disks in Theorem 1

For the other direction, let ϕ be a proper labelling of P. Let D' be the set of disks, whose centres are in ϕ. Since ϕ is a proper labelling in 1P model, the distance between any pair of points in P' is at least two. This implies that the disks corresponding to P' are non-intersecting. $\square$

Next, we show that labelling in the 2P model is as difficult as labelling in the 1P model.

Lemma 1 *Let ϕ be a proper labelling of a set of points P in the 2P model, and let ϕ' be the labelling obtained from ϕ by changing the labels present in ϕ to be anchored at their bottom end points. Then, ϕ' is a proper labelling in the 1P model.*

Proof: If ϕ' is not proper, there exists at least a pair of points p_i and p_j such that $\phi'(p_i)$ and $\phi'(p_j)$ intersect during rotation. Without loss of generality, suppose that $\ell_{p_i} \geq \ell_{p_j}$. Therefore, the distance between p_i and p_j is at most ℓ_{p_i}, implying that, at some point during rotation $\phi(p_i)$ intersects p_j (and thus $\phi(p_j)$), contradicting the assumption that ϕ is a proper labelling. $\square$

Theorem 2 *MRIS in the 2P model is as difficult as MRIS in the 1P model, and we can obtain a solution to one from that of the other.*

Proof: For a given set of points, let ϕ and ϕ' be the solutions to MRIS in the 1P and 2P models, respectively. Clearly, ϕ' is also a proper labelling in the 2P model. Therefore, we have $|\phi| \geq |\phi'|$. On the other hand, based on Lemma 1 we can obtain a proper labelling of the same size in the 1P model from ϕ. This implies that $|\phi'| \geq |\phi|$. Therefore, we have $\phi = \phi'$, and the solutions we obtain for one model from the solution of the other, are optimal. □

For the slider model, we now show that there is an optimal labelling, in which all labels are anchored at their midpoint.

Lemma 2 *Let ϕ be a proper labelling of a set of points in the slider model, and let ϕ' be the labelling obtained by changing all segments assigned in ϕ to be anchored at their midpoint. Then, the resulting labelling is also proper in the slider model.*

Proof: If ϕ' is not a proper labelling, there exist points p and q such that $\phi'(p)$ and $\phi'(q)$ intersect during rotation, which implies that the distance between p and q is at most $(\ell_p + \ell_q)/2$, where ℓ_p and ℓ_q are the lengths of the segments assigned to p and q, respectively. Let $\phi(p) = p'p''$ and $\phi(q) = q'q''$, in which p' and q' are the top end points of these segments. Obviously, $|p'p| + |pp''| = \ell_p$ and $|q'q| + |qq''| = \ell_q$. During the rotation, when q' is on segment pq, we have

$$|qq'| + |pp''| < |pq|.$$

Otherwise, the segments would intersect. Similarly, when q'' is on pq, we have

$$|qq''| + |pp'| < |pq|.$$

This, however, implies that

$$\ell_q + \ell_p = |qq'| + |qq''| + |pp'| + |pp''| < 2 \cdot |pq|.$$

This yields

$$|pq| > (\ell_p + \ell_q)/2;$$

a contradiction. Therefore, ϕ' is also a proper labelling. □

In the next section, we study the MRIS problem for segments in the 1P model, but by Lemma 1 our results also apply to the 2P model. Note that in the 1P model, labels centred at p and q remain disjoint during rotation if and only if disks D_p and D_q are disjoint.

3 A PTAS for MRIS

For point $p \in P$, let D_p denote the disk centred at p with radius ℓ_p. In a proper labelling ϕ in the 1P model, if points p and q are both present in ϕ, their corresponding disks, D_p and D_q, may intersect, but neither disk may contain the centre of the other. This is what makes MRIS for segments in the 1P model different from GMIS, in which the objects in the output must be disjoint. In this section, we first review some of the PTAS presented for GMIS, and adapt one of them for our problem.

Note that transforming an instance of MRIS to GMIS, based on the idea used in Theorem 1, does not work since the length of the segments (the radii of the disks) are not equal. To see this, consider two disks D_p and D_q of radius 2 and 8, respectively, in an instance of MRIS. To obtain an equivalent instance for GMIS, we replace each disk with a disk of half its radius, as in Theorem 1. Therefore, we have two disks D'_p and D'_q, corresponding to D_p and D_q, of radius 1 and 4, respectively. D'_p and D'_q may be a solution in the GMIS instance, but their corresponding disks may not be a solution in the MRIS instance (D'_p and D'_q may be disjoint but D_q may contain the centre of D_p); this is demonstrated in Figure 2.

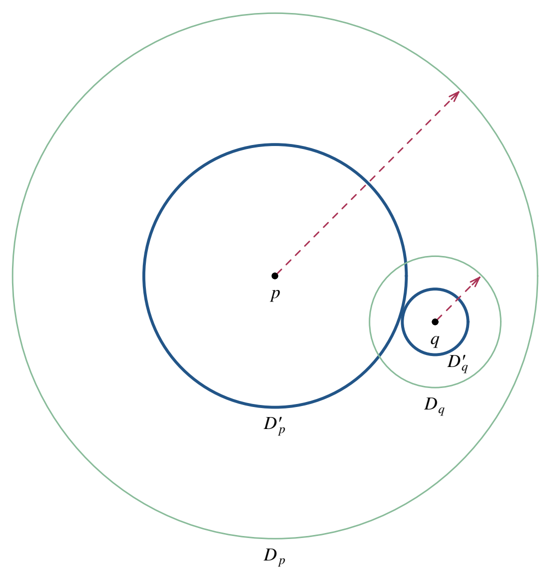

Figure 2: Disks D_p and D_q are disjoint but D'_p and D'_q are not centre-disjoint

For GMIS, Hochbaum and Maass [14] presented a PTAS for n unit disks or squares on the plane in $n^{O(1/\epsilon^2)}$ time. Agarwal et al. [1] improved this algorithm for unit-height objects using dynamic programming with time complexity $n^{O(1/\epsilon^{d-1})}$, in which d is the number of dimensions. Erlebach et al. [8] extended Hochbaum and Maass' algorithm using a multi-level grid to handle arbitrarily-sized but fat objects (informally, objects with bounded aspect ratio). Chan [5] presented a similar algorithm for fat objects using a quadtree instead of multi-level grids, improving the time complexity to $n^{O(1/\epsilon^{d-1})}$. In the same paper, Chan [5] also presented a divide-and-conquer algorithm based on the geometric separator theorem [16], that runs in time $n^{O(1/\epsilon^{d})}$ for unweighted and fat objects in the plane, improving the space complexity of the previous algorithm. More recently, Chan and Har-Peled [6] presented a PTAS with time complexity $n^{O(1/\epsilon^{d})}$ based on local search. They also presented a constant-factor approximation algorithm based on linear programming.

In the following section, we adapt Chan's [5] shifted quadtree algorithm for solving MRIS for segments in the 1P model. The main reason for preferring this algorithm to other PTAS for GMIS is its lower time complexity. The algorithm presented by Hochbaum and Maass [14] is simpler but cannot handle arbitrarily-sized segments in MRIS. To make it mostly self-contained, we repeat the necessary definitions and proofs of Chan [5], trying to simplify them where possible.

3.1 The Algorithm

For simplicity, we first map (using scaling and translation) all disks $D = \{D_1, D_2, \ldots, D_n\}$ to fit inside a unit square with its lower left corner at the origin, and store them in a quadtree. Let r_i be the radius of disk D_i after this mapping. A quadtree cell at depth d has side length 2^{-d}. Two disks are *centre-disjoint* if neither contains the centre of the other. A disk of radius ℓ is *k-aligned* if it is inside a quadtree cell of size at most $k\ell$.

The algorithm presented by Chan [5], which uses the shifting-quadtree technique, assumes fat input objects. Also, the objects in the output of GMIS must be disjoint. In MRIS, input objects are rotating segments, which are not fat, but the disks that results from their rotation are. However, the objects in the output of MRIS are centre-disjoint, but may not be disjoint. Therefore, the results of Chan [5] do not apply directly to MRIS. We modify Chan's algorithm to handle centre-disjoint objects with the aid of Lemmas 3 and 4.

Lemma 3 *Let p be a point on the plane. Any set S of centre-disjoint disks that contain p is of size at most 6.*

Proof: Consider the centres of the disks in S ordered radially around p. We show that for any two consecutive disks in this order with centres a and b (see Figure 3), we have $\angle apb \geq \frac{\pi}{3}$. To do so, we show that in triangle pab, ab is the longest side. Since p is inside and b is outside the disk centred at a, we have $|ab| > |ap|$. Using a similar argument for p and the disk centred at b, we have $|ab| > |bp|$. Therefore, ab is the longest side of triangle pab. □

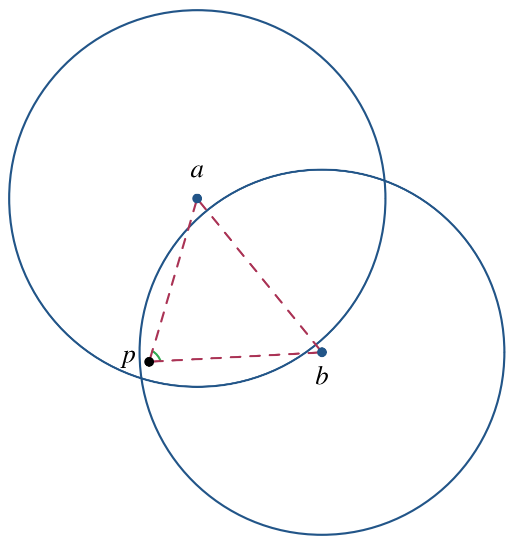

Figure 3: Centre-disjoint disks centred at a and b, containing point p (Lemma 3)

Lemma 4 *There is a constant c such that the size of any centre-disjoint set of k-aligned disks intersecting the boundary of any quadtree cell is bounded by ck.*

Proof: Let C be a quadtree cell at depth d, and let B be a set of centre-disjoint, k-aligned disks that intersect the boundary of C. If a k-aligned disk intersects C, its radius is at least $r := 2^{-d}/k$, based on the definition of k-aligned objects. Therefore, since all disks in B are k-aligned, the radius of any of them is at least r. Let C' and C'' be axis-aligned squares of side lengths $2^{-d} - r$ and $2^{-d} + r$, respectively, both centred at the centre of C.

Place a set X of points on C, C', and C'' with distance r, as shown in Figure 4. Since $2^{-d}/r = k$, we have $|X| \leq 12k$. Any disk of radius at least r that intersects the boundary of C contains at least one point from X. On the other hand, at most six centre-disjoint disks can contain any point in X due to Lemma 3. Therefore, the number of disks in B is at most $72k$. $\square$

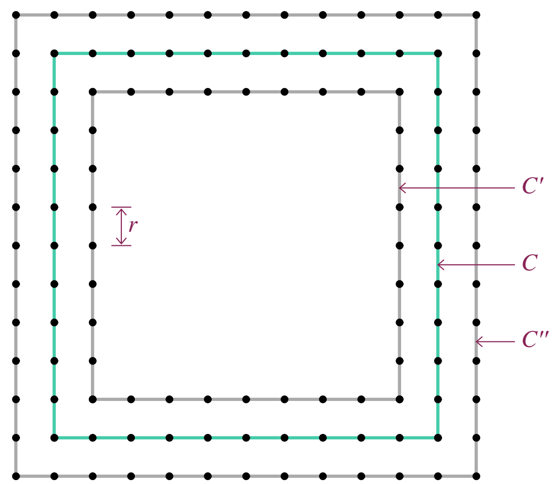

Figure 4: Points placed around the boundary of the grid cell C in Lemma 4

Let C be a quadtree cell, let B be a set of centre-disjoint disks intersecting its boundary, and let I be a set of disks inside C. Let $\mathrm{MRIS}(C, B, I)$ denote the maximum size of a centre-disjoint subset I' of I such that $I' \cup B$ is also centre-disjoint.

Lemma 5 *Let D be a set of disks that are stored in a quadtree and that are k-aligned. Let C be a quadtree cell, let B be a set of disks intersecting the boundary of C, and let I be the set of disks completely in the interior of C. The value of MRIS(C, B, I) can be computed in time $n^{O(k)}$ from the values of MRIS with the four children of C passed as its first argument.*

Proof: Let C_i for $1 \leq i \leq 4$ denote the child cells of C in the quadtree. For a set of disks X, let $C^{\mathrm{in}}(X)$ and $C^{\mathrm{on}}(X)$ denote the subset of X completely in the interior of C and the subset intersecting the boundary of C, respectively. Let B' be the subset of I with those disks that intersect the boundary of the children of C. For a centre-disjoint subset J of B', whose disks are also centre-disjoint from the disks of B, the value $\mathrm{MRIS}_J(C, B, I)$, denoting the maximum size of a centre-disjoint subset of $I \setminus B$ that includes J can be computed as follows:

$$\mathrm{MRIS}_J(C, B, I) = \sum_{i=1}^{4} \mathrm{MRIS}(C_i, C_i^{\mathrm{on}}(B \cup J), C_i^{\mathrm{in}}(I)) + |J|$$

To compute the value of $\mathrm{MRIS}(C, B, I)$, we find the maximum value of $\mathrm{MRIS}_J(C, B, I)$ for every centre-disjoint subset J of B'. Applying Lemma 4 to such subsets and the children of C, we can show that the size of any such subset is $O(k)$. Therefore, we can compare the value of $\mathrm{MRIS}_J(C, B, I)$ for every such subset (there are at most $n^{O(k)}$ of them) to find the value of $\mathrm{MRIS}(C, B, I)$ in time $n^{O(k)}$. □

Theorem 3 *A $(1-\epsilon)$-approximate solution to MRIS for segments in the 1P model can be computed in time $n^{O(1/\epsilon)}$, for any real constant ϵ with $0 < \epsilon < 1$.*

Proof: Let P and D be defined as in the beginning of this section. Let $k = 1/\epsilon$. We store D in a compressed quadtree (a quadtree in which nodes with only one non-empty child cell are merged, resulting in $O(n)$ nodes), and modify Lemma 5 to consider merged nodes. Let $\overline{C}$ be the root cell of this quadtree. We compute the value of $\mathrm{MRIS}(C, B, I)$ for every possible quadtree cell C, and inputs B and I in a bottom-up manner.

1. For each leaf cell C, $\mathrm{MRIS}(C, B, I)$ can be computed in $O(1)$ time slice. The set I has only one element.

2. For any other cell C, I is always $C^{\mathrm{in}}(D)$, and B is always a subset of $C^{\mathrm{on}}(D)$ (there are $n^{O(k)}$ subsets); computing $\mathrm{MRIS}(C, B, I)$ for every such input takes $n^{O(k)}$ time by Lemma 5.

Therefore, we can find the exact value of $\mathrm{MRIS}(\overline{C}, \emptyset, D)$, by computing MRIS for every node of the quadtree recursively in time $n^{O(1/\epsilon)}$, assuming that every disk is k-aligned.

Using the shifting technique of [5] (see also [14] and [8]), we can translate disks k times, such that in one of these translations at least an $(1-\epsilon)$-fraction of the disks in an optimal solution to MRIS for D are k-aligned. Computing $\mathrm{MRIS}(\overline{C}, \emptyset, D)$ after each such translation, and taking the maximum of these values, achieves the desired approximation factor. □

3.2 Placing Arbitrary Objects

The algorithm of Section 3.1 can be extended to work for a combination of vertical and horizontal segments (or of any orientation), or even for arbitrary objects in the FP model. For arbitrary objects of constant complexity, we similarly denote with D_i the disk that results by rotating the i-th object around its anchor point. Unlike vertical segments in the 1P model, two objects may intersect during rotation, even if the corresponding disks are centre-disjoint. To see this, consider a horizontal segment of unit length, anchored at its left endpoint at the origin, and a vertical segment of unit length, anchored at its bottom endpoint at $(1, -0.5)$; this is demonstrated in Figure 5.

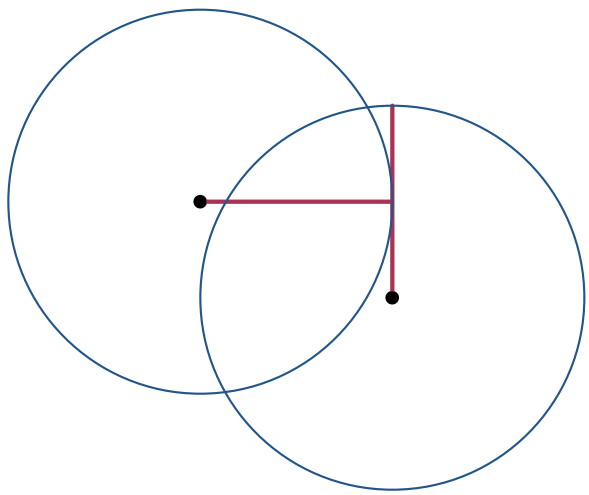

Figure 5: Horizontal and vertical bars that intersect during rotation

We modify the results of Section 3.1 as follows. Instead of finding a centre-disjoint subset of a set of disks, our goal is to find a subset of disks such that there exists a proper labelling that includes all of the objects that correspond to them. Lemma 4 can therefore be modified as follows.

Lemma 6 *There is a constant c such that for any quadtree cell C and any set of objects O whose corresponding disks are k-aligned and intersect the boundary of C, the size of any proper labelling of O is bounded by ck.*

Proof: Let C be a quadtree cell at depth d, let B be a set of k-aligned disks that intersect the boundary of C, and let O be the set of objects that correspond to the elements of B. Since the anchor of each object is inside that object, only objects whose disks are centre-disjoint can appear in ϕ (if an objects intersects the anchor of another object during rotation, their disks certainly intersect). Lemma 4 shows that the size of any centre-disjoint subset of B is bounded by ck for some constant c. Thus, the size of ϕ cannot be any greater. This implies the required upper bound. □

Using Lemma 6 in Lemma 5 and Theorem 3 implies the following corollary, after slight modifications.

Corollary 1 *For any real ϵ with $0 < \epsilon < 1$, a $(1-\epsilon)$-approximate solution to MRIS for arbitrary objects in the FP model can be computed in time $n^{O(1/\epsilon)}$.*

4 Labelling with Coalescing

In this section, we discuss a variant of the problem studied in the previous section for the 1P model, in which some of the labels may be coalesced. We give two motivating examples. Suppose that the label assigned to each feature point is horizontal and contains some text. If a label conflicts with another and cannot be visible during rotation, its text must be appended to another, nearby label, increasing the latter's length. This is demonstrated in Figure 6. The two left-most labels intersect during rotation, as shown in parts a and b; instead of rotating the plane counterclockwise, the labels are equivalently rotated counterclockwise. In part c, one of these labels are coalesced with the other. As another example, suppose that the height of the bar assigned to a feature point shows the quantity of a variable at that location. If a bar is not visible, it must be coalesced with another, preferably nearby bar, the length of which must be increased to represent the sum of their quantities.

In what follows, we consider two variants for this problem for the 1P model based on restricting the labels that can be coalesced.

4.1 Unrestricted Coalescing

Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of points, and let ℓ_{p_i} be the length of the vertical segment assigned to p_i. In a labelling, each point is either visible or is coalesced with another point. The *effective length* $\overline{\ell}_q$ of the segment assigned to a visible point q, is the sum of its length and the lengths of the segments coalesced with it. If C_i denotes the set of points coalesced with p_i,

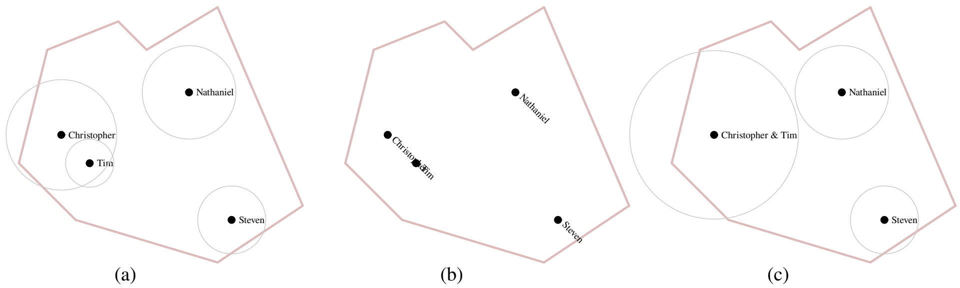

Figure 6: Merging labels of intersecting labels in a rotating map

we have

$$\overline{\ell}_{p_i} = \ell_{p_i} + \sum_{p_j \in C_i} \ell_{p_j}.$$

Point p_j may be coalesced with point p_i only if the distance between p_i and p_j is no greater than the effective length of the segment assigned to p_i after coalescing the points which are closer than p_j. In other words, the following condition is necessary for coalescing p_j with p_i:

$$d(p_i, p_j) \leq \ell_{p_i} + \sum_{p_k \in C_i \text{ and } d(p_i,p_k)<d(p_i,p_j)} \ell_{p_k}$$

In a proper labelling ϕ, none of the visible labels must intersect during the rotation of the plane. In MRIS with coalescing (MRISC), the goal is to find a proper labelling that minimises the maximum length of segments assigned to the points.

The decision version of MRISC with parameter x for a set of points P asks if there is a proper labelling, in which the maximum effective length of the segments assigned to points in P is bounded by x. Theorem 4, shows that MRISC is NP-hard in the 1P model, even if all input points are on a line.

Theorem 4 *MRISC is NP-hard in the 1P model, even if all input points are on a line.*

Proof: We reduce the Partition problem (a special case of Subset Sum) to MRISC. Given a multiset of numbers S, the goal in the Partition problem is to decide whether it is possible to partition S into subsets S_1 and S_2 such that the sums of the elements of these subsets are equal.

Let the multiset S be an instance of Partition of size n and sum w. Let $s_1, s_2, \ldots, s_n$ be the elements of S. We create an instance of MRISC as follows. For each element s_i of S, we place a point p_i the x-axis. Let p_1 be the origin. For $i \in \{1, \ldots, n-1\}$, p_{i+1} is placed to the right of p_i on the x-axis such that the distance between p_i and p_{i+1} is $s_i + s_{i+1}$. This is demonstrated for a multiset with elements 2, 1, 3, and 1 in Figure 7. Let ℓ_{p_i} be s_i for $1 \leq i \leq n$. We add points p_{n+1} and p_{n+2} at $-10w$ and $10w$, respectively. Let $\ell_{p_{n+1}} = \ell_{p_{n+2}} = 15w$. We show that we can obtain a solution to Partition from a solution to MRISC for $P = \{p_1, p_2, \ldots, p_{n+2}\}$ with $x = 15.5w$ as upper bound for the maximum effective segment length.

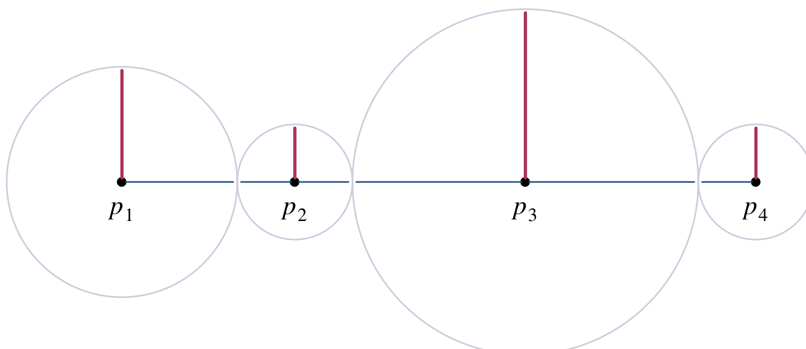

Figure 7: The placement of points in Theorem 4 for an example multiset

Point p_{n+1} cannot be coalesced with p_{n+2} because the maximum effective length of p_{n+2} is $16w$ (if all points except p_{n+1} are coalesced with p_{n+2}), whereas the distance between p_{n+1} and p_{n+2} is $20w$. Similarly, p_{n+2} cannot be coalesced with p_{n+1}. Neither can any two points p_i and p_j with $1 \leq i < j \leq n$ be coalesced with each other. This can be shown by contradiction, considering the vertex in C_i closest to p_i and showing that it cannot be coalesced with p_i.

Therefore, p_{n+1} and p_{n+2} are always visible, and the only possible coalescings are p_j with p_i for $i \in \{n+1, n+2\}$ and $j \in \{1, 2, \ldots, n\}$. On the other hand, p_j for $j \in \{1, 2, \ldots, n\}$ cannot be visible, since their segments intersect with those of p_{n+1} and p_{n+2} during rotation. Consequently, $\{p_1, p_2, \ldots, p_n\}$ needs to be partitioned into subsets C_{n+1} and C_{n+2} such that the maximum of $\ell_{p_{n+1}} + \sum_{p_j \in C_{n+1}} \ell_{p_j}$ and $\ell_{p_{n+2}} + \sum_{p_j \in C_{n+2}} \ell_{p_j}$ is $15.5w$, or equivalently, the sum of the lengths the elements of C_{n+1} and of those in C_{n+2} is $w/2$. From this partition we can obtain the subsets S_1 and S_2 for Partition. □

We now present an ILP formulation for MRISC.

Theorem 5 *There is an integer linear programme for solving MRISC for a set of n points, with $O(n^2)$ variables and $O(n^2)$ constraints.*

Proof: Let $P = \{p_1, \ldots, p_n\}$ be the given set of points. For $i \in \{1, \ldots, p_n\}$, we introduce a binary variable x_i that shows whether p_i is visible, and for $i, j \in \{p_1, \ldots, p_n\}$, we introduce a binary variable y_{ij} that shows whether p_i is coalesced with p_j:

$$\begin{array}{rr} \forall i \in \{1, 2, \ldots, n\} & x_i \in \{0, 1\} \\ \forall i, j \in \{1, 2, \ldots, n\} \quad i \neq j & y_{ij} \in \{0, 1\} \end{array}$$

Since each point can be coalesced with at most one point, we also have the following constraints:

$$\sum_{1 \leq j \leq n} y_{ij} \leq 1, \qquad \forall i \in \{1, 2, \ldots, n\}$$

A point that is coalesced with another point is not visible:

$$x_i = 1 - \sum_{1 \leq j \leq n} y_{ij}, \qquad \forall i \in \{1, 2, \ldots, n\}$$

We introduce an auxiliary variable $\overline{\ell}_{p_i}$ to express the effective length of the segment assigned to p_i. We have:

$$\overline{\ell}_{p_i} = x_i \ell_{p_i} + \sum_{1 \leq j \leq n} y_{ji} \ell_{p_j}, \qquad \forall i \in \{1, 2, \ldots, n\}$$

A point can be coalesced with another only if the distance between them is bounded as follows:

$$y_{ji} d(p_i, p_j) \leq \ell_{p_i} + \sum_{d(p_i, p_k) < d(p_i, p_j)} y_{ki} \ell_{p_k}, \qquad \forall i, j \in \{1, 2, \ldots, n\}$$

The objective is to minimize an auxiliary variable z that meets the following constraints:

$$z \geq \overline{\ell}_{p_i}, \qquad \forall i \in \{1, 2, \ldots, n\}$$

$\square$

4.2 Restricted Coalescing

Let $P = \{p_1, p_2, \ldots, p_n\}$ be a set of points and let ℓ_{p_i} be the length of the vertical segment assigned to p_i. Each point in P is either a *vital* point or an *optional* point. Let P_V and P_O denote the sets of vital and optional points, respectively; we have $P = P_V \cup P_O$. In a labelling, every vital point must be visible. Any optional point p may be visible or may be coalesced with a pre-specified vital point $vital(p)$. If p_i is coalesced with p_j, the length of the segment assigned to p_j increases by ℓ_{p_i}. In a proper labelling ϕ, none of the visible labels should intersect during the rotation of the plane. In MRIS with restricted coalescing (MRISRC), the goal is to find a proper labelling of maximum size.

The decision version of MRISRC asks if there is a proper labelling of size k. We now show that MRISRC in the 1P model is NP-hard.

Theorem 6 *MRISRC is NP-hard in the 1P model.*

Proof: We reduce from GMIS for unit disks to MRISRC. Let $D = \{d_1, d_2, \ldots, d_n\}$ be a set of n unit disks on the plane. Define a set of points $P = \{p_0, p_1, p_2, \ldots, p_n\}$ such that p_i is the centre of d_i for $1 \leq i \leq n$, and p_0 is a point of distance at least $2n$ from any disk in D. Also, let $\ell_{p_i} = 2$ for $0 \leq i \leq n$, $P_V = \{p_0\}$, $P_O = \{p_1, p_2, \ldots, p_n\}$, and $vital(p_i) = p_0$ for $0 \leq i \leq n$. We show that any independent subset of D of size k corresponds to a proper labelling of size $k+1$ for P, and vice versa. This would imply that there is an independent subset of D of size k if and only if there is a proper labelling of P of size $k+1$.

Consider any independent subset D' of D of size k. Let P' be the union of $\{p_0\}$ and the centres of the disks in D'. Since D' is independent, the distance between any pair of points in P' is at least 2, implying that the disks that result from rotating the labels of two elements of P' are centre-disjoint. Since $n-k$ points are coalesced with p_0, ℓ_{p_0} is $2(n-k+1)$. The distance of p_0 to any other point in P is at least $2n$. Thus, the rotation disk of p_0 is disjoint from that of every other element of P'. Therefore, a labelling of size $k+1$ containing P' is a proper labelling.

For the converse, let P' be the set of points in a proper labelling of P of size $k+1$. Let D' be the subset of D that consists of the disks whose centres are in P'. Since the rotation disks of the elements of P' are centre disjoint, the distance between any pair of points in P' is at least two. Therefore, D' is independent. $\square$

Theorem 6 also shows that even if at most one optional point is assigned to each vital point, the problem remains NP-hard. We now show that MRISRC can be formulated as an ILP.

Theorem 7 *There is an integer linear programme for solving MRISRC for a set of n points, with $O(n)$ variables and $O(n^2)$ constraints.*

Proof: Let $P = \{p_1, \ldots, p_n\}$ be the given set of points. For $i \in \{1, \ldots, n\}$, we introduce two binary variables: x_i that shows the visibility of point p_i, and $\overline{x}_i$ that shows its exclusion from the labelling. We have $\overline{x}_i = 1 - x_i$. For any vital point p_i, $x_i = 1$.

$$\begin{aligned} & x_i \in \{0, 1\}, && \forall i \in \{1, 2, \ldots, n\} \\ & \overline{x}_i = 1 - x_i, && \forall i \in \{1, 2, \ldots, n\} \\ & x_i = 1, && \forall p_i \in P_{\mathrm{V}} \end{aligned}$$

The auxiliary variable $\overline{\ell}_{p_i}$ indicates the effective length of the segment assigned to p_i. The effective length of each segment can be described using the coalescing relationships between the points as follows.

$$\begin{aligned} & \overline{\ell}_{p_i} = \ell_{p_i}, && \forall p_i \in P_{\mathrm{O}} \\ & \overline{\ell}_{p_i} = \ell_{p_i} + \sum_{vital(p_j) = p_i} x_j \ell_{p_j}, && \forall p_i \in P_{\mathrm{V}} \end{aligned}$$

To prevent the segments assigned to visible points from intersecting, we have the following constraint:

$$\overline{\ell}_{p_i} > (1 - \overline{x}_i - \overline{x}_j) \cdot d(p_i, p_j), \qquad \forall i, j \in \{1, 2, \ldots, n\}$$

where, $d(p_i, p_j)$ is the Euclidean distance between points p_i and p_j.

The objective is to maximize

$$\sum_{1 \le i \le n} x_i,$$

subject to the above constraints.

□

5 Concluding Remarks

We have studied the problem of labelling rotating maps. We have proved the problem to be NP-hard even for vertical labels of zero width, and modified a PTAS presented by Chan [5] for geometric maximum independent set to solve this problem. We also have studied a variant of this problem for vertical, zero-width labels, in which the labels that do not appear in the labelling must be coalesced with a label that does, and we have have presented ILPs for solving them.

Several problem remain for further investigation such as the following.

- Theorem 4 proves that MRISC is NP-hard. Is MRISC strongly NP-hard?
- Are approximation and randomized algorithms possible for MRISC and MRISRC?
- Other coalescing models are interesting to investigate. They may be more suitable for specific applications.

Acknowledgments

We are grateful to the anonymous reviewers for many suggestions to improve this paper.

References

[1] P. K. Agarwal, M. J. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Computational Geometry*, 11(3–4):209–218, 1998. doi:10.1016/S0925-7721(98)00028-5.

[2] K. Been, E. Daiches, and C.-K. Yap. Dynamic map labeling. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):773–780, 2006. doi:10.1109/TVCG.2006.136.

[3] K. Been, M. Nöllenburg, S.-H. Poon, and A. Wolff. Optimizing active ranges for consistent dynamic map labeling. *Computational Geometry*, 43(3):312–328, 2010. doi:10.1016/J.COMGEO.2009.03.006.

[4] R. G. Cano, C. C. de Souza, and P. J. de Rezende. Fast optimal labelings for rotating maps. In Sheung-Hung Poon, Md. Saidur Rahman, and Hsu-Chun Yen, editors, *11th International Conference and Workshops, WALCOM 2017*, volume 10167 of *Lecture Notes in Computer Science*, pages 161–173. Springer, 2017. doi:10.1007/978-3-319-53925-6_13.

[5] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *Journal of Algorithms*, 46(2):178–189, 2003. doi:10.1016/S0196-6774(02)00294-8.

[6] T. M. Chan and S. Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 2012. doi:10.1007/s00454-012-9417-5.

[7] T. Erlebach, T. Hagerup, K. Jansen, M. Minzlaff, and A. Wolff. Trimming of graphs, with application to point labeling. *Theory of Computing Systems*, 47(3):613–636, 2010. doi:10.1007/S00224-009-9184-8.

[8] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM Journal on Computing*, 34(6):1302–1323, 2005. doi:10.1137/S0097539702402676.

[9] M. Formann and F. Wagner. A packing problem with applications to lettering of maps. In Robert L. Scot Drysdale, editor, *Seventh Annual Symposium on Computational Geometry, SoCG 1991*, pages 281–288. ACM, 1991. doi:10.1145/109648.109680.

[10] R. J. Fowler, M. Paterson, and S. L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Information Processing Letters*, 12(3):133–137, 1981. doi:10.1016/0020-0190(81)90111-3.

[11] A. Gemsa, B. Niedermann, and M. Nöllenburg. Trajectory-based dynamic map labeling. In Leizhen Cai, Siu-Wing Cheng, and Tak Wah Lam, editors, *24th International Symposium on Algorithms and Computation, ISAAC 2013*, volume 8283 of *Lecture Notes in Computer Science*, pages 413–423. Springer, 2013. doi:10.1007/978-3-642-45030-3_39.

[12] A. Gemsa, M. Nöllenburg, and I. Rutter. Consistent labeling of rotating maps. *Computational Geometry*, 7(1):308–331, 2011. doi:10.20382/jocg.v7i1a15.

[13] A. Gemsa, M. Nöllenburg, and I. Rutter. Evaluation of labeling strategies for rotating maps. *ACM Journal of Experimental Algorithmics*, 21(1):1.4:1–1.4:21, 2016. doi:10.1145/2851493.

[14] D. S. Hochbaum and W. Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *Journal of the ACM*, 32(1):130–136, 1985. doi:10.1145/2455.214106.

[15] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182(1):105–142, 1999. doi:10.1007/BF02392825.

[16] W. D. Smith and N. C. Wormald. Geometric separator theorems and applications. In *39th Annual Symposium on Foundations of Computer Science, FOCS 1998*, pages 232–243. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743449.

[17] M. J. van Kreveld, T. Strijk, and A. Wolff. Point labeling with sliding labels. *Computational Geometry*, 13(1):21–47, 1999. doi:10.1016/S0925-7721(99)00005-X.

[18] Y. Yokosuka and K. Imai. Polynomial time algorithms for label size maximization on rotating maps. *Journal of Information Processing*, 25:572–579, 2017. doi:10.2197/ipsjjip.25.572.