The Online Broadcast Range-Assignment Problem

Mark de Berg 🆔
Department of Computing Science, TU Eindhoven, The Netherlands
m.t.d.berg@tue.nl

Aleksandar Markovic
Department of Computing Science, TU Eindhoven, The Netherlands
info@aleksmarkovic.com

Seeun William Umboh 🆔
School of Computer Science, The University of Sydney, Australia
william.umboh@sydney.edu.au

─── Abstract ───

Let $P = \{p_0, \ldots, p_{n-1}\}$ be a set of points in $\mathbb{R}^d$, modeling devices in a wireless network. A range assignment assigns a range $r(p_i)$ to each point $p_i \in P$, thus inducing a directed communication graph $\mathcal{G}_r$ in which there is a directed edge (p_i, p_j) iff $\mathrm{dist}(p_i, p_j) \leqslant r(p_i)$, where $\mathrm{dist}(p_i, p_j)$ denotes the distance between p_i and p_j. The range-assignment problem is to assign the transmission ranges such that $\mathcal{G}_r$ has a certain desirable property, while minimizing the cost of the assignment; here the cost is given by $\sum_{p_i \in P} r(p_i)^\alpha$, for some constant $\alpha > 1$ called the distance-power gradient.

We introduce the online version of the range-assignment problem, where the points p_j arrive one by one, and the range assignment has to be updated at each arrival. Following the standard in online algorithms, resources given out cannot be taken away – in our case this means that the transmission ranges will never decrease. The property we want to maintain is that $\mathcal{G}_r$ has a broadcast tree rooted at the first point p_0. Our results include the following.

- We prove that already in $\mathbb{R}^1$, a 1-competitive algorithm does not exist. In particular, for distance-power gradient $\alpha = 2$ any online algorithm has competitive ratio at least 1.57.
- For points in $\mathbb{R}^1$ and $\mathbb{R}^2$, we analyze two natural strategies for updating the range assignment upon the arrival of a new point p_j. The strategies do not change the assignment if p_j is already within range of an existing point, otherwise they increase the range of a single point, as follows: NEAREST-NEIGHBOR (NN) increases the range of $\mathrm{nn}(p_j)$, the nearest neighbor of p_j, to $\mathrm{dist}(p_j, \mathrm{nn}(p_j))$, and CHEAPEST INCREASE (CI) increases the range of the point p_i for which the resulting cost increase to be able to reach the new point p_j is minimal. We give lower and upper bounds on the competitive ratio of these strategies as a function of the distance-power gradient α. We also analyze the following variant of NN in $\mathbb{R}^2$ for $\alpha = 2$: 2-NEAREST-NEIGHBOR (2-NN) increases the range of $\mathrm{nn}(p_j)$ to $2 \cdot \mathrm{dist}(p_j, \mathrm{nn}(p_j))$,
- We generalize the problem to points in arbitrary metric spaces, where we present an $O(\log n)$-competitive algorithm.

31st International Symposium on Algorithms and Computation (ISAAC 2020).
Editors: Yixin Cao, Siu-Wing Cheng, and Minming Li; Article No. 60; pp. 60:1–60:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1   Introduction

Consider a collection of wireless devices, each with its own transmission range. The transmission ranges induce a directed communication network, where each device p_i can directly send a message to any device p_j in its transmission range. If p_j is not within range, a message from p_i can still reach p_j if there is a path from p_i to p_j in the communication network. The energy consumption of a device depends on its transmission range: the greater the range, the more power is needed. This leads to the range-assignment problem: assign transmissions ranges to the devices such that the resulting network has some desired connectivity property, while minimizing the total power consumption.

Mathematically we can model the problem as follows. Let $P = \{p_0, \ldots, p_{n-1}\}$ be a set of n points in $\mathbb{R}^d$. For an assignment $r : P \to \mathbb{R}_{\geqslant 0}$, let $\mathcal{G}_r$ be the directed graph on the vertex set P obtained by putting a directed edge from a vertex p_i to a vertex p_j iff $\mathrm{dist}(p_i, p_j) \leqslant r(p_i)$, where $\mathrm{dist}(p_i, p_j)$ denotes the distance between p_i and p_j. We call $\mathcal{G}_r$ the *communication graph* on P induced by the range assignment r. The *cost* of a range assignment r is defined as $\mathrm{cost}_\alpha(r) := \sum_{p_i \in P} r(p_i)^\alpha$, where $\alpha \geqslant 1$ is called the *distance-power gradient*. In practice, α typically varies from 1 to 6 [15]. We then want to find a range assignment that minimizes the cost while ensuring that $\mathcal{G}_r$ has some desired property. Properties that have been investigated in this context include strong connectivity [9, 14], h-hop strong connectivity [8, 10, 14], broadcast capability – here $\mathcal{G}_r$ must contain a broadcast tree (that is, an arborescence) rooted at the source point p_0 – , and h-hop broadcast capability [2, 13]; see the survey by Clementi et al. [6] for an overview of the various range-assignment problems. Most previous work considered the Euclidean setting. There has been some work on arbitrary metric spaces for the strong connectivity version [4, 12]. (Note that while the 2-dimensional version seems the most relevant setting, the distances may not be Euclidean due to obstacles that reduce the strength of the signal of a device.)

In this paper we focus on the broadcast version of the range-assignment problem. This version can be solved optimally in a trivial manner when $\alpha = 1$, by setting $r(p_0) := \max_{0 \leqslant i < n} \mathrm{dist}(p_0, p_i)$ and $r(p_i) := 0$ for $i > 0$. Clementi et al. [7] showed a polynomial time algorithm for the 1-dimensional problem when $\alpha \geqslant 2$. Moreover, Clementi et al. [5] showed the problem is NP-hard for any $\alpha > 1$ and any $d \geqslant 2$. Clementi et al. [7], Clementi et al. [5], and Wan et al. [17] also showed that the problem can be approximated within a factor $c \cdot 2^\alpha$ for any $\alpha \geqslant 2$ and for a certain constant c. Furthermore, Clementi et al. [5] showed that for any $d \geqslant 2$ and for any $\alpha \geqslant d$, there is a function $f : \mathbb{N} \times \mathbb{R} \to \mathbb{R}$ such that the problem can be approximated within a factor $f(d, \alpha)$ in the d-dimensional Euclidean space. Fuchs [11] showed that for $d = 2$, the problem remains NP-hard even for so-called well-spread instances for any $\alpha > 1$. In dimension $d \geqslant 3$, he also showed that the problem is NP-hard to approximate within a factor of $51/50$ when $\alpha > 1$; the result also holds for well-spread instances when $\alpha > d$.

Our contribution.   We study the online version of the broadcast range-assignment problem. Here the points $p_0, p_1, \ldots, p_{n-1}$ come in one by one, and the goal is to maintain a range assignment r such that $\mathcal{G}_r$ contains a broadcast tree on the currently inserted points, rooted at the first point p_0. Of course one can simply recompute the assignment from scratch, but in online algorithms one requires that resources that have been given out cannot be taken back. For the range assignment problem this means that we are not allowed to decrease the range of any point. In fact, our algorithms have the useful property that upon arrival of each point, we change the current range assignment only minimally: either we do not change it at

all – this happens when the newly arrived point is already within range of an existing point – or we increase the range of only a single point. Our goal is to obtain algorithms with a good competitive ratio.[1] As far as we know, the range-assignment problem has not been studied from the perspective of online algorithms.

We first prove a lower bound on the competitive ratio achievable by any online algorithm: even in $\mathbb{R}^1$ there is a constant $c_\alpha > 1$ (which depends on the power-distance gradient α) such that no online algorithm can be c_α-competitive. For $\alpha = 2$, we have $c_\alpha > 1.57$.

We then investigate the following two natural online algorithms for the broadcast range-assignment problem. Suppose the point p_j arrives. Our algorithms all set $r(p_j) := 0$ and, as mentioned, they do not change any of the ranges $r(p_0), \ldots, r(p_{j-1})$ if $|p_i p_j| \leqslant r(p_i)$ for some $0 \leqslant i < j$. When p_j is not within range of an already inserted point, the algorithms increase the range of one point, as follows. Let $\mathrm{nn}(p_j)$ denote the nearest neighbor of p_j in the set $\{p_0, \ldots, p_{j-1}\}$, with ties broken arbitrarily.

- Nearest-Neighbor (NN for short) increases the range of $\mathrm{nn}(p_j)$ to $\mathrm{dist}(p_j, \mathrm{nn}(p_j))$.
- Cheapest Increase (CI for short) increases the range of p_{i^*} to $\mathrm{dist}(p_{i^*}, p_j)$, where p_{i^*} is a point minimizing the cost increase of the assignment, which is $\mathrm{dist}(p_{i^*}, p_j)^\alpha - r(p_{i^*})^\alpha$ where $r(p_{i^*})$ denotes the current range of p_{i^*}.

■ **Table 1** Overview of results on NN and CI.

dimension	distance-power gradient	lower bound for NN	upper bound for NN and CI
$d = 1$	$\alpha = 2$	2	2
$d = 2$	$\alpha = 2$	≈ 7.61	322
			2-NN: 36
	$2 < \alpha < \alpha^* \approx 4.3$	$\approx 6(1 + 0.52^\alpha)$	$\alpha \frac{2^\alpha - 3}{2^{\alpha-1} - \alpha}$
	$\alpha \geqslant \alpha^* \approx 4.3$		≈ 12.94

The results are summarized in Table 1. Note the lower bounds hold only for NN, while the upper bounds hold for NN and CI; the exception is the third row, which is for 2-NN (see below). The lower bound of $6(1 + 0.52^\alpha)$ mentioned in the table – the exact bound is $6(1 + (\frac{\sqrt{6} - \sqrt{2}}{2})^\alpha)$ – applies to all $\alpha > 1$, and thus implies the given lower bound for $\alpha = 2$. Recall that for $d = 1$ and $\alpha = 2$, we also have a universal lower bound of 1.57 that holds for any online algorithm and, hence, also for CI. The exact value of α^* is $\alpha^* = \arg\min F_\alpha^*$, where $F_\alpha^* = \alpha \frac{2^\alpha - 3}{2^{\alpha-1} - \alpha}$.

As can be seen in the table NN is $O(1)$-competitive for $\alpha = 2$, but the competitive ratio is quite large, namely 322. We therefore also analyze the following variant of NN, which (if p_j is not yet within range of an existing point) proceeds as follows:

- 2-Nearest-Neighbor (2-NN for short) increases the range of $\mathrm{nn}(p_j)$ to $2 \cdot \mathrm{dist}(p_j, \mathrm{nn}(p_j))$.

We prove that the competitive ratio of 2-NN is at most 36 for $\alpha = 2$. Thus, while still rather large, the competitive ratio is a lot smaller than what we were able to prove for NN. It is interesting to note that both NN and 2-NN make decisions that are independent of α. Hence, NN obtains a solution that is simultaneously competitive for all $\alpha \geqslant 2$.

---

[1] The *competitive ratio* [16] of an online algorithm compares the cost of the solution it maintains to the cost achieved by the optimal offline algorithm. More precisely, an online algorithm ALG is *c-competitive* if there is a constant a such that for any instance I, the cost of ALG is at most $c \cdot \mathrm{OPT}(I) + a$. Note the the offline algorithm must still maintain the solution over time. Thus the offline problem is not the same as the static problem, where we only want a solution for the final configuration.

As a final contribution we generalize the broadcast problem to points in arbitrary metric spaces. Since to the best of our knowledge this version has not been studied before, we present in the appendix an approximation algorithm for the offline setting; its approximation ratio is 5^α. In this offline setting the algorithm must be what Boyar et al. [3] call an *incremental algorithm*: an algorithm that, even though it may know the future, maintains a feasible solution at any time. For the online setting (where the future is unknown) we obtain an $O(4^\alpha \log n)$-competitive algorithm.

Notation.   We let $P := p_0, \ldots, p_{n-1}$ denote the input sequence, where we assume without loss of generality that p_i is inserted at time i and that all p_i are distinct. Define $P_i := p_0, \ldots, p_i$, and denote the range of a point $p_i \in P_j$ just after the insertion of the point p_j by $r_j(p_i)$. Thus in the online version we require that $r_j(p_i) \leqslant r_{j+1}(p_i)$. For an algorithm ALG we use $\text{cost}_\alpha(\text{ALG}(P))$ to denote the cost incurred by ALG on input P for distance-power gradient α. Finally we denote the ball of radius ρ centered at a point p by $B(p, \rho)$; note that in $\mathbb{R}^1$ this is an interval of length 2ρ and in $\mathbb{R}^2$ it is a disk of radius ρ.

2   Online range-assignment in $\mathbb{R}^1$

In this section we prove that no online algorithm can have a competitive ratio arbitrarily close to 1, even in $\mathbb{R}^1$. We also prove bounds on the competitive ratio of NN and CI in $\mathbb{R}^1$.

A universal lower bound.   To prove the lower bound we consider an arbitrary online algorithm ALG. Our adversary then first presents the points $p_0 = 0$, $p_1 = x$, and $p_2 := \delta_\alpha \cdot x$. Depending on the range assignment ALG has done so far, the adversary either ends the instance or presents a fourth point $p_3 = -\delta_\alpha \cdot x$. By picking a suitable value for δ_α and making x sufficiently large, we can obtain a lower bound. This is made precise in the following theorem.

▶ **Theorem 2.1.** *For any distance-power gradient $\alpha > 1$, there is a constant $c_\alpha > 1$ such that any online algorithm for the range assignment problem in $\mathbb{R}^1$ has a competitive ratio of at least c_α. For $\alpha = 2$ this constant is $c_2 \approx 1.58$.*

Proof.   Let $\alpha > 1$ and let ALG be an algorithm with competitive ratio $c \geqslant 1$, i.e., there is a constant a such that the cost of ALG is upper bounded by $c \cdot \text{OPT} + a$. We also define

$$c_\alpha := \max_{\delta > 1} \min \left(\frac{\delta^\alpha}{1 + (\delta - 1)^\alpha}, \frac{\delta^\alpha + (\delta - 1)^\alpha}{\delta^\alpha}, \frac{1 + (\delta + 1)^\alpha}{\delta^\alpha} \right),$$

$$\text{and} \quad \delta_\alpha := \arg\max_{\delta > 1} \min \left(\frac{\delta^\alpha}{1 + (\delta - 1)^\alpha}, \frac{\delta^\alpha + (\delta - 1)^\alpha}{\delta^\alpha}, \frac{1 + (\delta + 1)^\alpha}{\delta^\alpha} \right).$$

We show that $c \geqslant c_\alpha$ by constructing the following families of instances consisting of, respectively, three and four points, and parametrized by the real number $x \geqslant 1$:

$$\mathcal{F}_1 := \{\{p_0 = 0, p_1 = x, p_2 = \delta_\alpha \cdot x\}\}$$

$$\text{and} \quad \mathcal{F}_2 := \{\{p_0 = 0, p_1 = x, p_2 = \delta_\alpha \cdot x, p_3 = -\delta_\alpha \cdot x\}\}.$$

Note that there is a one-to-one correspondence between the instances in both families: each instance of $\mathcal{F}_1$ is the beginning of exactly one instance of $\mathcal{F}_2$ and each instance of $\mathcal{F}_2$ starts like exactly one instance of $\mathcal{F}_1$.

For any x, depending on what ALG does after p_2 is inserted, we choose an instance from either the family $\mathcal{F}_1$ or the family $\mathcal{F}_2$ using the following rule: if after p_2 is inserted, ALG has a disk of radius at least $\delta_\alpha \cdot x$, we choose $\mathcal{F}_1$, otherwise we choose $\mathcal{F}_2$. In the former

case, ALG pays at least $\delta_\alpha^\alpha \cdot x^\alpha$ while the optimal solution would be to place a disk of radius x at p_0 and a disk of radius $(\delta_\alpha - 1) \cdot x$ at p_1 and pay $x^\alpha + (\delta_\alpha - 1)^\alpha \cdot x^\alpha$. Since the competitive ratio of ALG is c, we have that $\delta_\alpha^\alpha \cdot x^\alpha \leqslant c \cdot x^\alpha (1 + (\delta_\alpha - 1)^\alpha) + a$ and hence

$$c \geqslant \frac{\delta_\alpha^\alpha}{1 + (\delta_\alpha - 1)^\alpha} - \frac{a}{x^\alpha(1 + (\delta_\alpha - 1)^\alpha)}.$$

Since the second term can be made arbitrarily small by choosing x large enough, c must be at least $\frac{\delta_\alpha^\alpha}{(\delta_\alpha - 1)^\alpha}$.

In the latter case, ALG has one disk of radius at least x and one of radius at least $(\delta_\alpha - 1) \cdot x$ before p_3 is inserted. We split this case into two subcases: in the first one, ALG increases the radius of the disk at p_0 and in the second one, ALG increases the radius of the disk at p_1. The cost ALG has to pay after p_3 has been inserted is at least either $\delta_\alpha^\alpha \cdot x^\alpha + (\delta_\alpha - 1)^\alpha \cdot x^\alpha$ in the first subcase, or $x^\alpha + (\delta_\alpha + 1)^\alpha \cdot x^\alpha$ in the second, whereas the optimal solution for both subcases would be to place only one disk of radius $\delta_\alpha \cdot x$ at p_0 and pay $\delta_\alpha^\alpha \cdot x^\alpha$. Since the competitive ratio of ALG is c, we have that $\delta_\alpha^\alpha \cdot x^\alpha + (\delta_\alpha - 1)^\alpha \cdot x^\alpha \leqslant c \cdot \delta_\alpha^\alpha \cdot x^\alpha + a$ for the first subcase and hence

$$c \geqslant \frac{\delta_\alpha^\alpha + (\delta_\alpha - 1)^\alpha}{\delta_\alpha^\alpha} - \frac{a}{\delta_\alpha^\alpha \cdot x^\alpha};$$

and that $x^\alpha + (\delta_\alpha + 1)^\alpha \cdot x^\alpha \leqslant c\delta_\alpha^\alpha \cdot x^\alpha + a$ for the second subcase and hence

$$c \geqslant \frac{1 + (\delta_\alpha + 1)^\alpha}{\delta_\alpha^\alpha} - \frac{a}{\delta_\alpha^\alpha \cdot x^\alpha}.$$

Since, in both subcases, the second term can be made arbitrarily small by choosing x large enough, c must be at least $\frac{\delta_\alpha^\alpha + (\delta_\alpha - 1)^\alpha}{\delta_\alpha^\alpha}$ for the first subcase, and at least $\frac{1 + (\delta_\alpha + 1)^\alpha}{\delta_\alpha^\alpha}$, otherwise there is an infinite family of instances contradicting the competitive ratio for these two subcases.

Therefore, the competitive ratio of ALG must be at least the minimum of the competitive ratio between these cases, which is exactly c_α. Even though it is not clear how to compute the value of c_α for any fixed $\alpha > 1$, it is easy to see it is strictly bigger than 1. If $\alpha = 2$, we have

$$\begin{aligned}
c_2 :&= \max_{\delta > 1} \min \left(\frac{\delta^2}{1 + (\delta - 1)^2}, \frac{\delta^2 + (\delta - 1)^2}{\delta^2}, \frac{1 + (\delta + 1)^2}{\delta^2} \right) \\
&= \frac{1}{12} \left(4 + \sqrt[3]{496 - 24\sqrt{183}} + 2\sqrt[3]{62 + 3\sqrt{183}} \right) \\
&\approx 1.58
\end{aligned}$$

which is achieved for

$$\begin{aligned}
\delta_2 :&= \arg\max_{\delta > 1} \min \left(\frac{\delta^2}{1 + (\delta - 1)^2}, \frac{\delta^2 + (\delta - 1)^2}{\delta^2}, \frac{1 + (\delta + 1)^2}{\delta^2} \right) \\
&= \frac{1}{3} \left(5 + \sqrt[3]{62 - 3\sqrt{183}} + \sqrt[3]{62 + 3\sqrt{183}} \right) \\
&\approx 4.15.
\end{aligned}$$

◀

Bounds for NN **and** CI. We now prove bounds on the competitive ratio of the algorithms NN and CI explained in the introduction.

▶ **Theorem 2.2.** *Consider the range-assignment problem in* $\mathbb{R}^1$ *with distance-power gradient* α.

(i) *For any* $\alpha > 1$, *the competitive ratio of* CI *is at most* 2.

(ii) *For any* $\alpha > 1$, *the competitive ratio of* NN *is exactly* 2.

Proof. We first prove the upper bounds. Assume without loss of generality that $p_0 = 0$. We first prove that both NN and CI perform optimally for $\alpha > 1$ on any sequence $p_0, p_1, \ldots, p_{n-1}$ with $p_j \geqslant 0$ for all $1 \leqslant j < n$.

▷ **Claim.** Suppose $p_0 = 0$ and $p_j \geqslant 0$ for all $1 \leqslant j < n$. Then NN and CI are optimal.

Proof. We first observe that for any point p_j the following holds for the graph $\mathcal{G}_{r_j}$ that we have after the insertion of p_j: for any point p_i with $0 < i \leqslant j$ there is a path from the source p_0 to p_i that only uses edges directed from left to right, that is, edges $(p_{i'}, p_{i''})$ with $p_{i'} < p_{i''}$. Indeed, if the path uses an edge $(p_{i'}, p_{i''})$ with $p_{i'} > p_{i''}$ then the subpath from p_0 to $p_{i'}$ must contain an edge (p_s, p_t) with $p_s \leqslant p_{i''} \leqslant p_t$, and then we can go directly from p_s to $p_{i''}$. This observation implies that there exists an optimal strategy OPT such that the balls $B(p_i, r_j(p_i))$ of the currently inserted points never extend beyond the currently rightmost point, a property which holds for NN and CI as well. (Intuitively, the part of $B(p_i, r_j(p_i))$ to the right of the rightmost point is currently useless, and the part of $B(p_i, r_j(p_i))$ to the left of p_i is not needed because we never need edges going to the left. Hence, we decrease $r_j(p_i)$ until the right endpoint of $B(p_i, r_j(p_i))$ coincides with the currently rightmost point, and increase the range of p_i later, as needed.)

Now imagine running NN, CI, and OPT simultaneously on P. We claim that NN and CI do exactly the same, and that their cost increase after the insertion of any point p_j is at most the cost increase of OPT. To see this, let $p_{j'}$ be the rightmost point just before inserting p_j. If $p_j < p_{j'}$ then NN and CI do not increase any range – since $p_{j'}$ is reachable from p_0, the point p_j must already be reachable as well – and so the cost increase is zero. If $p_j > p_{j'}$ then NN and CI both increase the range of $p_{j'}$ from 0 to $p_j - p_{j'}$. For NN this is clear. For CI it follows from the fact that $\alpha > 1$. Indeed, increasing the range of some $p_i < p_{j'}$ gives a cost increase $(r_{j-1}(p_i) + x + (p_j - p_{j'}))^\alpha - (r_{j-1}(p_i))^\alpha$, for some $x \geqslant 0$. This is more than $(p_j - p_{j'})^\alpha$, since we must have $r_{j-1}(p_i) + x > 0$. By a similar reasoning, and using that the balls of OPT do not extend beyond $p_{j'}$, we conclude that the cost increase of OPT cannot be smaller than $(p_j - p_{j'})^\alpha$. Hence, NN and CI are optimal on a sequence of non-negative points.

◁

Next, we prove that the optimality for non-negative points gives a competitive ratio of at most 2 for any input sequence P. Let P^+ and P^- denote the subsequences of P consisting of the points with non-negative and non-positive points, respectively. Note that the source point $p_0 = 0$ is included in both subsequences. We claim that $\mathrm{cost}_\alpha(\mathrm{OPT}((P)) \geqslant \mathrm{cost}_\alpha(\mathrm{OPT}((P^+))$. Indeed, we can modify the optimal solution for P to a valid solution for P^+ whose cost is at most $\mathrm{cost}_\alpha(\mathrm{OPT}((P))$, as follows: whenever the range of a point $p_i \notin P^+$ is increased to reach a point $p_j \in P^+$, we instead increase the range of p_0 by the same amount. A similar argument gives $\mathrm{cost}_\alpha(\mathrm{OPT}((P)) \geqslant \mathrm{cost}_\alpha(\mathrm{OPT}((P^-))$.

We now argue that $\mathrm{cost}_\alpha(\mathrm{NN}(P)) \leqslant \mathrm{cost}_\alpha(\mathrm{NN}(P^+)) + \mathrm{cost}_\alpha(\mathrm{NN}(P^-))$ and, similarly, that $\mathrm{cost}_\alpha(\mathrm{CI}(P)) \leqslant \mathrm{cost}_\alpha(\mathrm{CI}(P^+)) + \mathrm{cost}_\alpha(\mathrm{CI}(P^-))$.

Imagine running NN simultaneously on P, on P^+ and on P^-. We claim that the increase of $\mathrm{cost}_\alpha(\mathrm{NN}(P))$ upon the arrival of a new point p_j is at most the increase of $\mathrm{cost}_\alpha(\mathrm{NN}(P^+))$ if $p_j > 0$, and at most the increase of $\mathrm{cost}_\alpha(\mathrm{NN}(P^-))$ if $p_j < 0$. To see this, assume without loss of generality that $p_j > 0$ and suppose the increase of $\mathrm{cost}_\alpha(\mathrm{NN}(P))$ is non-zero. Then p_j

lies to the right of the currently rightmost point, p_i. Both $\text{NN}(P)$ and $\text{NN}(P^+)$ then increase the range of p_i, and pay the same cost. The only exception is when $i = 0$, that is, p_j is the first point with $p_j > 0$. In this case $\text{NN}(P)$ may pay less than $\text{NN}(P^+)$, since $\text{NN}(P)$ could already have increased the range of p_0 due to arrivals of points to the left of p_0.

A similar argument works for CI. Indeed, $\text{CI}(P^+)$ and $\text{CI}(P^-)$ never extend a ball beyond the currently rightmost and leftmost point, respectively. Hence, when the new point p_j lies, say, to the right of the currently rightmost point p_i, then $\text{CI}(P^+)$ would pay $(\text{dist}(p_i, p_j))^\alpha$. Since $\text{CI}(P)$ also has the option to increase the range of p_i, it will never pay more.

Hence, for NN– a similar computation holds for CI– we get

$$\text{cost}_\alpha(\text{NN}(P)) \leqslant \text{cost}_\alpha(\text{NN}(P^+)) + \text{cost}_\alpha(\text{NN}(P^-)) \leqslant 2 \cdot \text{OPT}(P).$$

It remains to prove the lower bound for part (ii) of the theorem. Assume for a contradiction that there is a constant a such that for all inputs P we have $\text{cost}_\alpha(\text{NN}(P)) \leqslant (2 - \varepsilon) \cdot \text{cost}_\alpha(\text{OPT}(P)) + a$. Consider the input $p_0 = 0$, $p_1 = \delta x$, $p_2 = x$, and $p_3 = -x$, for some $\delta \in (0, 1]$ and $x > 0$ to be determined later. The optimal solution has $r_3(p_0) = x$ and $r_3(p_1) = r_3(p_2) = r_3(p_3) = 0$, while NN has $r_3(p_0) = x$ and $r_3(p_1) = (1 - \delta)x$ and $r_3(p_2) = r_3(p_3) = 0$. Hence, the competitive ratio that NN achieves on this instance is

$$c = \frac{((1 - \delta)^\alpha + 1)x^\alpha - a}{x^\alpha} = (1 - \delta)^\alpha + 1 - \frac{a}{x^\alpha},$$

which is larger than $2 - \varepsilon$ when we pick δ sufficiently small and x sufficiently large. ◀

3 Online range-assignment in $\mathbb{R}^2$

3.1 Bounds on the competitive ratio of NN and CI when $\alpha > 2$

As before, let $p_0, \ldots, p_{n-1}$ be the sequence of inserted points, with p_0 being the source point. Consider a fixed point p_i, and a disk D centered at p_i – the disk D need not have radius equal to the range of p_i. Define $S(p_i, D) := \{p_j : j \geqslant i \text{ and } p_j \in D\}$ to be the set containing p_i plus all points arriving after p_i that lie in D. For a point p_j, define $\text{cost}_\alpha(\text{NN}, p_j)$ to be the cost incurred by NN when p_j is inserted; in other words, $\text{cost}_\alpha(\text{NN}, p_j) := 0$ when p_j falls into an existing disk $B(p_i, r_{j-1}(p_i))$, and $\text{cost}_\alpha(\text{NN}, p_j) := (r_j(p_k))^\alpha - (r_{j-1}(p_k))^\alpha$ otherwise, where $p_k := \text{nn}(p_j)$. Define $\text{cost}_\alpha(\text{CI}, p_j)$ similarly for CI. Finally, for $p_j \in S(p_i, D)$ define

$$F_\alpha(p_j) = F_\alpha(p_j; p_i, D) := \min\{\text{dist}(p_j, p_k)^\alpha \mid p_k \in S(p_i, D) \text{ and } k < j\}.$$

The next lemma shows that we can use the function F_α to upper bound the cost of NN and CI. We later apply this lemma to all disks in an optimal solution to bound the competitive ratio. Note that $\text{cost}_\alpha(\text{NN}, p_j) \leqslant F_\alpha(p_j)$. Indeed, NN either pays zero (when p_j already lies inside a disk) or it expands the disk of p_j's nearest neighbor (which may or may not lie in D) which costs at most $F_\alpha(p_j)$. Similarly $\text{cost}_\alpha(\text{CI}, p_j) \leqslant F_\alpha(p_j)$. Hence we have:

▶ **Lemma 3.1.** *Let p_i be any input point and D a disk centered at p_i. Then for any subset $S(D) \subseteq S(p_i, D) \setminus \{p_i\}$ we have:*

$$\sum_{p_j \in S(D)} \text{cost}_\alpha(\text{NN}, p_j) \leqslant \sum_{p_j \in S(D)} F_\alpha(p_j) \qquad and \qquad \sum_{p_j \in S(D)} \text{cost}_\alpha(\text{CI}, p_j) \leqslant \sum_{p_j \in S(D)} F_\alpha(p_j).$$

Lemma 3.1 suggests the following strategy to bound the competitive ratio of NN (and CI). Consider, for each point p_i, the final disk D placed at p_i in an optimal solution, and let ρ be its radius. The cost of this disk is ρ^α. We charge the cost of the disks placed by NN (or CI) at points p_j inside D – this cost can be bounded using the function F_α, by Lemma 3.1 – to the cost of D. This motivates the following definition:

$$F_\alpha^* := \max \frac{1}{\rho^\alpha} \sum_{p_j \in S(D)} F_\alpha(p_j),$$

where the maximum is over any possible input instance P, any point $p_i \in P$, any disk D of radius ρ centered at p_i, and any subset $S(D) \subseteq S(p_i, D) \setminus \{p_i\}$. The value F_α^* bounds the maximum total charge to any disk D in the optimal solution, relative to D's cost ρ^α. The next lemma shows that for $\alpha > 2$, the value F_α^* is bounded by a constant (depending on α).

▶ **Lemma 3.2.** *We have that $F_\alpha^* \leqslant \alpha \frac{2^\alpha - 3}{2^{\alpha-1} - \alpha}$ for any $\alpha > 2$.*

The formal proof of the lemma is quite technical and can be found in the full version. Here we just sketch the intuition, also showing why the condition $\alpha > 2$ is needed. The quantity F_α^* can be thought of in the following way. Consider a disk D of radius ρ centered at p_i, and imagine the points in $S(D)$ arriving one by one. (The points in $S(p_i, D) \setminus S(D)$ are irrelevant.) Whenever a new point p_j arrives, then F_α^* increases by $\mathrm{dist}(p_j, \mathrm{nn}(p_j))^\alpha$, where $\mathrm{nn}(p_j)$ is p_j's nearest neighbor among the already arrived points from $S(D)$ including p_i. Since the more points arrive the distances to the nearest neighbor will decrease – more precisely, we cannot have many points whose nearest neighbor is at a relatively large distance – the hope is that the sum of these distance to the power α converges, and this is indeed what we can prove for $\alpha > 2$. For $\alpha = 2$ it does not converge, as shown by the following example.

Let D be a unit disk centered at p_i, and consider the inscribed square σ of D. Note that the radius of σ – the distance from its center to its vertices – is 1. We insert a set $S(D)$ of $n-1$ points in rounds, as follows. In the first round we partition σ into four squares of radius $1/2$, and we insert a point in the center of each of them. These four points all have p_i as nearest neighbor, and F_α^* increases by $4 \cdot (1/2)^\alpha = (1/2)^{\alpha-2}$. We recurse in each of the four squares. Thus in the k-th round, we have 4^{k-1} squares of radius $(1/2)^{k-1}$, each of which is partitioned into four squares of radius $(1/2)^k$, and we place a point inside each such subsquare. This increases F_α^* by $4^k \cdot (1/2^k)^\alpha = (1/2^{2-\alpha})^k$. The total cost is $\sum_{k=1}^{t}(1/2^{2-\alpha})^k$, where $t := \Theta(\log n)$ is the number of rounds.

Note that $1/2^{2-\alpha} = 1$ for $\alpha = 2$, giving $F_2^* = \Omega(\log n)$, while for $\alpha > 2$ the total cost converges. Also note that the example only gives a lower bound on F_2^*, it does not show that NN has unbounded competitive ratio for $\alpha = 2$. The reason is that NN actually pays less than F_2^*, since most points p_j are already within range of an existing point upon insertion, and so we do not have to pay $\mathrm{dist}(p_j, \mathrm{nn}(p_j))^\alpha$. Indeed, in the next section we prove, using a different argument, that NN is $O(1)$-competitive even for $\alpha = 2$.

Using Lemma 3.2 we can prove a bound on the competitive ratio of NN and CI.

▶ **Theorem 3.3.** *Let $F_\alpha^* := \alpha \frac{2^\alpha - 3}{2^{\alpha-1} - \alpha}$. For any $\alpha > 2$, the competitive ratio of NN and CI in $\mathbb{R}^2$ is at most $\min\{F_\beta^* \mid 2 < \beta \leqslant \alpha\}$. Hence, for $\alpha \leqslant \alpha^*$, where $\alpha^* = \arg\min F_\alpha^* \approx 4.3$, the competitive ratio is at most $\alpha \frac{2^\alpha - 3}{2^{\alpha-1} - \alpha}$, and for $\alpha > \alpha^*$ it is at most 12.94.*

Proof. Consider a sequence $p_0, p_1, \ldots, p_{n-1}$ of points in the plane. Let D_j be the disk centered at p_j in an optimal solution, after the last point p_{n-1} has been handled, and let ρ_j be its radius. Thus the cost of the optimal solution is $\mathrm{OPT} := \sum_{j=0}^{n-1} \rho_j^\alpha$. To bound the cost

of NN on the same sequence, we charge the cost of inserting p_i, with $0 < i \leqslant n - 1$, to a disk D_j such that $j < i$ and $p_i \in D_j$. Such a disk D_j exists, since after p_i's insertion, p_i is contained in a disk of an existing point p_j and so p_i will also be contained in D_j, the final disk of p_j. (If there are more such points, we take an arbitrary one.) Let $S(D_j)$ be the set of points that charge disk D_j. Note that $\{p_1, \ldots, p_{n-1}\} = \bigcup_{j=0}^{n-2} S(D_j)$. Hence, using Lemmas 3.1 and 3.2, for any $2 < \beta \leqslant \alpha$, for NN (and similarly for CI) we get:

$$
\begin{aligned}
\mathrm{cost}_\alpha(\mathrm{NN}) &= \textstyle\sum_{i=0}^{n-2} \sum_{p_j \in S(D_i)} \mathrm{cost}_\alpha(\mathrm{NN}, p_j) \\
&= \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha \sum_{p_j \in S(D_i)} \frac{\mathrm{cost}_\alpha(\mathrm{NN}, p_j)}{\rho_i^\alpha} \\
&\leqslant \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha \sum_{p_j \in S(D_i)} \frac{\mathrm{dist}(p_j, \mathrm{nn}(p_j))^\alpha}{\rho_i^\alpha} \\
&\leqslant \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha \sum_{p_j \in S(D_i)} \frac{\mathrm{dist}(p_j, \mathrm{nn}(p_j))^\beta}{\rho_i^\beta} && \text{because } \mathrm{dist}(p_j, \mathrm{nn}(p_j)) \leqslant \rho \\
&= \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha \sum_{p_j \in S(D_i)} \frac{F_\beta(p_j)}{\rho_i^\beta} && \text{by Lemma 3.1} \\
&\leqslant \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha F_\beta^* \\
&\leqslant \beta \frac{2^\beta - 3}{2^{\beta-1} - \beta} \textstyle\sum_{i=0}^{n-2} \rho_i^\alpha && \text{by Lemma 3.2} \\
&= \beta \frac{2^\beta - 3}{2^{\beta-1} - \beta} \mathrm{OPT}. && \blacktriangleleft
\end{aligned}
$$

The next theorem gives a lower bound on the competitive ratio of NN.

▶ **Theorem 3.4.** *For any $\alpha > 1$, NN has a competitive ratio of at least $6(1 + (\frac{\sqrt{6} - \sqrt{2}}{2})^\alpha) \approx 6(1 + 0.52^\alpha)$ in the plane. In particular, for $\alpha = 2$, we get a lower bound of 7.6 on the competitive ratio.*

Proof. Let p_0 be the source placed at the origin. The following construction is depicted in Figure 1. We place $p_1, \ldots, p_{18}$ in a disk of radius 1 around p_0 as explained next, such that a possible solution is to place that single disk and pay 1s. For simplicity, in the rest of the proof we use polar coordinates. Let $\varepsilon > 0$ be a positive number. Let then $p_1 = (\varepsilon, 0)$, $p_2(\varepsilon, \pi/3)$, ..., and $p_6 = (\varepsilon, 5\pi/3)$ be the next six points. NN places a disk of radius ε on p_0. Let further $p_7 = (1, 0)$, $p_8 = (1, \pi/3)$, ..., and $p_{12} = (1, 5\pi/3)$ be the next six points. Here NN places six disks of radius $1 - \varepsilon$ centered around $p_1, \ldots, p_6$, paying $6(1 - \varepsilon)^\alpha$. Finally, let $p_{13} = (1, \pi/6 - \varepsilon)$, $p_{14} = (1, 3\pi/6 - \varepsilon)$,..., and $p_{18} = (1, 11\pi/6 - \varepsilon)$ be the last six points. NN is now forced to place 6 disks of radius almost equal to the side of a 12-gon of radius 1, that is $2\sin(\pi/12) - \delta$ for some $\delta > 0$ that tends to 0 as ε tends to 0.

Thus, we have that for any $\varepsilon > 0$, there is an instance on which a solution of cost 1 exists, whereas NN is forced to pay $\varepsilon^\alpha + 6(1 - \varepsilon)^\alpha + 6(2\sin(\pi/12) - \delta)^\alpha$, where $\delta > 0$ tends to 0 as ε tends to 0. We can therefore conclude that NN has to pay at least $6(1 + (2\sin(\pi/12))^\alpha) = 6(1 + (2\frac{\sqrt{6} - \sqrt{2}}{4})^\alpha) = 6(1 + (\frac{\sqrt{6} - \sqrt{2}}{2})^\alpha) \approx 6(1 + 0.52^\alpha)$, whereas $\mathrm{OPT} \leqslant 1$.

We can then scale this construction and thus, there is no constant a such that $\mathrm{cost}(\mathrm{NN}) \leqslant c \cdot \mathrm{OPT} + a$ for $c < 6(1 + (\frac{\sqrt{6} - \sqrt{2}}{2})^\alpha)$. ◀

3.2 Bounds on the competitive ratio of NN and 2-NN when $\alpha = 2$

Above we proved upper bounds for NN and CI for $\alpha > 2$, and we gave a lower bound for NN for any $\alpha > 1$. We now study NN and 2-NN for the case $\alpha = 2$. Unfortunately, the arguments below do not apply to CI.

An upper bound on the competitive ratio of 2-NN for $\alpha = 2$. Let $P := p_0, p_1, \ldots, p_{n-1}$ be the input instance. Recall that $\mathrm{nn}(p_i)$ is the closest point to p_i among $p_0, \ldots, p_{i-1}$. Upon insertion of a point p_i, if p_i is not covered by the current set of balls $B(p_{i'}, r_{j-1}(p_{i'}))$
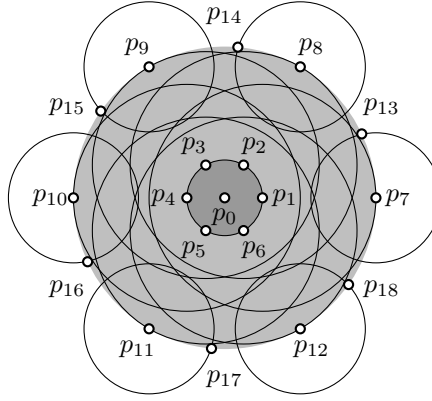
Figure 1 Lower bound on the competitive ratio of NN. The light gray disk (of radius 1) represents the optimal solution on the boundary of which points $p_7, \ldots, p_{18}$ are placed. Points $p_1, \ldots, p_6$ are placed on the boundary of a disk of radius ε (in dark gray). The algorithm NN is forced to place one first disk of radius ε around p_0, then six disks of radius $1 - \varepsilon$ around $p_1, \ldots, p_6$. And finally six disks of radius roughly 0.5 around $p_7, \ldots, p_{12}$.

with $i' < i$, then 2-NN increases the range of $\mathrm{nn}(p_i)$ to $2 \cdot \mathrm{dist}(p_i, \mathrm{nn}(p_i))$, and otherwise it does nothing. Suppose that upon the insertion of some point p_i, we increase the range of $\mathrm{nn}(p_i)$. We now define D_i^* as the disk centered at p_i (*not* at $\mathrm{nn}(p_i)$) and of radius $d_i/2$, where $d_i := \mathrm{dist}(p_i, \mathrm{nn}(p_i))$. We call D_i^* the *charging disk* of p_i. Note that the charging disk is a tool in the proof, it is not a disk used by the algorithm. If 2-NN did nothing upon insertion of p_i because p_i was already covered by a disk, we define $D_i^* := \emptyset$.

▶ **Lemma 3.5.** *For every pair of charging disks D_i^* and D_j^* with $j \neq i$, we have $D_i^* \cap D_j^* = \emptyset$.*

Proof. Without loss of generality we assume that $i < j$. Suppose for a contradiction that $D_i^* \cap D_j^* \neq \emptyset$. Let $p_{i'} := \mathrm{nn}(p_i)$ and $p_{j'} := \mathrm{nn}(p_j)$, and let $d_i := \mathrm{dist}(p_i, p_{i'})$ and $d_j := \mathrm{dist}(p_j, p_{j'})$. Since $i' < i < j$, we have $\mathrm{dist}(p_j, p_{i'}) > 2d_i$, otherwise p_j lies inside the disk of $p_{i'}$ when p_j is inserted and we would have $D_j^* = \emptyset$. On the other hand, $d_i/2 + d_j/2 \geqslant \mathrm{dist}(p_i, p_j)$ because $D_i^* \cap D_j^* \neq \emptyset$. Since $d_j \leqslant \mathrm{dist}(p_i, p_j)$, which is true because we assumed $i < j$, this implies $d_i \geqslant \mathrm{dist}(p_i, p_j)$. But then $\mathrm{dist}(p_j, p_{i'}) \leqslant d_i + \mathrm{dist}(p_i, p_j) \leqslant 2d_i$, a contradiction. ◀

▶ **Lemma 3.6.** *For any points p_i and p_j with $i < j$, let $D_j^{\mathrm{OPT}}(p_i)$ be the disk centered at p_i after p_j is inserted in an optimal solution and let $\rho_j(p_i)$ be its radius. Furthermore, let $D_j^{1.5\,\mathrm{OPT}}(p_i)$ be the disk centered at p_i of radius $1.5 \cdot \rho_j(p_i)$. Then, for every point p_k, there is a point p_i such that the charging disk D_k^* is contained in $D_k^{1.5\,\mathrm{OPT}}(p_i)$.*

Proof. Let p_i be such that p_k is contained in $D_k^{\mathrm{OPT}}(p_i)$. Upon insertion of p_k, we create the charging disk D_k^* of radius $\frac{1}{2} \mathrm{dist}(p_k, \mathrm{nn}(p_k)) \leqslant \frac{1}{2} \mathrm{dist}(p_i, p_k)$ centered at p_k. Therefore, the point of D_k^* furthest from p_i is at distance at most $\frac{3}{2} \mathrm{dist}(p_i, p_k)$. Thus $D_k^* \subset D_k^{1.5OPT}(p_i)$. ◀

Using these two lemmas, we can conclude the following.

▶ **Theorem 3.7.** *In $\mathbb{R}^2$ the strategy 2-NN is 36-competitive for $\alpha = 2$.*

Proof. Recall that the charging disk D_i^* has radius $\mathrm{dist}(p_i, \mathrm{nn}(p_i))/2$. Thus the cost incurred by 2-NN upon the insertion of p_i is at most $(2 \cdot \mathrm{dist}(p_i, \mathrm{nn}(p_i)))^2 \leqslant 16 \cdot \mathrm{radius}(D_i^*)^2$. By Lemma 3.5, the disks D_i^* are pairwise disjoint. Let $\mathcal{D}_{\mathrm{OPT}}$ denote the set of disks in an optimal solution, and let OPT be its cost. Then by Lemma 3.6 we have $\sum_{i=1}^{n-1} \mathrm{radius}(D_i^*)^2 \leqslant \sum_{D \in \mathcal{D}_{\mathrm{OPT}}}((3/2) \cdot \mathrm{radius}(D))^2 = \frac{9}{4}\,\mathrm{OPT}$. Hence the total cost incurred by 2-NN is bounded by $16 \cdot \sum_{i=1}^{n-1} \mathrm{radius}(D_i^*)^2 \leqslant 36\,\mathrm{OPT}$. ◀

Upper bound on the competitive ratio of NN **for** $\alpha = 2$**.** We now prove an upper bound on the competitive ratio of NN using a similar strategy as for 2-NN. The proof uses charging disks, as above. The main difference being how the charging disks are defined.

Suppose that NN increases the range of $\mathrm{nn}(p_i)$ upon the insertion of p_i. Then the *charging disk* D_i^* is the disk of radius $\gamma \cdot d_i$ that is centered on the midpoint of the segment $p_i \, \mathrm{nn}(p_i)$, where $d_i := \mathrm{dist}(p_i, \mathrm{nn}(p_i))$ and γ is a constant to be determined later. If NN did nothing upon insertion of p_i, we define $D_i^* := \emptyset$. As shown in the full version of the paper, the charging disks are disjoint if we pick γ suitably.

▶ **Lemma 3.8.** *Let $\gamma < \frac{3-\sqrt{7}}{4}$. Then for every pair D_i^*, D_j^* of charging disks with $i \neq j$, we have $D_i^* \cap D_j^* = \emptyset$.*

We also need the following lemma, whose proof is similar to that of Lemma 3.6.

▶ **Lemma 3.9.** *For any points p_i and p_j with $i < j$, let $D_j^{\mathrm{OPT}}(p_i)$ be the disk centered at p_i of radius $\rho_j(p_i)$ after p_j is inserted in an optimal solution and let $D_j^{(1.5+\gamma)\,\mathrm{OPT}}(p_i)$ be the disk centered at p_i of radius $(1.5 + \gamma)\rho_j(p_i)$. Then, for every point p_k, there is a point p_i such that the disk D_k^* is contained in the disk $D_k^{(1.5+\gamma)\,\mathrm{OPT}}(p_i)$.*

Putting everything together we obtain the following theorem.

▶ **Theorem 3.10.** *In $\mathbb{R}^2$ the strategy NN is 322-competitive for $\alpha = 2$.*

Proof. Recall that $\mathrm{radius}(D_i^*) = \gamma \cdot \mathrm{dist}(p_i, nn(p_i))$. Thus the cost incurred by NN upon the insertion of p_i is at most $\mathrm{dist}(p_i, nn(p_i))^2 \leqslant ((1/\gamma) \cdot \mathrm{radius}(D_i^*))^2$. By Lemma 3.8, the disks D_i^* are pairwise disjoint. If $\mathcal{D}_{\mathrm{OPT}}$ denotes the set of disks used in an optimal solution, then by Lemma 3.9 we have $\sum_{i=1}^{n-1} \rho(D_i^*)^2 \leqslant \sum_{D \in \mathcal{D}_{\mathrm{OPT}}} ((1.5 + \gamma) \cdot \rho(D))^2 = (1.5 + \gamma)^2 \, \mathrm{OPT}$, where OPT is the cost of an optimal solution. Hence the total cost incurred by NN is at most $\frac{1}{\gamma^2} \sum_{i=1}^{n-1} \rho(D_i^*)^2 = \frac{1}{\gamma^2} \cdot (1.5 + \gamma)^2 \, \mathrm{OPT}$. Since this holds for any value of $\gamma < \frac{3-\sqrt{7}}{4}$, we can conclude that the cost incurred by NN is at most $\frac{4^2}{(3-\sqrt{7})^2} \cdot (1.5 + \frac{3-\sqrt{7}}{4})^2 \, \mathrm{OPT} = (163 + 60\sqrt{7}) \, \mathrm{OPT} < 322 \, \mathrm{OPT}$. ◀

4 Online range-assignment in general metric spaces

In this section we consider the problem in general metric spaces. In the full version of the paper we also consider the offline variant of the problem; here we focus on the online variant, for which we give an $O(\log n)$-competitive algorithm. The key insight to our algorithms is to formulate the problem as a set-cover problem and apply linear-programming techniques. As we will see later, applying the online set cover algorithm of Alon et al. [1] yields a competitive ratio much worse than $O(\log n)$, so we need to exploit structural properties of the particular set cover instances arising from our problem.

4.1 A set cover formulation and its LP

Let $\mathcal{R}$ be the set of distances between pairs of points. Observe that we can restrict ourselves without loss of generality to only using ranges from $\mathcal{R}$. This allows us to formulate the problem in terms of set cover: The elements are the points $p_0, p_1, \ldots, p_{n-1}$, with p_0 being the source point, and for each $0 \leqslant i \leqslant n - 2$ and $r \in \mathcal{R}$ there is a set $S_{i,r} := \{p_j : j > i \text{ and } \mathrm{dist}(p_i, p_j) \leqslant r\}$ with cost r^α. (Note that $S_{i,r}$ is the set of points arriving after p_i that are within range r of p_i). In the following, we abuse notation and also write $j \in S_{i,r}$ for points $p_j \in S_{i,r}$. We also say that $S_{i,r}$ is *centered* at p_i.

Observe that a feasible range assignment corresponds to a feasible set cover. A set cover is *minimally feasible* if removing any set from it causes an element to be uncovered. Since a minimally feasible set cover picks at most one set $S_{i,r}$ for each i, it corresponds to a feasible range assignment. (Note that applying the online set cover algorithm of Alon et al. [1] only gives a competitive ratio of $O(\log^2 n / \log \log n)$ as our set cover instance has $n-1$ elements and $|\mathcal{R}|(n-1)$ sets.)

We can now formulate our problem as an integer linear program. To this end we introduce, for each range $r \in \mathcal{R}$ and each point p_i a variable $x_{i,r}$, where $x_{i,r} = 1$ indicates we choose the set $S_{i,r}$ (or, in other words, that we assign range r to p_i) and $x_{i,r} = 0$ indicates we do not choose $S_{i,r}$. Allowing the $x_{i,r}$ to take fractional values gives us the following relaxed LP.

$$
\begin{array}{lll}
\text{Minimize} & \displaystyle\sum_{0 \leqslant i \leqslant n-2} \sum_{r \in \mathcal{R}} x_{i,r} \cdot r^\alpha & \\[2em]
\text{Subject to} & \displaystyle\sum_{i,r: j \in S_{i,r}} x_{i,r} \geqslant 1 & \text{for all } 1 \leqslant j \leqslant n-1 \\[2em]
& x_{i,r} \geqslant 0 & \text{for all } (i,r) \text{ with } 0 \leqslant i \leqslant n-2 \text{ and } r \in \mathcal{R}
\end{array}
\tag{1}
$$

The dual LP corresponding to the LP above is as follows.

$$
\begin{array}{lll}
\text{Maximize} & \displaystyle\sum_{1 \leqslant j \leqslant n} y_j & \\[2em]
\text{Subject to} & \displaystyle\sum_{j \in S_{i,r}} y_j \leqslant r^\alpha & \text{for all } (i,r) \text{ with } 0 \leqslant i \leqslant n-2 \text{ and } r \in \mathcal{R} \\[2em]
& y_j \geqslant 0 & \text{for all } 1 \leqslant j \leqslant n-1
\end{array}
\tag{2}
$$

We say that the set $S_{i,r}$ is *tight* if the corresponding dual constraint is tight, that is, if $\sum_{j \in S_{i,r}} y_j = r^\alpha$.

4.2 The online algorithm and its analysis

Recall that in the online version, we are given the source p_0 and then the points $p_1, \dots, p_{n-1}$ arrive one-by-one. When a point p_i arrives, its distances to previous points and the source are revealed.

The algorithm. Let $\gamma > 1$ be a constant that we will set later. The basic idea of the algorithm is that when a point p_i arrives, we will raise its associated dual variable y_i until some set $S_{j,r}$ containing p_i is tight and then update the range of point p_j to be $r_i(p_j) := \gamma \max\{r : \sum_{k \in S_{j,r}: k \leqslant i} y_k = r^\alpha\}$. In other words, the range of p_j becomes γ times the largest radius of the tight sets centered at p_j.

Here is a more precise description of the algorithm. When p_i arrives, we initialize its dual variable $y_i := 0$. If $p_i \in S_{j,r}$ for some $j < i$ and range r with $\sum_{k \in S_{j,r}: k \leqslant i} y_k = r^\alpha$, then we set $r_i(p_j) := \gamma \max\{r : \sum_{k \in S_{j,r}: k \leqslant i} y_k = r^\alpha\}$ for one such j. (It can happen that some $S_{j,r}$ is tight but that $r_{i-1}(p_j)$ is still smaller than r, because when multiple sets become tight at the same time, we only increase the range of one point.) Otherwise, we increase y_i until for some $j < i$ and range r we have $\sum_{k \in S_{j,r}: k \leqslant i} y_k = r^\alpha$; we then set p_j's new range to $r_i(p_j) := \gamma r$ for one such j. In both cases, we only set p_j's range, the other ranges remain unchanged. Note that in the event that multiple sets centered at different points become tight simultaneously, we only update the range of one of them.

Analysis. We begin our analysis of the algorithm by showing the feasibility of the constructed dual solution y and the corresponding range assignment. For each point p_i, the algorithm stops raising y_i once some set $S_{j,r}$ containing p_i is tight and then updates p_j's radius to be $\gamma r > r$. This guarantees that no dual constraint is violated and that p_i is covered by p_j.

Next we analyze the cost of this algorithm. We use the shorthand r_j for the final range $r_{n-1}(p_j)$ of the point p_j. First, we argue that it suffices to bound the cost of the points whose ranges are large enough. Let $H = \{0 \leqslant i \leqslant n-2 : r_i \geqslant \max_{0 \leqslant j \leqslant n-2} r_j/n\}$. Then, the cost of the algorithm is

$$\sum_i r_i^\alpha = \sum_{i \in H} r_i^\alpha + \sum_{i \notin H} r_i^\alpha \leqslant \sum_{i \in H} r_i^\alpha + n(\max_j r_j/n)^\alpha \leqslant (1 + 1/n^{\alpha-1}) \sum_{i \in H} r_i^\alpha \leqslant 2 \sum_{i \in H} r_i^\alpha,$$

where the second last inequality is because $\sum_{i \in H} r_i^\alpha \geqslant \max_j r_j^\alpha$ and the last is because $\alpha > 1$. In the remainder of this section we will show that

$$\sum_{i \in H} r_i^\alpha \leqslant O(\log n) \cdot \sum_{1 \leqslant j \leqslant n-1} y_j. \tag{3}$$

The theorem then follows from the Weak Duality Theorem of Linear Programming which states that value of any feasible solution to the primal (minimization) problem is always greater than or equal to the value of any feasible solution to its associated dual problem.

For $0 \leqslant i \leqslant n-2$, our algorithm sets the final range r_i of point p_i such that $r_i = \gamma r$ for some $r \in \mathcal{R}$ such that $\sum_{k \in S_{i,r}} y_k = r^\alpha$. Thus, we get

$$\left(\frac{r_i}{\gamma}\right)^\alpha = \sum_{j \in S_{i,r_i/\gamma}} y_j,$$

and so

$$\sum_{i \in H} r_i^\alpha = \sum_{i \in H} \gamma^\alpha \left(\sum_{j \in S_{i,r_i/\gamma}} y_j\right) = \gamma^\alpha \sum_{1 \leqslant j \leqslant n-1} y_j \cdot \left|\{i \in H : j \in S_{i,r_i/\gamma}\}\right|,$$

where the last equality follows by interchanging the sums. Thus, to prove Inequality (3) it suffices to prove the following lemma.

▶ **Lemma 4.1.** *For every $1 \leqslant j \leqslant n-1$ and any fixed $\gamma > 3$, we have $|\{i \in H : j \in S_{i,r_i/\gamma}\}| = O(\gamma^\alpha \log n)$.*

Proof. Define $H_j = \{i \in H : j \in S_{i,r_i/\gamma}\}$. We will show that for every $i, i' \in H_j$, either $r_i > \frac{\gamma-1}{2} r_{i'}$ or $r_{i'} > \frac{\gamma-1}{2} r_i$. This implies that the t-th smallest range (among the points in H_j) is at least $((\gamma-1)/2)^t$ times the smallest range (among those points). Since $\frac{\max_{i \in H_j} r_i}{\min_{i \in H_j} r_i} \leqslant n$, this means that $|H_j| = O(\log_{(\gamma-1)/2} n) = O(\log n)$.

Suppose $i, i' \in H_j$. Let p_k be the last-arriving point that causes our algorithm to update r_i, and $p_{k'}$ be the last-arriving point that causes our algorithm to update $r_{i'}$. Since the arrival of any point causes at most one point's range to be updated, we have that $p_k \neq p_{k'}$. Suppose that p_k arrived before $p_{k'}$. By construction of $r_{i'}$, we have $\text{dist}(p_{k'}, p_{i'}) = r_{i'}/\gamma$. Moreover, since $i, i' \in H_j$, we have $\text{dist}(p_i, p_j) \leqslant r_i/\gamma$ and $\text{dist}(p_{i'}, p_j) \leqslant r_{i'}/\gamma$. Therefore, by the triangle inequality,

$$\text{dist}(p_i, p_{k'}) \leqslant \text{dist}(p_i, p_j) + \text{dist}(p_j, p_{i'}) + \text{dist}(p_{i'}, p_{k'}) \leqslant 2\frac{r_{i'}}{\gamma} + \frac{r_i}{\gamma}.$$

Since p_k arrived before $p_{k'}$ and $p_{k'}$ caused our algorithm to update $r_{i'}$, the point $p_{k'}$ must have been uncovered when it arrived, and so $\mathrm{dist}(p_i, p_{k'}) > r_i$. Therefore, we get

$$r_i < \mathrm{dist}(p_i, p_{k'}) \leqslant 2\frac{r_{i'}}{\gamma} + \frac{r_i}{\gamma}$$

and so $r_{i'} > \frac{\gamma-1}{2} r_i$ as desired. In the case that $p_{k'}$ arrived before p_k, a similar argument yields $r_i > \frac{\gamma-1}{2} r_{i'}$.   ◄

By setting $\gamma = 4$ we obtain the following theorem.

▶ **Theorem 4.2.** *For any power-distance gradient $\alpha > 1$, there is a $O(4^\alpha \log n)$-competitive algorithm for the online range assignment problem in general metric spaces.*

5   Concluding Remarks

We introduced the online version of the broadcast range-assignment problem, and we analyzed the competitive ratio of two natural algorithm, NN and CI, in $\mathbb{R}^1$ and $\mathbb{R}^2$ as a function of the power-distance gradient α. While NN is $O(1)$-competitive in $\mathbb{R}^2$ and for $\alpha = 2$ the best competitive ratio we can prove is quite large, namely 322. The variant 2-NN has a better ratio, namely 36, but this is still large. We conjecture that the actual competitive ratio of NN is actually much closer to the lower bound we proved, which is 7.61. We also conjecture that CI has a constant (and small) competitive ratio in $\mathbb{R}^2$. Another approach to getting better competitive ratios might be to develop more sophisticated algorithms. For the general (metric-space) version of the problem, the main question is whether an algorithm with constant competitive ratio is possible.

While the requirement that we cannot decrease the range of any point in the online setting is perhaps not necessary in practice, our algorithms have the additional benefit that they modify the range of at most one point. Thus it can also be seen as the first step in studying a more general version, where we are allowed to modify (increase or decrease) the range of, say, two points. In general, it is interesting to study trade-offs between the number of modifications and the competitive ratio. Studying deletions is then also of interest.

─── **References** ───

1   Noga Alon, Baruch Awerbuch, Yossi Azar, Niv Buchbinder, and Joseph Naor. The online set cover problem. *SIAM J. Comput.*, 39(2):361–370, 2009. `doi:10.1137/060661946`.

2   Christoph Ambühl, Andrea E. F. Clementi, Miriam Di Ianni, Nissan Lev-Tov, Angelo Monti, David Peleg, Gianluca Rossi, and Riccardo Silvestri. Efficient algorithms for low-energy bounded-hop broadcast in ad-hoc wireless networks. In *Proc. 21st Annual Symposium on Theoretical Aspects of Computer Science (STACS 2004)*, volume 2996 of *Lecture Notes in Computer Science*, pages 418–427. Springer, 2004. `doi:10.1007/978-3-540-24749-4_37`.

3   Joan Boyar, Stephan J. Eidenbenz, Lene M. Favrholdt, Michal Kotrbčík, and Kim S. Larsen. Online dominating set. *Algorithmica*, 81(5):1938–1964, 2019. `doi:10.1007/s00453-018-0519-1`.

4   Gruia Călinescu. Approximate min-power strong connectivity. *SIAM J. Discret. Math.*, 27(3):1527–1543, 2013. `doi:10.1137/100819540`.

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. Verho-even. Some ecent 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 assignmet 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. Springer, 1999. `doi:10.1007/978-3-540-48413-4_21`.

10 Andrea E. F. Clementi, Paolo Penna, and Riccardo Silvestri. On the power assignment problem in radio networks. *Mob. Networks Appl.*, 9(2):125–140, 2004. `doi:10.1023/B:MONE.0000013624.32948.87`.

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

12 Fabrizio Grandoni. On min-power steiner tree. In *Proc. 20th Annual European Symposium on Algorithms (ESA 2012)*, volume 7501 of *Lecture Notes in Computer Science*, pages 527–538. Springer, 2012. `doi:10.1007/978-3-642-33090-2_46`.

13 Sudipto Guha and Samir Khuller. Improved methods for approximating node weighted steiner trees and connected dominating sets. *Inf. Comput.*, 150(1):57–74, 1999. `doi:10.1006/inco.1998.2754`.

14 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`.

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

16 Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985. `doi:10.1145/2786.2793`.

17 Peng-Jun Wan, Gruia Călinescu, Xiang-Yang Li, and Ophir Frieder. Minimum-energy broadcasting in static ad hoc wireless networks. *Wirel. Networks*, 8(6):607–617, 2002. `doi:10.1023/A:1020381720601`.