Stable Approximation Algorithms for the Dynamic Broadcast Range-Assignment Problem

Mark de Berg ✉

Department of Mathematics and Computer Science,
Eindhoven University of Technology, The Netherlands

Arpan Sadhukhan ✉

Department of Mathematics and Computer Science,
Eindhoven University of Technology, The Netherlands

Frits Spieksma ✉

Department of Mathematics and Computer Science,
Eindhoven University of Technology, The Netherlands

Abstract

Let P be a set of points in $\mathbb{R}^d$ (or some other metric space), where each point $p \in P$ has an associated transmission range, denoted $\rho(p)$. The range assignment ρ induces a directed communication graph $\mathcal{G}_\rho(P)$ on P, which contains an edge (p, q) iff $|pq| \leqslant \rho(p)$. In the broadcast range-assignment problem, the goal is to assign the ranges such that $\mathcal{G}_\rho(P)$ contains an arborescence rooted at a designated root node and the cost $\sum_{p \in P} \rho(p)^2$ of the assignment is minimized.

We study the dynamic version of this problem. In particular, we study trade-offs between the stability of the solution – the number of ranges that are modified when a point is inserted into or deleted from P – and its approximation ratio. To this end we introduce the concept of *k-stable algorithms*, which are algorithms that modify the range of at most k points when they update the solution. We also introduce the concept of a *stable approximation scheme*, or *SAS* for short. A SAS is an update algorithm ALG that, for any given fixed parameter $\varepsilon > 0$, is $k(\varepsilon)$-stable and that maintains a solution with approximation ratio $1+\varepsilon$, where the stability parameter $k(\varepsilon)$ only depends on ε and not on the size of P. We study such trade-offs in three settings.

- For the problem in $\mathbb{R}^1$, we present a SAS with $k(\varepsilon) = O(1/\varepsilon)$. Furthermore, we prove that this is tight in the worst case: any SAS for the problem must have $k(\varepsilon) = \Omega(1/\varepsilon)$. We also present algorithms with very small stability parameters: a 1-stable $(6 + 2\sqrt{5})$-approximation algorithm – this algorithm can only handle insertions – a (trivial) 2-stable 2-approximation algorithm, and a 3-stable 1.97-approximation algorithm.
- For the problem in $\mathbb{S}^1$ (that is, when the underlying space is a circle) we prove that no SAS exists. This is in spite of the fact that, for the static problem in $\mathbb{S}^1$, we prove that an optimal solution can always be obtained by cutting the circle at an appropriate point and solving the resulting problem in $\mathbb{R}^1$.
- For the problem in $\mathbb{R}^2$, we also prove that no SAS exists, and we present a $O(1)$-stable $O(1)$-approximation algorithm.

Most results generalize to when the range-assignment cost is $\sum_{p \in P} \rho(p)^\alpha$, for some constant $\alpha > 1$. All omitted theorems and proofs are available in the full version of the paper [14].

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Computational geometry, online algorithms, broadcast range assignment, stable approximation schemes

Digital Object Identifier 10.4230/LIPIcs.SWAT.2022.15

Related Version *Full Version*: `https://doi.org/10.48550/arXiv.2112.05426`

Funding MdB, AS, and FS are supported by the Dutch Research Council (NWO) through Gravitation-grant NETWORKS-024.002.003.

18th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2022).
Editors: Artur Czumaj and Qin Xin; Article No. 15; pp. 15:1–15:21
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements We thank the reviewers of an earlier version of the paper for pointing us to some important references and for other helpful comments.

1 Introduction

The broadcast range-assignment problem. Let P be a set of points in $\mathbb{R}^d$, representing transmission devices in a wireless network. By assigning each point $p \in P$ a transmission range $\rho(p)$, we obtain a *communication graph* $\mathcal{G}_\rho(P)$. The nodes in $\mathcal{G}_\rho(P)$ are the points from P and there is a directed edge (p, q) iff $|pq| \leqslant \rho(p)$, where $|pq|$ denotes the Euclidean distance between p and q. The energy consumption of a device depends on its transmission range: the larger the range, the more energy it needs. More precisely, the energy needed to obtain a transmission range $\rho(p)$ is given by $\rho(p)^\alpha$, for some real constant $\alpha > 1$ called the *distance-power gradient*. In practice, α depends on the environment and ranges from 1 to 6 [18]. Thus the overall cost of a range assignment is $\text{cost}_\alpha(\rho(P)) := \sum_{p \in P} \rho(p)^\alpha$, where we use $\rho(P)$ to denote the set of ranges given to the points in P by the assignment ρ. The goal of the range-assignment problem is to assign the ranges such that $\mathcal{G}_\rho(P)$ has certain connectivity properties while minimizing the total cost [6]. Desirable connectivity properties are that $\mathcal{G}_\rho(P)$ is (h-hop) strongly connected [8, 9, 10, 16] or that $\mathcal{G}_\rho(P)$ contains a *broadcast tree*, that is, an arborescence rooted at a given source $s \in P$. The latter property leads to the *broadcast range-assignment problem*, which is the topic of our paper.

The broadcast range-assignment problem has been studied extensively, sometimes with the extra condition that any point in P is reachable in at most h hops from the source s. For $\alpha = 1$ the problem is trivial in any dimension: setting the range of the source s to $\max\{|sp| : p \in P\}$ and all other ranges to zero is optimal; however, for any $\alpha > 1$ the problem is NP-hard in $\mathbb{R}^d$ for $d \geqslant 2$ [5, 15]. Approximation algorithms and results on hardness of approximation are known as well [4, 7, 15]. Many of our results will be on the 1-dimensional (or: linear) broadcast range-assignment problem. Linear networks are important for modeling road traffic information systems [3, 17] and as such they have received ample attention. In $\mathbb{R}^1$, the broadcast range-assignment problem is no longer NP-hard, and several polynomial-time algorithms have been proposed, for the standard version, the h-hop version, as well as the weighted version [2, 4, 7, 11, 12]. The currently fastest algorithms for the (standard and h-hop) broadcast range-assignment problem run in $O(n^2)$ time [11].

All results mentioned so far are for the static version of the problem. Our interest lies in the dynamic version, where points can be inserted into and deleted from P (except the source, which should always be present). This corresponds to new sensors being deployed and existing sensors being removed, or, in a traffic scenario, cars entering and exiting the highway. Recomputing the range assignment from scratch when P is updated may result in all ranges being changed. The question we want to answer is therefore: is it possible to maintain a close-to-optimal range assignment that is relatively stable, that is, an assignment for which only few ranges are modified when a point is inserted into or deleted from P? And which trade-offs can be achieved between the quality of the solution and its stability?

To the best of our knowledge, the dynamic problem has not been studied so far. The online problem, where the points from P arrive one by one (there are no deletions) and it is not allowed to decrease ranges, is studied by De Berg et al. [13]. This restriction is arguably unnatural, and it has the consequence that a bounded approximation ratio cannot be achieved. Indeed, let the source s be at $x = 0$, and suppose that first the point $x = 1$ arrives, forcing us to set $\rho(s) := 1$, and then the points $x = i/n$ arrive for $1 \leqslant i < n$. In the optimal static solution at the end of this scenario all points, except the rightmost one, have range $1/n$; for $\alpha = 2$ this induces a total cost of $n \cdot (1/n)^2 = 1/n$. But if we are not allowed to decrease the range of s after setting $\rho(s) = 1$, the total cost will be (at least) 1, leading to an unbounded approximation ratio. Therefore, [13] analyze the competitive ratio: they compare the cost of their algorithm to the cost of an optimal offline algorithm (which knows the future arrivals, but must still maintain a valid solution at all times without decreasing any range). As we will see, by allowing to also decrease a few ranges, we are able to maintain solutions whose cost is close even to the static optimum.

Our contribution. Before we state our results, we first define the framework we use to analyze our algorithms. Let P be a dynamic set of points in $\mathbb{R}^d$, which includes a fixed source point s that cannot be deleted.

An update algorithm ALG for the dynamic broadcast range-assignment problem is an algorithm that, given the current solution (the current ranges of the points in the current set P) and the location of the new point to be inserted into P, or the point to be deleted from P, modifies the range assignment so that the updated solution is a valid broadcast range assignment for the updated set P. We call such an update algorithm *k-stable* if it modifies at most k ranges when a point is inserted into or deleted from P. Here we define the range of a point currently not in P to be zero. Thus, if a newly inserted point receives a positive range it will be counted as receiving a modified range; similarly, if a point with positive range is deleted then it will be counted as receiving a modified range. To get a more detailed view of the stability, we sometimes distinguish between the number of increased ranges and the number of decreased ranges, in the worst case. When these numbers are k^+ and k^-, respectively, we say that ALG is *(k^+, k^-)-stable*. This is especially useful when we separately report on the stability of insertions and deletions; often, when insertions are (k_1, k_2)-stable then deletions will be (k_2, k_1)-stable.

We are not only interested in the stability of our update algorithms, but also in the quality of the solutions they provide. We measure this in the usual way, by considering the approximation ratio of the solution. As mentioned, we are interested in trade-offs between the stability of an algorithm and its approximation ratio. Of particular interest are so-called stable approximation schemes, defined as follows.

▶ **Definition 1.** *A* stable approximation scheme*, or* SAS *for short, is an update algorithm* ALG *that, for any given yet fixed parameter $\varepsilon > 0$, is $k(\varepsilon)$-stable and that maintains a solution with approximation ratio $1 + \varepsilon$, where the stability parameter $k(\varepsilon)$ only depends on ε and not on the size of P.*

Notice that in the definition of a SAS we do not take the computational complexity of the update algorithm into account. We point out that, in the context of dynamic scheduling problems (where jobs arrive and disappear in an online fashion, and it is allowed to re-assign jobs), a related concept has been introduced under the name *robust PTAS*: a polynomial-time algorithm that, for any given parameter $\varepsilon > 0$, computes a $(1 + \varepsilon)$-approximation with re-assignment costs only depending on ε, see e.g. [19] and [20].

We now present our results. Recall that $\text{cost}_\alpha(\rho(P)) := \sum_{p \in P} \rho(p)^\alpha$, is the cost of a range assignment ρ, where $\alpha > 1$ is a constant. To make the results easier to interpret, we state the results for $\alpha = 2$; the dependencies of the bounds on the parameter α can be found in the theorems presented in later sections.

- In Section 3 we present a SAS for the broadcast range-assignment problem in $\mathbb{R}^1$, with $k(\varepsilon) = O(1/\varepsilon)$. We prove that this is tight in the worst case, by showing that any SAS for the problem must have $k(\varepsilon) = \Omega(1/\varepsilon)$.

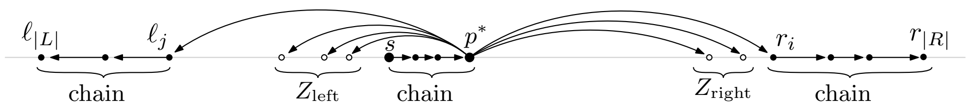

Figure 1 The structure of an optimal solution. The non-filled points are zero-range points, the solid black points all have a standard range (for $\ell_{|L|}$ and $r_{|R|}$ the standard range is zero), except for the root-crossing point which (in this example) has a long range.

- Our SAS (as well as some other algorithms) needs to know an optimal solution after each update. The fastest existing algorithms to compute an optimal solution in $\mathbb{R}^1$ run in $O(n^2)$ time. In Section 2 we show how to recompute an optimal solution in $O(n \log n)$ time after each update, which we believe to be of independent interest. As a result, our SAS also runs in $O(n \log n)$ time per update.
- There is a very simple 2-stable 2-approximation algorithm. We show that a 1-stable algorithm with bounded approximation ratio does not exist when both insertions and deletions must be handled. For the insertion-only case, however, we give a 1-stable $(6 + 2\sqrt{5})$-approximation algorithm. We have not been able to improve upon the approximation ratio 2 with a 2-stable algorithm, but we show that with a 3-stable we can get a 1.97-approximation. Due to lack of space, these results are mostly delegated to the appendix.
- Next we study the problem in $\mathbb{S}^1$, that is, when the underlying 1-dimensional space is circular. This version has, as far as we know, not been studied so far. We first prove that in $\mathbb{S}^1$ an optimal solution for the static problem can always be obtained by cutting the circle at an appropriate point and solving the resulting problem in $\mathbb{R}^1$. This leads to an algorithm to solve the static problem optimally in $O(n^2 \log n)$ time. We also prove that, in spite of this, a SAS does not exist in $\mathbb{S}^1$.
- Finally, we consider the problem in $\mathbb{R}^2$. Based on the no-SAS proof in $\mathbb{S}^1$, we show that the 2-dimensional problem does not admit a SAS either. In addition, we present an 17-stable 12-approximation algorithm for the 2-dimensional version of the problem.

All omitted results and proofs are there in the full version of the paper [14].

2 Maintaining an optimal solution in $\mathbb{R}^1$

Before we can present our stable algorithms for the broadcast range-assignment problem in $\mathbb{R}^1$, we first introduce some terminology and we discuss the structure of optimal solutions. We also present an efficient subroutine to maintain an optimal solution.

2.1 The structure of an optimal solution

Several papers have characterized the structure of optimal broadcast range assignments in $\mathbb{R}^1$, in a more or less explicit manner. We use the characterization by Caragiannis et al. [4], which is illustrated in Figure 1 and described next.

Let $P := L \cup \{s\} \cup R$ be a point set in $\mathbb{R}^1$. Here s is the designated source node, $L := \{\ell_1, \dots, \ell_{|L|}\}$ contains all points from P to the left of s, and $R := \{r_1, \dots, r_{|R|}\}$ contains all points to the right of s. The points in L are numbered in order of increasing distance from s, and the same is true for the points in R. The points $\ell_{|L|}$ and $r_{|R|}$ are called *extreme points*. In the following, and with a slight abuse of notation, we sometimes use p or q to refer a generic point from P – that is, a point that could be s, or a point from R, or a point from L. Furthermore, we will not distinguish between points in P and the corresponding nodes in the communication graph $\mathcal{G}_\rho(P)$.

For a non-extreme point $r_i \in R$, we define r_{i+1} to be its *successor*; similarly, ℓ_{i+1} is the successor of ℓ_i. The source s has (at most) two successors, namely r_1 and ℓ_1. The successor of a point p is denoted by $\text{succ}(p)$; for an extreme point p we define $\text{succ}(p) = \text{NIL}$. If $\text{succ}(p) = q \neq \text{NIL}$, the we call p the *predecessor* of q and we write $\text{pred}(q) = p$. A *chain* is a path in the communication graph $\mathcal{G}_\rho(P)$ that only consists of edges connecting a point to its successor. Thus a chain either visits consecutive points from $\{s\} \cup R$ from left to right, or it visits consecutive points from $\{s\} \cup L$ from right to left. It will be convenient to consider the empty path from s to itself to be a chain as well.

Consider a range assignment ρ. We say that a point $q \in P$ is *within reach* of a point $p \in P$ if $|pq| \leqslant \rho(p)$. Let $\mathcal{B}$ a broadcast tree in $\mathcal{G}_\rho(P)$ – that is, $\mathcal{B}$ is an arborescence rooted at s. A point in $R \cup L$ in $\mathcal{B}$ is called *root-crossing* in $\mathcal{B}$ if it has a child on the other side of s; the source s is root-crossing if it has a child in L and a child in R. The following theorem, which holds for any distance-power gradient $\alpha > 1$, is proven in [4].

▶ **Theorem 2** ([4]). *Let P be a point set in $\mathbb{R}^1$. If all points in $P \setminus \{s\}$ lie to the same side of the source s, then the optimal solution induces a chain from s to the extreme point in P. Otherwise, there is an optimal range assignment ρ such that $\mathcal{G}_\rho(P)$ contains a broadcast tree $\mathcal{B}$ with the following structure:*

- *$\mathcal{B}$ has a single root-crossing point, p^*.*
- *$\mathcal{B}$ contains a chain from s to p^*.*
- *All points within reach of p^*, except those on the chain from s to p^*, are children of p^*.*
- *Let r_i and ℓ_j be the rightmost and leftmost point within reach of p^*, respectively. Then $\mathcal{B}$ contains a chain from r_i to $r_{|R|}$, and a chain from ℓ_j to $\ell_{|L|}$.*

From now on, whenever we talk about optimal range assignments and their induced broadcast trees, we implicitly assume that the broadcast tree has the structure described in Theorem 2. Note that the communication graph $\mathcal{G}_\rho(P)$ induced by an optimal range assignment ρ can contain more edges than the ones belonging to the broadcast tree $\mathcal{B}$. Obviously, for ρ to be optimal it must be a minimum-cost assignment inducing $\mathcal{B}$.

Define the *standard range* of a non-extreme point $r_i \in R$ to be $|r_i r_{i+1}|$; the standard range of the extreme point $r_{|R|}$ is defined to be zero. The standard ranges of the points in L are defined similarly. The source s has two standard ranges, $|s\ell_1|$ and $|sr_1|$. A range assignment in which every point has a standard range is called a *standard solution*; a standard solution may or may not be optimal. Note that, in the static problem, it is never useful to give a point a non-zero range that is smaller than its standard range. Hence, we only need to consider three types of points: *standard-range points*, *zero-range points*, and *long-range points*. Here zero-range points are non-extreme points with a zero range, and a point is said to have a *long range* if its range is greater than its standard range. Theorem 2 implies that an optimal range assignment has the following properties; see also Figure 1.

- There is at most one long-range point.
- The set $Z \subset P$ of zero-range points (which may be empty) can be partitioned into two subsets, Z_{left} and Z_{right}, such that Z_{left} consists of consecutive points that lie to the left of the source s, and Z_{right} consists of consecutive points to that lie to the right of s.

2.2 An efficient update algorithm

Using Theorem 2 an optimal solution for the broadcast range-assignment problem can be computed in $O(n^2)$ time [11]. Below we show that maintaining an optimal solution under insertions and deletions can be done more efficiently than by re-computing it from scratch: using a suitable data structure, we can update the solution in $O(n \log n)$ time. This will also be useful in later sections, when we give algorithms that maintain a stable solution.

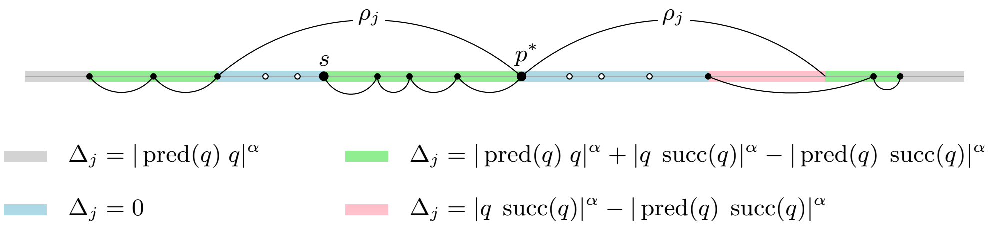

Figure 2 Various cases that can arise when a new point q is inserted into P. Open disks indicate zero-range points. The arcs indicate the ranges of the points before the insertion of q, where the range of the root-crossing point is drawn both to its right and to its left. The colored intervals relate the possible locations of q to the corresponding values Δ_j, where Δ_j refers to the (signed) difference of the cost of the range assignment before and after the insertion of q.

Recall that an optimal solution for a given point set P has a single root-crossing point, p^*. Once the range $\rho(p^*)$ is fixed, the solution is completely determined. Since $\rho(p^*) = |p^*p|$ for some point $p \neq p^*$, there are $n-1$ candidate ranges for a given choice of the root-crossing point p^*. The idea of our solution is to implicitly store the cost of the range assignment for each candidate range of p^* such that, upon the insertion or deletion of a point in P, we can find the best range for p^* in $O(\log n)$ time. By maintaining n such data structures $\mathcal{T}_{p^*}$, one for each choice of the root-crossing point p^*, we can then find the overall best solution.

The data structure for a given root-crossing point. Next we explain our data structure for a given candidate root-crossing point p^*. We assume without loss of generality that p^* lies to the right of the source point s; it is straightforward to adapt the structure to the (symmetric) case where p^* lies to the left of s, and to the case where $p^* = s$.

Let $\mathcal{R}_{p^*}$ be the set of all ranges we need to consider for p^*, for the current set P. The range of a root-crossing point must extend beyond the source point. Hence,

$$\mathcal{R}_{p^*} := \{|p^*p| : p \in P \text{ and } |p^*p| > |p^*s|\}.$$

Let $\lambda_1, \ldots, \lambda_m$ denote the sequence of ranges in $\mathcal{R}_{p^*}$, ordered from small to large. (If $\mathcal{R}_{p^*} = \emptyset$, there is nothing to do and our data structure is empty.) As mentioned, once we fix a range λ_j for the given root-crossing point p^*, the solution is fully determined by Theorem 2: there is a chain from s to p^*, a chain from the rightmost point within range of p^* to the right-extreme point, and a chain from the leftmost point within range of p^* to the left-extreme point. We denote the resulting range assignment[1] for P by $\Gamma(P, p^*, \lambda_j)$.

Our data structure, which implicitly stores the costs of the range assignments $\Gamma(P, p^*, \lambda_j)$ for all $\lambda_j \in \mathcal{R}_{p^*}$, is an augmented balanced binary search tree $\mathcal{T}_{p^*}$. The key to the efficient maintenance of $\mathcal{T}_{p^*}$ is that, upon the insertion of a new point p (or the deletion of an existing point), many of the solutions change in the same way. To formalize this, let Δ_j be the signed difference of the cost of the range assignment $\Gamma(P, p^*, \lambda_j)$ before and after the insertion of q, where Δ_j is positive if the cost increases. Figure 2 shows various possible values for Δ_j, depending on the location of the new point q with respect to the range λ_j. It follows from the figure that there are only four possible values for Δ_j. This allows us to design our data structure $\mathcal{T}_{p^*}$ such that it can be updated using $O(1)$ *bulk updates* of the following form:

[1] When P lies completely to one side of s, then the range assignment is formally not root-crossing. We permit ourselves this slight abuse of terminology because by considering s as root-crossing point, setting $\rho(s) := |s\,\mathrm{succ}(s)|$ and adding a chain from $\mathrm{succ}(s)$ to the extreme point, we get an optimal solution.

Given an interval I of range values and an update value Δ, add Δ to the cost of $\Gamma(P, p^*, \lambda_j)$ for all $\lambda_j \in I$.

In the full version of the paper [14]. we define the information stored in $\mathcal{T}_{p^*}$ and we show how bulk updates can be done in $O(\log n)$ time. We eventually obtain the following theorem.

▶ **Theorem 3.** *An optimal solution to the broadcast range-assignment problem for a point set P in $\mathbb{R}^1$ can be maintained in $O(n \log n)$ per insertion and deletion, where n is the number of points in the current set P.*

3 A stable approximation scheme in $\mathbb{R}^1$

In this section we use the structure of an optimal solution provided by Theorem 2 to obtain a SAS for the 1-dimensional broadcast range-assignment problem. Our SAS has stability parameter $k(\varepsilon) = O((1/\varepsilon)^{1/(\alpha-1)})$, which we will show to be asymptotically optimal.

The optimal range assignment can be very unstable. Indeed, suppose the current point set is $P := \{s, r_1, \ldots, r_n\}$ with $s = 0$ and $r_i = i$ ($1 \leqslant i \leqslant n$), and take any $\alpha > 1$. Then the (unique) optimal assignment $\rho_{\rm opt}$ has $\rho_{\rm opt}(s) = \rho_{\rm opt}(r_1) = \cdots = \rho_{\rm opt}(r_{n-1}) = 1$ and $\rho_{\rm opt}(r_n) = 0$. If now the point $\ell_1 = -n$ is inserted, then the optimal assignment becomes $\rho_{\rm opt}(s) = n$ and $\rho_{\rm opt}(r_1) = \cdots = \rho_{\rm opt}(r_n) = \rho_{\rm opt}(\ell_1) = 0$, causing n ranges to be modified.

Next, we will define a feasible solution, referred to as a *canonical range assignment* ρ_k that is more stable than an optimal assignment, while still having a cost close to the cost of an optimal solution. Here k is a parameter that allows a trade-off between stability and quality of the solution. The assignment ρ_k for a given point set P will be uniquely determined by the set P–it does not depend on the order in which the points have been inserted or deleted. This means that the update algorithm simply works as follows. Let $\rho_k(P)$ be the canonical range assignment for a point set P, and suppose we update P by inserting a point q. Then the update algorithm computes $\rho_k(P \cup \{q\})$ and it modifies the range of each point $p \in P \cup \{q\}$ whose canonical range in $\rho_k(P \cup \{q\})$ is different from its canonical range in $\rho_k(P)$. The goal is now to specify ρ_k such that (i) many ranges in $\rho_k(P \cup \{q\})$ are the same as in $\rho_k(P)$, (ii) the cost of $\rho_k(P)$ is close to the cost of $\rho_{\rm opt}(P)$.

The instance in the example above shows that there can be many points whose range changes from being standard to being zero (or vice versa) when preserving optimality of the consecutive instances. Our idea is therefore to construct solutions where the number of points with zero range is limited, and instead give many points their standard range; if we do this for points whose standard range is relatively small, then the cost of this solution remains bounded compared to the cost of an optimum solution. We now make this idea precise.

Consider a point set P and let $\rho_{\rm opt}$ be an optimal range assignment satisfying the structure described in Theorem 2. Assuming there are points in P on both sides of the source, $\rho_{\rm opt}$ induces a broadcast tree $\mathcal{B}$ with the structure depicted in Figure 1. Let $\rho_{\rm st}(p)$ be the standard range of a point p. The canonical range assignment ρ_k is now defined as follows.

- If all points from P lie to the same side of s, then $\rho_k(p) := \rho_{\rm opt}(p)$ for all $p \in P$. Note that in this case $\rho_k(p) = \rho_{\rm st}(p)$ for all $p \in P$.
- Otherwise, let Z be the set of zero-range points in $\rho_{\rm opt}(P)$. If $|Z| \leqslant k$ then let $Z_k := Z$; otherwise let $Z_k \subseteq Z$ be the k points from Z with the largest standard ranges, with ties broken arbitrarily. We define ρ_k as follows.
  - $\rho_k(p) := \rho_{\rm opt}(p)$ for all $p \in P \setminus Z$. Observe that this means that $\rho_k(p) = \rho_{\rm st}(p)$ for all $p \in P \setminus Z$ except (possibly) for the root-crossing point.
  - $\rho_k(p) := 0$ for all $p \in Z_k$.
  - $\rho_k(p) := \rho_{\rm st}(p)$ for all $p \in Z \setminus Z_k$.

Notice that ρ_k is a feasible solution since $\rho_k(p) \geqslant \rho_{\text{opt}}(p)$ for each $p \in P$. The next lemma analyzes the stability of the canonical range assignment ρ_k. Recall that for any range assignment ρ – hence, also for ρ_k – and any point q not in the current set P, we have $\rho(q) = 0$ by definition.

▶ **Lemma 4.** *Consider a point set P and a point $q \notin P$. Let $\rho_{\text{old}}(p)$ be the range of a point p in $\rho_k(P)$ and let $\rho_{\text{new}}(p)$ be the range of p in $\rho_k(P \cup \{q\})$. Then*

$$|\{p \in P \cup \{q\} : \rho_{\text{new}}(p) > \rho_{\text{old}}(p)\}| \leqslant k+3 \text{ and } |\{p \in P \cup \{q\} : \rho_{\text{new}}(p) < \rho_{\text{old}}(p)\}| \leqslant k+3.$$

Proof. The range of a point $p \in P \cup \{q\}$ can increase due to the insertion of q only if
(i) $p = q$ and $\rho_{\text{new}}(q) > 0$, or
(ii) p is a zero-range point in $\rho_k(P)$, or
(iii) p is the root-crossing point in $\rho_k(P \cup \{q\})$, or
(iv) the standard range of p increases due to the insertion of q, or
(v) $p = s$ and, out of the two standard ranges it has, s gets assigned a larger one in $\rho_k(P\cup\{q\})$ than in $\rho_k(P)$.

Recall that we defined ρ_k such that the number of zero-range points is at most k. Furthermore, at most one standard range can increase due to the insertion of q, namely, the standard range of a point that is extreme in P but not in $P \cup \{q\}$. When this happens, however, q is extreme in $P \cup \{q\}$ and so $\rho_{\text{new}}(q) = 0$; this implies that cases (i) and (iv) cannot both happen. Hence, $|\{p \in P \cup \{q\} : \rho_{\text{new}}(p) > \rho_{\text{old}}(p)\}| \leqslant k + 3$.

The range of a point p can decrease only if
(i) p is a zero-range point in $\rho_k(P \cup \{q\})$, or
(ii) p is the root-crossing point in $\rho_k(P)$, or
(iii) the standard range of p decreases due to the insertion of q, or
(iv) $p = s$ and, out of the two standard ranges it has, p gets a assigned a smaller one in $\rho_k(P \cup \{q\})$ than in $\rho_k(P)$.

Since the only point whose standard range decreases is the predecessor of q in P, we conclude that $|\{p \in P \cup \{q\} : \rho_{\text{new}}(p) < \rho_{\text{old}}(p)\}| \leqslant k + 3$. ◀

Next we bound the approximation ratio of ρ_k.

▶ **Lemma 5.** *For any set P and any $\alpha > 1$, we have $\text{cost}_\alpha(\rho_k(P)) \leqslant \left(1 + \frac{2^\alpha}{k^{\alpha-1}}\right) \cdot \text{cost}_\alpha(\rho_{\text{opt}}(P))$.*

Proof. If all points in P lie to the same side of s then $\rho_k(P) = \rho_{\text{opt}}(P)$, and we are done. Otherwise, let p^* be the root-crossing point. The only points receiving a different range in $\rho_k(P)$ when compared to $\rho_{\text{opt}}(P)$ are the points in $Z \setminus Z_k$; these points have $\rho_k(p) = \rho_{\text{st}}(p)$ while $\rho_{\text{opt}}(p) = 0$. This means we are done when $Z \setminus Z_k = \emptyset$. Thus we can assume that $|Z| > k$, so $Z \setminus Z_k \neq \emptyset$. Assume without loss of generality that $\rho_{\text{opt}}(p^*) = 1$. As each $p \in Z$ is within reach of p^*, we have $\sum_{p \in Z} \rho_{\text{st}}(p) \leqslant 2$. Since Z_k contains the k points with the largest standard ranges among the points in Z, we have $\max\{\rho_{\text{st}}(p) : p \in Z \setminus Z_k\} \leqslant 2/k$. Hence,

$$\sum_{p \in Z \setminus Z_k} \rho_k(p)^\alpha = \sum_{p \in Z \setminus Z_k} \rho_{\text{st}}(p)^\alpha = \sum_{p \in Z \setminus Z_k} \rho_{\text{st}}(p)^{\alpha-1} \cdot \rho_{\text{st}}(p) \leqslant \left(\frac{2}{k}\right)^{\alpha-1} \sum_{p \in Z \setminus Z_k} \rho_{\text{st}}(p) \leqslant \frac{2^\alpha}{k^{\alpha-1}}.$$

(The analysis can be made tighter by using that $\sum_{p \in Z \setminus Z_k} \rho_{\text{st}}(p) \leqslant 2 - k \max_{p \in Z \setminus Z_k} \rho_{\text{st}}(p)$, but this will not change the approximation ratio asymptotically.) We conclude that

$$\frac{\text{cost}_\alpha(\rho_k(P))}{\text{cost}_\alpha(\rho_{\text{opt}}(P))} \leqslant \frac{\sum_{p \in P \setminus (Z \setminus Z_k)} \rho_k(p)^\alpha + \sum_{p \in Z \setminus Z_k} \rho_k(p)^\alpha}{\sum_{p \in P \setminus (Z \setminus Z_k)} \rho_{\text{opt}}(p)^\alpha} \leqslant 1 + \frac{2^\alpha}{k^{\alpha-1}},$$

where the last inequality follows because we have $\rho_k(p) = \rho_{\text{opt}}(p)$ for all $p \in P \setminus (Z \setminus Z_k)$ and $\sum_{p \in P \setminus (Z \setminus Z_k)} \rho_{\text{opt}}(p)^\alpha \geqslant 1$. ◀

By maintaining the canonical range assignment ρ_k for $k = (2^\alpha/\varepsilon)^{1/(\alpha-1)} = O((1/\varepsilon)^{1/(\alpha-1)})$ we obtain the following theorem.

▶ **Theorem 6.** *There is a SAS for the dynamic broadcast range-assignment problem in $\mathbb{R}^1$ with stability parameter $k(\varepsilon) = O((1/\varepsilon)^{1/(\alpha-1)})$, where $\alpha > 1$ is the distance-power gradient. The time needed by the SAS to compute the new range assignment upon the insertion or deletion of a point is $O(n \log n)$, where n is the number of points in the current set.*

Proof. Our SAS maintains the canonical range assignment ρ_k for $k = (2^\alpha/\varepsilon)^{1/(\alpha-1)} = O((1/\varepsilon)^{1/(\alpha-1)})$. We then have $\text{cost}_\alpha(\rho_k(P)) \leqslant (1+\varepsilon) \cdot \rho_{\text{opt}}(P)$ by Lemma 5. Furthermore, the number of modified ranges when P is updated is $2k + 6$ by Lemma 4. To determine the assignment ρ_k, we need to know an optimal assignment ρ_{opt} with the structure from Theorem 2. Such an optimal assignment can be maintained in $O(n \log n)$ time per update, by Theorem 3. Once we have the new optimal assignment, the new optimal assignment can be determined in $O(n)$ time. ◀

Next we show that the stability parameter $k(\varepsilon)$ in our SAS is asymptotically optimal.

▶ **Theorem 7.** *Any SAS for the dynamic broadcast range-assignment problem in $\mathbb{R}^1$ must have stability parameter $k(\varepsilon) = \Omega((1/\varepsilon)^{1/(\alpha-1)})$, where $\alpha > 1$ is the distance-power gradient.*

Proof. Let ALG be a k-stable algorithm, where $k \geqslant 4$ and $k^{\alpha-1} \geqslant \frac{1}{2^{\alpha+1}(2^{\alpha-1}-1)}$ and k is even, and let ρ_{alg} be the range assignment it maintains. Note that the condition on k is satisfied for k large enough. We will show that the approximation ratio of ALG is at least $1 + \frac{1}{2^{\alpha+2}k^{\alpha-1}}$. Since a SAS has approximation ratio $1+\varepsilon$, this implies that the stability parameter $k(\varepsilon)$ of ALG must satisfy $k(\varepsilon) = \Omega((1/\varepsilon)^{1/(\alpha-1)})$.

Consider the point set $P := \{s, r_1, r_2, \dots r_{2k}\}$, where $s = 0$ and $r_i = i/(2k)$ for $i = 1, 2, \dots, 2k$. We consider two cases.

Case I: The number of zero-range points in $\rho_{\text{alg}}(P)$ is at least $k/2$, where we assume without loss of generality that all points with range less than $1/(2k)$ actually have range zero. The cheapest possible solution in this case is to have exactly $k/2$ zero-range points, k points with range $1/(2k)$, and $k/2$ points with range $1/k$, for a total cost of

$$\text{cost}_\alpha(\rho_{\text{alg}}(P)) \geqslant k \cdot \left(\frac{1}{2k}\right)^\alpha + \frac{k}{2} \cdot \left(\frac{1}{k}\right)^\alpha = \left(1 + \frac{2^{\alpha-1}-1}{2}\right) \cdot 2k \left(\frac{1}{2k}\right)^\alpha .$$

An optimal solution has cost $2k \cdot (1/(2k))^\alpha$, and so the approximation ratio of ALG in Case I is at least $1 + \frac{2^{\alpha-1}-1}{2}$, which is at least $1 + \frac{1}{2^{\alpha+2}k^{\alpha-1}}$ since $k^{\alpha-1} \geqslant \frac{1}{2^{\alpha+1}(2^{\alpha-1}-1)}$.

Case II: The number of zero-range points $\rho_{\text{alg}}(P)$ is less than $k/2$. Now suppose the point $\ell_1 = -1$ arrives. Since $\rho_{\text{alg}}(P)$ had less than $k/2$ zero-range points and ALG can modify at most k ranges, $\rho_{\text{alg}}(P \cup \{\ell_1\})$ has less than $3k/2$ zero-range points. Hence, at least $k/2$ points in $P \cup \{\ell_1\}$ have a range that is at least $1/(2k)$, one of which must have a range at least 1. This implies that $\text{cost}_\alpha(\rho_{\text{alg}}(P \cup \{\ell_1\})) \geqslant 1 + (k/2 - 1) \cdot \left(\frac{1}{2k}\right)^\alpha \geqslant 1 + \frac{1}{2^{\alpha+2}k^{\alpha-1}}$, where the last inequality holds since $k/2 - 1 \geqslant k/4$ (because $k \geqslant 4$). An optimal range assignment on $P \cup \{\ell_1\}$ has $\rho_{\text{opt}}(s) = 1$ and all other ranges equal to zero, for a total cost of 1, and so the approximation ratio of ALG in Case II is at least $1 + \frac{1}{2^{\alpha+2}k^{\alpha-1}}$ as well. ◀

Table 1 An overview of the approximation ratio of 1-stable, 2-stable and 3-stable algorithms.

ℓ-stable algorithm	Approximation Ratio	Remarks
$\ell = 1$	$6 + 2\sqrt{5} \approx 10.47$	$\alpha = 2$, insertions only
$\ell = 2$	2	for any $\alpha > 1$
$\ell = 3$	1.97	$\alpha = 2$

4 1-Stable, 2-Stable, and 3-Stable Algorithms in $\mathbb{R}^1$

In Section 3 we presented a $(2k + 6)$-stable algorithm with approximation ratio $1 + 2^\alpha/k^{\alpha-1}$, which provided us with a SAS. For small k the algorithm is not very good: the most stable algorithm we can get is 6-stable, by setting $k = 0$. A careful analysis shows that the approximation ratio of this 6-stable algorithm is 3, for $\alpha = 2$. Below we briefly discuss the results we obtained for more stable algorithms; details are available in the full version [14].

We give a 1-stable $O(1)$-approximation algorithm; obviously, this is the best we can do in terms of stability. This algorithm can only handle insertions. We also show that this is necessarily the case: a 1-stable algorithm that can handle insertions as well as deletions cannot have bounded approximation ratio. We then present a straightforward 2-stable 2-approximation algorithm, which simply gives every point its standard range. Finally, we study 3-stable algorithms: we show that using a 3-stable algorithm it is possible to get an approximation ratio strictly below 2. See Table 1 for an overview of results. We now briefly sketch our 3-stable algorithm.

A 3-stable algorithm with approximation ratio less than 2. Given the simplicity of our 2-stable 2-approximation algorithm, it is surprisingly difficult to obtain an approximation ratio strictly smaller than 2. In fact, we have not been able to do this with a 2-stable algorithm. Below we show this is possible with a 3-stable algorithm, at least for the case $\alpha = 2$, which we assume from now on.

Recall that for any set P with points on both sides of the source point s, there is an optimal range assignment inducing a broadcast tree with a single root-crossing point; see Figure 1. Unfortunately the root-crossing point may change when P is updated. This may cause many changes if we maintain a solution with a good approximation ratio and the same root-crossing point as the optimal solution. We therefore restrict ourselves to *source-based range assignments*, where s is the root-crossing point. The main question is then how large the range of s should be, and which points within range of s should be zero-range points.

We now define our source-based range assignment, which we denote by ρ_{sb}, more precisely. It will be uniquely defined by the set P; it does not depend on the order in which points have been inserted or deleted. Let δ be a parameter with $1/2 < \delta < 1$; later we will choose δ such that the approximation ratio of our algorithm is optimized. We call a point $p \in P \setminus \{s\}$ *expensive* if $\text{succ}(p) \neq \text{NIL}$ and $|p\, \text{succ}(p)| > \delta \cdot |s\, \text{succ}(p)|$, and we call it *cheap* otherwise. The source s is defined to be always expensive. (This is consistent in the sense that for $p = s$ the condition $|p\, \text{succ}(p)| > \delta \cdot |s\, \text{succ}(p)|$ holds for both successors, since $\delta < 1$.) We denote the set of all expensive points in P by P_{exp} and the set of all cheap points by P_{cheap}. Define $d_{\max} := \max\{|s\, \text{succ}(p)| : p \in P_{\text{exp}}\}$, that is, $d_{\max}$ is the maximum distance from s to the successor of any expensive point. We say that a point $p \in P_{\text{exp}}$ is *crucial* if $|s\, \text{succ}(p)| = d_{\max}$. Typically there is a single crucial point, but there can also be two: one on the left and one on the right of s. Our source-based range assignment ρ_{sb} is now defined as follows.

- $\rho_{\text{sb}}(s) := d_{\max}$,
- $\rho_{\text{sb}}(p) := 0$ for all $p \in P_{\text{exp}} \setminus \{s\}$, and
- $\rho_{\text{sb}}(p) := \rho_{\text{st}}(p)$ for all $p \in P_{\text{cheap}}$, where $\rho_{\text{st}}(p)$ denotes the standard range of a point.

It is easily checked that we can maintain this range assignment with a 3-stable algorithm. The challenge is to analyze its approximation ratio. In the full version of the paper [14]. we show that, for a suitable choice of δ, the approximation ratio is strictly smaller than 2, leading to the following theorem.

▶ **Theorem 8.** *There exists a 3-stable 1.97-approximation algorithm for the dynamic broadcast range-assignment problem in $\mathbb{R}^1$ for $\alpha = 2$.*

5 The problem in $\mathbb{S}^1$

We now turn to the setting where the underlying space is $\mathbb{S}^1$, that is, the points in P lie on a circle and distances are measured along the circle. In Section 5.1, we prove that the structure of an optimal solution in $\mathbb{S}^1$ is very similar to the structure of an optimal solution in $\mathbb{R}^1$ as formulated in Theorem 2. In spite of this, and contrary to the problem in $\mathbb{R}^1$, we prove in Section 5.2 that no SAS exists for the problem in $\mathbb{S}^1$.

Again, we denote the source point by s. The clockwise distance from a point $p \in \mathbb{S}^1$ to a point $q \in \mathbb{S}^1$ is denoted by $d_{\text{cw}}(p, q)$, and the counterclockwise distance by $d_{\text{ccw}}(p, q)$. The actual distance is then $d(p, q) := \min(d_{\text{cw}}(p, q), d_{\text{ccw}}(p, q))$. The closed and open clockwise interval from p to q are denoted by $[p, q]^{\text{cw}}$ and $(p, q)^{\text{cw}}$, respectively.

5.1 The structure of an optimal solution in $\mathbb{S}^1$

Here we prove that the structure of an optimal solution in $\mathbb{S}^1$ is very similar to the structure of an optimal solution in $\mathbb{R}^1$. The heart of this proof is the following lemma 9. Define the *covered region* of P with respect to a range assignment ρ, denoted by $\text{cov}(\rho, P)$, to be the set of all points $r \in \mathbb{S}^1$ such that there exists a point $p \in P$ with $\rho(p) \geqslant d(p, r)$.

▶ **Lemma 9.** *Let P be a point set in $\mathbb{S}^1$ with $|P| > 2$ and let ρ_{opt} be an optimal range assignment for P. Then there exists a point $r \in \mathbb{S}^1$ such that $r \notin \text{cov}(\rho_{\text{opt}}, P)$.*

Lemma 9 implies that an optimal solution for an instance in $\mathbb{S}^1$ corresponds to an optimal solution for an instance in $\mathbb{R}^1$ derived as follows. For a point $r \in \mathbb{S}^1$, define the mapping $\mu_r : P \to \mathbb{R}^1$ such that $\mu_r(s) := 0$, and $\mu_r(p) := d_{\text{cw}}(s, p)$ for all $p \in [s, r]^{\text{cw}}$, and $\mu_r(p) := -d_{\text{ccw}}(s, p)$ for all $p \in [r, s]^{\text{cw}}$. Let $\mu_r(P)$ denote the resulting point set in $\mathbb{R}^1$.

▶ **Theorem 10.** *Let P be an instance of the broadcast range-assignment problem in $\mathbb{S}^1$. There exists a point $r \in \mathbb{S}^1$ such that an optimal range assignment for $\mu_r(P)$ in $\mathbb{R}^1$ induces an optimal range assignment for P. Moreover, we can compute an optimal range assignment for P in $O(n^2 \log n)$ time, where n is the number of points in P.*

Proof. Let $r \in \mathbb{S}^1$ be a point such that $r \notin \text{cov}(\rho_{\text{opt}}, P)$, which exists by Lemma 9. Consider the mapping μ_r. Any feasible range assignment for $\mu_r(P)$ induces a feasible range assignment for P in $\mathbb{S}^1$, since $d(p, q) \leqslant |\mu_r(p)\mu_r(q)|$ for any two points $p, q \in P$. Conversely, an optimal range assignment for P induces a feasible range assignment for $\mu_r(P)$, since the point r is not covered in the optimal solution. This proves the first part of the theorem.

Now let $P := \{s, p_1, \ldots, p_n\}$, where the points p_i are ordered clockwise from s. For $0 \leqslant i \leqslant n$, let r_i be a point in $(p_i, p_{i+1})^{\text{cw}}$, where $p_0 = p_{n+1} = s$. Since $\mu_{r_i} = \mu_r$ for any $r \in (p_i, p_{i+1})^{\text{cw}}$, an optimal solution can be computed by finding the best solution over all mappings μ_{r_i}. The only difference between μ_{r_i} and $\mu_{r_{i+1}}$ is the location that p_{i+1} is mapped to, so after computing an optimal solution for $\mu_1(P)$ in $O(n^2 \log n)$ time, we can go through the mappings $\mu_2, \dots, \mu_n$ and update the optimal solution in $O(n \log n)$ time using Theorem 3. Hence, an optimal range assignment for P can be computed in $O(n^2 \log n)$ time. ◀

Next we prove Lemma 9. Without loss of generality we identify $\mathbb{S}^1$ with a circle of perimeter 1. Let ρ_{opt} be a fixed optimal range-assignment on P. We will need the following lemma.

▶ **Lemma 11.** *If $|P| > 2$ then $\rho_{\text{opt}}(p) < \frac{1}{2}$ for all $p \in P$.*

Proof. Note that setting $\rho(s) = \frac{1}{2}$ and $\rho(p) = 0$ for all $p \in P \setminus \{s\}$ gives a feasible solution. Since $\rho(s) > 0$ in any feasible solution, this means that $\rho_{\text{opt}}(p) < \frac{1}{2}$ for all $p \neq s$. Hence, it suffices to show that $\rho_{\text{opt}}(s) < \frac{1}{2}$. If there is no point $p \in P$ which is diametrically opposite s then clearly $\rho_{\text{opt}}(s) < \frac{1}{2}$. Now suppose some point $p \in P$ lies diametrically opposite s. Let $q \in P \setminus \{s, p\}$ be a point that maximizes the distance from s among all points in $P \setminus \{s, p\}$. The point q exists since $|P| > 2$. Note that $d(s, q) + d(q, p) = \frac{1}{2}$. Hence, setting $\rho(s) = d(s, q)$ and $\rho(q) = d(q, p)$ (and keeping all other ranges zero) gives a solution of cost $d(s, q)^\alpha + d(q, p)^\alpha$, which is less than $\left(\frac{1}{2}\right)^\alpha$ since $\alpha > 1$. Thus $\rho_{\text{opt}}(s) < \frac{1}{2}$, which finishes the proof. ◀

Before we proceed, we introduce some more notation.

For two points $p, q \in \mathbb{S}^1$, we let $[p, q]^{\text{cw}} \subset \mathbb{S}^1$ denote the closed clockwise interval from p to q. In other words, $[p, q]^{\text{cw}}$ is the clockwise arc along $\mathbb{S}^1$ from p to q, including its endpoints. Furthermore, we define $(p, q)^{\text{ccw}}$ to be the open clockwise interval from p to q. The intervals $[p, q]^{\text{ccw}}$ and $(p, q)^{\text{ccw}}$ are defined similarly, but for the counterclockwise direction. Now consider a directed edge (p, q) in a communication graph $\mathcal{G}_\rho(P)$. We say that (p, q) is a *clockwise edge* if $\rho(p) \geqslant d_{\text{cw}}(p, q)$, and we say that it is a *counterclockwise edge* if $\rho(p) \geqslant d_{\text{ccw}}(p, q)$. Lemma 11 implies that an edge cannot be both clockwise and counterclockwise in an optimal range assignment, assuming $|P| > 2$. Finally, we define the *covered region* of a subset $Q \subseteq P$ with respect to a range assignment ρ to be the set of all points $r \in \mathbb{S}^1$ such that there exists a point $p \in Q$ such that $\rho(p) \geqslant d(p, r)$. We denote this region by $\text{cov}(\rho, Q)$. Furthermore, the *counterclockwise covered region* of Q, denoted by $\text{cov}_{\text{ccw}}(\rho, Q)$, is the set of all points $r \in \mathbb{S}^1$ such that there exists a point $p \in Q$ such that $\rho(p) \geqslant d_{\text{ccw}}(p, r)$. The *clockwise covered region* of Q, denoted by $\text{cov}_{\text{cw}}(\rho, Q)$, is defined similarly.

We can now state the main lemma of this section.

▶ **Lemma 9.** *Let P be a point set in $\mathbb{S}^1$ with $|P| > 2$ and let ρ_{opt} be an optimal range assignment for P. Then there exists a point $r \in \mathbb{S}^1$ such that $r \notin \text{cov}(\rho_{\text{opt}}, P)$.*

Proof. Let $d_{\text{hop}}(p, q)$ denote the hop distance from p to q in the communication graph $\mathcal{G}_{\rho_{\text{opt}}}(P)$. Let $\mathcal{B}$ broadcast tree rooted at s in $\mathcal{G}_{\rho_{\text{opt}}}(P)$ with the following properties.

- $\mathcal{B}$ is a shortest-path tree in terms of hop distance, that is, the hop-distance from s to any point p in $\mathcal{B}$ is equal to $d_{\text{hop}}(s, p)$.
- Among all such shortest-path trees, $\mathcal{B}$ maximizes the number of clockwise edges.

For two points $p, q \in P$, let $\pi(p, q)$ denote the path from p to q in $\mathcal{B}$, and let $|\pi(p, q)|$ be its length, that is, the number of edges on the path. Note that $|\pi(s, p)| = d_{\text{hop}}(s, p)$ for any $p \in P$. Let $\text{pa}(p)$ denote the parent of a point p in $\mathcal{B}$ and define

$$S_{\text{cw}} = \{p \in P \setminus \{s\} : \ (\text{pa}(p), p) \text{ is a clockwise edge}\}$$

and

$$S_{\text{ccw}} = \{p \in P \setminus \{s\} : \ (\text{pa}(p), p) \text{ is a counterclockwise edge}\}.$$

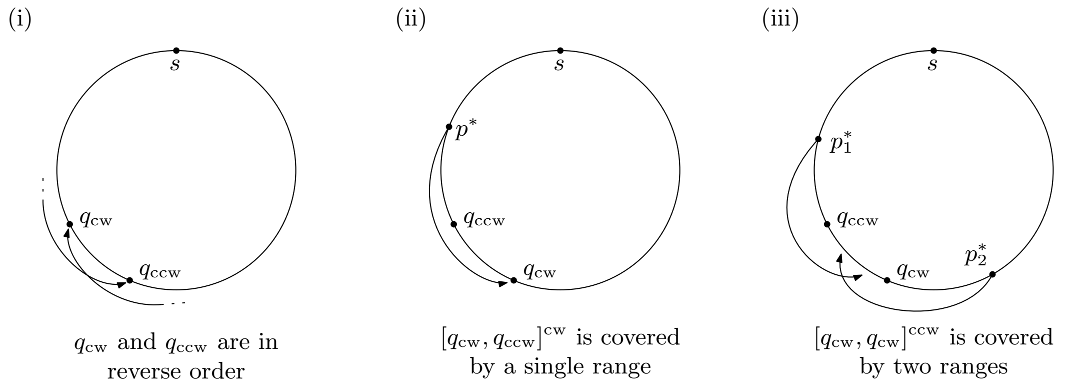

Figure 3 Illustration for the proof of Lemma 9. Note that the point p^* in part (ii) of the figure could also lie in $[s, q_{\text{cw}}]^{\text{cw}}$. Similarly, in part (iii) the points p_1^* and p_2^* could lie on "the other side" of s.

Note that $S_{\text{cw}} \cup S_{\text{ccw}} = P \setminus \{s\}$. Now define

$$q_{\text{cw}} = \text{the point from } S_{\text{cw}} \text{ that maximizes } d_{\text{cw}}(s,p),$$

where $q_{\text{ccw}} = s$ if $S_{\text{cw}} = \emptyset$. Similarly, define

$$q_{\text{ccw}} = \text{the point from } S_{\text{ccw}} \text{ that maximizes } d_{\text{ccw}}(s,p),$$

where $q_{\text{ccw}} = s$ if $S_{\text{ccw}} = \emptyset$. Let $\text{anc}(p)$ be the set of ancestors in $\mathcal{B}$ of a point $p \in P$, that is, $\text{anc}(p)$ contains the points of $\pi(s,p)$ excluding the point p. The following observation will be used repeatedly in the proof.

▷ Observation. If $(\text{pa}(p), p)$ is a clockwise edge, then $[s,p]^{\text{cw}} \subset \text{cov}(\rho_{\text{opt}}, \text{anc}(p))$. Similarly, if $(\text{pa}(p), p)$ is a counterclockwise edge, then $[s,p]^{\text{ccw}} \subset \text{cov}(\rho_{\text{opt}}, \text{anc}(p))$.

Proof. Assume $(\text{pa}(p), p)$ is a clockwise edge; the proof for when $(\text{pa}(p), p)$ is a counterclockwise edge is similar. If $s \in [\text{pa}(p), p)]^{\text{cw}}$ – this includes the case where $\text{pa}(p) = s$ – then the statement obviously holds, so assume $\text{pa}(p) \in [s,p]^{\text{cw}}$. Since $(\text{pa}(p), p)$ is a clockwise edge, it then suffices to prove that $[s, \text{pa}(p)]^{\text{cw}} \subset \text{cov}(\rho_{\text{opt}}, \text{anc}(p))$. Note that $\text{cov}(\rho_{\text{opt}}, \text{anc}(p))$ is connected, because the points in $\text{anc}(p)$ form a path, namely $\pi(s, \text{pa}(p))$. Since $\pi(s,p)$ is shortest path, $p \notin \text{cov}(\rho_{\text{opt}}, \text{anc}(\text{pa}(p))$, which implies that $[s, \text{pa}(p)]^{\text{cw}} \subset \text{cov}(\rho_{\text{opt}}, \text{anc}(\text{pa}(p))) \subset \text{cov}(\rho_{\text{opt}}, \text{anc}(p))$. ◁

We now proceed to show that q_{ccw} must lie clockwise from q_{cw}, as seen from s, that is, the situation shown in Fig. 3(i) cannot happen.

▷ Claim. $d_{\text{cw}}(s, q_{\text{cw}}) + d_{\text{ccw}}(s, q_{\text{ccw}}) < 1$.

Proof. Note that $d_{\text{cw}}(s, q_{\text{cw}}) + d_{\text{ccw}}(s, q_{\text{ccw}}) \neq 1$, since otherwise $q_{\text{cw}} = q_{\text{ccw}}$ which cannot happen since $S_{\text{cw}} \cap S_{\text{ccw}} = \emptyset$.

Now assume for a contradiction that $d_{\text{cw}}(s, q_{\text{cw}}) + d_{\text{ccw}}(s, q_{\text{ccw}}) > 1$, which means that $q_{\text{ccw}} \in [s, q_{\text{cw}}]^{\text{cw}}$. Since q_{cw} is reached from its parent by a clockwise edge, this implies that $q_{\text{ccw}} \in \text{cov}(\rho_{\text{opt}}, \text{anc}(q_{\text{cw}}))$ by the observation above. Hence, $d_{\text{hop}}(s, q_{\text{cw}}) \geqslant d_{\text{hop}}(s, q_{\text{ccw}})$. An analogous argument shows that $d_{\text{hop}}(s, q_{\text{ccw}}) \geqslant d_{\text{hop}}(s, q_{\text{cw}})$. Hence, $d_{\text{hop}}(s, q_{\text{ccw}}) = d_{\text{hop}}(s, q_{\text{cw}})$. This implies that the edge $(\text{pa}(q_{\text{cw}}), q_{\text{cw}})$ passes over q_{ccw}, otherwise some other edge of $\pi(s, q_{\text{cw}})$ would pass over q_{ccw} and we would have $d_{\text{hop}}(s, q_{\text{ccw}}) < d_{\text{hop}}(s, q_{\text{cw}})$. But then we also have a shortest path from s to q_{ccw} whose last edge is a clockwise edge, contradicting the definition of $\mathcal{B}$. ◁

So we can assume that $d_{\mathrm{cw}}(s, q_{\mathrm{cw}}) + d_{\mathrm{ccw}}(s, q_{\mathrm{ccw}}) < 1$ or, in other words, that q_{ccw} lies clockwise from q_{cw}, as seen from s. Clearly no point from P lies in $(q_{\mathrm{cw}}, q_{\mathrm{ccw}})^{\mathrm{cw}}$. If we have $(q_{\mathrm{cw}}, q_{\mathrm{ccw}})^{\mathrm{cw}} \not\subset \mathrm{cov}(\rho_{\mathrm{opt}}, P)$ then we are done, so assume for a contradiction that $(q_{\mathrm{cw}}, q_{\mathrm{ccw}})^{\mathrm{cw}} \subset \mathrm{cov}(\rho_{\mathrm{opt}}, P)$. This can happen in three ways, each of which will lead to a contradiction.

Case I: There exists a point $p^* \in \mathcal{B}$ such that $q_{\mathrm{cw}} \in \mathrm{cov}_{\mathrm{ccw}}(\rho_{\mathrm{opt}}, \{p^*\})$.

See Fig. 3(ii) for an illustration of the situation. If $p^* = s$ then $d_{\mathrm{hop}}(s, q_{\mathrm{cw}}) = 1$. Since $q_{\mathrm{cw}} \in S_{\mathrm{cw}}$ this means that q_{cw} must also have an incoming clockwise edge from s. But then $\rho_{\mathrm{opt}}(s) \geqslant \frac{1}{2}$, which contradicts Lemma 11. So $p^* \neq s$. Now note that p^* must have an outgoing clockwise edge in $\mathcal{B}$, else we can reduce the range of p^* to $d_{\mathrm{ccw}}(p^*, q_{\mathrm{ccw}})$, which is smaller than $d_{\mathrm{ccw}}(p^*, q_{\mathrm{cw}})$, and still get a feasible solution. Observe that $p^* \notin \pi(s, q_{\mathrm{cw}})$; otherwise we must have $p^* = \mathrm{pa}(q_{\mathrm{cw}})$ (since q_{cw} lies in the range of p^*) which contradicts that $q_{\mathrm{cw}} \in S_{\mathrm{cw}}$. So for any point from P in the region $[s, q_{\mathrm{cw}}]^{\mathrm{cw}}$ there exists a path from s in the communication graph induced by ρ_{opt} that does not use p^*. We now have two subcases.

If $p^* \in [s, q_{\mathrm{ccw}}]^{\mathrm{ccw}}$ then clearly $p^* \in S_{\mathrm{ccw}}$ (otherwise the definition of q_{cw} is contradicted). Hence, each point from P in the region $[s, p^*]^{\mathrm{ccw}}$ has a path from s that does not use p^*. This implies that can reduce the range of p^* to $d_{\mathrm{ccw}}(p^*, q_{\mathrm{ccw}})$ and still get a feasible solution.

If $p^* \in [s, q_{\mathrm{cw}}]^{\mathrm{cw}}$ then obviously we can also reduce the range of p^* to $d_{\mathrm{ccw}}(p^*, q_{\mathrm{ccw}})$ and still get a feasible solution.

So both subcases lead to the desired contradiction.

Case II: There exists a point $p^* \in \mathcal{B}$ such that $q_{\mathrm{ccw}} \in \mathrm{cov}_{\mathrm{cw}}(\rho_{\mathrm{opt}}, \{p^*\})$.

In the proof of Case I we never used that $\mathcal{B}$ maximizes the number of clockwise edges. Hence, a symmetric argument shows that Case II also leads to a contradiction.

Case III: There are two points $p_1^*, p_2^* \in P$ such that $[q_{\mathrm{cw}}, q_{\mathrm{ccw}}]^{\mathrm{cw}} \subseteq \mathrm{cov}_{\mathrm{ccw}}(\rho_{\mathrm{opt}}, \{p_1^*\}) \cup \mathrm{cov}_{\mathrm{cw}}(\rho_{\mathrm{opt}}, \{p_2^*\})$.

See Fig. 3(iii) for an illustration of the situation. We can assume that $q_{\mathrm{cw}} \notin \mathrm{cov}_{\mathrm{ccw}}(\rho_{\mathrm{opt}}, \{p_1^*\})$ and $q_{\mathrm{ccw}} \notin \mathrm{cov}_{\mathrm{cw}}(\rho_{\mathrm{opt}}, \{p_2^*\})$, otherwise we are in Case I or Case II. Now either $p_2^* \notin \pi(s, p_1^*)$ or $p_1^* \notin \pi(s, p_2^*)$ or both. Without loss of generality, assume $p_2^* \notin \pi(s, p_1^*)$. Then $p_2^* \neq s$ and all points from P in the region $[s, q_{\mathrm{ccw}}]^{\mathrm{ccw}}$ have a path from s in the communication graph $\mathcal{G}_{\rho_{\mathrm{opt}}}(P)$ that does not use p_2^*. The point p_2^* must have an outgoing counterclockwise edge, else we can reduce the range of p_2^* to $d_{\mathrm{cw}}(p_2^*, q_{\mathrm{cw}})$ and still get a feasible solution. We have two subcases.

If $p_2^* \in [s, q_{\mathrm{ccw}}]^{\mathrm{ccw}}$ then by reducing the range of p_2^* to $d_{\mathrm{cw}}(p_2^*, q_{\mathrm{cw}})$ we still get a feasible solution.

If $p_2^* \in [s, q_{\mathrm{cw}}]^{\mathrm{cw}}$ then p_2^* must be reached by a clockwise edge from its parent in $\mathcal{B}$, otherwise the definition of q_{ccw} would be contradicted. Hence, for each point from P in the region $[s, p_2^*]^{\mathrm{cw}}$ there is a path from s that does not use p_2^*. So again we can reduce the range of p_2^* to $d_{\mathrm{cw}}(p_2^*, q_{\mathrm{cw}})$ we still get a feasible solution.

Thus both subcases lead to a contradiction.

This finishes the proof of the lemma. ◀

5.2 Non-existence of a SAS in $\mathbb{S}^1$

We have seen that an optimal solution for a set P in $\mathbb{S}^1$ can be obtained from an optimal solution in $\mathbb{R}^1$, if we cut $\mathbb{S}^1$ at an appropriate point r. It is a fact however that the insertion of a new point into P may cause the location of the cutting point r to change drastically. Next we show that this means that the dynamic problem in $\mathbb{S}^1$ does not admit a SAS.

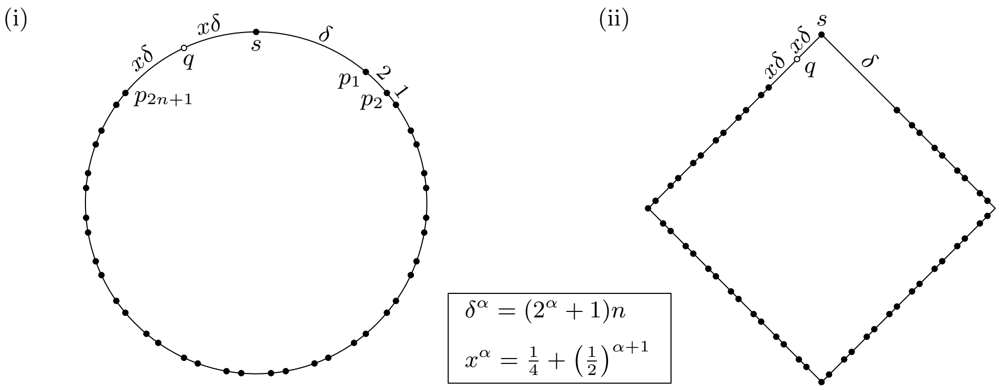

Figure 4 (i) The instance showing that there is no SAS in $\mathbb{S}^1$. (ii) The instance in $\mathbb{R}^2$.

▶ **Theorem 12.** *The dynamic broadcast range-assignment problem in $\mathbb{S}^1$ with distance power gradient $\alpha > 1$ does not admit a SAS. In particular, there is a constant $c_\alpha > 1$ such that the following holds: for any n large enough, there is a set $P := \{s, p_1, \dots, p_{2n+1}\}$ and a point q in $\mathbb{S}^1$ such that any update algorithm* ALG *that maintains a c_α-approximation must modify more than $2n/3 - 1$ ranges upon the insertion of q into P.*

The rest of this section is dedicated to proving Theorem 12. We will prove the theorem for

$$c_\alpha := \min\left(1 + 2^{\alpha-4} - \frac{1}{8},\ 1 + \frac{2^{\alpha-1} - 1}{3 \cdot 2^\alpha + 2},\ 1 + \frac{\min\left(2^\alpha - 1, \frac{3^\alpha - 2^\alpha - 1}{2}, \frac{4^\alpha - 2^\alpha - 2}{3}\right)}{4(2^\alpha + 1)}\right).$$

Note that each term is a constant strictly greater than 1 for any fixed constant $\alpha > 1$. In particular, for $\alpha = 2$ we have $c_\alpha = 1 + \frac{1}{14}$.

Let $P := \{s, p_1, \dots, p_{2n+1}\}$, where $d_{\mathrm{cw}}(p_i, p_{i+1}) = 2$ for odd i and $d_{\mathrm{cw}}(p_i, p_{i+1}) = 1$ for even i; see Fig. 4(i). Let $d_{\mathrm{cw}}(s, p_1) = \delta$, where $\delta^\alpha = (2^\alpha + 1)n$. Finally, let $d_{\mathrm{cw}}(p_{2n+1}, q) = d_{\mathrm{cw}}(q, s) = x\delta$, where $x^\alpha = \frac{1}{4} + \left(\frac{1}{2}\right)^{\alpha+1}$. Note that $(1/2)^\alpha < x^\alpha < 1/2$ for any $\alpha > 1$.

Let $\rho(p)$ denote the range given to a point p by ALG. A directed edge (p, p') in the communication graph induced by ρ is called a *clockwise edge* if $\rho(p) \geqslant d_{\mathrm{cw}}(p, p')$, and it is called a *counterclockwise edge* if $\rho(p) \geqslant d_{\mathrm{ccw}}(p, p')$. Observe that we may assume that no edge (p, p') is both clockwise and counterclockwise, because otherwise $\rho(p) \geqslant (\delta + 3n + 2x\delta)/2$, which is much too expensive for an approximation ratio of at most c_α. Define the range $\rho(p)$ of a point in P to be CW-*minimal* if $\rho(p)$ equals the distance from p to its clockwise neighbor in P. Similarly, $\rho(p)$ is CCW-*minimal* if $\rho(p)$ equals the distance from p to its counterclockwise neighbor. The idea of the proof is to show that before the insertion of q, most of the points $s, p_1, \dots, p_{2n+1}$ must have a CW-minimal range, while after the insertion most points must have a CCW-minimal range. This will imply that many ranges must be modified from being CW-minimal to being CCW-minimal.

Before the insertion of q, giving every point a CW-minimal range leads to a feasible assignment of total cost $\delta^\alpha + (2^\alpha + 1)n = 2\delta^\alpha$. After the insertion of q, giving every point a CCW-minimal range leads to a feasible assignment of total cost $2(x\delta)^\alpha + (2^\alpha + 1)n = (2x^\alpha + 1)\delta^\alpha$. Hence, if OPT$(\cdot)$ denotes the cost of an optimal range assignment, then we have:

▶ **Observation 13.** OPT$(P) \leqslant 2\delta^\alpha$ *and* OPT$(P \cup \{q\}) \leqslant (2x^\alpha + 1)\delta^\alpha < 2\delta^\alpha$.

We first prove a lower bound on the total cost of the points $p_1, \ldots, p_{2n+1}$. Intuitively, only $o(n)$ of those points can be reached from s or q (otherwise the range of s or q would be too expensive) and the cheapest way to reach the remaining points will be to use only CW-minimal or CCW-minimal ranges.

▶ **Lemma 14.** *$\sum_{i=1}^{2n+1} \rho(p_i)^\alpha \geqslant (2^\alpha + 1)n - o(n)$, both before and after the insertion of q.*

Proof. By Observation 13, we have $\rho(p)^\alpha \leqslant c_\alpha \cdot 2\delta^\alpha$ and, hence, $\rho(p) \leqslant (2c_\alpha)^{1/\alpha} \cdot \delta < 3\delta$, for any point p. Consider the interval $I = [y_1, y_2]^{\text{cw}}$ where $d_{\text{cw}}(s, y_1) = 3\delta$ and $d_{\text{ccw}}(q, y_2) = 3\delta$. All the points in $I \cap P$ are at distance more than 3δ from s or q and hence $I \cap P \subseteq \text{cov}(\rho_{\text{opt}}, P \setminus \{s, q\})$. Let $p_i \in I \cap P$ be the point whose clockwise distance from s is minimum, and let $p_j \in I \cap P$ be the point whose counterclockwise distance from q is minimum. Then the cost of covering all the points in $I \cap P$ using the points in $P \setminus \{s, q\}$ is at least $\sum_{t=i}^{j-1} d_{\text{cw}}(p_t, p_{t+1})^\alpha - 2^\alpha$, where the term -2^α is because the covered region may leave one interval $[p_t, p_{t+1}]^{\text{cw}}$ uncovered. Recall that the cost of assigning all the points in $P \setminus \{s, q\}$ a CW-minimal range is $(2^\alpha + 1)n$. Note that $i = O(\delta)$ since $d_{\text{cw}}(s, p_i) \leqslant 3\delta + 2$ and $(2n + 1) - j = O(\delta)$ since $d_{\text{cw}}(p_j, q) \leqslant 3\delta + 2$. Hence,

$$\sum_{i=1}^{2n+1} \rho(p_i)^\alpha \geqslant (2^\alpha + 1)n - O(\delta) \cdot 2^\alpha \geqslant (2^\alpha + 1)n - o(n),$$

since $\delta = ((2^\alpha + 1)n)^{1/\alpha} = o(n)$. ◀

The following lemma gives a key property of the construction.

▶ **Lemma 15.** *The point p_{2n+1} cannot have an incoming counterclockwise edge before q is inserted, and the point p_1 cannot have an incoming clockwise edge after q has been inserted.*

Proof. Suppose before insertion of q the point p_{2n+1} has an incoming counterclockwise edge. The cheapest incoming counterclockwise edge would be from s and this is already too expensive. Indeed, if $\rho(s) \geqslant 2x\delta$ then by Lemma 14 the total cost of the range assignment by ALG is at least

$$\begin{aligned}
(2x\delta)^\alpha + (2^\alpha + 1)n - o(n) &= \left(2^\alpha \cdot \left(\frac{1}{4} + \left(\frac{1}{2}\right)^{\alpha+1}\right) + 1\right) \cdot \delta^\alpha - o(n) \\
&= \left(1 + \left(2^{\alpha-3} - \frac{1}{4}\right)\right) \cdot 2\delta^\alpha - o(n) \\
&> \left(1 + \frac{1}{2} \cdot \left(2^{\alpha-3} - \frac{1}{4}\right)\right) \cdot 2\delta^\alpha \qquad \text{for } n \text{ sufficiently large} \\
&\geqslant c_\alpha \cdot \text{OPT}(P) \qquad \text{by definition of } c_\alpha \text{ and Observation 13.}
\end{aligned}$$

This contradicts the approximation ratio of ALG, proving the first part of the lemma.

Now suppose after the insertion of q the point p_1 has an incoming clockwise edge. The cheapest way to achieve this is with $\rho(s) = \delta$, which is too expensive. Indeed, by Lemma 14 the total cost of the range assignment is then at least

$$
\begin{aligned}
\delta^\alpha + (2^\alpha+1)n - o(n) &= \frac{2\delta^\alpha}{(2x^\alpha+1)\delta^\alpha}\cdot(2x^\alpha+1)\delta^\alpha - o(n) \\
&\geqslant \left(1+\frac{1}{2}\cdot\left(\frac{2\delta^\alpha}{(2x^\alpha+1)\delta^\alpha}-1\right)\right)\cdot \textsc{opt}(P\cup\{q\}) \quad \text{for } n \text{ sufficiently large} \\
&= \left(1+\frac{2-(2x^\alpha+1)}{2(2x^\alpha+1)}\right)\cdot \textsc{opt}(P\cup\{q\}) \\
&= \left(1+\frac{1-\left(\frac{1}{2}+\frac{1}{2^\alpha}\right)}{2\left(\frac{1}{2}+\frac{1}{2^\alpha}+1\right)}\right)\cdot \textsc{opt}(P\cup\{q\}) \quad \text{since } 2x^\alpha = \tfrac{1}{2}+\tfrac{1}{2^\alpha} \\
&= \left(1+\frac{2^{\alpha-1}-1}{3\cdot 2^\alpha+2}\right)\cdot \textsc{opt}(P\cup\{q\}) \\
&\geqslant c_\alpha\cdot \textsc{opt}(P\cup\{q\}) \qquad \text{by definition of } c_\alpha \text{ and Observation 13.}
\end{aligned}
$$

This contradicts the approximation ratio of ALG, proving the second part of the lemma. ◀

We are now ready to prove that many edges must change from being CW-minimal to being CCW-minimal when q is inserted. Observe that before and after the insertion of a point q, the distance between any two points is either 1, 2 or at least 3. Hence, in the following lemma we may assume that $\rho(p) \in \{0,1,2\}\cup[3,\infty)$ for any point $p \in P\cup\{q\}$.

▶ **Lemma 16.** *Before the insertion of q, at least $4n/3+1$ of the points from $\{s,p_1,\dots,p_{2n}\}$ have a* CW*-minimal range and after the insertion of q at least $4n/3+1$ of the points from $\{q,p_1,\dots,p_{2n}\}$ have a* CCW*-minimal range.*

Proof. We prove the lemma for the situation before q is inserted; the proof for the situation after the insertion of q is similar. It will be convenient to define $p_0 := s$ (although we may still use s if we want to stress that we are talking about the source). Recall that p_{2n+1} does not have an incoming counterclockwise edge in the communication graph $\mathcal{G}_\rho(P)$ before the insertion of q. Let π^* be a minimum-hop path from s to p_{2n+1} in $\mathcal{G}_\rho(P)$. Since p_{2n+1} does not have an incoming counterclockwise edge and π^* is a minimum-hop path, all edges in π are clockwise. We assign each point p_j with $1 \leqslant j \leqslant 2n+1$ to the edge (p_i,p_t) in π^* such that $i+1 \leqslant j \leqslant t$, and we define $A(p_i,p_t) := \{p_{i+1},\dots,p_t\}$ to be the set of all points assigned to (p_i,p_t). We define the *excess* of a point $p_j \in A(p_i,p_t)$ to be

$$\mathrm{excess}(p_j) := \frac{1}{|A(p_i,p_t)|}\cdot\left(\rho(p_i)^\alpha - \sum_{p_\ell\in A(p_i,p_t)} d(p_{\ell-1},p_\ell)^\alpha\right).$$

We say that an edge (p_i,p_t) in π^* is CW-minimal if p_i has a CW-minimal range. Note that if a point p_j is assigned to a CW-minimal edge, then this is the edge (p_{j-1},p_j) and $\mathrm{excess}(p_j)=0$. Intuitively, $\mathrm{excess}(p_j)$ denotes the additional cost we pay for reaching p_j compared to reaching it by a CW-minimal edge, if we distribute the additional cost of a non-CW-minimal edge over the points assigned to it. Because each of the points $p_1,\dots,p_{2n+1}$ is assigned to exactly one edge on the path π^*, we have

$$\sum_{p_i\in\pi^*}\rho(p)^\alpha \geqslant \sum_{j=1}^{2n+1} d(p_{j-1},p_j)^\alpha + \sum_{j=1}^{2n+1}\mathrm{excess}(p_j) \geqslant \textsc{opt}(P) + \sum_{j=1}^{2n+1}\mathrm{excess}(p_j) \qquad (1)$$

where the second inequality follows from Observation 13 and because $p_0 = s$. The following claim is proved in the full version. (Essentially, the smallest possible excess is obtained when $|A(p_i,p_t)| \in \{1,2,3\}$; the three terms in the claim correspond to these cases.)

▷ Claim. If p_j is not assigned to a CW-minimal edge then $\text{excess}(p_j) \geqslant c'_\alpha$, where $c'_\alpha = \min\left(2^\alpha - 1, \frac{3^\alpha - 2^\alpha - 1}{2}, \frac{4^\alpha - 2^\alpha - 2}{3}\right)$.

Now suppose for a contradiction that less than $4n/3 + 1$ points from $\{s, p_1, \ldots, p_{2n+1}\}$ have a CW-minimal range. Then at least $2n/3 + 1$ points p_j have $\text{excess}(p_j) \geqslant c'_\alpha$ by the claim above. By Inequality (1) the total cost incurred by ALG is therefore more than

$$\begin{aligned} \text{OPT}(P) + c'_\alpha \cdot (2n/3) &= \text{OPT}(P) + \frac{c'_\alpha}{3(2^\alpha + 1)} \cdot 2(2^\alpha + 1)n & (2)\\ &> \left(1 + \frac{\min\left(2^\alpha - 1, \frac{3^\alpha - 2^\alpha - 1}{2}, \frac{4^\alpha - 2^\alpha - 2}{3}\right)}{4(2^\alpha + 1)}\right) \cdot \text{OPT}(P) & (3)\\ &\geqslant c_\alpha \cdot \text{OPT}(P) & (4) \end{aligned}$$

which contradicts the approximation ratio achieved by ALG. ◀

Lemma 16 implies that at least $4n/3$ of the points $p_1, \ldots, p_{2n+1}$ have a CW-minimal range before q is inserted, and at least $4n/3$ of those points have a CCW-minimal range after the insertion. Hence, at least $2n + 1 - 2 \cdot (2n/3 + 1) = 2n/3 - 1$ points must change from being CW-minimal to being CCW-minimal, thus finishing the proof of Theorem 12.

The claim in the proof of Lemma 16

▷ Claim. If p_j is not assigned to a CW-minimal edge then $\text{excess}(p_j) \geqslant c'_\alpha$, where $c'_\alpha = \min\left(2^\alpha - 1, \frac{3^\alpha - 2^\alpha - 1}{2}, \frac{4^\alpha - 2^\alpha - 2}{3}\right)$.

Proof. Consider a non-CW-minimal edge (p_i, p_t). First suppose only a single point p_j is assigned to (p_i, p_t). Then $t = i + 1$ and $p_j = p_t$. Hence, $\rho(p_i) \geqslant d(p_{j-1}, p_j) + 1$ because we assumed $\rho(p_i) \in \{0, 1, 2\} \cup [3, \infty)$. Thus when $|A(p_i, p_t)| = 1$ then

$$\text{excess}(p_j) \geqslant (d(p_{j-1}, p_j) + 1)^\alpha - d(p_{j-1}p_j)^\alpha \geqslant 2^\alpha - 1 \geqslant c'_\alpha.$$

Now suppose $|A(p_i, p_t)| > 1$. Let z_1 be the number of points $p_j \in A(p_i, p_t)$ with $d(p_{j-1}, p_j) = 1$, and let z_2 be the number of points $p_j \in A(p_i, p_t)$ with $d(p_{j-1}, p_j) = 2$. Since $|A(p_i, p_t)| > 1$ we have $z_1 \geqslant 1$ and $z_2 \geqslant 1$ and $|z_1 - z_2| \leqslant 1$. When $|A(p_i, p_t)| = 2$ then $z_1 = z_2 = 1$, and we are distributing the cost of an edge of length at least 3, minus the costs of edges of length 2 and 1, over two points. Thus in this case we have

$$\text{excess}(p_j) \geqslant \frac{3^\alpha - 2^\alpha - 1}{2}.$$

Similarly, when $|A(p_i, p_t)| = 3$ then $z_1 = 2$ and $z_2 = 1$ (or vice versa, but that will only lead to a larger excess), and we have

$$\text{excess}(p_j) \geqslant \frac{4^\alpha - 2^\alpha - 2}{3}.$$

It remains to argue that we do not get a smaller excess when $|A(p_i, p_t)| \geqslant 4$. To see this, we compare the excess we get when (p_i, p_t) is an edge of π with the excesses we would get when, instead of (p_i, p_t), the edges (p_i, p_{i+2}) and (p_{i+2}, p_t) would be in π^*. Note that

$$d(p_i, p_t)^\alpha = \Big(d(p_i, p_{i+2}) + d(p_{i+2}, p_t)\Big)^\alpha > d(p_i, p_{i+2})^\alpha + d(p_{i+2}, p_t)^\alpha$$

since $\alpha > 1$. Hence,

$$\begin{aligned}
&\frac{d(p_i,p_t)^\alpha - \sum_{\ell=i+1}^{t} d(p_{\ell-1},p_\ell)^\alpha}{t-i} \\
&\quad > \quad \frac{\left(d(p_i,p_{i+2})^\alpha - \sum_{\ell=i+1}^{i+2} d(p_{\ell-1},p_\ell)^\alpha\right) + \left(d(p_{i+2},p_t)^\alpha - \sum_{\ell=i+3}^{t} d(p_{\ell-1},p_\ell)^\alpha\right)}{t-i} \\
&\quad \geqslant \quad \frac{d(p_i,p_{i+2})^\alpha - \sum_{\ell=i+1}^{i+2} d(p_{\ell-1},p_\ell)^\alpha}{2} + \frac{d(p_{i+2},p_t)^\alpha - \sum_{\ell=i+3}^{t} d(p_{\ell-1},p_\ell)^\alpha}{t-i-2}
\end{aligned}$$

where the last inequality uses that $\frac{a_1+a_2}{b_1+b_2} \geqslant \min\left(\frac{a_1}{b_1}, \frac{a_2}{b_2}\right)$ for any $a_1, a_2, b_1, b_2 > 0$. Thus the excess we get for (p_i, p_t) is at least the minimum of the excesses we would get for (p_i, p_{i+2}) and (p_{i+3}, p_t). More generally, when $|A(p_i, p_t)| > 4$ then we can compare the excess for (p_i, p_t) with the excesses we get when we would replace (p_i, p_t) with a path of smaller edges, each being assigned two or three points. The excess for (p_i, p_{i+2}) is at least the minimum of the excesses for these shorter edges. (Reducing to edges that are assigned a single point is not useful, since these may be CW-minimal and have zero excess.) This finishes the proof of the claim. ◁

6 The 2-dimensional problem

The broadcast range-assignment problem is NP-hard in $\mathbb{R}^2$, so we cannot expect a characterization of the structure of an optimal solution similar to Theorem 2. Using a similar construction as in $\mathbb{S}^1$ we can also show that the problem in $\mathbb{R}^2$ does not admit a SAS.

▶ **Theorem 17.** *The dynamic broadcast range-assignment problem in $\mathbb{R}^2$ with distance power gradient $\alpha > 1$ does not admit a SAS. In particular, there is a constant $c_\alpha > 1$ such that the following holds: for any n large enough, there is a set $P := \{s, p_1, \ldots, p_{2n+1}\}$ and a point q in $\mathbb{R}^2$ such that any update algorithm* ALG *that maintains a c_α-approximation must modify at least $2n/3 - 1$ ranges upon the insertion of q into P.*

Proof. We use the same construction as in $\mathbb{S}^1$, where we embed the points on a square and the distances used to define the instance are measured along the square; see Fig. 4(ii). We now discuss the changes needed in the proof to deal with the fact that distances in $\mathbb{R}^2$ between points from $P \cup \{q\}$ may be smaller than when measured along the square. With a slight abuse of terminology, we will still refer to an edge (p, p') that was clockwise in $\mathbb{S}^1$ as a clockwise edge, and similarly for counterclockwise edges.

Note that Observation 13 still holds. Now consider Lemma 14. The proof used that the points p_i at distance more than 3δ from s or q must be covered by the ranges of the points $p_1, \ldots, p_{2n+1}$. We now restrict our attention to the points that are also at distance more than 3δ from a corner of the square. Each such point p_i must be covered by the range of some point p_j on the same edge of the square. Hence, the distance in $\mathbb{R}^2$ of from p_j to p_i is the same as the distance in $\mathbb{S}^1$, so we can use the same reasoning as before. Thus the exclusion of the points that are at distance at most 3δ from a corner of the square only influences the constant in the $o(n)$ term in the lemma. Hence, Lemma 14 still holds.

The proof of Lemma 15 still holds, since the cheapest counterclockwise edge to p_{2n+1} before the insertion of q is still from s (and the distance from s to p_{2n+1} did not change), and the cheapest clockwise edge to p_1 after the insertion of q is still from s (and the distance from s to p_1 did not change).

It remains to check Lemma 16. The proof still holds, except that the claim that $\text{excess}(p_j) \geqslant c'_\alpha$ may not be true for the given value of c'_α when p_j is near a corner of the square, because the distances between points on different edges of the square do not correspond to the distances in $\mathbb{S}^1$. To deal with this, we simply ignore the excess of any point within distance 3δ from a corner. This reduces the total excess by $o(n)$. It is easily verified that this does not invalidate the rest of the proof: we have to subtract $o(n)$ from the formulae in Equality (2), but this is still larger than $c_\alpha \cdot \text{OPT}(P)$.

We conclude that all lemmas still hold, which proves Theorem 17. ◀

Although the problem in $\mathbb{R}^2$ does not admit a SAS, there is a relatively simple $O(1)$-stable $O(1)$-approximation algorithm for $\alpha \geqslant 2$. The algorithm is based on a result by Ambühl [1], who showed that a minimum spanning tree (MST) on P gives a 6-approximation for the static broadcast range-assignment problem: turn the MST into a directed tree rooted at the source s, and assign as a range to each point $p \in P$ the maximum length of any of its outgoing edges. To make this stable, we set the range of each point to the maximum length of any of its incident edges (not just the outgoing ones). Because an MST in $\mathbb{R}^2$ has maximum degree 6, this leads to 17-stable 12-approximation algorithm; see the full version [14].

7 Concluding remarks

We studied the dynamic broadcast range-assignment problem from a stability perspective, introducing the notions of k-stable algorithms and stable approximation schemes (SASs). Our results provide a fairly complete picture of the problem in $\mathbb{R}^1$, in $\mathbb{S}^1$, and in $\mathbb{R}^2$. In particular, we presented a SAS in $\mathbb{R}^1$ that has an asymptotically optimal stability parameter, and showed that the problem does not admit a SAS in $\mathbb{S}^1$ and $\mathbb{R}^2$. Future work can focus on improving the (the upper and/or lower bounds for) approximation ratios that we have obtained for algorithms with constant stability parameter.

References

1 Christoph Ambühl. An optimal bound for the MST algorithm to compute energy efficient broadcast trees in wireless networks. In *Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP 2005)*, volume 3580 of *Lecture Notes in Computer Science*, pages 1139–1150, 2005. `doi:10.1007/11523468_92`.

2 Mohammad R. Ataei, Amir H. Banihashemi, and Thomas Kunz. Low-complexity energy-efficient broadcasting in one-dimensional wireless networks. *IEEE Trans. Veh. Technol.*, 61(7):3276–3282, 2012. `doi:10.1109/TVT.2012.2204077`.

3 Mostafa A. Bassiouni and Chun-Chin Fang. Dynamic channel allocation for linear macrocellular topology. In *Proc. 1999 ACM Symposium on Applied Computing*, pages 382–388, 1999. `doi:10.1145/298151.298391`.

4 Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. New results for energy-efficient broadcasting in wireless networks. In *Proc. 13th International Symposium on Algorithms and Computation (ISAAC 2002)*, volume 2518 of *Lecture Notes in Computer Science*, pages 332–343, 2002. `doi:10.1007/3-540-36136-7_30`.

5 Andrea E. F. Clementi, Pierluigi Crescenzi, Paolo Penna, Gianluca Rossi, and Paola Vocca. On the complexity of computing minimum energy consumption broadcast subgraphs. In *Proc. 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, volume 2010 of *Lecture Notes in Computer Science*, pages 121–131. Springer, 2001. `doi:10.1007/3-540-44693-1_11`.

6 Andrea E. F. Clementi, Gurvan Huiban, Paolo Penna, Gianluca Rossi, and Yann C. Verhoeven. Some recent theoretical advances and open questions on energy consumption in ad-hoc wireless networks. In *Proc. 3rd Workshop on Approximation and Randomization Algorithms in Communication Networks (ARACNE 2002)*, 2002.

7 Andrea E. F. Clementi, Miriam Di Ianni, and Riccardo Silvestri. The minimum broadcast range assignment problem on linear multi-hop wireless networks. *Theor. Comput. Sci.*, 299(1-3):751–761, 2003. doi:10.1016/S0304-3975(02)00538-8.

8 Andrea E. F. Clementi, Paolo Penna, Afonso Ferreira, Stephane Perennes, and Riccardo Silvestri. The minimum range assignment problem on linear radio networks. *Algorithmica*, 35(2):95–110, 2003. doi:10.1007/s00453-002-0985-2.

9 Andrea E. F. Clementi, Paolo Penna, and Riccardo Silvestri. Hardness results for the power range assignment problem in packet radio networks. In *Proc. 3rd International Workshop on Randomization and Approximation Techniques in Computer Science, and 2nd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (RANDOM-APPROX'99)*, volume 1671 of *Lecture Notes in Computer Science*, pages 197–208, 1999. doi:10.1007/978-3-540-48413-4_21.

10 Andrea E. F. Clementi, Paolo Penna, and Riccardo Silvestri. The power range assignment problem in radio networks on the plane. In *Proc. 17th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1770 of *Lecture Notes in Computer Science*, pages 651–660, 2000. doi:10.1007/3-540-46541-3_54.

11 Gautam K. Das, Sandip Das, and Subhas C. Nandy. Range assignment for energy efficient broadcasting in linear radio networks. *Theor. Comput. Sci.*, 352(1-3):332–341, 2006. doi:10.1016/j.tcs.2005.11.046.

12 Gautam K. Das and Subhas C. Nandy. Weighted broadcast in linear radio networks. *Inf. Process. Lett.*, 106(4):136–143, 2008. doi:10.1016/j.ipl.2007.10.016.

13 Mark de Berg, Aleksandar Markovic, and Seeun William Umboh. The online broadcast range-assignment problem. In *Proc. 31st International Symposium on Algorithms and Computation (ISAAC)*, volume 181 of *LIPIcs*, pages 60:1–60:15, 2020. doi:10.4230/LIPIcs.ISAAC.2020.60.

14 Mark de Berg, Arpan Sadhukhan, and Frits C. R. Spieksma. Stable approximation algorithms for the dynamic broadcast range-assignment problem. *CoRR*, abs/2112.05426, 2021. arXiv:2112.05426.

15 Bernhard Fuchs. On the hardness of range assignment problems. *Networks*, 52(4):183–195, 2008. doi:10.1002/net.20227.

16 Lefteris M. Kirousis, Evangelos Kranakis, Danny Krizanc, and Andrzej Pelc. Power consumption in packet radio networks. *Theor. Comput. Sci.*, 243(1-2):289–305, 2000. doi:10.1016/S0304-3975(98)00223-0.

17 Rudolf Mathar and Jürgen Mattfeldt. Optimal transmission ranges for mobile communication in linear multihop packet radio networks. *Wireless Networks*, 2(4):329–342, 1996. doi:10.1007/BF01262051.

18 Kaveh Pahlavan and Allen H. Levesque. *Wireless information networks, Second Edition*. Wiley series in telecommunications and signal processing. Wiley-VCH, 2005.

19 Peter Sanders, Naveen Sivadasan, and Martin Skutella. Online scheduling with bounded migration. *Math. Oper. Res.*, 34(2):481–498, 2009. doi:10.1287/moor.1090.0381.

20 Martin Skutella and José Verschae. A robust PTAS for machine covering and packing. In *Proc. 18th Annual European Symposium (ESA 201)*, volume 6346 of *Lecture Notes in Computer Science*, pages 36–47, 2010. doi:10.1007/978-3-642-15775-2_4.