Boundary Labeling in a Circular Orbit

Annika Bonerath
University of Bonn, Germany

Martin Nöllenburg
Algorithms and Complexity Group, TU Wien, Austria

Soeren Terziadis
TU Eindhoven, The Netherlands

Markus Wallinger
Chair for Efficient Algorithms, TU Munich, Germany

Jules Wulms
TU Eindhoven, The Netherlands

Abstract

Boundary labeling is a well-known method for displaying short textual labels for a set of point features in a figure alongside the boundary of that figure. Labels and their corresponding points are connected via crossing-free leaders. We propose *orbital boundary labeling* as a new variant of the problem, in which (i) the figure is enclosed by a circular contour and (ii) the labels are placed as disjoint circular arcs in an annulus-shaped orbit around the contour. The algorithmic objective is to compute an orbital boundary labeling with the minimum total leader length. We identify several parameters that define the corresponding problem space: two leader types (straight or orbital-radial), label size and order, presence of candidate label positions, and constraints on where a leader attaches to its label. Our results provide polynomial-time algorithms for many variants and NP-hardness for others, using a variety of geometric and combinatorial insights.

2012 ACM Subject Classification Human-centered computing → Visualization; Theory of computation → Computational geometry

Keywords and phrases External labeling, Orthoradial drawing, NP-hardness, Polynomial algorithms

Digital Object Identifier 10.4230/LIPIcs.GD.2024.22

Related Version A preliminary abstract of this paper has been presented at EuroCG 2024.
Full Version: `https://arxiv.org/abs/2403.19052` [8]

Funding *Annika Bonerath*: German Research Foundation under Germany's Excellence Strategy, EXC-2070 – 390732324 – PhenoRob.
Martin Nöllenburg: Vienna Science and Technology Fund (WWTF) [10.47379/ICT19035].
Soeren Terziadis: Vienna Science and Technology Fund (WWTF) [10.47379/ICT19035] and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101034253.
Markus Wallinger: Vienna Science and Technology Fund (WWTF) [10.47379/ICT19035].

1 Introduction

Labeling spatial data on a map is a well-studied topic in computational geometry [1, 10, 17]. Commonly the feature points are annotated with labels that display the names or additional descriptions, ensuring non-overlapping labels to guarantee full readability. The labels are placed either next to the feature points [16] (*internal* label positions), or remotely along the contour of a bounding shape such that the feature points are connected to their labels by crossing-free leaders (*external* labeling models) [5]. Often, for high feature-point densities, the external labeling model is more advantageous since the background map is not obscured by

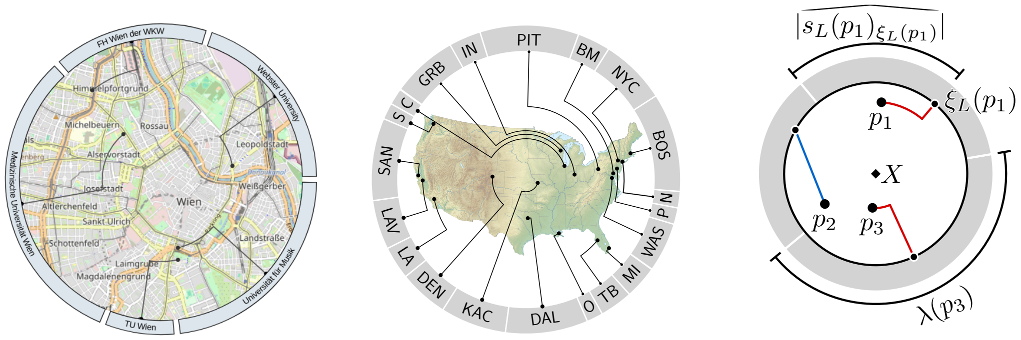

Figure 1 An orbital labeling highlighting points of interest in a map section (left) and an orbital labeling of NFL-Championship winning teams, where the label sizes are scaled with the number of won titles (center). Our notation (right) with an SL-leader in blue and OR-leaders in red.

annotations. A special case of external labeling is boundary labeling [3,4] where the labels are attached to the (mostly rectangular) boundary of the map. The interest in visualizing data on round displays, e.g., on smartwatch faces (see Figure 1) or on round displays in cockpits, is growing, as discussed by Islam et al. [14] in their recent survey; but, from a visualization perspective, it is still an under-explored topic compared to traditional rectangular displays. The design space description of Islam et al. [14] includes geospatial data representations as well as placement of text labels and icons on a round display. With this in mind, we initiate the investigation of boundary labeling for maps with circular boundaries, surrounded by a peripheral fixed-width ring, in which the labels are placed. We call these labels *orbital*. We assume that the lengths of the orbital labels are normalized, s.t., they sum up to at most the perimeter of the boundary of the map. Orbital labels can also be used for other purposes, e.g., to display donut charts representing statistical data values of different feature points within the map, such that the label sizes are proportional to the data values (Figure 1). Previous research on circular map displays considered either multirow circular labels where the sum of label lengths does not equate to the map's boundary length [12], radial labels [2,9], or horizontal labels [9,13,15]. The latter two settings are relevant on rectangular displays but not suitable for circular displays with a narrow annulus reserved for labels. Furthermore, these settings differ in their geometric properties and hence their labeling algorithms do not immediately generalize to orbital labels.

Formally, we assume that we are given a disk D in the plane $\mathbb{R}^2$ centered at a point X. The disk contains n points $P = \{p_1, \ldots, p_n\}$. We call the set P of points *features* and we refer to the boundary of the disk as the *boundary* B. The feature closest to X is denoted by $p_{\min}$. We may assume that D has a radius of 1. Throughout this paper, unless otherwise specified, the *angle between* two points p, q refers to the smallest angle with the center X of D, i.e., $\min(\measuredangle pXq, \measuredangle qXp)$. Every feature $p \in P$ has an associated *label* representing additional information to be placed along a circular arc on the boundary starting at a point $b_1 \in B$ and ending at a point $b_2 \in B$. The circular arc along B is denoted as $\widehat{b_1 b_2}$. Usually, the start and endpoint of the label are not fixed in the input, however, the length of the arc is part of the input. We represent the associated label simply as a number $\lambda(p)$, which indicates the length of the associated label. We assume that $\sum_{i=1}^{n} \lambda(p_i)$ is equal to the circumference of D, i.e., all labels can be placed in a non-overlapping way without gaps between the arcs on B.

In a *labeling* L, every feature $p \in P$ is assigned a label with starting point $s_L(p) \in B$ and an endpoint $e_L(p) \in B$, s.t., $|\widehat{s_L(p)e_L(p)}| = \lambda(p)$. We require that all labels in L are pairwise non-overlapping. Additionally, every feature p is connected to its label via a curve called a *leader*. We denote the length of a leader ψ using $|\psi|$. We call the point on a label arc where the leader connects to B the *port* $\xi_L(p)$. We represent a port by the *port ratio* $\rho(p) = \frac{|\widehat{s_L(p)\xi_L(p)}|}{\lambda(p)}$ which is the ratio of the arc from the starting point $s_L(p)$ to the port $\xi_L(p)$ and the arc from the start-point to the end-point. $\varepsilon \leq |\widehat{s_L(p)\xi_L(p)}|$ and $\varepsilon \leq |\widehat{\xi_L(p)e_L(p)}|$, i.e., there is always at least a small distance between the ends and the port of the label. A distance of ε between two points on B should be interpreted as a distance along B. Now, we define the generic orbital labeling problem.

▶ **Problem 1** (ORBITAL BOUNDARY LABELING). *Given a disk D, containing n feature points P with their labels $\lambda(p)$ for $p \in P$, compute a labeling L, in which all leaders are pairwise interior-disjoint and where the sum of leader lengths is minimal.*[1]

A labeling in which the sum over all leader lengths is minimum is also called a *leader length minimal labeling*. We consider two *leader types* in this paper. A *straight-line* leader or SL-*leader* is simply a straight-line segment starting at p and ending at $\xi_L(p)$. Its *length* is the Euclidean distance between p and $\xi_L(p)$. An *orbital-radial* leader or OR-*leader* consists of two parts: a (possibly empty) orbital circular arc with center point X starting at the feature p and ending at a bend point q, and a radial segment that connects q to $\xi_L(p)$; see Figure 1. We call the line through X and $\xi_L(p)$ the *supporting line* of the radial part. Note that for any pair of feature and port, there are exactly two possible OR-leaders. We call an OR-leader leaving its feature in clockwise direction a *clockwise leader* and analogously define *counter-clockwise leaders*. We will also refer to the OR-leader whose orbital part spans an angle larger than π as the feature's *long* and to the other one as its *short* leader. For a feature p_i let r_i be the radius of the circle concentric with D containing p_i. The length of the OR-leader can be expressed as a function $g : P \times [0, 2\pi] \to \mathbb{R}$ with $g(p_i, \theta) = r_i\theta + (1 - r_i)$, where $0 \leq \theta < 2\pi$ describes the angle spanned by the orbital part of a leader connected to p_i.

▶ **Observation 2.** *For a fixed feature p, the function $g(p, \theta)$ is continuous and linear in θ.*

Based on this problem description, we delineate the space of the possible problem variants and a suitable naming scheme for such variants in the following section.

2 Problem Space

In this paper, we refer to variants of ORBITAL BOUNDARY LABELING based on the six-dimensional T-$COSA$-ORBITAL BOUNDARY LABELING scheme introduced in this section. The first dimension, denoted by T, determines the leader type, whereas the other dimensions, denoted by variables $COSA$, characterize properties of the labels (A encodes two dimensions). We use OR and SL as substitutes for T in the T-$COSA$ scheme. Without the T- prefix, we refer to both OR-$COSA$ and SL-$COSA$ (leader types are not mixed in a labeling). We mostly focus on OR-leaders, while still discussing which of the results extend to SL-leaders.

For the five dimensions regarding the labels, we use each letter of $COSA$ to describe the variants for the respective dimension.

- **[C] Candidate port positions on the boundary.** If we are given a set C of candidate positions on B, we require in any valid labeling L that the set Ξ_L of all ports in L is a subset of C, we say *the port candidates are locked* (and use the symbol $C^{🔒}$) otherwise they are *free* ($C^{🔑}$). For variants with $C^{🔒}$ we assume that sufficiently many, but no more than linearly many candidates are specified ($n \leq |C|$). Otherwise, $C^{🔑}$ is the more reasonable choice.

[1] This definition does allow a leader to contain another feature, i.e., the endpoint of another leader.

Table 1 An overview of the problem space and our results. Only **locked** port ratios ($A^{🔒}$) are shown. The abbreviation "w. NP-h" denotes weakly NP-hard. Red cells represent polytime, due to the reduction to the algorithm of Benkert et al. [7]. Light red cells are covered by the reduction but are superseded by faster dedicated approaches. Blue cells with a question mark are conjectures. Cells indicate the section containing the relevant result; SL-leader results are in the full version [8].

			$C^{🔒}$		$C^{📌}$	
			$A^{🔒}_{\equiv}$	$A^{🔒}_{\not\equiv}$	$A^{🔒}_{\equiv}$	$A^{🔒}_{\not\equiv}$
OR	$O^{🔒}$	$S_{\equiv}$	$O(\|C\|n^2)$ [S. 4.1]	$O(\|C\|n^2)$ [S. 4.1]	$O(n^2)$ [S. 4.3]	$O(n^2)$ [S. 4.3]
		$S_{\not\equiv}$	$O(\|C\|n^2)$ [S. 4.1]	$O(\|C\|n^2)$ [S. 4.1]	$O(n^2)$ [S. 4.3]	$O(n^2)$ [S. 4.3]
	$O^{📌}$	$S_{\equiv}$	$O(\|C\|^2n)$ [S. 4.2]	$O(\|C\|n^3)$ [S. 3]	$O(n^5)$ [S. 3]	$O(n^5)$ [S. 3]
		$S_{\not\equiv}$	$O(\|C\|^4)$ [S. 3]	$O(\|C\|^4)$ [S. 3]	w. NP-h [S. 5]	w. NP-h [S. 5]
SL	$O^{🔒}$	$S_{\equiv}$	$O(\|C\|n^2)$ [8, B.1]	$O(\|C\|n^2)$ [8, B.1]]	$O(n^2)$? [8, B.3]	$O(n^2)$? [8, B.3]
		$S_{\not\equiv}$	$O(\|C\|n^2)$ [8, B.1]	$O(\|C\|n^2)$ [8, B.1]	$O(n^2)$? [8, B.3]	$O(n^2)$? [8, B.3]
	$O^{📌}$	$S_{\equiv}$	$O(\|C\|^2n)$ [8, B.2]			
		$S_{\not\equiv}$			w. NP-h [8, B.4]	w. NP-h [8, B.4]

- **[O] Order.** Consider the cyclic order of labels around B. If a certain label order is pre-specified, we say *the label order is locked* ($O^{🔒}$); otherwise, for the unconstrained setting, we say *the label order is free* ($O^{📌}$).
- **[S] Size of labels.** We distinguish the setting where $\lambda(p) = \frac{2\pi}{n}$ for all $p \in P$, in which case we say that *the label size is uniform* ($S_{\equiv}$), otherwise *the label size is non-uniform* ($S_{\not\equiv}$).
- **[A] Port position on labels.** We differentiate between *uniform port ratios*, where $\rho(p) = \rho(q)$ for all $p, q \in P$ ($A_{\equiv}$), and *non-uniform port ratios* ($A_{\not\equiv}$). Additionally the port ratios can be *fixed* as part of the input ($A^{🔒}$) or can be *free* to be chosen ($A^{📌}$).

To specify a T-$COSA$ variant we substitute C, O, S and A with $C^{📌}/C^{🔒}$, $O^{📌}/O^{🔒}$, $S_{\equiv}/S_{\not\equiv}$ and $A^{📌}_{\equiv}/A^{📌}_{\not\equiv}/A^{🔒}_{\equiv}/A^{🔒}_{\not\equiv}$, respectively. Whenever a statement applies to all variants along a $COSA$ dimension, we drop the sub- or superscript. For example, $C^{📌}O^{🔒}SA^{🔒}_{\equiv}$ refers to the variants where the leader style is either OR or SL, the port candidates are free ($C^{📌}$), the order is locked ($O^{🔒}$), the label sizes could be fixed to be uniform or they could be non-uniform (S) and all port ratios are fixed ($A^{🔒}$) to the same ($A_{\equiv}$) given value.

Contributions. The remainder of the paper is structured as follows (see also Table 1). Our focus is OR-$COSA^{🔒}$ and we show how results for these variants extend to SL-$COSA^{🔒}$ and eventually to all $COSA$ variants. Section 3 presents a reduction from OR-$COSA$ to a problem called BOUNDARY LABELING [7], for which polynomial time algorithms are known. This approach is applicable to a number of OR-$COSA$ variants. In Section 4 we introduce dedicated approaches for some variants, which improve the runtime of the reduction approach. In Section 5 we prove that OR-$C^{📌}O^{📌}S_{\not\equiv}A^{🔒}$ variants are weakly NP-hard. Section 6 outlines the extensions to straight-line leaders and free port ratios (details in Appendices B and C of the full version [8], respectively), while Section 7 provides some concluding remarks.

3 Reduction to Boundary Labeling

We begin by investigating the relation between ORBITAL BOUNDARY LABELING and an already established related problem called BOUNDARY LABELING. In an instance of BOUNDARY LABELING, we are given a set of points which are entirely left of a vertical line ℓ. Additionally, we have a set of n disjoint rectangular labels, which are placed to the right of ℓ. The goal is

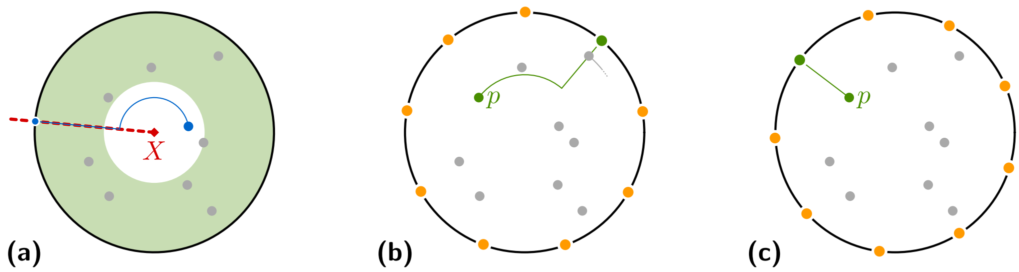

Figure 2 All leaders of gray points in **(a)** are entirely contained within the green annulus. The radial part of the blue leader spans the entire intersection between its supporting line and the green annulus. In the labeling shown in **(b)**, a clockwise rotation of the ports is not possible without introducing a crossing between the green leader and the one in gray (indicated with a dashed line). In **(c)** any rotation clockwise or counterclockwise increases the length of the green leader. Both **(b)** and **(c)** show the implied ports in yellow.

to find a set of parallel-orthogonal (po) leaders, which consist of a line segment parallel to ℓ starting at a feature point and a second line segment orthogonal to ℓ connecting to a port on a label. If the label sizes are interchangeable, i.e., every point can be connected to any label, there is an $O(n^3)$ dynamic program [6, 7], which minimizes an arbitrary badness function for a leader (under the assumption that the function can be computed in linear time). The dynamic program, which first splits the instance into horizontal strips, defined by points and the boundaries of the rectangles, then assigns a specific horizontal strip to the feature farthest from ℓ and then recurses on the two sub-instances defined by this assignment. One instance contains all features and rectangles above the assigned horizontal strip, the second everything below.

To create an instance of Boundary Labeling from an instance of Orbital Boundary Labeling, we consider D as an annulus (by introducing a small circle at the center, which does not contain any features), find a line segment s orthogonal to B from the boundary of the small circle to a point on B and transform the annulus into a rectangle, turning B into a straight line. The basis for this reduction is the following observation, which – intuitively speaking – prove that there exists a radial line, which we can use to cut and unroll our instance to remove the cyclic nature of Orbital Boundary Labeling.

▶ **Observation 3.** *In a crossing-free leader-length minimal labeling for problems in* OR-$COSA$*, the supporting line of the radial part for the leader of* p_{min} *intersect no other leader.*

Observation 3 is also illustrated in Figure 2a. Based on this observation we know that in any crossing free labeling, there is always a radial line, which does not intersect a leader other than the one of $p_{\min}$. Moreover it is sufficient to check all possibilities for the port of $p_{\min}$ to obtain such a line. Next, we need to show that there is only a polynomial number of possibilities for the port of $p_{\min}$. The exact number depends on the variant we are considering. If we are considering a labeling for a problem in the set OR-$C^{🔒}OSA$, we are given the set C of possible ports, leading to the following observation.

▶ **Observation 4.** *For problems in* $C^{🔒}OSA$ *there are only* $|C|$ *possibilities for* $\xi_L(p_{min})$*.*

It is less obvious if and how we can discretize these options for problems without a fixed candidate set for the ports. If we are considering a problem in the set $C^{✎}O^{🔒}SA$, we can reduce the number of relevant options using the following lemma.

▶ **Lemma 5.** *For problems in* OR-$C^{📍}O^{🔒}SA$ *there are only* n^2 *possibilities for* $\xi_L(p_{min})$.

Proof. Let S be the set of the n intersection points between B and any ray starting at X through a feature. We prove that there is a leader-length minimal labeling, in which at least one port is in S. Assume that L is a leader-length minimal labeling, where no port is in S. Consider a small rotation of all ports clockwise. If such a rotation is not possible without introducing a crossing, the radial segment of a clockwise leader already contains another feature, and therefore its port was in S (see Figure 2b). Otherwise, a small enough rotation neither changes the order of labels nor does it introduce any intersections and therefore results in a new valid labeling. If this rotation decreases the total leader length, then L was not optimal, which contradicts our assumption. If the total leader length stays the same, we can continue the rotation until either the orbital segment of a counter-clockwise leader reaches length 0 or the radial segment of a clockwise leader hits another feature. In both cases, its port is in S. If the leader length increases, we instead rotate counter-clockwise. Again, if the total leader length decreases or stays the same, the arguments above apply. Assume therefore that the total leader length again increases. Since by Observation 2 the change in leader length is linear in the angle by which we rotate, there must be a single leader that increases its length in both rotation directions, which implies that its orbital part has length 0 in L (see Figure 2c).

Therefore in any optimal labeling, at least one port is in S. Since the order is fixed, choosing a specific label to have its port at an element in S fixes all other ports and specifically the location of the port of $p_{\min}$. Therefore every element of S induces n possible choices for this port which results in n^2 choices in total. ◀

In general, the previous method does not extend to the problems in OR-$C^{📍}O^{📍}SA$, since choosing one element in S still leaves $(n-1)!$ possible orders of the remaining labels. However, in the special case where both the label sizes and the port ratios are uniform, we know that all ports in any valid labeling are distributed equally along B, which implies that again every element in S induces exactly n possible labels (even if we do not know which label is connected to which port). Therefore, we state the following observation.

▶ **Observation 6.** *For problems in* OR-$C^{📍}O^{📍}S_{≡}A_{≡}$ *there are only* n^2 *possibilities for* $\xi_L(p_{min})$.

We can now proceed to describe how we create an instance of Boundary Labeling based on a given instance of Orbital Boundary Labeling.

▶ **Lemma 7.** *Given a port* ξ *for* p_{min}*, we can reduce any problem in* OR-$CO^{🔒}SA^{🔒}$ *to* Boundary Labeling *with po-leaders [7].*

Proof. We first map all features to points in the plane (also shown in Figure 3). Let X be the center of D. For any feature $p_i \in P$, let α_i be the angle between ξ and p_i (recall that this is defined as the smaller angle formed between these two points at X) and let $r_i = |Xp_i|$. Using polar coordinates we then create a point $q_i = (r_i, \alpha_i)$ for every point $p_i \in P$. We now place a vertical line ℓ at $x = 1$. Recall that the radius of D is 1 and therefore all points $q_1, \ldots, q_n$ are left of ℓ. Since the port ratio is locked as part of the input, the fixed position of ξ also fixes the exact position of the label of $p_{\min}$. If the problem is in the set OR-$CO^{🔒}SA^{🔒}$, the order of labels is fixed, and therefore ξ also fixes the position of all other labels. For every point p_i we place a rectangle of height $\lambda(p_i) - 2\varepsilon$ in the order of the labels (recall that any port has to have a minimal distance of ε to the boundary of the label), s.t., the lowest point of the lowest rectangle is at height $|\widehat{s_L(p_{\min})\xi}| + \varepsilon$.

(a) Instance of ORBITAL BOUNDARY LABELING. **(b)** Instance of BOUNDARY LABELING.

Figure 3 An instance **(a)** of ORBITAL BOUNDARY LABELING together with a fixed port for $p_{\min}$. We use an appropriate mapping to construct an instance of BOUNDARY LABELING, whose solution **(b)** corresponds to the solution of ORBITAL BOUNDARY LABELING. This maps the angle σ between two features in **(a)** to their horizontal distance between the mapped features in **(b)** and the distance ρ of a feature to B in **(a)** to the distance of the mapped feature to ℓ in **(b)**. Note that only the remaining non-fixed labels are part of the new instance.

Now, the length of the radial segment of an orbital-radial leader connecting p_i to a point $b \in B$ in our problem is equal to the length of the orthogonal linepart of the *po*-leader connecting q_i to a point v on ℓ. The relation between the length of the parallel part of the *po*-leader and the length of the orbital segment of the OR-leader is more complicated. The length of the parallel part of the *po*-leader is simply the difference in y-coordinate between q_i and v. Note that we mapped the clockwise angle of a point q_i relative to ξ to the y-coordinate of q_i. However, the length of the orbital part of the OR-leader is dependent on the distance of p_i to X, i.e., two OR-leaders whose orbital part span the same angle can have different lengths. Specifically, the length of an orbital part in an OR-leader of a feature p_i is exactly $r_i \cdot \alpha_i$. Therefore, we define our badness function simply as $\mathrm{bad}(q_i, v) = 1 - r_i + r_i(|\alpha_i - y(q_i)|)$, where $y(q)$ is the y-coordinate of q_i. Note that the restriction of port placement on ℓ to a specific range (e.g., a point in a candidate set C or a point within a label corresponding to a fixed port ratio) can be encoded in this badness function too, by setting the value of $\mathrm{bad}(q_i, v) = \infty$ if v lies outside of the permitted range. Finally, also note that we can check in $O(n)$ time if the split into sub-instances induced by the combination of q and v respects the fixed order of the labels and, if it does not, also set $\mathrm{bad}(q_i, v) = \infty$. This completes the reduction. ◀

If we do not have a fixed order of labels, we have to pay attention to the order in which the rectangles are placed. However, if the labels have uniform size, any order will result in the same set of rectangles resulting in the following observation.

▶ **Observation 8.** *Given a port ξ for p_{min}, we can reduce any problem in* OR-$COSA$ *to po-*BOUNDARY LABELING *[7].*

Lastly, if the label sizes are not uniform, and if a candidate set C is fixed, we can adapt the dynamic program slightly to leverage this fact, by placing smaller rectangles at every possible candidate position and avoid overlap by restricting the sub-instances to appropriate ranges. The proof of the following lemma can be found in Appendix A of the full version [8].

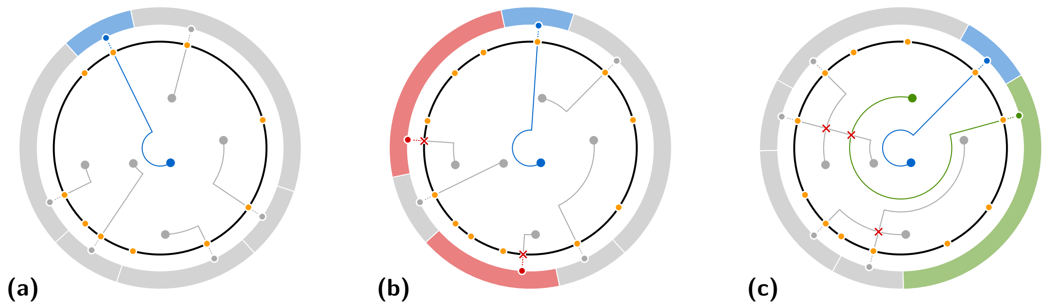

Figure 4 Three rotations for case OR-C🔒O🔒SA🔒. Candidates are shown in yellow. The blue feature is the first that is placed and we iteratively test every port candidate. Due to the fixed order, the other leaders are directly obtained. In a valid labeling **(a)** all ports (obtained due to a fixed port ratio) coincide with a candidate and no two leaders cross. A labeling is invalid if ports do not coincide with a candidate – highlighted red in **(b)** – or the obtained leaders contain crossings between themselves, which is shown with the red crosses in **(c)**. In **(c)** the green leader changed from a clockwise to a counter-clockwise leader to avoid crossing the blue leader.

▶ **Lemma 9.** *Given a port ξ for p_{min}, any problem in* OR-C🔒O📌$S_{\equiv}A$🔒 *can be solved in $O(|C|^3)$ time by adapting the dynamic program of Benkert et al. [7].*

With all the pieces assembled, we obtain the runtime of solving a number of problem variants via this reduction. We present the relevant runtimes in the following theorem.

▶ **Theorem 10.** *For any problem P in* OR-CO🔒SA🔒, OR-C🔒OSA🔒 *and* OR-C📌O📌SA🔒 *let A_P be the size of the set of possible ports for p_{min}, s.t., the set can be computed in $O(A_P)$ time and let the time required to solve the created instance of* Boundary Labeling *be $O(B_P)$. Then P can be solved in time $O(A_P \cdot B_P)$.*

Proof. We compute the set of possible ports for $p_{\min}$ as stated in Observation 4, Lemma 5 or Observation 6. For any element of these sets, we create an instance of Boundary Labeling and solve it as described in Lemma 7, Observation 8 and Lemma 9. We obtain a *po*-labeling minimizing the badness function, which has a one-to-one correspondence to a leader length minimal OR-labeling of our original instance. ◀

We remark that Benkert et al. [7] point out that their algorithm can be changed to handle non-uniform labels leading to a pseudopolynomial time algorithm. By Theorem 10 this would result in pseudopolynomial time algorithms for the problems in OR-C📌O🔒$S_{\equiv}A$🔒, however, these are superseded by a dedicated approach in Section 4, so we omit further details here.

4 Improvement via dedicated approaches

Having established the reduction as a general approach baseline, we will now present a variety of bespoke approaches, which improve the runtime implied by Theorem 10.

4.1 Locked Port Candidates and Locked Order

We begin by investigating the problem set OR-C🔒O🔒SA🔒. Recall that we are given a set C of candidate positions for the ports. The placement of the label of $p_{\min}$ determines the position of all other labels. The following lemma (which is applicable to a larger set of problems) shows that the same is true for all OR-leaders.

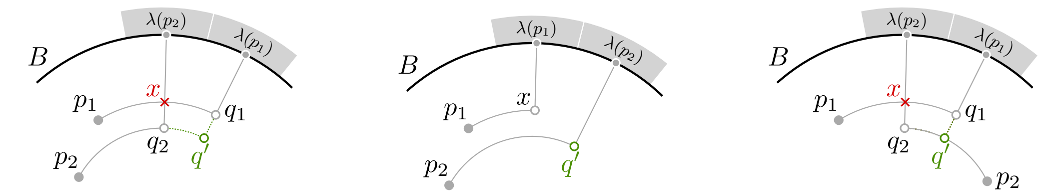

(a) Case 1. **(b)** Rerouted leaders. **(c)** Case 2.

Figure 5 Given a free label order O📌 we can reroute the leaders to arrive at a crossing-free solution with a shorter total leader length.

▶ **Lemma 11.** *For* OR-CO🔒SA🔒*, the choice of a port for p_{min} uniquely determines the* OR*-leader of any other feature p', which does not cross the* OR*-leader of p_{min}.*

Proof. Note that the orbital parts of the two possible OR-leaders connecting p' to any point on B form a circle concentric with B. Therefore one of the two orbital parts crosses the radial part of the leader of $p_{\min}$ (we also refer again to Figure 2a). The lemma follows. ◀

By Lemma 11 it is sufficient to place $\lambda(p_{\min})$ with $\xi_L(p)$ coinciding with one candidate ($O(|C|)$ possibilities). Then we check in $O(n)$ time if the ports of the remaining labels in the correct order also coincide with candidates, in $O(n^2)$ time that no two leaders cross, and finally we compute in $O(n)$ time the total leader length, leading to a total runtime of $O(|C|n^2)$ (see Figure 4) and to the following theorem.

▶ **Theorem 12.** *The problems* OR-C🔒O🔒SA🔒 *can be solved in $O(|C|n^2)$ time.*

4.2 Locked Port Candidates, Free Order and Uniform Port Distribution

If the order can freely be chosen, then a uniform port ratio together with uniform labels guarantee that any two labels can be exchanged without creating overlap between two labels. We can thus utilize a matching algorithm to solve the problems OR-C🔒O📌$S_{\equiv}A_{\equiv}$🔒. To obtain a crossing-free labeling we first prove the following lemma showing that a solution minimizing total leader length naturally does not contain any crossings.

▶ **Lemma 13.** *Given an instance of a problem variant in* OR-CO📌$S_{\equiv}A_{\equiv}$ *every leader-length minimal labeling L is crossing-free.*

Proof. Assume a leader-length minimal labeling L contains two crossing leaders γ_1 and γ_2 connecting p_1 to its port $\xi_L(p_1) = \xi_1$ and p_2 to its port $\xi_L(p_2) = \xi_2$, respectively. Both leaders begin with an orbital segment $\widehat{p_1 q_1}$ (or $\widehat{q_1 p_1}$) and $\widehat{p_2 q_2}$ (or $\widehat{q_2 p_2}$), respectively, connecting to their bend points q_1 and q_2, followed by their radial straight-line segment $q_1\xi_1$ and $q_2\xi_2$. Clearly, the crossing x occurs between the radial segment of the point closer to the center of D and the orbital segment of the point closer to B. We assume, w.l.o.g., that $x = \widehat{p_1 q_1} \cap q_2\xi_2$, and that p_1 is on the counter-clockwise end of $\widehat{p_1 q_1}$. Let q' be the intersection of the supporting line of $q_1\xi_1$ and the circle containing $\widehat{q_2 p_2}$. There are two cases, as shown in Figure 5.

In the first case (Figure 5a), we can replace γ_1 with a curve consisting of $\widehat{p_1 x}$ and $x\xi_2$ and γ_2 with curve consisting of $\widehat{p_2 q'}$ and $q'\xi_1$ (Figure 5b). Since $|q_2 x| = |q' q_1|$ and $|\widehat{x q_1}| > |\widehat{q_2 q'}|$, the total leader length has decreased. Note that the rerouting might have introduced new crossings, but since this method reduces the total leader length, we can iteratively apply this procedure and will never obtain an already seen labeling. Since there is a finite number of

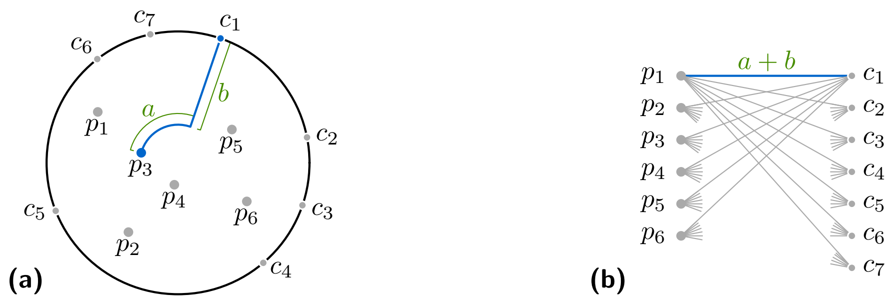

Figure 6 Case $C^{\text{🔒}}O^{\text{🔑}}S_{\equiv}A^{\text{🔒}}_{\equiv}$. Each feature p_1, …, p_6 and port candidate $c_1, \ldots, c_7$ **(a)** introduces a vertex in the weighted complete bipartite graph **(b)**. An edge in the bipartite graph corresponds to a leader and is weighted with the leader's length.

possible solutions, we have to arrive at a solution, which does not contain crossings anymore (otherwise we could apply the procedure infinitely many times contradicting the finite number of possible solutions). We arrive at a labeling, that has a smaller total leader length, which is a contradiction to L being optimal. While the second case (Figure 5c) looks different geometrically, we can resolve the crossing identically to the first case to again reduce the sum of leader lengths. In the special case where both features have the same distance to X, it is again obvious that by simply switching the ports the newly obtained leaders are a subset of the old leaders, removing the overlap and reducing the total leader length. In all three cases, we arrived at a labeling that is better than L which is a contradiction, concluding the proof. ◀

Next, we show that we can always use the shorter of the two OR-leaders.

▶ **Lemma 14.** *Given an instance of a problem variant in* OR-$CO^{\text{🔑}}S_{\equiv}A_{\equiv}$ *any leader-length minimal labeling L uses only the shorter of the two possible* OR*-leaders for any point.*

Proof. Assume L contains a long OR-leader for a point p. Replacing the leader with the short leader between p and its port yields a labeling L' with a shorter total leader length. If L' is crossing-free this is a contradiction to the optimality of L, otherwise we can iteratively apply the uncrossing procedure of the proof of Lemma 13, which, by Lemma 13, results in a crossing free labeling with the same or smaller total leader length compared to L' again contradicting the optimality of L. ◀

With Lemmas 13 and 14 we know that every combination of feature and port defines a unique leader and the shortest possible set of such leaders is crossing free. This leads naturally to a formulation of the problem as finding a minimum-weight matching between features and ports. However, we still have to guarantee that the ports we select in such a matching are equally distributed around B; a requirement stemming from uniform label-sizes and fixed uniform port ratios. To this end, we state the following observation.

▶ **Observation 15.** *Given a set C of candidate ports, we can partition C into at most $k = \frac{|C|}{n}$ subsets $C_1, \ldots, C_k$ of size n, s.t., all candidates in one set C_i are equally distributed around B. This can be done in $O(|C|^2)$ time.*

Now we can state the central result of this subsection.

▶ **Theorem 16.** *The problems in* OR-$C^{\text{🔒}}O^{\text{🔑}}S_{\equiv}A^{\text{🔒}}_{\equiv}$ *can be solved in $O(|C|^2 n)$ time.*

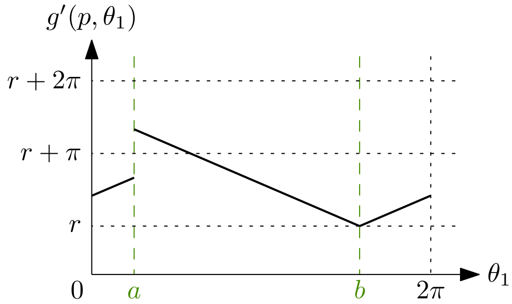

Figure 7 An example of the function $g'(p, \theta_1)$ for a point p with a distance of $1 - r$ to X. the points a and b are defined as in the proof of Lemma 17.

Proof. We begin by creating the k subsets (Observation 15). For each subset C_i we let G_i be a weighted complete bipartite graph between the set of features P and C_i using the length of the (short) leader between a feature $p \in P$ and a potential port $c \in C$ as the weight of the edge (p, c); see Figure 6. For OR-leaders it is by Lemma 14 sufficient to use the length of the shorter of the two leaders. Now a minimum weight bipartite matching in G corresponds to a leader-length minimal labeling. We know by Lemma 13, that such a labeling is crossing-free and an optimal solution (for C_i). Such a matching in a bipartite graph with $|V|$ vertices and $|E|$ edges can be computed in $O(|V|^2 \log |V| + |V||E|)$ time [11]. In our case $|V| = 2n$ and $|E| = n^2$. Therefore the runtime for one subset is $O(n^3)$ and since we run this algorithm for at most $\frac{|C|}{n}$ subsets, we arrive at a final runtime of $O(|C|^2 + |C|n^2)$. ◀

It is important to note that the total runtime of all iterations of the matching algorithm is strictly better than the runtime of $O(|C|n^3)$ yielded by Theorem 10, due to the preprocessing runtime of $O(|C|^2)$. Theorem 16 only guarantees an improvement for $|C| \in o(n^3)$.

4.3 Free Candidates and Locked Order

Now we turn to the problem set OR-$COSA$. Intuitively, these are problem variants, in which we can rotate the ports continuously around B to obtain other labelings as long as we do not change the label order or introduce any crossings. To solve these problems, we formulate a univariate piece-wise linear function of bounded complexity, whose global minimum corresponds to an optimal labeling.

Let $\omega(i)$ be the index of the label of p_i in the fixed order (assuming $\omega(1) = 1$). For any feature p_i let θ_i^L be the angle between p_i and $\xi_L(p_i)$ in a specific labeling L. Recall that by Lemma 11 the choice of the port for the innermost feature in L also fixes all other ports and leaders. We will therefore replace θ_i^L with $\Theta_i(\theta_1^L)$, i.e., a function which simply returns the value of θ_i^L implied by the value of θ_1^L. Since the order of labels is fixed, it is not guaranteed that every feature is connected to its port using the short OR-leader. To do so we define a Boolean variable $cw(i, \theta_1^L)$. Let s_i be the intersection of B and a ray starting at X through p_i. Then $cw(i, \theta_1^L)$ is true if (in the labeling L implied by the value θ_1) starting at s_i and traversing B clockwise we encounter $\xi_L(p_i)$ before $\xi_L(p_1)$ and false otherwise.

With this we define the function $g' : P \times [0, 2\pi] \to \mathbb{R}$ as

$$g'(p_i, \theta_1) = \begin{cases} g(p_i, \Theta_i(\theta_1)) & \text{if } cw(i, \theta_1) \text{ or } i = 1 \\ g(p_i, 2\pi - \Theta_i(\theta_1)) & \text{otherwise} \end{cases} \quad (1)$$

We now provide a lemma bounding the complexity of these functions.

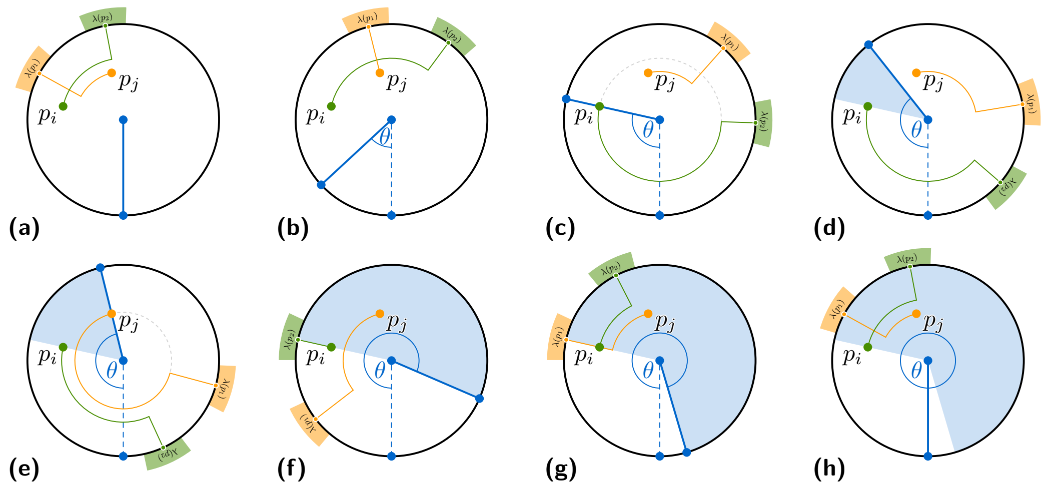

Figure 8 Illustration of the interval (blue area) of θ in which the two leaders of the points p_i and p_j do not cross, while increasing θ from 0 **(a)** to 2π **(h)**. The heavy blue line is the supporting line for the radial part of the leader of $p_{\min}$. The endpoints are the value of θ for which the leader of p_i changes from clockwise to counter-clockwise to avoid crossing the leader of the innermost point **(c)** and the value for which the radial part of the leader of p_j crosses p_i **(g)**.

▶ **Lemma 17.** *Any function $g'(p_i, \theta_1)$ consists of at most three continuous linear parts in the interval $0 \leq \theta_1 < 2\pi$.*

Proof. Assume $cw(i, 0)$ is true, i.e., we encounter $\xi_L(p_i)$ before $\xi_L(p_1)$ when traversing B clockwise starting at s_i. While increasing θ_1 the value of $cw(i, \theta_1)$ changes to false exactly when $\xi_L(p_1) = s_i$ and back to true exactly when $\xi_{L'}(p_i) = s_i$, where L and L' are labelings implied by some values $\theta_1 = a$ and $\theta_1 = b$, respectively, s.t., $a < b$ (see also Figure 7). Since for $\theta_1[0, a) \cup [b, 2\pi)$ we have $g'(p_i, \theta_1) = g(p_i, \Theta_i(\theta_1))$ and for $\theta_1 \in [a, b)$ we have $g'(p_i, \theta_1) = g(p_i, 2\pi - \Theta_i(\theta_1))$ and by Observation 2 the function g is continuous and linear in these intervals. The case for $cw(i, 0)$ being false is symmetrical and the lemma follows. ◀

We can now create a function $h(\theta_1) = \sum_{i=1}^{n} g'(p_i, \theta_1)$, which exactly captures the total leader length of a labeling implied by θ_1, which adheres to the order ω. However, it is important to observe that, while such a labeling does not contain any crossing between the leader of p_1 and any other leader, crossings between any other pair of leaders are still possible. To avoid these, we restrict θ_1 to $O(n^2)$ intervals which capture exactly the values, in which the labeling implied by θ_1 does not contain any crossing. We define a new Boolean variable $cr(i, j, \theta_1)$, which is true if the two OR-leaders of p_i and p_j that are implied by θ_1, cross. Let p_i be the feature closer to the center of D than p_j. Similar to the proof of Lemma 17, we can find two values θ_a, θ_b, s.t., $cr(i, j, x) = cr(i, j, x')$ for any value $x, x' \in [\theta_a, \theta_b)$ and $cr(i, j, x) = cr(i, j, x')$ for any value $x, x' \in [0, \theta_a) \cup [\theta_b, 2\pi)$.

▶ **Lemma 18.** *Let $\mathcal{I}$ be the set of intervals, s.t., θ_1 implies a labeling in which all leaders are crossing-free if and only if $\theta_1 \in I$ for some $I \in \mathcal{I}$. Then $|\mathcal{I}| \leq n^2$.*

Proof. For any pair of features, there is exactly one interval in which their leaders do not cross and one interval in which they do (considering the intervals modulo 2π). An example is shown in Figure 8. We prove the lemma by induction. The base case is a single feature, which is always crossing free, therefore $|\mathcal{I}| = 1^2$. Assume that there are at most n^2 such intervals for n features. We add the $(n + 1)$-th feature p. For any of the existing n features,

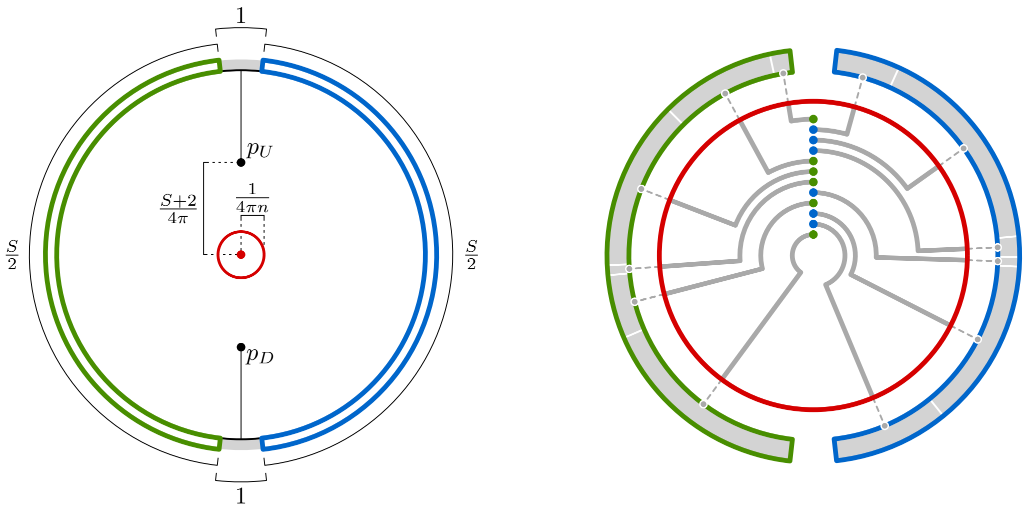

(a) Placement of two blocker features p_U and p_D. **(b)** Detail of the placement of $p_1, \ldots, p_n$.

Figure 9 Visualization of the reduction. The two points p_U and p_D are placed close to the boundary **(a)** and all points $p_1, \ldots, p_n$ are placed in the very small red circle. A zoomed-in picture of the red circle is shown in **(b)**.

p defines a single interval in which their leaders do not cross. Consider the possibilities for adding one of these intervals I' to $\mathcal{I}$. Let $I \in \mathcal{I}$. If $I \cap I' = \emptyset$, we remove I, if $I \cap I' = I$, we retain I and if $I \cap I' = I'$, we remove I and add I'. Otherwise, I contains one or two endpoints of I' but not the entire interval. If I contains only one endpoint of I', we remove I and add $I \cap I'$. Since I' is an interval, this can happen with at most two other intervals. Both times we remove one interval and add a new one, thereby not changing the size of $\mathcal{I}$. Finally, if I contains both endpoints of I' but not the entire interval – intuitively the intervals wrap around B in two different directions – then we remove I and add the two continuous parts of $I \cap I'$ increasing $|\mathcal{I}|$ by one. In this case all other intervals of $\mathcal{I}$ are either entirely contained in I or disjoint from I and therefore $|\mathcal{I}|$ does not increase any further. We iteratively add all n new intervals and increase $|\mathcal{I}|$ by a total of at most n. Therefore $|\mathcal{I}| \leq n^2 + n < n^2 + 2n + 1 = (n+1)^2$. ◀

With this, we have everything in place to prove the following theorem.

▶ **Theorem 19.** *The problem variants in* OR-$COSA$ *can be solved in* $O(n^2)$ *time.*

Proof. We set up the function $g'(p_i, \theta_1)$ as in Equation 1 and let $h(\theta_1) = \sum_{i=1}^{n} g'(p_i, \theta_1)$ as above. Then we restrict θ_1 to the set $\mathcal{I}$ of n^2 intervals (Lemma 18) in which the implied labelings do not contain any crossings. Since by Lemma 17 there are at most $O(1)$ continuous linear pieces for any function g', we conclude that h consists of $O(n)$ continuous linear pieces and thus also has at most $O(n)$ local minima. To find the global minimum of h it is sufficient to check all $O(n)$ local minima as well as the $O(n^2)$ endpoints of the intervals in $\mathcal{I}$. ◀

5 Free Candidates, Free Order and Non-uniform labels are NP-hard

The results presented so far cover most of the top half of Table 1. It remains to address the problems in OR-$COS_{\equiv}A$. Due to the flexibility of non-uniform labels, as well as the free order and lack of a discrete candidate set, these problems turn out to be weakly NP-hard.

▶ **Theorem 20.** *Given an instance of* OR-$COS_{\equiv}A$ *together with $k \in \mathbb{R}$ it is (weakly)* NP*-hard to decide whether there exists a labeling L with a total leader length of less than k.*

Proof. For the purpose of this proof, we will relax the requirement that D has a radius of 1. The final construction can be scaled down to meet that requirement. The reduction (visualized in Figure 9) is from the weakly NP-hard problem Partition, where we are given a set $\mathcal{X}$ of n integers with $S = \sum_{x \in \mathcal{X}} x$ and need to decide if $\mathcal{X}$ can be partitioned into two sets $\mathcal{X}_1$ and $\mathcal{X}_2$, s.t., $\sum_{x \in \mathcal{X}_1} x = \sum_{x \in \mathcal{X}_2} x = S/2$. For the reduction we place for every $x_i \in \mathcal{X}$ a feature $p_i = \left(0, \frac{i}{4\pi n^2}\right)$. Additionally we place two features $p_U = (0, r)$ and $p_D = (0, -r)$, where $r > \frac{S+2}{4\pi}$. We define $\lambda(p_i) = x_i$ for all $1 \leq i \leq n$ and $\lambda(p_U) = \lambda(p_D) = 1$. Note that $\sum_{i=1}^{n} \lambda(p_i) + \lambda(p_U) + \lambda(p_D) = S + 2$ and the radius of the enclosing disk is therefore $\frac{S+2}{2\pi}$.

Any feature p_i, s.t., $1 \leq i \leq n$ is contained in a disk of radius $\frac{1}{4\pi n}$ and circumference $\frac{1}{2n}$ centered at the origin. Let $o(i)$ and $r(i)$ be the orbital and radial part of $\gamma_L^{p_i}$, respectively. Note that the sum over all $r(i)$ is equal in all labelings. Let this sum be equal to L_{radial}. Further note that for any p_i, $o(i) < \frac{1}{2n}$. Therefore the sum $\sum_{i=1}^{n} o(i) < \frac{1}{2}$.

For the problem variants OR-$COS_{\equiv}A_{\equiv}$, in any labeling L the port ratios $\rho(p_U)$, and $\rho(p_D)$ are necessarily equal. For the variant OR-$COS_{\equiv}A_{\equiv}$ port ratios are described as part of the input and we define them, s.t., $\rho(p_U) = \rho(p_D)$. Finally we set $k = 1/2 + L_{\text{radial}}$.

If there exists a partition of $\mathcal{X}$ into two sets $\mathcal{X}_1, \mathcal{X}_2$, s.t., $\sum_{x \in \mathcal{X}_1} x = \sum_{x \in \mathcal{X}_2} x$, then we can make three observations. First, there exists a labeling L in which the length of the orbital part of $\gamma_L^{p_U}$ and $\gamma_L^{p_D}$ is equal to 0 and therefore $\gamma_L^{p_U}$ and $\gamma_L^{p_D}$ are straight lines. Second, in L both spaces between the labels of p_U and p_D are equally spaced, i.e., $|\overparen{e_L(p_U) s_L(p_D)}| = |\overparen{e_L(p_D) s_L(p_U)}|$, since $\rho(p_U) = \rho(p_D)$. Third, in a labeling L', in which the length of the orbital part of $\gamma_L^{p_U}$ or $\gamma_L^{p_D}$ is not equal to 0, the sum of the length of the orbital parts of $\gamma_{L'}^{p_U}$ and $\gamma_{L'}^{p_D}$ (and therefore the sum over the lengths of all orbital parts of leaders in L') is at least $\frac{2\pi}{S+2} \cdot \frac{S+2}{4\pi} = \frac{1}{2}$. This is because the difference between $\overparen{e_{L'}(p_U) s_{L'}(p_D)}$ and $\overparen{e_{L'}(p_D) s_{L'}(p_U)}$ is at least 1 (since the label sizes are integers). Therefore the sum over all leader lengths in L is less than $1/2 + L_{\text{radial}}$, while in L' it is at least $1/2 + L_{\text{radial}}$ and L' can never be optimal.

Assume now that $\mathcal{X}$ can be partitioned into two subsets $\mathcal{X}_1, \mathcal{X}_2$, s.t., $\sum_{x \in \mathcal{X}_1} x = \sum_{x \in \mathcal{X}_2} x$. Then the labels can be equally partitioned and the leaders of p_U and p_D can be straight lines. Therefore the total sum of leader lengths is less than k. Conversely, assume no such partition exists. Then the leaders of p_U and p_D must together contain orbital segments of length at least $1/2$ and the total sum of leader lengths is at least k, concluding the proof. ◀

6 Extensions to SL-Leaders and Free Port Ratios

So far we considered problems in OR-$COSA$, i.e., settings which use OR-leaders and have a fixed port ratio. Some results can be translated (with small to medium effort adaptions) to the settings using SL-leaders and/or free port ratios. In this section, we will give a high-level overview of which results can be adapted and how. The set of results for free port ratios is shown in Table 2; results for SL-leaders are included as the bottom part of both Table 1 and 2.

Some results (e.g., Theorem 12) are independent of the leader length (beyond computation of leader length). Using a matching algorithm as stated in Theorem 16 also extends to the SL variant (although it requires proving that here also leader-length minimal labelings are crossing-free). Extending Lemma 19 is harder, since the number of possible local minima of a function describing the sum of SL leader lengths is not obvious. We include a conjecture for the settings SL-$COSA$ one of which extends to a setting with free port ratio due to

Table 2 A tabular overview of the problem space and our results. Only **free** port ratios are shown. The abbreviation "w. NP-h" stands for weakly NP-hard. Blue cells are conjectures. These results can be found in the full version [8]; cells indicate the section containing the relevant result.

			C🔒		C📍	
			$A_{\equiv}$📍	$A_{\equiv}$📍	$A_{\equiv}$📍	$A_{\equiv}$📍
OR	O🔒	$S_{\equiv}$	$O(\lvert C\rvert n^2)$ [8, C.1]	$O(\lvert C\rvert^2 n^3)$ [8, C.2]	$O(n^2)$ [8, C.1]	$O(n^6)$ [8, C.2]
		$S_{\equiv}$		$O(\lvert C\rvert^5 n)$ [8, C.2]		$O(n^7)$ [8, C.2]
	O📍	$S_{\equiv}$	$O(\lvert C\rvert^2 n)$ [8, C.1]	$O(\lvert C\rvert^2 n^3)$ [8, C.2]	$O(n^5)$ [8, C.1]	$O(n^6)$ [8, C.2]
		$S_{\equiv}$			w. NP-h [8, C.3]	
SL	O🔒	$S_{\equiv}$	$O(\lvert C\rvert n^2)$ [8, C.1]		$O(n^2)$? [8, C.1]	
		$S_{\equiv}$				
	O📍	$S_{\equiv}$	$O(\lvert C\rvert^2 n)$ [8, C.1]			
		$S_{\equiv}$			w. NP-h [8, C.3]	

the equivalence of labelings. The construction of our NP-hardness reduction can be reused for the SL-leader case, however, it requires new arguments that it still works as intended. Details can be found in Section B of the full version [8].

While OR-leaders are related to *po*-leaders, no such immediate relation exists for SL-leaders, one reason being that OR-leaders monotonically increase their distance to X, which is not true for SL-leaders. Therefore the reduction to Boundary Labeling explained in Section 3 does not extend to SL-leaders. It also does not work as is for free port ratios, since we relied on the fact that a position of the port of $p_{\min}$ fixes the position of its label, which is not true for free port ratios. For some instances, we can leverage candidate positions or a fixed order to prove again that only a certain set of linear or quadratic size needs to be considered for such a label. Therefore we obtain similar runtimes with an additional linear or quadratic factor. However this reduction only works for non-uniform port ratios, since otherwise the subproblems in the dynamic program of Benkert et al. [7] are not independent. Detailed results for free port ratios can be found in Section C of the full version [8].

7 Conclusion

We have introduced orbital labeling as a variant of boundary labeling for circular boundaries, in which labels are placed as circular arcs in an annulus along the boundary. We provided a broad overview of problem variants, based on five different parameters of this labeling problem. We showed that from an algorithmic point of view, the different parameter combinations lead to distinctively different computational problems. In general, it appears that (unsurprisingly) the non-uniform label setting is computationally harder than the uniform setting. Similarly computing the layout for a fixed order is easier than for a free order of labels. The fixed candidate setting discretizes the problem and allows for an exhaustive search through all possible solutions, in contrast to the free candidate setting. In opposition to leader length, non-uniform port ratios seem to make the problem more approachable if the port ratios are also free, since free but uniform port ratios introduce a property, which has to be fulfilled globally, while free non-uniform port ratios can be fixed locally. Concerning leader types, the linear behaviour of the length of an OR-leader relative to the angle of their port with the x-axis allows for some approaches, which are not (or at least not immediately) applicable to SL-leaders.

It would be interesting to further investigate orbital labeling through a visualization lens as well: the variants have distinct visual styles and portray varying levels of visual complexity and uniformity. We are conducting user experiments to determine whether certain variants are superior to others.

References

1 Pankaj K. Agarwal, Marc van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. *Computational Geometry: Theory and Applications*, 11(3-4):209–218, 1998. doi:10.1016/S0925-7721(98)00028-5.

2 Sarah E. Battersby, John E. Stewart, Ana Lopez-De Fede, Kevin C. Remington, and Kathy Mayfield-Smith. Ring maps for spatial visualization of multivariate epidemiological data. *Journal of Maps*, 7(1):564–572, 2011. doi:10.4113/jom.2011.1182.

3 Michael A. Bekos, Michael Kaufmann, Antonios Symvonis, and Alexander Wolff. Boundary labeling: Models and efficient algorithms for rectangular maps. In János Pach, editor, *Graph Drawing (GD'04)*, volume 3383 of *LNCS*, pages 49–59. Springer, 2004. doi:10.1007/978-3-540-31843-9_7.

4 Michael A Bekos, Michael Kaufmann, Antonios Symvonis, and Alexander Wolff. Boundary labeling: Models and efficient algorithms for rectangular maps. *Computational Geometry*, 36(3):215–236, 2007. doi:10.1016/J.COMGEO.2006.05.003.

5 Michael A. Bekos, Benjamin Niedermann, and Martin Nöllenburg. *External Labeling: Fundamental Concepts and Algorithmic Techniques.* Synthesis Lectures on Visualization. Morgan & Claypool Publishers, 2021. doi:10.2200/S01115ED1V01Y202107VIS014.

6 Marc Benkert, Herman Haverkort, Moritz Kroll, and Martin Nöllenburg. Algorithms for multi-criteria one-sided boundary labeling. In Seok-Hee Hong and Takao Nishizeki, editors, *Graph Drawing (GD'07)*, volume 4875 of *LNCS*, pages 243–254. Springer, 2008. doi:10.1007/978-3-540-77537-9_25.

7 Marc Benkert, Herman J Haverkort, Moritz Kroll, and Martin Nöllenburg. Algorithms for multi-criteria boundary labeling. *Journal of Graph Algorithms and Applications*, 13(3):289–317, 2009. doi:10.7155/jgaa.00189.

8 Annika Bonerath, Martin Nöllenburg, Soeren Terziadis, Markus Wallinger, and Jules Wulms. On orbital labeling with circular contours, 2024. arXiv:2403.19052, doi:10.48550/arXiv.2403.19052.

9 Martin Fink, Jan-Henrik Haunert, André Schulz, Joachim Spoerhase, and Alexander Wolff. Algorithms for labeling focus regions. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2583–2592, 2012. doi:10.1109/TVCG.2012.193.

10 Michael Formann and Frank Wagner. A packing problem with applications to lettering of maps. In *Symposium on Computational Geometry (SoCG'91)*, pages 281–288. ACM, 1991. doi:10.1145/109648.109680.

11 Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987. doi:10.1145/28869.28874.

12 Andreas Gemsa, Jan-Henrik Haunert, and Martin Nöllenburg. Multirow boundary-labeling algorithms for panorama images. *ACM Transactions on Spatial Algorithms and Systems*, 1(1):1–30, 2015. doi:10.1145/2794299.

13 Jan-Henrik Haunert and Tobias Hermes. Labeling circular focus regions based on a tractable case of maximum weight independent set of rectangles. In Falko Schmid, Christian Kray, and Holger Fritze, editors, *SIGSPATIAL International Workshop on Interacting with Maps (MapInteract@GIS'14)*, pages 15–21, 2014. doi:10.1145/2677068.2677069.

14 Alaul Islam, Tingying He, Anastasia Bezerianos, Bongshin Lee, Tanja Blascheck, and Petra Isenberg. Visualizing information on smartwatch faces: A review and design space. *CoRR*, abs/2310.16185, 2023. doi:10.48550/arXiv.2310.16185.

15 Benjamin Niedermann and Jan-Henrik Haunert. Focus+context map labeling with optimized clutter reduction. *International Journal of Cartography*, 5(2-3):158–177, 2019. doi:10.1080/23729333.2019.1613072.

16 Marc van Kreveld, Tycho Strijk, and Alexander Wolff. Point labeling with sliding labels. *Computational Geometry: Theory and Applications*, 13(1):21–47, 1999. doi:10.1016/S0925-7721(99)00005-X.

17 Alexander Wolff. Graph drawing and cartography. In Roberto Tamassia, editor, *Handbook on Graph Drawing and Visualization*, pages 697–736. Chapman and Hall/CRC, 2013.