Tight Bounds for Online TSP on the Line

Antje Bjelde[*1], Yann Disser[2], Jan Hackfeld[†1], Christoph Hansknecht[3], Maarten Lipmann[4], Julie Meißner[‡1], Kevin Schewior[§5], Miriam Schlöter[1], and Leen Stougie[6]

[1]TU Berlin, Institute of Mathematics, Germany.
{bjelde,hackfeld,jmeiss,schloeter}@math.tu-berlin.de
[2]TU Darmstadt, Institute of Mathematics, Germany.
disser@mathematik.tu-darmstadt.de
[3]TU Braunschweig, Institute for Mathematical Optimization, Germany.
c.hansknecht@tu-braunschweig.de
[4]Maarten Lipmann, Amsterdam, Netherlands. Maarten.lipmann@gmail.com
[5]Universidad de Chile, Santiago, Chile. kschewior@gmail.com
[6]Vrije Universiteit, Department of Econometrics and Operations Research & CWI, Amsterdam, Netherlands. l.stougie@vu.nl

Abstract

We consider the online traveling salesperson problem (TSP), where requests appear online over time on the real line and need to be visited by a server initially located at the origin. We distinguish between closed and open online TSP, depending on whether the server eventually needs to return to the origin or not. While online TSP on the line is a very natural online problem that was introduced more than two decades ago, no tight competitive analysis was known to date. We settle this problem by providing tight bounds on the competitive ratios for both the closed and the open variant of the problem. In particular, for closed online TSP, we provide a 1.64-competitive algorithm, thus matching a known lower bound. For open online TSP, we give a new upper bound as well as a matching lower bound that establish the remarkable competitive ratio of 2.04.

Additionally, we consider the online Dial-A-Ride problem on the line, where each request needs to be transported to a specified destination. We provide an improved non-preemptive lower bound of 1.75 for this setting, as well as an improved preemptive algorithm with competitive ratio 2.41.

Finally, we generalize known and give new complexity results for the underlying offline problems. In particular, we give an algorithm with running time $\mathcal{O}(n^2)$ for closed offline TSP on the line with release dates and show that both variants of offline Dial-A-Ride on the line are NP-hard for any capacity $c \geq 2$ of the server.

1 Introduction

In the online Traveling Salesperson Problem (TSP) on the line, we consider a server initially located at the origin of the real line that has to serve requests that appear over time. The server has unit speed and serves requests (in any order) by moving to the position of the corresponding request at some time after its release. The objective in online TSP on the line is to minimize the makespan, i.e., the time until all requests have been served. In the *closed* variant of the problem, the server needs to return to the origin after serving all requests, while the *open* variant has no such requirement.

[*]Supported by Einstein Foundation Berlin in the framework of Matheon
[†]Supported by DFG Priority Programme 1736 Algorithms for Big Data.
[‡]Supported by Einstein Foundation Berlin in the framework of Matheon and by the German Science Foundation (DFG) under contract ME 3825/1.
[§]Partially supported by the Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003.

Online TSP is a natural online problem similar to the classical k-server problem [22]. In the latter, the order in which requests need to be served is prescribed, and the problem thus becomes trivial on the line for $k = 1$ server. In contrast, online TSP on the line is a non-trivial problem that arises in 1-dimensional collection/delivery problems. Examples include robotic welding/screwing/depositing material, horizontal/vertical item delivery systems, and the collection of objects from mass storage shelves. The online Dial-A-Ride problem additionally allows transportation requests that specify a source and destination that need to be visited by the server in this order. If the capacity of the server is finite, it limits the number of requests that can be transported simultaneously. The online Dial-A-Ride problem on the line arises, e.g., when controlling industrial or personal elevators.

While both online TSP and online Dial-A-Ride on the line are among the most natural online problems and have been studied extensively over the last two decades [2, 4, 5, 6, 7, 10, 15, 17, 18, 19], no satisfactory (tight) analysis was known for either problem in terms of competitive ratios. We address this shortcoming for TSP on the line by providing a tight upper bound for the closed variant, as well as tight bounds for the open variant. We emphasize that our results for the open and closed variant of the problem are independent and require substantially different approaches. Aside from our results for online TSP, we narrow the gaps for online Dial-A-Ride on the line by giving improved bounds. In addition to online results, we study the computational complexity of the underlying offline problems.

1.1 Our results

We have the following results[1] (cf. Tables 1 and 2) :

Tight bounds for online TSP on the line. Our main results are best-possible online algorithms for both the open and closed variant of online TSP on the line, as well as a new (tight) lower bound for the open variant. Our algorithm for the closed variant has a competitive ratio of $(9 + \sqrt{17})/8 \approx 1.64$, matching a lower bound of Ausiello et al. [6] and improving on their 1.75-competitive algorithm. For open TSP on the line, we give a lower bound of 2.04 on the competitive ratio, which is the first bound strictly greater than 2. We also provide an optimal online algorithm matching this bound and improving on the 2.33-competitive algorithm by Ausiello et al. [6]. Our results settle online TSP on the line from the perspective of competitive analysis.

Improved bounds for online Dial-A-Ride on the line. Our lower bounds for online TSP on the line immediately apply to preemptive and non-preemptive online Dial-A-Ride on the line. In particular, our lower bound of 2.04 is the first bound greater than 2 for the open variant of the problem. Additionally, we provide a simple preemptive 2.41-competitive algorithm, which improves a (non-preemptive) 3.41-competitive algorithm by Krumke [17]. For the closed Dial-A-Ride variant, the lower bound of 1.64 by Ausiello et al. [6] was improved for one server with unit capacity without preemption to 1.71 by Ascheuer et al. [2]. We improve this bound further to 1.75 for any finite capacity $c \geq 1$. The best known algorithm for closed Dial-A-Ride on the line for finite capacity $c \geq 1$ is 2-competitive and was given by Ascheuer et al. [2].

New offline complexity results. Regarding offline TSP on the line with release times, Psaraftis et al. [23] showed a dynamic program that solves the open variant in quadratic time. We refute their claim that all optimal closed tours have a very simple structure with a counterexample, and we adapt their algorithm to find an optimal closed tour in quadratic time. For the non-preemptive offline Dial-A-Ride problem on the line, results have previously been obtained for the closed variant without release times. For capacity $c = 1$ Gilmore and Gomory [13] and Atallah and Kosaraju [3] gave polynomial time algorithms, and Guan [14] proved hardness for the case $c = 2$. We show that both the open and closed variant of the problem are NP-hard for *any* capacity $c \geq 2$. Additionally, we show that the case with release times and any $c \geq 1$ is NP-hard. The complexity of offline Dial-A-Ride on the line with unbounded capacity remains open.

[1] Parts of our results were already claimed in [21], but mostly with weaker bounds and without a conclusive proof. Nevertheless, some of our ideas are inspired by the approaches described in [21].

Table 1: Overview of our results for online TSP on the line and online Dial-A-Ride on the line.

Online	closed: lower bound	closed: upper bound	open: lower bound	open: upper bound
online TSP on the line				
new	1.64 [5, 6]	**1.64** (Th. 3)	**2.04** (Th. 4)	**2.04** (Th. 10)
old		1.75 [5, 6]	2 [4, 6]	2.33 [4, 6]
Dial-A-Ride on the line				
preemptive	1.64 [5, 6]	2 [2]	**2.04** (Th. 4)	**2.41** (Th. 12)
non-preemptive	**1.75** (Th. 13)			3.41 [17]

Table 2: Overview of our results for offline TSP and Dial-A-Ride on the line with release times.

Offline	closed	open
TSP on the line	$\mathcal{O}(n^2)$ (Th. 15)	$\mathcal{O}(n^2)$ [23]
Dial-A-Ride on the line		
non-preemptive	NP-hard (Th. 16)	NP-hard (Th. 16)

1.2 Further related work

For the online TSP problem in general metric spaces, Ausiello et al. [6] show a lower bound of 2 on the competitive ratio for the open version and a 1.64 lower bound for the closed version, both bounds being achieved on the real line. For the open online TSP, they present a 2.5-competitive algorithm, and for the closed version they give a 2-competitive algorithm. Jaillet and Wagner [15] give 2-competitive algorithms for the closed version that can additionally deal with precedence constraints or multiple servers. Blom et al. [7] consider the closed online TSP problem on the non-negative part of the real line and present a best possible algorithm with competitive ratio 1.5. They also study a "fair" setting where the optimum does not travel outside the convex hull of the known requests, and they derive an algorithm for the real half-line with a better competitive ratio of 1.28 for this setting. Krumke et al. [19] show that there cannot be a competitive algorithm for open online TSP with the objective of minimizing the maximum flow time instead of minimizing the makespan. For the real line they define a fair setting and give a competitive algorithm for it.

The *online repairperson problem* is the open online TSP problem with the objective of minimizing the weighted sum of completion times. Feuerstein and Stougie [10] show a lower bound of 5.83 on the best-possible competitive ratio for this problem and provide a 9-competitive algorithm for the real line. Krumke et al. [18] give a best-possible online algorithm with competitive ratio 5.83 for general metric spaces.

For the the closed online Dial-A-Ride problem without preemption, Feuerstein and Stougie [10] show a lower bound of 2 for the competitive ratio in general, and present an algorithm with a best-possible competitive ratio of 2 for the case that the server has infinite capacity. Ascheuer et al. [2] analyze different algorithms for the same setting and present a 2-competitive algorithm for any finite capacity $c \geq 1$. For minimizing the sum of completion times instead of the makespan, Feuerstein and Stougie [10] further show a lower bound of 3 for a server with unit capacity and a lower bound of 2.41 independent of the capacity. Moreover, they provide a 15-competitive algorithm for the real line and unlimited capacity. For the same objective function, Krumke et al. [18] present an algorithm with a competitive ratio of 5.83 for a server with unit capacity in an arbitrary metric space.

The offline version of the TSP problem is a well-studied NP-hard problem (e.g., see [20]). Afrati et al. [1] show that the offline traveling repairperson problem is NP-hard in general, but can be solved in time $\mathcal{O}(n^2)$ for the real line and unweighted sum of completion times objective. There are many offline variants of the Dial-A-Ride problem, differing in capacities, the underlying metric space, release times and deadlines, open versus closed tours, and in whether preemption is allowed (e.g., see [9]). The special case without release times and unit capacity is known as the *stacker crane problem*. Attalah and Kosaraju [3] present a polynomial algorithm for the closed, non-preemptive stacker crane problem on the real line. Frederickson and

Guan [11] show that this problem is NP-complete on trees. Guan [14] shows that the Dial-A-Ride problem remains easy on the line with capacities larger than one if preemption is allowed, and that it remains hard on trees. Finally, Charikar and Raghavachari [8] give a $\mathcal{O}(\sqrt{c}\log n \log\log n)$-approximation for the closed Dial-A-Ride problem in metric spaces with n points and without preemption. In the same paper they claim a 2-approximation for the problem on the line, however this result seems to be incorrect (personal communication).

2 Problem definition and notation

We consider a server that moves along the real line with (at most) unit speed. We let p_t denote the position of the server at time $t \geq 0$ and assume (without loss of generality) that $p_0 = 0$. With this notation, the speed limitation of the server can equivalently be expressed via $|p_t - p_{t'}| \leq |t - t'|$ for all $t, t' \geq 0$. A series of requests $\sigma_1, \dots, \sigma_n$ arrives over time with $\sigma_i = (a_i, b_i; t_i)$, where $t_i \geq 0$ denotes the release time of the request and $a_i, b_i \in \mathbb{R}$ denote its source and target position, respectively. For TSP, we have $a_i = b_i$ and write $\sigma_i = (a_i; t_i)$. If not stated otherwise, we assume $t_1 \leq t_2 \leq \ldots \leq t_n$. Moreover, we assume without loss of generality that $t_i \geq |a_i|$ holds because the server can not reach σ_i before time $|a_i|$ and it only helps the algorithm to know a request earlier. We further use the notation $A^{\mathrm{R}} := \max_{i=1,\dots,n}\{a_i, b_i, 0\}$ to denote the rightmost point that needs to be visited by the server, and similarly $A^{\mathrm{L}} := \min_{i=1,\dots,n}\{a_i, b_i, 0\}$. Here and throughout we refer to the negative direction of the real line as *left* and the positive direction as *right*.

In both TSP and Dial-A-Ride on the line, all requests need to be *served*. For TSP, we consider a request served if $p_t = a_i$ for some time $t \geq t_i$. For Dial-A-Ride, the server may collect request σ_i at time $t \geq t_i$ if $p_t = a_i$. In the preemptive Dial-A-Ride problem, the server can drop off any request it is carrying at its current location at any time. If request σ_i is dropped off at point p at time t, we consider it to be modified to the new request $(p, b_i; t)$. In the non-preemptive Dial-A-Ride problem, the server may only drop off a request at its target location. We consider a request served if it is ever dropped off at its target location.

In TSP on the line, the behavior of the server in our algorithms at time t will mostly depend on so-called *extreme requests*. For $t \geq 0$, we denote by $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}(t); t^{\mathrm{R}}(t))$ the unserved request that is rightmost of the position of the server p_t, provided such a request exists, i.e., the unserved request $\sigma = (a; t')$ with $t' \leq t$, $a > p_t$, and maximizing a. Analogously, $\sigma^{\mathrm{L}}(t) = (a^{\mathrm{L}}(t); t^{\mathrm{L}}(t))$ denotes unserved request that is leftmost of the position of the server p_t. If there is more than one right-most (left-most) request, we choose the one with the largest release time.

If the server has finite capacity $c \geq 1$, it can carry at most c requests at any time. We assume that no time is needed for picking up and dropping off requests, so that the server can pick up and drop off any number of requests at the same time, as long as its capacity is not exceeded.

We refer to a valid trajectory of the server together with the description of when it picks up and drops requests as a tour T. If the tour ends at $p_0 = 0$, we call it *closed*, otherwise it is *open*. We denote the makespan of the tour T by $|T|$. The objective in the open (closed) version of both TSP and Dial-A-Ride is to find an open (closed) tour T that serves all requests and minimizes $|T|$.

In the *offline* setting, we assume all requests to be known from the start. We let T^{Opt} denote an optimal offline tour. In the *online* setting, we assume that request σ_i is revealed at its release time t_i, at which point the tour of the server until time t_i must already have been fixed irrevocably. Additionally, we assume that the total number n of requests is unknown. We measure the quality of an online algorithm via its competitive ratio, i.e., the maximum over all sequences of requests of the ratio between the makespan of the tour it produces and $|T^{\mathrm{Opt}}|$.

In order to describe the trajectory of the server, we use the notation "move(a)" for the tour that moves the server from its current position with unit speed to the point $a \in \mathbb{R}$ and the notation "waituntil(s)" for the tour that keeps the server stationary until time s. We use the operator $\oplus$ to concatenate tours. For example, if T_0 is a tour of the server that ends at time t_0 at position p_{t_0}, then $T_0 \oplus \text{move}(a)$ describes the tour that ends at time $t_0 + |a - p_{t_0}|$, is identical to the tour T_0 until time t_0 and satisfies $p_t = p_{t_0} + (a - p_{t_0})(t - t_0)$ for $t_0 \leq t \leq t_0 + |a - p_{t_0}|$. Similarly, $T_0 \oplus \text{waituntil}(s)$ is the tour that ends at time $\max\{t_0, s\}$, is identical to the tour T_0 until time t_0 and that satisfies $p_{t_0} = p_t$ for all $s \in [t_0, s]$. For TSP on the line, we do not explicitly specify when a request is served, but we assume that the server serves a request whenever possible, i.e., whenever the server passes the location of a request that is already released and not yet served.

We skip the proofs of our results as they are very technical. They can be found in the full version of the paper.

3 Algorithm for closed online TSP

In this section we consider the closed online TSP problem and describe a best-possible algorithm with competitive ratio $\rho = (9+\sqrt{17})/8 \approx 1.64$, where ρ is the nonnegative root of the polynomial $4x^2 - 9x + 4$.

We start by developing some intuition for our algorithm. In the following $T^{\textsc{Alg}}$ is the tour derived by an algorithm $\textsc{Alg}$. Observe that the decision of how to move the server at time t only depends on its position p_t and the location of the left- and rightmost extreme requests $\sigma^{\mathrm{L}}(t) = (a^{\mathrm{L}}; t^{\mathrm{L}}(t))$ and $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}; t^{\mathrm{R}}(t))$: All other requests can be served during any tour serving $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$. We will show that in this setting we can assume that $a^{\mathrm{L}}(t) \le 0$ and $a^{\mathrm{R}}(t) \ge 0$, provided these extremes exist. If $\sigma^{\mathrm{R}}(t)$ (resp. $\sigma^{\mathrm{L}}(t)$) does not exist we set $a^{\mathrm{R}}(t) = t^{\mathrm{R}}(t) = 0$ (resp. $a^{\mathrm{L}}(t) = t^{\mathrm{L}}(t) = 0$). Thus, in contrast to the initial definition of extreme requests, in our setting a leftmost extreme is always left and a rightmost extreme is always right of the origin. When both extremes exists, we have, on a high level, three possible courses of action at time t. Either we immediately decide to serve $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$ in one of the two possible orders, or we wait for some time for additional information to make a more informed decision. Intuitively, the critical case for our competitiveness is the case where we decide to serve $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$ in a different order than $T^{\textsc{Opt}}$. Let $T_{\mathrm{RL}}(t)$ and $T_{\mathrm{LR}}(t)$ be the tours that start at the origin at time 0 and then move as follows,

$$
\begin{aligned}
T_{\mathrm{RL}}(t) &= \mathrm{waituntil}(t^{\mathrm{R}}(t) - |a^{\mathrm{R}}(t)|) \oplus \mathrm{move}(a^{\mathrm{R}}(t)) \\
&\quad \oplus \mathrm{move}(a^{\mathrm{L}}(t)) \oplus \mathrm{move}(p_0), \\
T_{\mathrm{LR}}(t) &= \mathrm{waituntil}(t^{\mathrm{L}}(t) - |a^{\mathrm{L}}(t)|) \oplus \mathrm{move}(a^{\mathrm{L}}(t)) \\
&\quad \oplus \mathrm{move}(a^{\mathrm{R}}(t)) \oplus \mathrm{move}(p_0).
\end{aligned}
$$

Note that $|T_{\mathrm{RL}}(t)|$ (resp. $|T_{\mathrm{LR}}(t)|$) is a lower bound for the length of the shortest tour serving $\sigma^{\mathrm{R}}(t)$ before $\sigma^{\mathrm{L}}(t)$ (resp. $\sigma^{\mathrm{L}}(t)$ before $\sigma^{\mathrm{R}}(t)$). Say that, at time t, both extremes exist and we greedily decide to immediately start serving the extremes in the same order as $T_{\mathrm{LR}}(t)$. To see how this can fail, assume that $T^{\textsc{Opt}}$ initially follows the tour $T_{\mathrm{RL}}(t)$, but continues to move to the left after serving $\sigma^{\mathrm{L}}(t)$. The time when $T^{\textsc{Opt}}$ reaches $a^{\mathrm{L}}(t)$ is $t' = t^{\mathrm{R}}(t) + |a^{\mathrm{R}}(t)| + |a^{\mathrm{L}}(t)|$, since $t^{\mathrm{R}}(t) \ge |a^{\mathrm{R}}(t)|$ by assumption. Let t_0 be the time when we reach the origin p_0 after serving $\sigma^{\mathrm{L}}(t)$, and assume that $t' \le t_0$. Now a new request $\sigma' = (p'; t_0)$ may arrive at time t_0 and position $p' = -|a^{\mathrm{L}}(t)| - (t_0 - t') = -t_0 + t^{\mathrm{R}}(t) + |a^{\mathrm{R}}(t)|$, that the optimum can serve immediately at time t_0. We then have

$$
\begin{aligned}
|T^{\textsc{Opt}}| &= t_0 + |p'| = t_0 + |a^{\mathrm{L}}(t)| + (t_0 - t') \\
&= 2t_0 - t^{\mathrm{R}}(t) - |a^{\mathrm{R}}(t)|.
\end{aligned}
$$

Our algorithm still needs to serve σ' and σ^{R} at time t_0, and hence

$$
\begin{aligned}
|T^{\textsc{Alg}}| &= t_0 + 2|p'| + 2|a^{\mathrm{R}}(t)| \\
&= 3t_0 - 2t^{\mathrm{R}}(t).
\end{aligned}
$$

For the algorithm to be ρ-competitive, we need $|T^{\textsc{Alg}}|/|T^{\textsc{Opt}}| \le \rho$, and we thus obtain a condition on the earliest time t_0 we may return to the origin.

Fact 1. If at time t a ρ-competitive algorithm serves the leftmost extreme $\sigma^{\mathrm{L}}(t)$ first and t_0 denotes the first time the server returns to the origin after having served $\sigma^{\mathrm{L}}(t)$, then

$$
t_0 \ge t_0^{\mathrm{L}}(t) := \frac{\rho|a^{\mathrm{R}}(t)| - (2-\rho)t^{\mathrm{R}}(t)}{2\rho - 3}. \tag{1}
$$

A symmetric statement with $t_0^{\mathrm{R}}(t)$ holds if the algorithm serves $\sigma^{\mathrm{R}}(t)$ first.

Fact 1 illustrates that waiting is sometimes necessary in order to be competitive. On the other hand, we can obviously not afford to wait too long. To quantify this, we introduce a lower bound on the length of $T^{\textsc{Opt}}$.

Definition 1. If both extreme request exist at time t we define the *greedy tour* $T^{\text{greedy}}(t)$ at time t as

$$T^{\text{greedy}}(t) := \begin{cases} T_{\text{LR}}(t), & \text{if } |T_{\text{LR}}(t)| \le |T_{\text{RL}}(t)|, \\ T_{\text{RL}}(t), & \text{else.} \end{cases}$$

If only $\sigma^{\text{L}}(t)$ exists, we set $T^{\text{greedy}}(t) := T_{\text{LR}}(t)$ and we set $T^{\text{greedy}}(t) := T_{\text{RL}}(t)$ if only $\sigma^{\text{R}}(t)$ exists.

Observation 1. We have

$$\begin{aligned} |T_{\text{RL}}(t)| &= t^{\text{R}}(t) + |a^{\text{R}}(t)| + 2|a^{\text{L}}(t)|, \\ |T_{\text{LR}}(t)| &= t^{\text{L}}(t) + |a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \end{aligned}$$

and $|T^{\text{greedy}}(t)| = \min\{|T_{\text{LR}}(t)|, |T_{\text{RL}}(t)|\} \le |T^{\text{OPT}}|$ if both extremes exist at time t.

Assume we are still waiting at the origin at time t, i.e. $p_t = 0$. From Observation 1, we conclude that if $t \le \rho|T^{\text{greedy}}(t)| - 2|a^{\text{R}}(t)| - 2|a^{\text{L}}(t)|$, we can wait until time $\rho|T^{\text{greedy}}(t)| - 2|a^{\text{R}}(t)| - 2|a^{\text{L}}(t)|$, and then still serve $\sigma^{\text{R}}(t)$ and $\sigma^{\text{L}}(t)$ and return to the origin p_0 until time $\rho|T^{\text{greedy}}| \le \rho|T^{\text{OPT}}|$, i.e., we can stay ρ-competitive. Formally, we make the following definition.

Definition 2. Let the *safe tour* $T^{\text{safe}}(t)$ at time t be defined as

$$T^{\text{safe}}(t) := \begin{cases} T^{\text{wait}} \oplus T_{\text{LR}}(t), & \text{if } |a^{\text{L}}(t)| \ge |a^{\text{R}}(t)|, \\ T^{\text{wait}} \oplus T_{\text{RL}}(t), & \text{else,} \end{cases}$$

with

$$T^{\text{wait}} := \text{waituntil}(\rho|T^{\text{greedy}}(t)| - 2|a^{\text{R}}(t)| - 2|a^{\text{L}}(t)|).$$

We still have to ensure that the definition of $T^{\text{safe}}(t)$ is compatible with the requirement from Fact 1 regarding the time t_0 when the tour first returns to the origin. In case $|a^{\text{L}}(t)| \ge |a^{\text{R}}(t)|$, we get $t_0 \ge t_0^{\text{L}}(t)$ with $t_0 = \rho|T^{\text{greedy}}(t)| - 2|a^{\text{R}}(t)| \ge (4\rho - 2)|a^{\text{R}}(t)|$. Symmetrically, for $|a^{\text{L}}(t)| < |a^{\text{R}}(t)|$, we get $(4\rho - 2)|a^{\text{L}}(t)| \ge t_0^{\text{R}}(t)$. In either case, we can derive the following condition on ρ.

Fact 2. The safe tour $T^{\text{safe}}(t)$ fulfills inequality (1) in Fact 1 if and only if $2 \ge \rho \ge \frac{9+\sqrt{17}}{8}$.

We are now ready to describe our algorithm (cf. Algorithm 1). We argued that it is a safe option to

---

Algorithm 1: $\textsc{Update}(t, \sigma^{\text{L}}(t), \sigma^{\text{R}}(t), p_t)$ for the closed online TSP Problem

---

this function is called upon release of a new extreme request
Input: time t, unserved extreme requests $\sigma^{\text{R}}(t)$ and $\sigma^{\text{L}}(t)$, position p_t of the server
Output: closed Online TSP tour serving all unserved requests
$A \leftarrow \text{argmax}_{g \in \{a^{\text{R}}(t), a^{\text{L}}(t)\}} |g|$; $a \leftarrow \text{argmin}_{g \in \{a^{\text{R}}(t), a^{\text{L}}(t)\}} |g|$
if $T^{\text{greedy}}(t) = T_{\text{LR}}(t)$ **then**
  $(a_1(t); t_1) \leftarrow \sigma^{\text{L}}(t)$; $(a_2(t); t_2) \leftarrow \sigma^{\text{R}}(t)$
else
  $(a_1(t); t_1) \leftarrow \sigma^{\text{R}}(t)$; $(a_2(t), t_2) \leftarrow \sigma^{\text{L}}(t)$
if $t^{\text{wait}} := \rho|T^{\text{greedy}}(t)| - (|p_t - A| + |A| + 2|a|) \ge t$ **then** // $T^{\text{safe}}(t)$ can be reached
(A) $T^{\text{ALG}} \leftarrow \text{waituntil}(t^{\text{wait}}) \oplus \text{move}(A) \oplus \text{move}(a) \oplus \text{move}(p_0)$
else if $\text{sign}(p_t) = \text{sign}(a_1(t))$ **or** $t + |p_t - a_2(t)| + |a_2(t)| < \frac{\rho|a_1(t)| - (2-\rho)t_1}{2\rho - 3}$ **then**
(B1,B2) $T^{\text{ALG}} \leftarrow \text{move}(a_1(t)) \oplus \text{move}(a_2(t)) \oplus \text{move}(p_0)$
else
(C) $T^{\text{ALG}} \leftarrow \text{move}(a_2(t)) \oplus \text{move}(a_1(t)) \oplus \text{move}(p_0)$
return T^{ALG}

---

follow $T^{\text{safe}}(t)$ in order to stay ρ-competitive, provided no further requests appear. It will turn out that it is indeed always good enough to follow $T^{\text{safe}}(t)$, if possible. However, at time t, we may be too far from the extreme that $T^{\text{safe}}(t)$ serves first in order to catch up with the safe tour, in which case we have to resort to secondary strategies. If at time t the server cannot reach $T^{\text{safe}}(t)$, it instead bases its behavior on the greedy tour as an estimate for T^{OPT}. Surprisingly, this estimate turns out to be sufficient to obtain an optimal online algorithm. There are three situations that can occur if the safe tour cannot be reached at time t. If the online server is on the same side of the origin as the extreme that $T^{\text{greedy}}(t)$ serves first, our algorithm decides to follow the greedy tour. If the online server is on the other side of the origin than the extreme that $T^{\text{greedy}}(t)$ serves first, we have to ensure that the condition of Fact 1 is not violated. If the tour serving the nearer extreme first satisfies (1), the algorithm serves this extreme first, i.e., it serves the extremes in a different order than $T^{\text{greedy}}(t)$. Otherwise, we can deduce from (1) being violated that we can afford to serve the opposite extreme first, i.e., to follow $T^{\text{greedy}}(t)$.

Theorem 3. *There is a $(9 + \sqrt{17})/8 \approx 1.64$-competitive algorithm for closed online* TSP *on the line.*

We obtain the main result of this section by analyzing each of the above cases.

4 Lower Bound for open online TSP

In this section, we consider open online TSP on the line and give a tight lower bound on the best-possible competitive ratio. Note that a lower bound of 2 is obvious: At time 1, we present a request either at -1 or 1, whichever is further away from the online server. The online tour has length at least 2 while the optimum tour has length 1. Remarkably, we are able to show a slightly larger bound that turns out to be tight.

Theorem 4. *Let $\rho \approx 2.04$ be the second-largest root (out of the four real roots) of $9\rho^4 - 18\rho^3 - 78\rho^2 + 210\rho - 107$. There is no $(\rho - \varepsilon)$-competitive algorithm for open TSP on the line for any $\varepsilon > 0$.*

In the following, we fix any online algorithm ALG and $\rho' \in (2, \rho)$ and describe an adversarial strategy that forces $|T^{\text{ALG}}|$ to be larger than $|T^{\text{OPT}}|$ by a factor of at least ρ'. After the first request σ_0^{R}, which is to the right[2] of the origin, we alternatingly present leftmost and rightmost extreme requests, in the i-th iteration called σ_i^{L} and σ_i^{R}, respectively, depending on ALG's behavior. Roughly, a new leftmost request σ_i^{L} appears whenever the last rightmost request σ_{i-1}^{R} is served, and a new rightmost request σ_i^{R} appears when ALG has moved close enough to σ_i^{L}. Importantly, we will show that some pair $(\sigma_i^{\text{L}}, \sigma_i^{\text{R}})$ is *critical*, in the following sense.

Definition 5. We call the last two requests $\sigma_0^\star = (a_0^\star; |a_0^\star|)$ and $\sigma_1^\star = (a_1^\star; |a_1^\star|)$ of a request sequence with $\text{sign}(a_0^\star) \neq \text{sign}(a_1^\star)$ and $0 < |a_0^\star| \leq |a_1^\star|$ *critical* for ALG if the following conditions hold:

(i) Both tours $\text{move}(a_0^\star) \oplus \text{move}(a_1^\star)$ and $\text{move}(a_1^\star) \oplus \text{move}(a_0^\star)$ serve all the requests presented until time $|a_1^\star|$.

(ii) ALG serves both $\sigma_0^\star$ and $\sigma_1^\star$ after time $|a_1^\star|$, and $p_{|a_1^\star|}$ lies between $a_0^\star$ and $a_1^\star$.

(iii) Let $k \in \{0, 1\}$ be such that ALG serves $\sigma_k^\star$ before $\sigma_{1-k}^\star$. Then ALG serves $\sigma_k^\star$ no earlier than $t^\star := (2\rho' - 2) \cdot |a_{1-k}^\star| + (\rho' - 2) \cdot |a_k^\star|$.

(iv) It holds that $|a_{1-k}^\star| / |a_k^\star| \leq 2$.

Indeed, we have the following lemma.

Lemma 6. *If there is a request sequence with two critical requests for* ALG*, we can release additional requests such that* ALG *is not $(\rho - \varepsilon)$-competitive on the resulting instance.*

In the proof, we use the notation from Definition 5. We assume that $\text{sign}(a_k^\star) \geq 0$; the other case is symmetric. For the sake of readability, we define $\sigma^{\text{L}} = (a^{\text{L}}; -a^{\text{L}}) := \sigma_{1-k}^\star$ and $\sigma^{\text{R}} = (a^{\text{R}}; a^{\text{R}}) := \sigma_k^\star$.

Conceptually, we want to present additional requests after $|a_1^*|$ so that ALG serves σ^{R} before σ^{L}. However, it will turn out that serving σ^{R} first is a mistake for ALG, compared with using the tour $T_{\text{LR}} := \text{move}(\sigma^{\text{L}}) \oplus \text{move}(\sigma^{\text{R}})$. Roughly, we make OPT follow the tour T_{LR} and then let it continue moving to the right until all requests are served by ALG. Accordingly, we will ensure that all additional requests we introduce coincide with OPT's position at their release time.

[2] We assume $p_1 \leq 0$; the other case is symmetrical.

Assume that we could force $\textsc{Alg}$ to serve σ^{L} immediately after σ^{R}, before serving any additional requests. In this case, we could simply introduce another request at a^{R} at time $|T_{\mathrm{LR}}|$, and, by Definition 5 (iii), we would have

$$\begin{aligned} |T^{\textsc{Alg}}| &\geq t^\star + 2(|a^{\mathrm{R}}| + |a^{\mathrm{L}}|) \\ &= (2\rho' - 2) \cdot |a^{\mathrm{L}}| + (\rho' - 2) \cdot |a^{\mathrm{R}}| \\ &\quad + 2(|a^{\mathrm{R}}| + |a^{\mathrm{L}}|) \\ &= \rho'(2|a^{\mathrm{L}}| + |a^{\mathrm{R}}|) = \rho'|T^{\textsc{Opt}}|, \end{aligned} \tag{2}$$

as claimed.

In general, however, $\textsc{Alg}$ may not serve σ^{L} immediately after σ^{R}, for example by waiting for a while at a^{R} – which forces us to postpone the release of additional requests. Of course, $\textsc{Alg}$ needs to start moving towards σ^{L} at some point to stay competitve if no new requests appear. Our goal is to balance these two effects by introducing one or two new requests.

The additional requests we use depend on the tour that $\textsc{Alg}$ takes after time $|a_1^\star|$. Towards this, let $t^{\star\star}$ be the earliest possible time that a server starting in a^{R} at time $t^\star$ could serve σ^{L}, that is,

$$\begin{aligned} t^{\star\star} &:= t^\star + |a^{\mathrm{R}}| + |a^{\mathrm{L}}| \\ &= (2\rho' - 2) \cdot |a^{\mathrm{L}}| + (\rho' - 2) \cdot |a^{\mathrm{R}}| + |a^{\mathrm{R}}| + |a^{\mathrm{L}}| \\ &= (2\rho' - 1) \cdot |a^{\mathrm{L}}| + (\rho' - 1) \cdot |a^{\mathrm{R}}|. \end{aligned}$$

We characterize the trajectory of $\textsc{Alg}$ at time $t \geq |a_1^\star|$ by the difference between $t^{\star\star}$ and the earliest possible time that $\textsc{Alg}$ can still serve σ^{L}, if it aborts its tour at time t and takes the shortest tour serving σ^{R} (if needed) and then σ^{L}. Formally, for $t \geq |a_1^\star|$, we define

$$\operatorname{delay}(t) := \begin{cases} t + 2|a^{\mathrm{R}}| + |a^{\mathrm{L}}| - p_t - t^{\star\star}, \\ \qquad\qquad \text{if } \sigma^{\mathrm{R}} \text{ not served at } t, \\ t + |a^{\mathrm{L}}| + p_t - t^{\star\star}, \\ \qquad\qquad \text{if } \sigma^{\mathrm{R}} \text{ served at } t \text{ but } \sigma^{\mathrm{L}} \text{ not}, \\ \text{undefined, else.} \end{cases}$$

It is easy to see that the following properties hold.

Fact 3. Consider some t such that $\operatorname{delay}(t)$ is defined. Let $\mathcal{T}$ be the set of tours that start in position p_t at time t and, if $\textsc{Alg}$ has not served σ^{R} at time t, that do not visit a^{L} before a^{R}. The following is true:

(i) There is no tour $T \in \mathcal{T}$ that arrives at a^{L} earlier than $t^{\star\star} + \operatorname{delay}(t)$.

(ii) There is a tour $T \in \mathcal{T}$ that arrives at a^{L} at time $t^{\star\star} + \operatorname{delay}(t)$.

The following two lemmata state useful properties of the delay function that will be used to define the additional requests.

Lemma 7. *There exists $W \geq 0$ with*

$$\operatorname{delay}\left(|T_{\mathrm{LR}}| + \frac{W}{\rho' - 1}\right) = W. \tag{3}$$

Lemma 8. *With W as in Lemma 7,* $\textsc{Alg}$ *serves σ^{R} no later than time $|T_{\mathrm{LR}}| + \frac{W}{\rho'-1}$.*

We will show that if we present an additional request at time $|T_{\mathrm{LR}}| + W/(\rho' - 1)$ (at a distance of $W/(\rho' - 1)$ to the right of σ^{R}) and $\textsc{Alg}$ decides to serve σ^{L} before the new request, the ratio between $\textsc{Alg}$'s and $\textsc{Opt}$'s additional costs (Inequality (2)) is at least ρ'. If $\textsc{Alg}$ can save time by serving the new request first and does so, we need to present yet another request.

It remains to show that we can define a request sequence (depending on $\textsc{Alg}$) that ends with a pair of critical requests. We use the following strategy:

- W.l.o.g. $p_1 \leq 0$. The first request is $\sigma_0^{\mathrm{R}} := (1, 1)$.

- Whenever, at some time, called t_i^{L} in the following, a request at $\sigma_{i-1}^{\mathrm{R}}$ gets served, we present the new request $\sigma_i^{\mathrm{L}} := (a_i^{\mathrm{L}} = -t_i^{\mathrm{L}}; t_i^{\mathrm{L}})$. Based on t_i^{L}, we define for $t \geq t_i^{\mathrm{L}}$ the two functions

$$\ell_i^{\mathrm{L}}(t) := (2\rho' - 3) \cdot t - (3 - \rho') \cdot t_i^{\mathrm{L}},$$
$$\ell_i^{\mathrm{R}}(t) := (4 - \rho') \cdot t - (2\rho' - 2) \cdot t_i^{\mathrm{L}},$$

  which can as well be viewed as lines in the path-time diagram.

- If at some time t_i^{R} after t_i^{L} ALG crosses $\ell_i^{\mathrm{L}}(t)$ or $\ell_i^{\mathrm{R}}(t)$, we present the request $\sigma_i^{\mathrm{R}} := (a_i^{\mathrm{R}} = t_i^{\mathrm{R}}; t_i^{\mathrm{R}})$.

- We stop the procedure when one of the following cases occurs. (The pair $(\sigma_i^{\mathrm{L}}, \sigma_i^{\mathrm{R}})$ will be shown to be critical in these cases.)

  Case 1: ALG serves σ_i^{L} before σ_i^{R} if no new requests appear.
  Case 2: ALG serves σ_i^{R} not before time $(2\rho' - 2) \cdot t_i^{\mathrm{L}} + (\rho' - 2) \cdot t_i^{\mathrm{R}}$ if no new requests appear.

The intuition behind the lines ℓ_i^{L} and ℓ_i^{R} is the following: Suppose the position of ALG at t_i^{R} is on or to the right of ℓ_i^{L} and ALG decides to serve σ_i^{L} before σ_i^{R} in case no new requests appear after t_i^{R}. Then the pair $(\sigma_i^{\mathrm{L}}, \sigma_i^{\mathrm{R}})$ satisfies Definition 5 (iii). The symmetric statement holds for ℓ_i^{R}. The following lemma ensures that, in each iteration, we obtain a (non necessarily critical) pair of unserved requests $(\sigma_i^{\mathrm{L}}, \sigma_i^{\mathrm{R}})$ and that Definition 5 (iv) is fulfilled.

Lemma 9. *Let $i \geq 1$. At time t_i^{L},* ALG *is to the right of ℓ_i^{L} and ℓ_i^{R} and crosses one of them after t_i^{L} and before it serves σ_i^{L}. We also have that $t_i^{\mathrm{R}} \leq 2t_i^{\mathrm{L}}$.*

In the proof of Theorem 4, we show that Case 1 or 2 eventually occurs, and we formalize the above intuition to show along with Lemma 9 that the requests $(\sigma_i^{\mathrm{L}}, \sigma_i^{\mathrm{R}})$ are indeed critical.

5 Algorithm for open online TSP

In Algorithm 2, we describe an algorithm for the open online TSP on the line which achieves a competitive ratio of $\rho \approx 2.04$, matching the lower bound presented in Section 4.

Theorem 10. *Algorithm 2 is ρ-competitive with $\rho \approx 2.04$ being the second-largest root of the polynomial $9\rho^4 - 18\rho^3 - 78\rho^2 + 210\rho - 107$.*

In the following, we discuss the different cases that can occur in Algorithm 2 and give an intuition of their interplay. Algorithm 2 is called every time a new extreme request is released. It then computes a new tour serving the current extreme requests and thus also all other requests between these extremes. It is important that we wait in certain cases to protect against the release of new extremes that force us to serve the extremes in a different order than we initially chose. Algorithm 2 lets the server return to the origin and wait there whenever possible, and moves to the extremes as late as possible. Intuitively, staying as close to the origin as possible has the benefit that the algorithm can delay the choice in which order to serve the extremes for as long as possible. We observe in the lower bound construction (Section 4), that a ρ-competitive algorithm may not serve a request too early, i.e., $|p_t|/t$ is bounded when a request is served. The exact bound on this ratio is computed in the proof of Lemma 9 and coincides with the bound in Lemma 11 for Algorithm 2.

Lemma 11. *The position of the server in a tour computed by Algorithm 2 satisfies*

$$\frac{|p_t|}{t} \leq \frac{3\rho - 5}{-3\rho^2 + 9\rho - 4} \approx 0.58$$

for all times $t \geq 0$.

We decide the order in which we serve the extremes and possibly wait based on the current position p_t of the server, the current time t and the release times and positions of the single or the two extremes. The server always tries to move back to the origin and wait there as long as possible, i.e., it follows the tour $T_0 := \mathrm{move}(0) \oplus \mathrm{waituntil}(\infty)$ for as long as possible. We use the notation "$T_0.\,\mathrm{until}(\mathrm{condition}(\tau))$" to denote

Algorithm 2: For the open online TSP problem on the line.

This function is called every time a new extreme request is released.
Input: Current time t, current extremes, current position p_t
Output: The next part of the tour for the server
$T_0 := \text{move}(0) \oplus \text{waituntil}(\infty)$
if $\exists$ *only single extreme* $\sigma_1 = (a_1; t_1)$ **then**
(P1),(E1) **return** $T_0.\,\text{until}(\tau + |p_\tau - a_1| \geq \rho t_1) \oplus \text{move}(a_1)$
if $0 \notin [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ **then**
  $(\sigma_1 = (a_1; t_1),\ \sigma_2 = (a_2; t_2)) \leftarrow$ extremes such that $|a_1| \leq |a_2|$
(O) **return** $\text{move}(a_1) \oplus T_0.\,\text{until}(\tau + |p_\tau - a_2| \geq \rho t_2) \oplus \text{move}(a_2)$
else
  $(\sigma_1 = (a_1; t_1),\ \sigma_2 = (a_2; t_2)) \leftarrow$ extremes such that $t_1 \leq t_2$
  if $t + |p_t - a_1| \leq L^{\sigma_1, \sigma_2}$ **then**
(P2) **return** $T_0.\,\text{until}(\tau + |p_\tau - a_1| = L^{\sigma_1, \sigma_2}) \oplus \text{move}(a_1) \oplus \text{move}(a_2)$
  else if $t + |p_t - a_2| \leq L^{\sigma_2, \sigma_1}$ **and** $|a_2| \leq \frac{3\rho - 5}{(2\rho-2)(7-3\rho)}(\rho t_1 + (\rho - 2)|a_1|)$ **then**
(A2) **return** $T_0.\,\text{until}(\tau + |p_\tau - a_2| = L^{\sigma_2, \sigma_1}) \oplus \text{move}(a_2) \oplus \text{move}(a_1)$
  else
(E2) **return** $\text{move}(a_1) \oplus \text{move}(a_2)$

the tour that follows T_0 until the first time t_0 that satisfies "condition(t_0)". Note that this may happen before the server reaches the origin.

We now discuss the different cases of the algorithm step-by-step. Let us first consider the simplest case where only one extreme request $\sigma_1 = (a_1; t_1)$ is present (cases (P1) and (E1) in Algorithm 2). The offline optimum Opt obviously cannot finish before time t_1 in this case. In order to guarantee ρ-competitiveness it is therefore sufficient to serve σ_1 at time ρt_1. Hence, we can afford to move to the origin and wait until the equation $t + |p_t - a_1| = \rho t_1$ is satisfied for the current time t and position p_t. This ensures that $|p_t|/t$ is bounded as stated in Lemma 11 (this is not clear and formally proved in the full version of the paper) and that we have the option to change our tour without much additional cost if an extreme on the other side is released later on. We call the resulting tour in this case, the *preferred tour* (case (P1)). It can happen, however, that we are only able to serve σ_1 later than time ρt_1, i.e. $t + |p_t - a_1| > \rho t_1$. This means that the until-condition is already satisfied and the server serves σ_1 immediately. We call this tour the *enforced tour* (case (E1)). The makespan of Algorithm 2 in this case is $|T^{\mathrm{Alg}}| = t_1 + |p_{t_1} - a_1|$. But we have $|T^{\mathrm{Opt}}| \geq t_1$ and also $|T^{\mathrm{Opt}}| \geq |p_t - a_1|$ as Algorithm 2 only visits points on the real line that must also be visited by Opt at some time. Overall, we have $|T^{\mathrm{Alg}}| \leq 2 \cdot \mathrm{Opt}$, implying ρ-competitiveness.

Now, consider the case where two extremes $\sigma_1 = (a_1; t_1)$ and $\sigma_2 = (a_2; t_2)$ are present. The two extremes define an interval

$$[a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)] = [\min\{a_1, a_2\}, \max\{a_1, a_2\}].$$

Note that $a^{\mathrm{L}}(t) < p_t < a^{\mathrm{R}}(t)$ holds by definition of extreme requests.

If $0 \notin [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ (case (O) in Algorithm 2), we let σ_1 denote the request closer to the origin, i.e., $|a_1| \leq |a_2|$. We then immediately serve σ_1 in order to ensure that Lemma 11 holds and $|p_t|/t$ stays small. We know that the offline optimum Opt cannot finish before time t_2. After serving σ_1 it is therefore safe to return to the origin and wait as long as we can to reach a_2 at time ρt_2. We can thus follow the tour T_0 after serving σ_1 at time $t + |p_t - a_2| \geq \rho t_2$. Possibly, this equation is already satisfied from the start and we serve σ_2 immediately.

Next, consider the case that $0 \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ (cases (P2), (A2) and (E2)). Let now σ_1 be the request released first, i.e. $t_1 \leq t_2$, and $t = t_2$ be the current time. We have a lower bound of $|T^{\mathrm{OPT}}| \geq t_1 + |a_1| + |a_2|$, irrespective of whether Opt serves σ_1 or σ_2 first. If we serve σ_1 first, which we call the *preferred tour* (case (P2)), we ensure that the tour produced by Algorithm 2 is not longer than $\rho|T^{\mathrm{Opt}}|$ by satisfying the inequality

$$t + |p_t - a_1| + |a_1| + |a_2| \leq \rho(t_1 + |a_1| + |a_2|). \tag{4}$$

Now assume that Opt serves σ_2 first and a new request $\sigma_1' = (a_1; t_2 + |a_1| + |a_2|)$ appears at the same position as σ_1 when Opt arrives there. The new request does not increase the cost for Opt, which is still only lower bounded by $|T^{\text{OPT}}| \geq t_2 + |a_1| + |a_2|$. But our algorithm, which served σ_1 first, may be closer to σ_2 at the time when this new request appears and now has to go all the way back to the position of a_1 after serving σ_2. The intuition why it is sufficient to protect against this worst-case is that if σ_1' appears at a position further away from the origin, then this additional distance has to be traveled by Opt as well, and if σ_1' appears closer to the origin, it only benefits our algorithm. In order to ensure ρ-competitiveness in this scenario, the following inequality has to also be satisfied:

$$\begin{aligned} &t + |p_t - a_1| + 2(|a_1| + |a_2|) \\ &\quad \leq \rho(t_2 + |a_1| + |a_2|). \end{aligned} \tag{5}$$

If we now define

$$\begin{aligned} L^{\sigma_1,\sigma_2} := \min\{&\rho t_1 + (\rho - 1)|a_1| + (\rho - 1)|a_2|, \\ &\rho t_2 + (\rho - 2)|a_1| + (\rho - 2)|a_2|\}, \end{aligned}$$

then Inequalities (4) and (5) are simultaneously satisfied if and only if $t + |p_t - a_1| \leq L^{\sigma_1,\sigma_2}$. If this latter inequality is satisfied with some slack, we again follow the tour T_0 until it becomes tight, so that $|p_t|/t$ stays small and we are more flexible to change our tour when new requests appear.

In case the conditions for the preferred tour are not satisfied, we try to serve σ_2 first. This is called the *anticipated tour* (case (A2)). The inequalities that need to be satisfied in this case are the same as those for the preferred tour with σ_1 and σ_2 exchanged. Moreover, to ensure that $|p_t|/t$ is bounded as claimed in Lemma 11, the following inequality also needs to be satisfied if Algorithm 2 serves σ_2 first:

$$\begin{aligned} |a_2| &\leq \frac{3\rho - 5}{(2\rho - 2)(7 - 3\rho)} \left(\rho t_1 + (\rho - 2)|a_1|\right) \\ &\approx 1.21 \cdot t_1 + 0.02 \cdot |a_1|. \end{aligned} \tag{6}$$

Intuitively, the inequality ensures that σ_2 is not too far from the origin compared to σ_1, as otherwise we would need to start too early to serve the extreme σ_2 and would violate the bound on $|p_t|/t$.

Lastly, in case neither the conditions for the preferred tour nor the anticipated tour are met, we immediately serve the earlier request σ_1 first and σ_2 directly afterwards. This is called the *enforced tour* (case (E2)). The main challenge in showing ρ-competitiveness in the analysis of the algorithm is to derive a better lower bound on Opt in this case by considering extremes that must have been released before.

6 Online Dial-A-Ride on the line

In this section we give a $(1 + \sqrt{2})$-competitive algorithm for (preemptive) open online Dial-A-Ride on the line. Let T_S^{OPT} denote the optimal tour over a set of requests S, starting at position $p_0 = 0$, and let R_t denote the set of released but not yet delivered requests at a time t. Our algorithm works as follows (cf. Algorithm 3): The server stays at position p_0 until the first request arrives. With every new arriving request, the server stops its current tour and returns to p_0, or stays at p_0 if it is already at the origin. The server starts following the tour $T_{R_t}^{\text{OPT}}$ at time $\sqrt{2} \cdot |T_{R_t}^{\text{OPT}}|$. By "unload" we denote the operation of unloading all requests the server is currently carrying at the current position. Formally, each such request $(a, b; t')$ is changed to $(p_t, b; t')$.

Algorithm 3: For the open online Dial-A-Ride problem with $c \geq 0$ fixed.

this function is called upon receiving a new request;
Input: new request σ, current position p_t, unserved requests R_t
Output: an open tour starting at p_t and serving all requests in R_t
return unload $\oplus$ move(p_0)
$\oplus$ waituntil$(\sqrt{2} \cdot |T_{R_t}^{\text{OPT}}|) \oplus T_{R_t}^{\text{OPT}}$

We show that we get back to the origin in time whenever a new request is released.

Theorem 12. *Algorithm 3 is $(1+\sqrt{2}) \approx 2.41$-competitive for the preemptive open online* Dial-A-Ride *problem with capacity $c \geq 1$.*

Note. Algorithm 3 can easily be modified to solve open Dial-A-Ride in general metric spaces and upholds the same competitive ratio. The proof remains the same.

We also provide a lower bound for non-preemptive closed Dial-A-Ride on the line that improves the lower bound of 1.70 from [2].

Theorem 13. *No algorithm for the non-preemptive closed* Dial-A-Ride *problem on the line with fixed capacity $c \geq 1$ has competitive ratio lower than $\rho = 1.75$.*

7 The offline problem

Psaraftis et al. [23] show that open offline TSP on the line with release dates can be solved in quadratic time. For the closed variant they claim that the optimal tour has the structure [23, pp. 215–216]:

$$\text{waituntil}() \oplus \text{move}(A^{\mathrm{R}}) \oplus \text{move}(A^{\mathrm{L}}) \oplus \text{move}(p_0)$$

or

$$\text{waituntil}() \oplus \text{move}(A^{\mathrm{L}}) \oplus \text{move}(A^{\mathrm{R}}) \oplus \text{move}(p_0).$$

Here the waiting time at the origin is chosen maximally such that all requests are still served. We contradict this claim by showing that an optimal server tour may need to turn around arbitrarily many times.

Theorem 14. *For every $k \in \mathbb{N}$, there is an instance of closed* TSP *on the line such that any optimal solution turns around at least $2k$ times.*

Next, we show a dynamic program that solves closed offline TSP on the line in quadratic time. It is inspired by the one of Psaraftis et al. [23] for the open variant and a dynamic program of Tsitsiklis [24] for the same problem with deadlines instead of release times.

The algorithm (cf. Algorithm 4) relies on the fact that an optimal server tour has a 'zig-zag shape' with decreasing amplitude.

Algorithm 4: Dynamic Program for Closed Offline TSP on the line.

Input: A set of requests $\sigma_i = (a_i; t_i)$ with $t_i \geq |a_i|$ for $-\ell \leq i \leq r$ and $a_i < a_{i+1}$ for $-\ell \leq i < r$.
Output: The minimum completion time $C_{\max}$ of a tour.
$C^{+}_{-\ell-1,r} \leftarrow t_r$.
$C^{-}_{-\ell,r+1} \leftarrow t_{-\ell}$.
for $i = -\ell, ..., r$ **do**
  $C^{-}_{i,r+1} \leftarrow \max\{t_i, C^{-}_{i-1,r+1} + a_i - a_{i-1}\}$.
for $j = r, .., \ell$ **do**
  $C^{+}_{-\ell-1,j} \leftarrow \max\{t_j, C^{+}_{\ell-1,j+1} + a_{j+1} - a_j\}$.
for $d = r + \ell, ..., 0$ **do**
  for $i = -\ell, .., r - d$ **do**
    $j \leftarrow i + d$.
    $C^{-}_{i,j} \leftarrow \max\{t_i, \min\{C^{+}_{i-1,j} + a_j - a_i, C^{-}_{i-1,j} + a_i - a_{i-1}\}\}$.
    $C^{+}_{i,j} \leftarrow \max\{t_j, \min\{C^{+}_{i,j+1} + a_{j+1} - a_j, C^{-}_{i,j+1} + a_j - a_i\}\}$.
return $C_{\max} = C^{+}_{0,0}$.

We index the requests by increasing positions. Then we compute for each index pair $i < j$ the completion time of two tours serving all requests to positions smaller than position a_i and all requests to positions larger than position a_j. We use $C^+_{i,j}$ for the best such tour ending at position a_j and $C^-_{i,j}$ for the best one ending at a_i. Starting with the largest difference $j - i$ and iteratively decreasing it, we recursively compute $C^+_{i,j}$ and $C^-_{i,j}$.

Theorem 15. *Algorithm 4 computes the minimum completion time of a server tour for offline TSP on the line in time $O(n^2)$.*

For the non-preemptive Dial-A-Ride problem on the line we show that the open and closed variant with release times are NP-hard. Without release times, we prove they are NP-hard for capacity $c \geq 2$. Our reductions are from the circular arc coloring problem, which is also used in a reduction for minimizing the sum of completion times of Dial-A-Ride on the line with capacity $c = 1$ [9].

Theorem 16. *The non-preemptive open and closed offline* Dial-A-Ride *problem on the line are NP-complete. For capacity $c \geq 2$ this even holds when all release times are 0.*

In the classification of closed dial-a-ride problems in [9], offline TSP on the line is the problem $1|s = t, d_j|\text{line}|C_{\max}$. De Paepe et al. [9] claim Tsitsiklis [24] shows a polynomial algorithm, but this is for the open version. We solve the closed variant by giving a polynomial algorithm (Theorem 15) as well as a counterexample to the algorithm of Psaraftis et al. [23]. Our Theorem 16 shows the problem $1, \text{cap}1|d_j|\text{line}|C_{\max}$ in the same classification scheme is NP-hard. While this is implicitly claimed in [9], no proof is given. For the generalization to arbitrary capacity but without release dates, $1||C_{\max}$, Guan [14] showed hardness for capacity $c = 2$ and our new hardness proof handles any capacity $c \geq 2$.

References

[1] Foto Afrati, Stavros Cosmadakis, Christos H. Papadimitriou, George Papageorgiou, and Nadia Papakostantinou. The complexity of the travelling repairman problem. *Informatique Theorique et Applications*, 20(1):79–87, 1986.

[2] Norbert Ascheuer, Sven Oliver Krumke, and Jörg Rambau. Online Dial-a-Ride problems: Minimizing the completion time. *Proceedings of the 17th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 639–650, 2000.

[3] Mikhail J. Atallah and S. Rao Kosaraju. Efficient solutions to some transportation problems with applications to minimizing robot arm travel. *SIAM Journal on Computing*, 17(5):849–869, 1988.

[4] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Serving requests with on-line routing. In *Proceedings of the 4th Scandinavian Workshop on Algorithm Theory Aarhus on Algorithm Theory (SWAT)*, pages 37–48, 1994.

[5] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Competitive algorithms for the on-line traveling salesman. In *Proceedings of the 4th International Workshop on Algorithms and Data Structures (WADS)*, pages 206–217, 1995.

[6] Giorgio Ausiello, Esteban Feuerstein, Stefano Leonardi, Leen Stougie, and Maurizio Talamo. Algorithms for the on-line travelling salesman. *Algorithmica*, 29(4):560–581, 2001.

[7] Michiel Blom, Sven O. Krumke, Willem de Paepe, and Leen Stougie. The online TSP against fair adversaries. *INFORMS Journal on Computing*, 13(2):138–148, 2001.

[8] Moses Charikar and Balaji Raghavachari. The finite capacity dial-a-ride problem. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 458–467, 1998.

[9] Willem E. de Paepe, Jan Karel Lenstra, Jiri Sgall, René A. Sitters, and Leen Stougie. Computer-aided complexity classification of Dial-a-Ride problems. *INFORMS Journal on Computing*, 16(2):120–132, 2004.

[10] Esteban Feuerstein and Leen Stougie. On-line single-server Dial-a-Ride problems. *Theoretical Computer Science*, 268(1):91–105, 2001.

[11] Greg N. Frederickson and Dih Jiun Guan. Nonpreemptive ensemble motion planning on a tree. *Journal of Algorithms*, 15(1):29–60, 1993.

[12] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, USA, 1979.

[13] Paul C. Gilmore and Ralph E. Gomory. Sequencing a one state-variable machine: A solvable case of the traveling salesman problem. *Operations Research*, 12(5):655–679, 1964.

[14] Dih Jiun Guan. Routing a vehicle of capacity greater than one. *Discrete Applied Mathematics*, 81(1-3):41–57, 1998.

[15] Patrick Jaillet and Michael R. Wagner. Generalized online routing: New competitive ratios, resource augmentation, and asymptotic analyses. *Operations Research*, 56(3):745–757, 2008.

[16] Richard M. Karp. Two combinatorial problems associated with external sorting. In *Combinatorial Algorithms, Courant Computer Science Symposium*, pages 17–29, 1972.

[17] Sven O. Krumke. Online optimization competitive analysis and beyond, 2001. Habilitation thesis.

[18] Sven O. Krumke, Willem E. de Paepe, Diana Poensgen, and Leen Stougie. News from the online traveling repairman. *Theoretical Computer Science*, 295(1-3):279–294, 2003.

[19] Sven O. Krumke, Luigi Laura, Maarten Lipmann, Alberto Marchetti-Spaccamela, Willem de Paepe, Diana Poensgen, and Leen Stougie. Non-abusiveness helps: An O(1)-competitive algorithm for minimizing the maximum flow time in the online traveling salesman problem. In *Proceedings of the 5th International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX)*, pages 200–214, 2002.

[20] Eugene L. Lawler, Jan Karel Lenstra, Alexander H. G. Rinnooy Kan, and David B. Shmoys, editors. *The Traveling Salesman Problem; A Guided Tour of Combinatorial Optimization*. Wiley, Chichester, 1985.

[21] Maarten Lipmann. *On-Line Rounting*. PhD thesis, Technical University Eindhoven, 2003.

[22] Mark S. Manasse, Lyle A. McGeoch, and Daniel D. Sleator. Competitive algorithms for server problems. *Journal of Algorithms*, 11(2):208–230, 1990.

[23] Harilaos N. Psaraftis, Marius M. Solomon, Thomas L. Magnanti, and Tai-Up Kim. Routing and scheduling on a shoreline with release times. *Management Science*, 36(2):212–223, 1990.

[24] John N. Tsitsiklis. Special cases of traveling salesman and repairman problems with time windows. *Networks*, 22(3):263–282, 1992.

A Proofs of Section 3

Algorithm 5: $\textsc{Update}(t, \sigma^{\mathrm{L}}(t), \sigma^{\mathrm{R}}(t), p_t)$ for the closed online TSP Problem

this function is called upon release of a new extreme request
Input: time t, unserved extreme requests $\sigma^{\mathrm{R}}(t)$ and $\sigma^{\mathrm{L}}(t)$, position p_t of the server
Output: closed Online TSP tour serving all unserved requests
$A \leftarrow \operatorname{argmax}_{g\in\{a^{\mathrm{R}}(t),a^{\mathrm{L}}(t)\}} |g|$; $a \leftarrow \operatorname{argmin}_{g\in\{a^{\mathrm{R}}(t),a^{\mathrm{L}}(t)\}} |g|$
if $T^{\text{greedy}}(t) = T_{\mathrm{LR}}(t)$ **then**
  $(a_1(t); t_1) \leftarrow \sigma^{\mathrm{L}}(t)$; $(a_2(t); t_2) \leftarrow \sigma^{\mathrm{R}}(t)$
else
  $(a_1(t); t_1) \leftarrow \sigma^{\mathrm{R}}(t)$; $(a_2(t), t_2) \leftarrow \sigma^{\mathrm{L}}(t)$
if $t^{\text{wait}} := \rho|T^{\text{greedy}}(t)| - (|p_t - A| + |A| + 2|a|) \geq t$ **then** // $T^{\text{safe}}(t)$ can be reached
(A) $T^{\textsc{Alg}} \leftarrow \text{waituntil}(t^{\text{wait}}) \oplus \text{move}(A) \oplus \text{move}(a) \oplus \text{move}(p_0)$
else if $\operatorname{sign}(p_t) = \operatorname{sign}(a_1(t))$ **or** $t + |p_t - a_2(t)| + |a_2(t)| < \frac{\rho|a_1(t)|-(2-\rho)t_1}{2\rho-3}$ **then**
(B1,B2) $T^{\textsc{Alg}} \leftarrow \text{move}(a_1(t)) \oplus \text{move}(a_2(t)) \oplus \text{move}(p_0)$
else
(C) $T^{\textsc{Alg}} \leftarrow \text{move}(a_2(t)) \oplus \text{move}(a_1(t)) \oplus \text{move}(p_0)$
return $T^{\textsc{Alg}}$

Lemma 17. *We can assume that the input of the algorithm fulfills the following properties without loss of generality.*

1. *All requests are extreme requests at the time they are released.*
2. $p_t \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ *for all t at which both extreme requests exists. Moreover we can assume $a^{\mathrm{L}}(t) < 0$ if the leftmost extreme exists and $a^{\mathrm{R}}(t) > 0$ if the rightmost extreme exist.*

Proof. We can assume the first property without loss of generality as all requests with a location in the interval $(a^{\mathrm{L}}(t), a^{\mathrm{R}}(t))$ will be served while the algorithm serves the two extreme requests $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$. For the second property note that it is $p_t \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ by definition of extreme requests. Assume, we have $a^{\mathrm{L}}(t) > 0$, then $\sigma^{\mathrm{L}}(t)$ will be served while the server returns to the origin after having served $\sigma^{\mathrm{R}}(t)$. If $a^{\mathrm{R}}(t) < 0$, we can argue similarly that $\sigma^{\mathrm{R}}(t)$ is served while the server returns to the origin after having served $\sigma^{\mathrm{L}}(t)$. ☐

Fact 1. If at time t a ρ-competitive algorithm serves the leftmost extreme $\sigma^{\mathrm{L}}(t)$ first and t_0 denotes the first time the server returns to the origin after having served $\sigma^{\mathrm{L}}(t)$, then

$$t_0 \geq t_0^{\mathrm{L}}(t) := \frac{\rho|a^{\mathrm{R}}(t)| - (2-\rho)t^{\mathrm{R}}(t)}{2\rho - 3}. \tag{1}$$

A symmetric statement with $t_0^{\mathrm{R}}(t)$ holds if the algorithm serves $\sigma^{\mathrm{R}}(t)$ first.

Proof. Recall the situation we constructed in Section 3. Assume at time some time t at which both extreme requests exist, the server greedily decides to serve the extremes in the same order as $T_{\mathrm{LR}}(t)$. Also assume that $T^{\textsc{Opt}}$ initially follows $T_{\mathrm{RL}}(t)$, but continues to move to the left after serving $\sigma^{\mathrm{L}}(t)$. The time when $T^{\textsc{Opt}}$ reaches $a^{\mathrm{L}}(t)$ is $t' = t^{\mathrm{R}}(t) + |a^{\mathrm{R}}(t)| + |a^{\mathrm{L}}(t)|$, since $t^{\mathrm{R}}(t) \geq |a^{\mathrm{R}}(t)|$ by assumption. Let t_0 be the time when we reach the origin p_0 after serving $\sigma^{\mathrm{L}}(t)$, and assume that $t' \leq t_0$. Now a new request $\sigma' = (t_0; p')$ may arrive at time t_0 and position $p' = -|a^{\mathrm{L}}| - (t_0 - t') = -t_0 + t^{\mathrm{R}}(t) + |a^{\mathrm{R}}(t)|$, that the optimum can serve immediately at time t_0. We then have $|T^{\textsc{Opt}}| = t_0 + |p'| = t_0 + |a^{\mathrm{L}}| + (t_0 - t') = 2t_0 - t^{\mathrm{R}}(t) - |a^{\mathrm{R}}(t)|$. On the other hand, our algorithm still needs to serve σ' and σ^{R} at time t_0, and hence $|T^{\textsc{Alg}}| = t_0 + 2|p'| + 2|a^{\mathrm{R}}(t)| = 3t_0 - 2t^{\mathrm{R}}(t)$. For the algorithm to be ρ-competitive, we need $|T^{\textsc{Alg}}|/|T^{\textsc{Opt}}| \leq \rho$. It holds that

$$\frac{|T^{\textsc{Alg}}|}{|T^{\textsc{Opt}}|} = \frac{3t_0 - 2t^{\mathrm{R}}(t)}{2t_0 - t^{\mathrm{R}}(t) - |a^{\mathrm{R}}(t)|} \leq \rho \Leftrightarrow t_0 \geq \frac{\rho|a^{\mathrm{R}}(t)| - (2-\rho)t^{\mathrm{R}}(t)|}{2\rho - 3}.$$ ☐

Fact 2. The safe tour $T^{\text{safe}}(t)$ fulfills inequality (1) in Fact 1 if and only if $2 \geq \rho \geq \frac{9+\sqrt{17}}{8}$.

Proof. Assume that the server is able to follow the safe tour at time t and assume the safe tour serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t)|$. This means that the server is able to reach the tour that waits at the origin till time $\rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{R}}(t)| - 2|a^{\mathrm{L}}(t)|$ and then moves as follows $\text{move}(a^{\mathrm{L}}(t)) \oplus \text{move}(a^{\mathrm{R}}(t))$. Hence, if the server follows the safe tour it returns to the origin after serving $\sigma^{\mathrm{L}}(t)$ at time $t_0 := \rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{R}}(t)|$. It is

$$\begin{aligned} t_0 = \rho T^{\text{greedy}}(t) - 2|a^{\mathrm{R}}(t)| &\geq \rho(2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) - 2|a^{\mathrm{R}}(t)| \\ &\geq (4\rho - 2)|a^{\mathrm{R}}(t)|. \end{aligned}$$

As a necessary criterion for ρ-competitiveness we need that inequality (1) is fulfilled. First note that $\frac{(2\rho-2)|a^{\mathrm{R}}(t)|}{2\rho-3} \geq \frac{\rho|a^{\mathrm{R}}(t)|-(2-\rho)t^{\mathrm{R}}(t)}{2\rho-3}$ if and only if $\rho \leq 2$ because $t^{\mathrm{R}}(t) \geq |a^{\mathrm{R}}(t)|$. Also note

$$(4\rho - 2)|a^{\mathrm{R}}(t)| \geq \frac{(2\rho-2)|a^{\mathrm{R}}(t)|}{2\rho-3} \Leftrightarrow (4\rho^2 - 9\rho + 4) \geq 0.$$

Hence, inequality (1) is fulfilled if and only if $2 \geq \rho \geq \frac{9+\sqrt{17}}{8}$. If the safe tour at time t serves $\sigma^{\mathrm{R}}(t)$ first, the claim can be proved analogously. $\square$

During the following analysis of the algorithm we need the following observation which is a direct consequence from the proof above:

Observation 2. If $\rho = \frac{9+\sqrt{17}}{8}$ we have

$$(4\rho - 2)|a^{\mathrm{R}}(t)| \geq \frac{\rho|a^{\mathrm{R}}(t)| - (2-\rho)t^{\mathrm{R}}(t)}{2\rho-3}.$$

Analysis of Algorithm 1 In the following we will show that Algorithm 1 is ρ-competitive where ρ is the nonnegative root of the polynomial $4x^2 - 9x + 4$. Recall that $T^{\text{greedy}}(t)$ denotes the fastest (offline) tour serving the unserved requests at time t. Also recall, $|T^{\text{greedy}}(t)| = \min\{|T_{\mathrm{LR}}(t)|, |T_{\mathrm{RL}}(t)|\}$ (Observation 1). For the proof we need the formal definition of being able to reach the safe tour $T^{\text{safe}}(t)$.
Recall, that the safe tour was defined as the tour that waits at the origin until time $\rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t)|$ and then without interruption first serves the larger of both extreme and the opposite extreme after that. Intuitively, we are able to reach the safe tour if at time t the position p_t of the server is not too far away from the extreme which the safe tour serves first.

Definition 18. We say that we can *reach* or *recover* $T^{\text{safe}}(t)$ at time t if

$$\rho|T^{\text{greedy}}(t)| \geq \begin{cases} t + |p_t - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|, & \text{if } |a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t)|, \\ t + |p_t - a^{\mathrm{R}}(t)| + |a^{\mathrm{R}}(t)| + 2|a^{\mathrm{L}}(t)|, & \text{otherwise.} \end{cases}$$

We also need the following notation during the proof. Let $A^{\mathrm{L}}(t) = \min_{t' \leq t}\{a^{\mathrm{L}}(t')\}$ and $A^{\mathrm{R}}(t) = \max_{t' \leq t}\{|a^{\mathrm{R}}(t')|\}$. Let $\sigma_1 = (a_1; t_1)$, $\sigma_2 = (a_2; t_2)$ and $\sigma = (a; t)$. We define by $\tau(t, \sigma_1, \sigma_2)$ the point in time at which the tour that starts at time t a position p_t and then moves as follows, $\text{move}(a_1) \oplus \text{move}(a_2) \oplus \text{move}(p_0)$, returns to the origin. Hence, we have

$$\tau(t, \sigma_1, \sigma_2) = t + |a^{\mathrm{L}}(t) - p_t| + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|$$

Similarly $\tau(p_t, \sigma)$ is defined as the point in time at which the tour that starts at time t at position p_t and then moves as follows, $\text{move}(a) \oplus \text{move}(p_0)$ returns to the origin. Hence, it is

$$\tau(t, \sigma) = t + |a^{\mathrm{L}}(t) - p_t| + |a^{\mathrm{L}}(t)|.$$

The following lemmas will be needed during the analysis of the algorithm.

Lemma 19. *If at time $t = t^{\mathrm{L}}(t)$ a new leftmost request $\sigma^{\mathrm{L}}(t)$ is released, $|a^{\mathrm{R}}(t)| > 0$ and $T^{\text{greedy}}(t) = T_{\mathrm{LR}}(t)$, then* Alg *is able to reach $T^{\text{safe}}(t)$. Symmetrically, if at time t a new rightmost request $\sigma^{\mathrm{R}}(t)$ is released, $|a^{\mathrm{L}}(t)| > 0$ and $T^{\text{greedy}}(t) = T_{\mathrm{RL}}(t)$, then* Alg *is able to reach $T^{\text{safe}}(t)$.*

Proof. We only give a proof for the first statement. A symmetric argument proves the second statement. Assume at time t the safe tour $T^{\text{safe}}(t)$ serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| \geq a^{\mathrm{R}}(t)$. In this case to be able to reach $T^{\text{safe}}(t)$ means that $\tau(t, \sigma^{\mathrm{L}}(t), \sigma^{\mathrm{R}}(t)) = t + |p_t - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \leq \rho|T^{\text{greedy}}(t)|$ (Definition 18) or equivalently $\tau(t, \sigma^{\mathrm{L}}(t)) + 2|a^{\mathrm{R}}(t)| \leq \rho|T^{\text{greedy}}(t)|$ By Lemma 17 we have $p_t \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$, hence $|p_t - a^{\mathrm{L}}(t)| \leq |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|$ and by assumption, $t^{\mathrm{L}}(t) \geq |a^{\mathrm{L}}(t)|$. Together with $\rho \geq 1.5$ and $T^{\text{greedy}}(t) = T_{\mathrm{LR}}(t)$ we have

$$\begin{aligned}
\tau(t, \sigma^{\mathrm{L}}(t)) &= t + |p_t - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t)| \\
&\leq t + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\
&= \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) + |a^{\mathrm{L}}(t)| \\
&\quad + (1-\rho)(t + |a^{\mathrm{L}}(t)|) + (1-2\rho)|a^{\mathrm{R}}(t)| \\
&\overset{\rho \geq 1.5}{\leq} \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) + |a^{\mathrm{L}}(t)| - |a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t)| \\
&= \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) - 2|a^{\mathrm{R}}(t)| \\
&\overset{\text{Obs. 1}}{\leq} \rho|T_{\mathrm{LR}}(t)| - 2|a^{\mathrm{R}}(t)| \\
&= \rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{R}}(t)|.
\end{aligned}$$

Thus, Alg can reach the safe tour $T^{\text{safe}}(t)$ at time t. Note that in the second inequality we used that $\rho \geq 1.5$ and $t \geq |a^{\mathrm{L}}(t)|$. We also used that $\rho|T_{\mathrm{LR}}(t)| = \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|)$ (Observation 1).
Assume now that $T^{\text{safe}}(t)$ serves $\sigma^{\mathrm{R}}(t)$ first, i.e., $|a^{\mathrm{R}}(t)| \geq |a^{\mathrm{L}}(t)$—. Using Lemma 17 we can again deduce that $|p_t - a^{\mathrm{R}}(t)| \leq |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \leq 2|a^{\mathrm{R}}(t)|$. Again we use $\rho \geq 1.5$, $T^{\text{greedy}}(T) = T_{\mathrm{LR}}(t)$ and $t \geq |a^{\mathrm{L}}(t)|$ to obtain

$$\begin{aligned}
\tau(t, a^{\mathrm{R}}(t)) &= t + |p_t - a^{\mathrm{R}}(t)| + |a^{\mathrm{R}}(t)| \\
&\leq t + 3|a^{\mathrm{R}}(t)| \\
&= \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) + (3-2\rho)|a^{\mathrm{R}}(t)| + (1-\rho)t - \rho|a^{\mathrm{L}}(t)| \\
&\overset{\rho \geq 1.5}{\leq} \rho(t + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) - 2|a^{\mathrm{L}}(t)| \\
&\overset{\text{Obs.1}}{=} \rho|T_{\mathrm{LR}}(t)| - 2|a^{\mathrm{L}}(t)| \\
&\leq \rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{L}}(t)|.
\end{aligned}$$

Hence, in this case Alg is also able to reach the safe tour at time t. ☐

Lemma 20. *Assume that at time t a new request is released, the safe tour $T^{\text{safe}}(t)$ serves $\sigma^{\mathrm{L}}(t)$ first, $T^{\text{greedy}}(t) = T_{\mathrm{RL}}(t)$ and $p_t \leq 0$. Then* Alg *starts a tour in the direction of $\sigma^{\mathrm{L}}(t)$ at time t. The symmetric statement exchanging $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$, and "leftmost" and "rightmost" also holds.*

Proof. That $T^{\text{safe}}(t)$ serves $\sigma^{\mathrm{L}}(t)$ first implies $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t)|$. If Algorithm 1 can reach the safe tour at time t (Case A), then by definition Alg follows the safe tour $T^{\text{safe}}(t)$ after possibly waiting some time. This implies that in this case Alg starts at tour in the direction of $\sigma^{\mathrm{L}}(t)$ at time t.
Assume that the algorithm is not able to reach $T^{\text{safe}}(t)$ at time t. Recall, that by our assumptions $p_t \leq 0$ and $T^{\text{greedy}}(t) = T_{\mathrm{RL}}(t)$. In the notation of the algorithm we have $a_1(t) = a^{\mathrm{R}}(t)$ and thus $\operatorname{sign}(p_1) \neq \operatorname{sign}(a_1(t))$. This means that the current position of the server is not on the same side of the origin as the extreme which $T^{\text{greedy}}(t)$ serves first. By definition, our algorithm will only start a tour serving $\sigma^{\mathrm{L}}(t)$ first at time t if and only if inequality (1) is not violated when doing so (Case C). We use that Alg is not able to reach the safe tour and Observation 2 to show that in this situation inequality (1) is always fulfilled when the server follows the tour first serving $\sigma^{\mathrm{L}}(t)$. Not being able to reach $T^{\text{safe}}(t)$ means that (we use Definition 18 and

$T^{\text{greedy}}(t) = T_{\text{RL}}(t)$)

$$\begin{aligned}\tau(t, a^{\text{L}}(t)) &= t + |p_t - a^{\text{L}}(t)| + |a^{\text{L}}(t)| \\ &\overset{\text{Def. 18}}{>} \rho|T^{\text{greedy}}(t)| - 2|a^{\text{R}}(t)| \\ &\geq \rho(t^{\text{R}}(t) + |a^{\text{R}}(t)| + 2|a^{\text{L}}(t)|) - 2|a^{\text{R}}(t)| \\ &\geq (4\rho - 2)|a^{\text{R}}(t)|.\end{aligned}$$

In the last inequality we used that $|a^{\text{L}}(t)| \geq |a^{\text{R}}(t)|$ and $t^{\text{R}}(t) \geq |a^{\text{R}}(t)|$ by assumption. By Observation 2 we have

$$\tau(t, a^{\text{L}}(t)) \overset{\text{Obs.2}}{\geq} (4\rho - 2)|a^{\text{R}}(t)| \geq \frac{\rho|a^{\text{R}}(t)| - (2-\rho)t^{\text{R}}(t)}{2\rho - 3}.$$

Hence, immediately serving $\sigma^{\text{L}}(t)$ satisfies inequality (1) in Fact 1 and thus the requirements for Case C are fulfilled because $\text{sign}(p_t) \neq \text{sign}(a_1(t))$ by assumption. Hence, Alg starts a tour in the direction of $\sigma^{\text{L}}(t)$. □

Theorem 3. *There is a $(9 + \sqrt{17})/8 \approx 1.64$-competitive algorithm for closed online* TSP *on the line.*

Proof. Assume that Algorithm 1 derived some tour for an initial sequence of requests. At time t a new extreme request is released and thus an updated tour is computed. We show that the resulting tour is ρ-competitive. During the analysis we assume without loss of generality that the extreme request released at time t is the rightmost request $\sigma^{\text{R}}(t)$. Thus, $t^{\text{R}}(t) = t$. Assume at time t, the server is able to reach the safe tour. By definition of the safe tour (Definition 2), we then have $|T^{\text{Alg}}| = \rho|T^{\text{greedy}}(t)| \leq \rho|\text{Opt}|$ by Observation 1. Thus, in this case the tour derived at time t is ρ-competitive. Hence, we will assume that at time t the server is not able to reach the safe tour.
Assume at first that $|a^{\text{L}}(t)| = 0$. Lower bounds for the offline optimum are $|T^{\text{Opt}}| \geq t + |a^{\text{R}}(t)|$ and $|T^{\text{Opt}}| \geq 2|A^{\text{R}}(t)| + 2|A^{\text{L}}(t)|$. If $|a^{\text{L}}(t)| = 0$ at time t, then, because of $|p_t - a^{\text{R}}(t)| \leq |A^{\text{L}}(t)| + |A^{\text{R}}(t)| \leq T^{\text{Opt}}/2$, we have

$$|T^{\text{Alg}}| = t + |p_t - a^{\text{R}}(t)| + |a^{\text{R}}(t)| \leq 3/2|T^{\text{Opt}}| < \rho|T^{\text{Opt}}|$$

and the tour is ρ-competitive. Thus, we can in the following assume that $|a^{\text{L}}(t)| > 0$.
If at time t the greedy tour serves $\sigma^{\text{R}}(t)$ before $\sigma^{\text{L}}(t)$ and $|a^{\text{L}}(t)| > 0$, then Lemma 19 implies that Alg is able to reach the safe tour at time t, contradicting our assumption. In the following, we will thus assume that the greedy tour at time t serves the leftmost extreme first, i.e. $T^{\text{greedy}}(t) = T_{\text{LR}}(t)$. For the analysis of the algorithm we have to consider the whole timeframe between $t^{\text{L}}(t)$ and the release of $\sigma^{\text{R}}(t)$ at time t. In particular, we have to take the behaviour of the server at time $t^{\text{L}} := t^{\text{L}}(t)$ into account. Note, that, by assumption, no new leftmost extreme is released during $(t^{\text{L}}, t]$.

We now give the remaining proof of the theorem. The proof is divided into the two major cases **(Case 1)** and **(Case 2)** which are again divided into two cases. The main part of the proof can be found in the Lemmas 21 through 35.

(Case 1) *The server does not serve a rightmost extreme during $[t^{\text{L}}, t)$.*

- **(Case 1.1)** *At some time $t^{\text{safe}} \in [t^{\text{L}}, t)$,* Alg *is able to recover the safe tour.*
  If at time t the server is able to reach a safe tour, then the resulting tour is ρ-competitive because by Definition 18 we then have $|T^{\text{Alg}}| = \rho|T^{\text{greedy}}(t)| \leq \rho|T^{\text{Opt}}|$ by Observation 1. In Lemma 23, Lemma 25 and Lemma 27 we show that it is sufficient for the tour derived at time t to be ρ-competitive if at some time during $[t^{\text{L}}, t)$ the server can reach a safe tour.
- **(Case 1.2)** Alg *does* not *reach a safe tour during $[t^{\text{L}}, t]$.*
  This case is treated in Lemma 28

(Case 2) Alg *serves a rightmost extreme during $[t^{\text{L}}, t)$* .
Assume that at time $t' \in [t^{\text{L}}, t)$ the last request before Alg serves the rightmost extreme $\sigma^{\text{R}}(t')$ is released. The request $\sigma^{\text{L}}(t)$ is still unserved at time t. Thus, at time t' the server starts a tour in the direction of $\sigma^{\text{R}}(t')$ and does not turn around before it is served. Denote by t_0 the point in time at which the tour derived

at time t' returns to the origin after having served $\sigma^{\mathrm{R}}(t')$. Lemma 29 provides us with a lower bound for t_0 which holds in all but one special case. This Lemma is an integral part to show ρ-competitiveness in **(Case 2.1)**. In the following we will distinguish between the cases in which the lower bound on t_0 holds in general and the case where it does not.

(Case 2.1) *At time t' it is $T^{\text{greedy}}(t') = T_{\mathrm{LR}}(t')$ or $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ while the $T^{\text{safe}}(t')$ serves $\sigma^{\mathrm{R}}(t')$ first.*
This case is treated in Lemma 29, Lemma 30 and Lemma 32.

(Case 2.2) *At time t' it is $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and $T^{\text{safe}}(t')$ serves $\sigma^{\mathrm{L}}(t')$ first.*
This case is treated in Lemma 33, Lemma 34 and Lemma 35.

□

General lemmas needed for the proof of Theorem 3

Lemma 21. *Assume at some time $\tilde{t}$ a new leftmost request is released, $|a^{\mathrm{R}}(\tilde{t})| > 0$, $|a^{\mathrm{L}}(\tilde{t})| \geq |a^{\mathrm{R}}(\tilde{t})|$ and the server is not able to reach the safe tour at time $\tilde{t}$. Then it holds that*

$$\tilde{t} + |p_{\tilde{t}}| > (2\rho - 2)(|a^{\mathrm{L}}(\tilde{t})| + |a^{\mathrm{R}}(\tilde{t})|)$$

and

$$2|a^{\mathrm{R}}(\tilde{t})| + |a^{\mathrm{L}}(\tilde{t})| - |p_{\tilde{t}}| < (\rho - 1)(\tilde{t} + |a^{\mathrm{L}}(\tilde{t})| + 2|a^{\mathrm{R}}(\tilde{t})|).$$

The symmetric statement exchanging a^{L} and a^{R}, and "leftmost" and "rightmost" also holds.

Proof. Since $|a^{\mathrm{L}}(\tilde{t})| \geq |a^{\mathrm{R}}(\tilde{t})|$ the safe tour at time $\tilde{t}$ serves $\sigma^{\mathrm{L}}(\tilde{t})$ first. Because at time $\tilde{t}$ the server could not reach the safe tour and using $\tilde{t} + |a^{\mathrm{L}}(\tilde{t}) - p_{\tilde{t}}| \leq \tilde{t} + |a^{\mathrm{L}}(\tilde{t})| + |p_{\tilde{t}}|$, we have

$$\begin{aligned} \tilde{t} + |p_{\tilde{t}}| &\overset{\text{Def. 18}}{>} \rho|T^{\text{greedy}}(\tilde{t})| - 2|a^{\mathrm{L}}(\tilde{t})| - 2|a^{\mathrm{R}}(\tilde{t})| \\ &\geq \rho(2|a^{\mathrm{L}}(\tilde{t})| + 2|a^{\mathrm{R}}(\tilde{t})|) - 2|a^{\mathrm{L}}(\tilde{t})| - 2|a^{\mathrm{R}}(\tilde{t})| \\ &= (2\rho - 2)(|a^{\mathrm{L}}(\tilde{t})| + |a^{\mathrm{R}}(\tilde{t})|). \end{aligned} \tag{7}$$

Using this together with $4\rho^2 - 5\rho - 2 \geq 0$ by our choice of ρ and $|a^{\mathrm{L}}(\tilde{t})| \geq |a^{\mathrm{R}}(\tilde{t})|$ we get that

$$\begin{aligned} &2|a^{\mathrm{R}}(\tilde{t})| + |a^{\mathrm{L}}(\tilde{t})| - |p_{\tilde{t}}| \\ &\leq 2|a^{\mathrm{R}}(\tilde{t})| + |a^{\mathrm{L}}(\tilde{t})| - |p_{\tilde{t}}| + (2 - \rho)|p_{\tilde{t}}| + (4\rho^2 - 5\rho - 2)|a^{\mathrm{R}}(\tilde{t})| \\ &= |a^{\mathrm{L}}(\tilde{t})| + (1 - \rho)|p_{\tilde{t}}| + (2\rho^2 - 3\rho)|a^{\mathrm{R}}(\tilde{t})| + (2\rho^2 - 2\rho)|a^{\mathrm{R}}(\tilde{t})| \\ &\leq (2\rho^2 - 3\rho + 1)|a^{\mathrm{L}}(\tilde{t})| + (2\rho^2 - 2\rho)|a^{\mathrm{R}}(\tilde{t})| + (1 - \rho)|p_{\tilde{t}}| \\ &= (\rho - 1)((2\rho - 2)(|a^{\mathrm{L}}(\tilde{t})| + |a^{\mathrm{R}}(\tilde{t})|) - |p_{\tilde{t}}| + |a^{\mathrm{L}}(\tilde{t})| + 2|a^{\mathrm{R}}(\tilde{t})|) \\ &\overset{(7)}{<} (\rho - 1)(\tilde{t} + |a^{\mathrm{L}}(\tilde{t})| + 2|a^{\mathrm{R}}(\tilde{t})|). \end{aligned}$$

□

For the following lemma we also need these additional assumption:

- At time t a new rightmost extreme $\sigma^{\mathrm{R}}(t)$ is released.
- At time t the safe tour cannot be reached.
- $|a^{\mathrm{L}}(t)| > 0$ and $t^{\mathrm{L}} := t^{\mathrm{L}}(t)$.
- $T^{\text{greedy}}(T) = T_{\mathrm{LR}}(t)$.

Lemma 22. *Assume that the above assumptions holds and that during $[t^{\mathrm{L}}, t]$ no safe tour can be reached. If at some time $t' \in [t^{\mathrm{L}}, t)$ the server starts a nonsafe tour in the direction of $\sigma^{\mathrm{L}}(t)$, then the server does not change its direction more than once during $[t', t]$ before having served any extreme. The symmetric statement of the lemma exchanging $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$, and "leftmost" and "rightmost" also holds.*

Proof. At first note that during $(t^{\mathrm{L}}, t]$ only rightmost extremes are released, as $\sigma^{\mathrm{L}}(t) = \sigma^{\mathrm{L}}(t^{\mathrm{L}})$ is still an unserved extreme request at time t. Assume that at some time $t'' \in (t', t]$ a new rightmost request $\sigma^{\mathrm{R}}(t'')$ is released which forces Alg to turn around before having served $\sigma^{\mathrm{L}}(t)$. By assumption no safe tour can be reached during $[t^{\mathrm{L}}, t]$ and we have $|a^{\mathrm{L}}(t)| > 0$. At $t'' \in (t', t]$ a new rightmost extreme is released. Thus, Lemma 19 implies $T^{\text{greedy}}(t'') = T_{\text{LR}}(t'')$. In the notation of the algorithm this means $a_1(t'') = a^{\mathrm{L}}(t)$. Hence, the algorithm (which is not in Case A by assumption) can only start a tour in the direction of $\sigma^{\mathrm{R}}(t'')$ at time t'' if $\operatorname{sign}(p_{t''}) \neq \operatorname{sign}(a^{\mathrm{L}}(t)) = \operatorname{sign}(a_1(t))$, i.e., $p_{t''} \geq 0$ (Case C). Thus, in the following we assume $p_{t''} \geq 0$ We will now show by contradiction that the algorithm cannot be forced to turn around again during $[t', t]$ before having served $\sigma^{\mathrm{R}}(t'')$. Assume that at a time $t''' \in (t'', t]$, before Alg could serve $\sigma^{\mathrm{R}}(t'')$, a new rightmost request $\sigma^{\mathrm{R}}(t''')$ is released that forces Alg to start a tour first serving $\sigma^{\mathrm{L}}(t)$. We have $|a^{\mathrm{R}}(t''')| > |a^{\mathrm{R}}(t'')|$, because $\sigma^{\mathrm{R}}(t'')$ is still unserved when $\sigma^{\mathrm{R}}(t''')$ is released. Again we can deduce by Lemma 19, that $T^{\text{greedy}}(t''') = T_{\text{LR}}(t''')$. We assumed $p_{t''} \geq 0$ and between t'' and t''' the server only moves to the right. Hence, it holds $p_{t'''} \geq 0$. Thus by Lemma 20 the server can only start a tour first serving $\sigma^{\mathrm{L}}(t)$ if $T^{\text{safe}}(t''')$ serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t''')|$. Hence, we can assume $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t''')|$. In the notation of the algorithm we have $a_1(t''') = a^{\mathrm{L}}(t)$ and $a_2 = a^{\mathrm{R}}(t''')$. Thus, it holds that $\operatorname{sign}(p_{t'''}) \neq \operatorname{sign}(a_1(t''')) = \operatorname{sign}(a^{\mathrm{L}}(t))$, i.e. the algorithm is not in Case B1 at time t''' (it is also not in Case A by assumption). Thus the server will start a tour first serving $\sigma^{\mathrm{L}}(t''')$ at time t''' if and only if inequality (1) in Fact 1 is not satisfied when serving $\sigma^{\mathrm{R}}(t) = \sigma_2(t)$ first (Case C). Thus, we additionally assume that at time t' inequality (1) is not satisfied when Alg serves $\sigma^{\mathrm{R}}(t')$ first, i.e.,

$$\begin{aligned} \tau(t', a^{\mathrm{R}}(t')) &= t' + 2|a^{\mathrm{R}}(t')| - |p_{t'}| \\ &< \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}. \end{aligned} \tag{8}$$

At time t'' we have two possible scenarios. First assume $T^{\text{safe}}(t'')$ serves $\sigma^{\mathrm{L}}(t'') = \sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t'')|$. By assumption the server cannot reach the safe tour at time t''. Also recall that $p_{t''} \geq 0$. Using this together with Observation 2 yields

$$\begin{aligned} \tau(t'', a^{\mathrm{R}}(t'')) &= t'' + 2|a^{\mathrm{R}}(t'')| - |p_{t''}| \\ &\overset{\text{Def.18}}{>} \rho|T^{\text{greedy}}(t)| - 2|a^{\mathrm{L}}(t)| \\ &\geq (4\rho - 2)|a^{\mathrm{L}}(t)| \\ &\overset{\text{Obs.2}}{\geq} \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}. \end{aligned} \tag{9}$$

At time t''' the server directly starts to move to left. Hence in this case, we have $|t''' - t''| = |p_{t''} - p_{t'}|$. Because of $p_{t'''} \geq p_{t''} \geq 0$, it also holds $|p_{t''} - p_{t'''}| = |p_{t'''}| - |p_{t''}|$. This implies $t''' + 2|a^{\mathrm{R}}(t'')| - |p_{t'''}| = t'' + 2|a^{\mathrm{R}}(t'')| - |p_{t''}|$. Together with $|a^{\mathrm{R}}(t''')| > |a^{\mathrm{R}}(t'')|$ we have

$$t''' + 2|a^{\mathrm{R}}(t''')| - |p_{t'''}| > t'' + 2|a^{\mathrm{R}}(t'')| - |p_{t''}| \overset{(9)}{\geq} \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3},$$

a contradiction with (8). If the safe tour at time t'' serves $\sigma^{\mathrm{R}}(t'')$ first, we immediately get $|a^{\mathrm{L}}(t)| < |a^{\mathrm{R}}(t'')| < |a^{\mathrm{R}}(t''')|$, a contradiction because we assumed that $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t''')|$. Hence, the algorithm cannot be forced to turn around again after the release of $\sigma^{\mathrm{R}}(t'')$. ☐

Lemmas regarding (Case 1)

Setting A. Consider the following setting of requests of an instance for Alg.

- At time t a new rightmost extreme $\sigma^{\mathrm{R}}(t)$ is released.

- At time t the safe tour cannot be reached.
- $|a^{\mathrm{L}}(t)| > 0$ and $t^{\mathrm{L}} := t^{\mathrm{L}}(t)$
- $T^{\mathrm{greedy}}(t) = T_{\mathrm{LR}}(t)$
- During $[t^{\mathrm{L}}, t]$ no rightmost extreme is served.

Lemma 23. *Assume the the server is in Setting A. If at some time $\hat{t} \in (t^{\mathrm{L}}, t]$ at which a new rightmost request is released the greedy tour serves $\sigma^{\mathrm{L}}(t)$ first, we have that*

$$\rho|T^{\mathrm{greedy}}(\hat{t})| = \rho|T_{\mathrm{LR}}(\hat{t})| > \rho|T^{\mathrm{greedy}}(\tilde{t})| + 2|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})|$$

for all $\tilde{t} \in [t^{\mathrm{L}}, \hat{t})$.

Proof. Recall that

$$|T_{\mathrm{LR}}(\hat{t})| \overset{\text{Obs.1}}{=} t^{\mathrm{L}}(\hat{t}) + |a^{\mathrm{L}}(\hat{t})| + 2|a^{\mathrm{R}}(\hat{t})|,$$
$$|T_{\mathrm{LR}}(\tilde{t})| \overset{\text{Obs.1}}{=} t^{\mathrm{L}}(\tilde{t}) + |a^{\mathrm{L}}(\tilde{t})| + 2|a^{\mathrm{R}}(\tilde{t})|.$$

Also, note that during $(t^{\mathrm{L}}, t]$ no new leftmost extreme is released and hence $t^{\mathrm{L}}(\hat{t}) = t^{\mathrm{L}}(\tilde{t}) = t^{\mathrm{L}}$ and $a^{\mathrm{L}}(\hat{t}) = a^{\mathrm{L}}(\tilde{t})$. Thus,

$$|T_{\mathrm{LR}}(\hat{t})| = |T_{\mathrm{LR}}(\tilde{t})| + 2|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})|.$$

This implies $\rho|T_{\mathrm{LR}}(\hat{t})| = \rho|T_{\mathrm{LR}}(\tilde{t})| + 2\rho|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})|$ and since $|T^{\mathrm{greedy}}(\tilde{t})| \leq |T_{\mathrm{LR}}(\tilde{t})|$, it holds that

$$\begin{aligned}
\rho|T^{\mathrm{greedy}}(\hat{t})| &= \rho|T_{\mathrm{LR}}(\hat{t})| \\
&= \rho|T_{\mathrm{LR}}(\tilde{t})| + 2\rho|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})| \\
&\geq \rho|T^{\mathrm{greedy}}(\tilde{t})| + 2\rho|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})| \\
&> \rho|T^{\mathrm{greedy}}(\tilde{t})| + 2|a^{\mathrm{R}}(\hat{t}) - a^{\mathrm{R}}(\tilde{t})|.
\end{aligned}$$

□

Note that the following symmetric version of Lemma 23 also holds.

Corollary 24. *Assume the following holds:*

- *At time t^{L} a new leftmost extreme $\sigma^{\mathrm{L}}(t)$ is released.*
- *At time t^{L} the safe tour cannot be reached.*
- $|a^{\mathrm{R}}(t^{\mathrm{L}})| > 0$ *and* $t^{\mathrm{R}} := t^{\mathrm{R}}(t^{\mathrm{L}})$
- $T^{\mathrm{greedy}}(t^{\mathrm{L}}) = T_{\mathrm{RL}}(t^{\mathrm{L}})$
- *During $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ no leftmost extreme is served.*

If at some time $\hat{t} \in (t^{\mathrm{R}}, t^{\mathrm{L}}]$ at which a new leftmost extreme is released the greedy tour serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first, we have that

$$\rho|T^{\mathrm{greedy}}(\hat{t})| = \rho|T_{\mathrm{RL}}(\hat{t})| > \rho|T^{\mathrm{greedy}}(\tilde{t})| + 2|a^{\mathrm{L}}(\hat{t}) - a^{\mathrm{L}}(\tilde{t})|$$

for all $\tilde{t} \in [t^{\mathrm{R}}, \hat{t})$.

Lemma 25. *Assume the algorithm is in Setting A and that the server can reach a safe tour at some time during $[t^{\mathrm{L}}, t)$. Let $t^{\mathrm{safe}} \in [t^{\mathrm{L}}, t)$ be the* last *time in this interval at which the server can recover the save tour. Then it holds that $p_t \leq 0$, the safe tour at time t^{safe} serves $\sigma^{\mathrm{L}}(t)$ first and the server does does not turn around after time t^{safe} before $\sigma^{\mathrm{L}}(t)$ is served.*

Proof. We will show that all cases except for the one in the claim lead to a contradiction. By assumption, the server does not serve any rightmost extremes during $[t^{\mathrm{L}}, t)$. Also, $\sigma^{\mathrm{L}}(t)$ is not served before time t. This means that a new rightmost request $\sigma^{\mathrm{R}}(t')$ with $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{R}}(t^{\mathrm{safe}})|$ has to be released at some time $t' \in (t^{\mathrm{safe}}, t]$ before ALG can serve any extreme. Since $t' > t^{\mathrm{safe}}$ and by definition of t^{safe}, the release of $\sigma^{\mathrm{R}}(t')$ has the effect, that ALG will now follow a nonsafe tour. Since at time t' a new rightmost extreme is released and $|a^{\mathrm{L}}(t)| > 0$ we can as a consequence of Lemma 19 assume that $T^{\mathrm{greedy}}(t') = T_{\mathrm{LR}}(t')$, because otherwise the server would be able to reach the safe tour at time $t' > t^{\mathrm{safe}}$, contradicting our assumption. In fact using Lemma 19 we can assume the during the whole interval $(t^{\mathrm{safe}}, t]$ we have $T^{\mathrm{greedy}} = T_{\mathrm{LR}}$ at times when a new request is released. Thus, with $\hat{t} = t'$ and $\tilde{t} = t^{\mathrm{safe}}$ the assumption for Lemma 23 are fulfilled which implies that the following holds:

$$\rho|T^{\mathrm{greedy}}(t')| = \rho|T_{\mathrm{LR}}(t')| > \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t') - a^{\mathrm{R}}(t^{\mathrm{safe}})|. \tag{10}$$

Assume, that $T^{\mathrm{safe}}(t^{\mathrm{safe}})$ serves the rightmost request first, i.e., $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{R}}(t^{\mathrm{safe}})| \geq |a^{\mathrm{L}}(t^{\mathrm{safe}})| = |a^{\mathrm{L}}(t)|$. The server can be in three situations at time t'. Either ALG is already moving along the safe tour towards $\sigma^{\mathrm{R}}(t^{\mathrm{safe}})$, it is still waiting or it is moving towards the safe tour to catch up with $T^{\mathrm{safe}}(t^{\mathrm{safe}})$. Recall, that by assumption $\sigma^{\mathrm{R}}(t^{\mathrm{safe}})$ is still unserved at time t'. In in every of the three situations, we still have by Definition 18

$$\begin{aligned}
&\tau(t', a^{\mathrm{R}}(t^{\mathrm{safe}}), a^{\mathrm{L}}(t')) \\
&\quad = t' + |p_{t'} - a^{\mathrm{R}}(t^{\mathrm{safe}})| + |a^{\mathrm{R}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{L}}(t^{\mathrm{safe}})| \\
&\quad \leq \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})|.
\end{aligned} \tag{11}$$

We will now aim for a contradiction and show that under this conditions at time $t' > t^{\mathrm{safe}}$ the server is again able to reach the safe tour. By using (10), (11) and $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{R}}(t^{\mathrm{safe}})|$ we get that

$$\begin{aligned}
&\tau(t', a^{\mathrm{L}}(t), a^{\mathrm{R}}(t')) \\
&\quad = \; t' + |p_{t'} - a^{\mathrm{R}}(t')| + |a^{\mathrm{R}}(t')| + 2|a^{\mathrm{L}}(t)| \\
&\quad = \; t' + |p_{t'} - |a^{\mathrm{R}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t') - |a^{\mathrm{R}}(t^{\mathrm{safe}})| + |a^{\mathrm{R}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{L}}(t)| \\
&\quad \overset{(11)}{\leq} \; \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t') - |a^{\mathrm{R}}(t^{\mathrm{safe}})| \\
&\quad \overset{(10)}{<} \; \rho|T^{\mathrm{greedy}}(t')|.
\end{aligned}$$

Thus, by Definition 18 the server can reach the safe tour at time $t' > t^{\mathrm{safe}}$, a contradiction.
Assume now that $T^{\mathrm{safe}}(t^{\mathrm{safe}})$ serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| = |a^{\mathrm{L}}(t^{\mathrm{safe}})| \geq |a^{\mathrm{R}}(t^{\mathrm{safe}})|$. If also $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t')|$ holds, then $T^{\mathrm{safe}}(t')$ also serves $\sigma^{\mathrm{L}}(t)$ first. By the same argumentation as above using (10) we get that

$$\begin{aligned}
&\tau(t', a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)) \\
&\quad = \; t' + |p_{t'} - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t')| \\
&\quad = \; t' + |p_{t'} - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t') - a^{\mathrm{R}}(t^{\mathrm{safe}})| \\
&\quad \overset{(10)}{\leq} \; \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t') - a^{\mathrm{R}}(t^{\mathrm{safe}})| \\
&\quad < \; \rho|T^{\mathrm{greedy}}(t')|.
\end{aligned}$$

Hence, ALG is again able to reach the safe tour at time t', a contradiction.
We can therefore assume that $|a^{\mathrm{L}}(t)| < |a^{\mathrm{R}}(t')|$. To analyze this case, we have to take the position of the server at time t' into account. Recall that the greedy tor at time t' serves $\sigma^{\mathrm{L}}(t')$ first. Thus, in the notation of the algorithm it is $a_1(t') = a^{\mathrm{L}}(t')$. If $p_{t'} < 0$, i.e., $\mathrm{sign}(a_1(t')) = \mathrm{sign}(p_{t'})$, the algorithm is in Case B1 at time t'. This means the server immediately starts a tour first serving $\sigma^{\mathrm{L}}(t)$ at time t'. Whenever at some time $t'' \in (t', t]$ a new request is released, we have $T^{\mathrm{greedy}}(t'') = T_{\mathrm{LR}}(t'')$ and $p_{t''} \leq 0$. Thus, at time t'' it holds $\mathrm{sign}(a_1(t')) = \mathrm{sign}(a^{\mathrm{L}}(t)) = \mathrm{sign}(p_{t''})$. Hence, whenever during $(t', t]$ a new request is released, the algorithm is in (Case B1) and start a tour in the direction of $\sigma^{\mathrm{L}}(t)$. Thus, the algorithm does not turn around before having served $\sigma^{\mathrm{L}}(t)$ and it also holds $p_t \leq p_{t'} \leq 0$. Thus, the assertion of the Lemma is fulfilled in this case.

If $p_{t'} \geq 0$, then since at time t^{safe} the safe tour serves $\sigma^{\text{L}}(t)$ first and $T^{\text{greedy}}(t^{\text{safe}}) = T_{\text{LR}}(t^{\text{safe}})$, again as above using (10)

$$\begin{aligned}
&\tau(t', a^{\text{L}}(t), a^{\text{R}}(t)) \\
&\quad = t' + |p_{t'}| + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t')| \\
&\quad = t' + |p_{t'}| + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t^{\text{safe}})| + 2|a^{\text{R}}(t') - a^{\text{R}}(t^{\text{safe}})| \\
&\quad \overset{(11)}{\leq} \rho|T^{\text{greedy}}(t^{\text{safe}})| + 2|a^{\text{R}}(t') - a^{\text{R}}(t^{\text{safe}})| \\
&\quad \overset{(10)}{<} \rho|T^{\text{greedy}}(t')|.
\end{aligned}$$

Thus, again ALG can recover the safe tour at time $t' > t^{\text{safe}}$, a contradiction. ☐

Note that with an analogous proof as above the following symmetric version of Lemma 25 also holds.

Corollary 26. *Assume the following holds:*

- *At time t^{L} a new leftmost extreme $\sigma^{\text{L}}(t^{\text{L}})$ is released.*
- *At time t^{L} the safe tour cannot be reached.*
- $|a^{\text{R}}(t^{\text{L}})| > 0$ *and* $t^{\text{R}} := t^{\text{R}}(t^{\text{L}})$
- $T^{\text{greedy}}(t^{\text{L}}) = T_{\text{RL}}(t^{\text{L}})$
- *During $[t^{\text{R}}, t^{\text{L}}]$ no leftmost extreme is served.*

also assume that the server can reach a safe tour at some time during $[t^{\text{R}}, t^{\text{L}})$. Let $t^{\text{safe}} \in [t^{\text{R}}, t^{\text{L}})$ be the last *time in this interval at which the server can recover the save tour. Then it holds that $p_{t^{\text{L}}} \geq 0$, the safe tour at time t^{safe} serves $\sigma^{\text{R}}(t^{\text{L}})$ first and the server does does not turn around after time t^{safe} before $\sigma^{\text{R}}(t^{\text{L}})$ is served.*

Lemma 27. *Assume the algorithm is in Setting A. If the server can reach the safe tour at some time during $[t^{\text{L}}, t)$, the tour derived at time t is ρ-competitive.*

Proof. Let $t^{\text{safe}} \in [t^{\text{L}}, t)$ be the *last* time between t^{L} and t at which the server can recover the safe tour, i.e., at time t^{safe} an extreme is released such that the algorithm is in Case A. By assumption, the server does not serve any rightmost extremes during $[t^{\text{L}}, t)$. Note that the requirements for Lemma 25 are fulfilled. Thus, it holds that $p_t \leq 0$, the safe tour at time t^{safe} serves $\sigma^{\text{L}}(t)$ first and the server does not turn around after time t^{safe} before $\sigma^{\text{L}}(t)$ is served. Using this and the fact that at time t the request $\sigma^{\text{L}}(t)$ is still unserved, we get

$$|T^{\text{ALG}}| \leq t + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| - |p_t|. \tag{12}$$

Using the fact that at time t^{safe} the safe tour first serving $\sigma^{\text{L}}(t)$ can be reached and that the server only moves to the left during $[t^{\text{safe}}, t]$ we get with Definition 18

$$t \leq \rho|T^{\text{greedy}}(t^{\text{safe}})| - 2|a^{\text{L}}(t^{\text{safe}})| - 2|a^{\text{R}}(t^{\text{safe}})| + |p_t|. \tag{13}$$

Also note that with $\hat{t} = t$ and $\tilde{t} = t^{\text{safe}}$ the requirements for Lemma 23 are fulfilled and it holds that

$$\rho|T^{\text{greedy}}(t)| = \rho|T_{\text{LR}}(t)| > \rho|T^{\text{greedy}}(t^{\text{safe}})| + 2|a^{\text{R}}(t) - a^{\text{R}}(t^{\text{safe}})|. \tag{14}$$

Thus, all in all we get

$$\begin{aligned}
|T^{\text{ALG}}| &\overset{(12)}{\leq} t + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| - |p_t| \\
&\overset{(13)}{\leq} \rho|T^{\text{greedy}}(t^{\text{safe}})| - 2|a^{\text{L}}(t)| - 2|a^{\text{R}}(t^{\text{safe}})| + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \\
&= \rho|T^{\text{greedy}}(t^{\text{safe}})| + 2|a^{\text{R}}(t) - a^{\text{R}}(t^{\text{safe}})| \\
&\overset{(14)}{<} \rho|T^{\text{greedy}}(t)| \leq \rho|T^{\text{OPT}}|.
\end{aligned}$$

☐

Observation 3. The statement of Lemma 27 is still true is we relax our assumptions a bit only assume that during $[t^{\text{safe}}, t]$ no rightmost extreme is served instead of assuming that this is true for $[t^{\text{L}}, t]$.

Lemma 28. *Assume the algorithm is in Setting A. If the server does **not** reach a safe tour during $[t^{\text{L}}, t]$, the tour derived at time t is ρ-competitive.*

Proof. To analyze this case we have to take the behaviour of ALG at time t^{L} into account. If we assume that at time $t'' \in [t^{\text{L}}, t]$ a request is released, the released request has to be a rightmost extreme, because $\sigma^{\text{L}}(t^{\text{L}})$ is still unserved at time t. For every $t'' \in (t^{\text{L}}, t]$ we have $|a^{\text{L}}(t'')| = |a^{\text{L}}(t)| > 0$ by assumption. Using Lemma 19 we can then assume that at time t'' we have it is $T^{\text{greedy}}(t'') = T_{\text{LR}}(t'')$ as otherwise ALG would be able to reach the safe tour at time t''.

(Case 1) *At time t^{L} the algorithm starts a (nonsafe) tour in the direction of $\sigma^{\text{L}}(t)$.*

(Case 1.1) *The server is **not** forced to turn around before serving $\sigma^{\text{L}}(t)$*

Using $|T^{\text{OPT}}| \geq t^{\text{L}} + |a^{\text{L}}(t)| + 2|a^{\text{R}}(t)|$ and $|T^{\text{OPT}}|/2 \geq |A^{\text{L}}(t)| + |A^{\text{R}}(t)| \geq |p_{t^{\text{L}}} - a^{\text{L}}(t)|$, we get

$$|T^{\text{ALG}}| = t^{\text{L}} + |p_{t^{\text{L}}} - a^{\text{L}}(t)| + |a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \leq 3/2|T^{\text{OPT}}|,$$

hence ALG is ρ-competitive.

(Case 1.2) *The server is forced to turn around before having served $\sigma^{\text{L}}(t)$*

Assume now that at some time $t' \in (t^{\text{L}}, t]$ a new rightmost request $\sigma^{\text{R}}(t')$ is released which forces ALG to turn around before having served $\sigma^{\text{L}}(t)$. We have $T^{\text{greedy}}(t') = T_{\text{LR}}(t')$. In the notation of the algorithm this means $a_1(t') = a^{\text{L}}(t)$. As the safe tour cannot be reached by assumption, the algorithm can only start a tour in the direction of $\sigma^{\text{R}}(t^{\text{newR}})$ at time t' if $\text{sign}(p_{t'}) \neq \text{sign}(a_1(t'))$, i.e., $p_{t'} \geq 0$ (Case C). Lemma 22 implies that the server cannot be forced to turn around again during $[t^{\text{L}}, t]$ before having served an extreme. Hence, using $p_{t^{\text{L}}} > p_{t'} \geq 0$ and the fact that $\sigma^{\text{R}}(t)$ must be released before $\sigma^{\text{R}}(t^{\text{newR}})$ is served by assumption, we get that

$$|T^{\text{ALG}}| \leq t^{\text{L}} + |p_{t^{\text{L}}}| + 2|a^{\text{R}}(t)| + 2|a^{\text{L}}(t)|.$$

Again using $t^{\text{L}} + |a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \leq |T^{\text{OPT}}|$ and $|p_{t^{\text{L}}}| + |a^{\text{L}}(t)| \leq |A^{\text{R}}(t)| + |A^{\text{L}}(t)| \leq |T^{\text{OPT}}|/2$, we get

$$|T^{\text{ALG}}| \leq 3|T^{\text{OPT}}|/2,$$

hence ALG is ρ-competitive.

(Case 2) *At time t^{L} the server starts a (nonsafe) tour in the direction of $\sigma^{\text{R}}(t^{\text{L}})$.* Note, that we only need to regard this case, if $|a^{\text{R}}(t^{\text{L}})| > 0$. Then we can apply Lemma 19 and deduce that $T^{\text{greedy}}(t^{\text{L}}) = T_{\text{RL}}(t^{\text{L}})$, because otherwise the server would be able to reach the safe tour at time t^{L}.

(Case 2.1) *The server is **not** forced to turn around before serving $\sigma^{\text{R}}(t^{\text{L}})$*

By assumption $\sigma^{\text{R}}(t)$ has to be released before $\sigma^{\text{R}}(t^{\text{L}})$ is served. Thus using $|T^{\text{OPT}}| \leq t^{\text{L}} + |a^{\text{L}}(t)| + 2|a^{\text{R}}(t)|$

$$\begin{aligned}|T^{\text{ALG}}| &= \tau(t^{\text{L}}, a^{\text{R}}(t), a^{\text{L}}(t))| \\ &\leq t^{\text{L}} + 2|a^{\text{R}}(t)| + |p_{t^{\text{L}}}| + 2|a^{\text{L}}(t)| \\ &\leq |T^{\text{OPT}}| + |a^{\text{L}}(t)| + |p_{t^{\text{L}}}|.\end{aligned}$$

If $|p_{t^{\text{L}}}| \leq |a^{\text{R}}(t)|$, we are done because then $|a^{\text{L}}(t)| + |p_{t^{\text{L}}}| \leq |A^{\text{L}}(t)| + |A^{\text{R}}(t)| \leq |T^{\text{OPT}}|/2$ by using $p_t \in [a^{\text{L}}(t), a^{\text{R}}(t)]$ (Lemma 17). If $p_{t^{\text{L}}} \geq 0$, then $|p_{t^{\text{L}}}| \leq |a^{\text{R}}(t)|$ follows again by Lemma 17. Assume $p_{t^{\text{L}}} < 0$. We have $T^{\text{greedy}}(t^{\text{L}}) = T_{\text{RL}}(t^{\text{L}})$. In the notation of the algorithm this means that $a_1(t^{\text{L}}) = a^{\text{R}}(t^{\text{L}})$ at time t^{L} and $\text{sign}(p_{t^{\text{L}}}) \neq \text{sign}(a^{\text{R}}(t^{\text{L}}))$. By assumption, the safe tour cannot be reached at time t^{L}. Thus, ALG can only start a tour in the direction $\sigma^{\text{R}}(t^{\text{L}})$ if serving $\sigma^{\text{L}}(t)$ first implies that inequality (1) in Fact 1 is not satisfied (Case C) which by Lemma 20 can only happen if $T^{\text{safe}}(t^{\text{L}})$ serves $\sigma^{\text{R}}(t^{\text{L}})$ first, i.e., $|a^{\text{R}}(t^{\text{L}})| > |a^{\text{L}}(t)| \geq |p_{t^{\text{L}}}|$. As the server does not serve a rightmost extreme during $[t^{\text{L}}, t)$ we get $|a^{\text{R}}(t)| \geq |a^{\text{R}}(t^{\text{L}})| \geq |p_{t^{\text{L}}}|$. Thus, we can apply the same reasoning as before to deduce that $|a^{\text{L}}(t)| + |p_{t^{\text{L}}}| \leq |T^{\text{OPT}}|/2$.

(Case 2.2) *The server is forced to turn around before having served* $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$

We will in the following show that this case cannot occur. Assume that the server turns around before having served $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. This can only happen if at some point in time $t' \in (t^{\mathrm{L}}, t]$ before ALG could serve $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ a new rightmost request with $|a^{\mathrm{R}}(t')| > |a^{\mathrm{R}}(t^{\mathrm{L}})|$ is released. Recall, that as a consequence of Lemma 19 we have $T^{\text{greedy}}(t') = T_{\mathrm{LR}}(t')$ as otherwise we would be able to reach the safe tour at time t' which contradicts our assumptions. Assume the safe tour at time t' serves $\sigma^{\mathrm{R}}(t')$ first, i.e., $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{L}}(t)|$. We then have

$$\begin{aligned}\rho|T_{\mathrm{LR}}(t')| &\geq \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t')|) \\ &\geq t^{\mathrm{L}} + (\rho-1)|a^{\mathrm{L}}(t)| + \rho|a^{\mathrm{L}}(t)| + 2\rho|a^{\mathrm{R}}(t')| \\ &= t^{\mathrm{L}} + 2\rho(|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t')|) - |a^{\mathrm{L}}(t)|.\end{aligned}$$

In the second inequality we used that $t^{\mathrm{L}} \geq |a^{\mathrm{L}}(t)|$ by assumption. Using $c > 1.5$, $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{L}}(t)|$ and $p_{t^{\mathrm{L}}} \in [a^{\mathrm{L}}(t^{\mathrm{L}}), a^{\mathrm{R}}(t^{\mathrm{L}})]$, this implies that

$$\begin{aligned}&\rho|T^{\text{greedy}}(t')| - 2|a^{\mathrm{L}}(t)| - |a^{\mathrm{R}}(t')| \\ &\quad= \rho|T_{\mathrm{LR}}(t')| - 2|a^{\mathrm{L}}(t)| - |a^{\mathrm{R}}(t')| \\ &\quad\geq t^{\mathrm{L}} + (2\rho-3)(|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t')|) + 2|a^{\mathrm{R}}(t')| \\ &\quad\geq t^{\mathrm{L}} + (2\rho-3)(|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t')|) + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t')| \\ &\quad> t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t')| \\ &\quad\geq t^{\mathrm{L}} + |p_{t^{\mathrm{L}}} - a^{\mathrm{R}}(t')| \\ &\quad= t' + |p_{t'} - a^{\mathrm{R}}(t')|.\end{aligned}$$

Hence, ALG can recover a safe tour at time t', contradiction our assumption that no safe tour can be reached during $[t^{\mathrm{L}}, t]$.

Now assume that at time t' the safe tour serves $\sigma^{\mathrm{L}}(t)$ first, then $|a^{\mathrm{R}}(t^{\mathrm{L}})| < |a^{\mathrm{R}}(t')| < |a^{\mathrm{L}}(t)|$. Recall, that at time t^{L} the greedy tour serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first, i.e., $T^{\text{greedy}}(t^{\mathrm{L}}) = T_{\mathrm{RL}}(t^{\mathrm{L}})$. If $p_{t^{\mathrm{L}}} < 0$ we get in the notation of the algorithm $\operatorname{sign}(p_{t^{\mathrm{L}}}) \neq \operatorname{sign}(a_1(t^{\mathrm{L}})) = \operatorname{sign}(a^{\mathrm{R}}(t^{\mathrm{L}}))$. Thus, as the safe tour cannot be reached at time t^{L}, the server will only start a tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ at time t^{L} if serving $\sigma^{\mathrm{L}}(t)$ first results in inequality (1) from Fact 1 not being satisfied. As a consequence of Lemma 20 this can only happen if the safe tour at time t^{L} served $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first, a contradiction.

Hence, we can assume that $p_{t^{\mathrm{L}}} \geq 0$. By assumption the server is not able to reach the safe tour at time t^{L} and it is $|a^{\mathrm{L}}(t)| > |a^{\mathrm{R}}(t^{\mathrm{L}})|$. Thus at time t^{L}, all conditions for Lemma 21 are fulfilled and we get that

$$(\rho-1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|) > 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}|. \tag{15}$$

At time t' the server turns around and can reach the origin before time $t^{\mathrm{R}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}|$, because at time t' the request $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is still unserved. Thus $t^{\mathrm{L}} + |p_{t^{\mathrm{L}}}| \leq t^{\mathrm{R}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}|$ Using this, we get that

$$\begin{aligned}&\rho|T^{\text{greedy}}(t')| - 2|a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t')| \\ &\quad= \rho|T_{\mathrm{LR}}(t')| - 2|a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t')| \\ &\quad= \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t')|) - 2|a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t')| \\ &\quad\overset{(15)}{>} (2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}|) + t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| - 2|a^{\mathrm{L}}(t)| - 2|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &\quad= t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| \\ &\quad\geq t' + |p_{t'}|\end{aligned}$$

Thus, ALG can again reach the safe tour at time t', contradicting our assumption.

□

Lemmas regarding (Case 2)

Setting B. Consider the following setting of requests of an instance for ALG:

- At time t a new rightmost extreme $\sigma^{\mathrm{R}}(t)$ is released.
- At time t the safe tour cannot be reached.
- We have $|a^{\mathrm{L}}(t)| > 0$ and $t^{\mathrm{L}} := t^{\mathrm{L}}(t)$.
- A rightmost extreme is served during $[t^{\mathrm{L}}, t)$.
- The last extreme before the server serves the rightmost extreme is released at time $t' \in [t^{\mathrm{L}}, t)$. In particular $|a^{\mathrm{R}}(t')| > 0$.
- t_0 is the time when the tour derived at time t' returns to the origin.

Lemma 29. *Assume the algorithm is in Setting B. Then, if at time t' we have $T^{\text{greedy}}(t') = T_{\mathrm{LR}}(t')$, or $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and the safe tour $T^{\text{safe}}(t')$ serves $\sigma^{\mathrm{R}}(t')$ first, it holds that*

$$t_0 \geq \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}.$$

Proof. At time t' the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t')$ and does not turn around before it is served. Assume at time t' we have $T^{\text{greedy}}(t') = T_{\mathrm{LR}}(t')$. In the notation of the algorithm this means $a_1(t') = a^{\mathrm{L}}(t')$. In this setting the server starts a tour in the direction of the rightmost extreme if $T^{\text{safe}}(t')$ serves the rightmost extreme first and the server can reach the safe tour (Case A) or if $\operatorname{sign}(p_{t'}) \neq \operatorname{sign}(a_1(t')) = \operatorname{sign}(a^{\mathrm{L}}(t'))$ and inequality (1) in Fact 1 is fulfilled when serving the rightmost extreme first (Case B2). By Fact 2, the safe tour fulfills inequality (1). Hence, if Case A occurs at time, the following inequality holds and in Case B2 it holds by definition,

$$t_0 \geq \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}. \tag{16}$$

Next assume that at time t' we have $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and that $T^{\text{safe}}(t')$ also serves $\sigma^{\mathrm{R}}(t')$ first, i.e., $|a^{\mathrm{R}}(t')| \geq |a^{\mathrm{L}}(t)|$. Thus, when the online server also serves $\sigma^{\mathrm{R}}(t')$ first, it is either on the safe tour (Case A) or it follows the safe tour without being able to catch up with it (Case B1 and Case C). This implies together with our assumptions on ρ and Observation 2,

$$t_0 \overset{\text{Def.18}}{\geq} \rho|T^{\text{greedy}}(t')| - 2|a^{\mathrm{L}}(t)| \geq (4\rho - 2)|a^{\mathrm{L}}(t)| \overset{\text{Obs.2}}{\geq} \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3},$$

hence (16) also holds in this case.

□

Lemma 30. *Assume the algorithm is in Setting B. Then, if at time t' we have $T^{\text{greedy}}(t') = T_{\mathrm{LR}}(t')$, or $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and the safe tour $T^{\text{safe}}(t')$ serves $\sigma^{\mathrm{R}}(t')$ first, it holds that*

$$2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_t| \leq (\rho - 1)|T^{\text{OPT}}|.$$

Proof. By Lemma 32 we know that

$$t_0 \geq \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}. \tag{17}$$

Assume at first that

$$|a^{\mathrm{R}}(t)| \leq \frac{3-\rho}{2\rho - 3}|a^{\mathrm{L}}(t)| + \frac{1-\rho}{2\rho - 3}t^{\mathrm{L}} + |p_t|. \tag{18}$$

Then using $1.640 \leq \rho \leq 2$ and inequality (17) we get

$$\begin{aligned}
&2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_t| \\
&\overset{(18)}{\leq} (2-\rho)\left(\frac{3-\rho}{2\rho-3}|a^{\mathrm{L}}(t)| + \frac{1-\rho}{2\rho-3}t^{\mathrm{L}} + |p_t| - |a^{\mathrm{R}}(t)|\right) \\
&\quad + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_t| \\
&= (\rho-1)\left(\frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho-3} - |p_t| + |a^{\mathrm{R}}(t)|\right) \\
&\overset{(17)}{\leq} (\rho-1)(t_0 - |p_t| + |a^{\mathrm{R}}(t)|) \\
&\leq (\rho-1)(t + |a^{\mathrm{R}}(t)|) \\
&\leq (\rho-1)|T^{\mathrm{OPT}}|.
\end{aligned}$$

If we assume

$$|a^{\mathrm{R}}(t)| > \frac{3-\rho}{2\rho-3}|a^{\mathrm{L}}(t)| + \frac{1-\rho}{2\rho-3}t^{\mathrm{L}} + |p_t|, \tag{19}$$

we obtain

$$\begin{aligned}
&2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_t| \\
&= ((3-\rho)|a^{\mathrm{L}}(t)| + (1-\rho)t^{\mathrm{L}}) - (2\rho-3)|a^{\mathrm{R}}(t)| - |p_t| \\
&\quad + (\rho-1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) \\
&\overset{(19)}{<} -(2\rho-3)|p_t| - |p_t| + (\rho-1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) \\
&\leq (\rho-1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) \\
&\leq (\rho-1)|T^{\mathrm{OPT}}|.
\end{aligned}$$

□

Note, that with an analogous proof as above the following symmetric statement of Lemma 30 also holds:

Corollary 31. *Assume the following assumptions hold:*

- *A time t^{L} a new leftmost extreme $\sigma^{\mathrm{L}}(t^{\mathrm{L}})$ is released.*
- *At time t^{L} the safe tour cannot be reached.*
- *We have $|a^{\mathrm{R}}(t^{\mathrm{L}})| > 0$ and $t^{\mathrm{R}} := t^{\mathrm{R}}(t^{\mathrm{L}})$.*
- *A leftmost extreme is served during $[t^{\mathrm{R}}, t^{\mathrm{L}})$*
- *The last extreme before the server serves the leftmost extreme is released at time $t' \in [t^{\mathrm{R}}, t^{\mathrm{L}})$. In particular $|a^{\mathrm{L}}(t')| > 0$.*

If at time t' we have $T^{\mathrm{greedy}}(t') = T_{\mathrm{RL}}(t')$, or $T^{\mathrm{greedy}}(t') = T_{\mathrm{LR}}(t')$ and the safe tour $T^{\mathrm{safe}}(t')$ serves $\sigma^{\mathrm{L}}(t')$ first, it holds that

$$2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| \leq (\rho-1)|T^{\mathrm{OPT}}|.$$

Lemma 32. *Assume the algorithm is in Setting B. Then, if at time t' we have $T^{\mathrm{greedy}}(t') = T_{\mathrm{LR}}(t')$, or $T^{\mathrm{greedy}}(t') = T_{\mathrm{RL}}(t')$ and the safe tour $T^{\mathrm{safe}}(t')$ serves $\sigma^{\mathrm{R}}(t')$ first, the tour derived at time t is ρ-competitive.*

Proof. By Lemma 29 we know that

$$t_0 \geq \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho-3}. \tag{20}$$

It also holds that $t_0 \leq t + |p_t|$. We have $|a^{\mathrm{L}}(t)| > 0$ and a rightmost extreme is released at time t. Thus, we can deduce by Lemma 19 that $T^{\text{greedy}}(t) = T_{\mathrm{LR}}(t)$ as otherwise ALG would be able to reach the safe tour at time t. We now show that after time t the server does not reach the origin before having served $\sigma^{\mathrm{L}}(t)$ or $\sigma^{\mathrm{R}}(t)$ assuming no further request appear.
It $p_t \leq 0$ the online server always follows the greedy tour in the direction of $\sigma^{\mathrm{L}}(t)$, because then we have $a_1(t) = a^{\mathrm{L}}(t)$ and thus $\operatorname{sign}(p_t) = \operatorname{sign}(a^{\mathrm{L}}(t))$, which means that the algorithm is in Case B1. If $p_t \geq 0$, we have $\operatorname{sign}(p_t) \neq \operatorname{sign}(a^{\mathrm{L}}(t)) = \operatorname{sign}(a_1(t))$. Hence, the server can only start a tour in the direction of $\sigma^{\mathrm{L}}(t)$ if $\tau(t, a^{\mathrm{R}}(t)) = t + 2|a^{\mathrm{R}}(t)| - |p_t| < \frac{\rho|a^{\mathrm{L}}(t)|-(2-\rho)t^{\mathrm{L}}}{2\rho-3}$ holds (Case C). By Lemma 20 this can only be case if the safe tour at time t serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t)|$. But because of $p_t \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ by Lemma 17, it holds that

$$t + 2|a^{\mathrm{R}}(t)| - |p_t| \geq t + |p_t| \geq t_0 \geq \frac{\rho|a^{\mathrm{L}}(t)| - (2-\rho)t^{\mathrm{L}}}{2\rho - 3}$$

which implies that ALG is not in Case C, i.e., it serves $\sigma^{\mathrm{R}}(t)$ first. Hence, we get the following upper bound for $|T^{\text{ALG}}|$

$$|T^{\text{ALG}}| \leq t + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_t|.$$

With $t + |a^{\mathrm{R}}(t)| \leq |T^{\text{OPT}}|$, it suffices to prove that

$$2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_t| \leq (\rho - 1)|T^{\text{OPT}}| \tag{21}$$

which follows from Lemma 21. □

Lemma 33. *Assume the algorithm is in Setting B and at time t' we have $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ while the safe tour $T^{\text{safe}}(t')$ serves $\sigma^{\mathrm{L}}(t')$ first. Then we have $t' = t^{\mathrm{L}}$, $p_{t^{\mathrm{L}}} \geq 0$ and T^{safe} cannot be reached at t^{L}. Moreover we can assume that during $(t^{\mathrm{L}}, t]$ the safe tour cannot be reached and that it holds that $T^{\text{greedy}}(\hat{t}) = T_{\mathrm{LR}}(\hat{t})$ for all $\hat{t} \in (t^{\mathrm{L}}, t]$ at which a rightmost extreme is released.*

Proof. By definition the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t')$ at time t'. This in particular means, that the server cannot reach the safe tour at time t' as the safe tour serves $\sigma^{\mathrm{L}}(t')$ by assumption. Assume $p_{t'} < 0$. Then together with $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and $|a^{\mathrm{L}}(t')| \geq |a^{\mathrm{R}}(t')|$, i.e., the safe tour at time t' serves $\sigma^{\mathrm{L}}(t)$ first, the requirements for Lemma 20 are fulfilled, which implies that the server will start a tour in the direction of $\sigma^{\mathrm{L}}(t')$ at time t', a contradiction. Hence, assume $p_{t'} \geq 0$. If $t' > t$, then at time t' a new rightmost extreme is released. As $T^{\text{greedy}}(t') = T_{\mathrm{RL}}(t')$ and $|a^{\mathrm{L}}(t')| = |a^{\mathrm{L}}(t)| > 0$ by assumption, the requirements for Lemma 19 are fulfilled, which implies that the server can reach the safe tour at time t', a contradiction. Hence, $t' = t^{\mathrm{L}}$. We can thus in the following assume $t' = t^{\mathrm{L}}$, $p_{t^{\mathrm{L}}} \geq 0$ and that the safe tour cannot be reached at time t^{L}.
It remains to be proven that during (t^{L}, t) the safe tour cannot be reached. Assume at some time $t^{\text{safe}} \in (t^{\mathrm{L}}, t)$ the server is able to reach the safe tour and that t^{safe} is the last time during this interval at which the safe tour can be reached.
Assume first that during $[t^{\text{safe}}, t]$ no rightmost extreme is served. Note that the algorithm is in Setting A except for the assumption that during $[t^{\mathrm{L}}, t]$ no rightmost extreme is served but we assumed that during $[t^{\text{safe}}, t]$ no rightmost extreme is served. With Observation 3 this implies that we can apply Lemma 27 which implies that the tour derived at time t is ρ-competitive.
Assume after time $t^{\text{safe}} > t^{\mathrm{L}}$ a rightmost extreme is served. The served rightmost extreme is released at some time $\tilde{t} \in [t^{\text{safe}}, t^{\mathrm{L}}]$, in particular $\tilde{t} > t^{\mathrm{L}}$. This means that at time $\tilde{t}$ we have $T^{\text{greedy}}(\tilde{t}) = T_{\mathrm{LR}}(\tilde{t})$, or $T^{\text{greedy}}(\tilde{t}) = T_{\mathrm{RL}}(\tilde{t})$ and the safe tour at time $\tilde{t}$ serves $\sigma^{\mathrm{L}}(\tilde{t})$ first, because the other remaining case can only occur at time $\tilde{t} = t^{\mathrm{L}}$ by our argument above. Thus, at time $\tilde{t}$ the requirements for Lemma 32 are fulfilled which implies that the tour derived at time t is ρ-competitive. Thus, we can assume that during (t^{L}, t) the safe tour cannot be reached. Assume at time $\hat{t} \in [t^{\mathrm{L}}, t)$ a rightmost extreme is released. Using Lemma 19 we can deduce that $T^{\text{greedy}}(\hat{t}) = T_{\mathrm{LR}}(\hat{t})$ as otherwise the server would be able to reach the safe tour at time $\hat{t}$. Thus $T^{\text{greedy}}(\hat{t}) = T_{\mathrm{LR}}(\hat{t})$, for all $\hat{t} \in (t^{\mathrm{L}}, t]$ at which a rightmost extreme appears. □

Lemma 34. *Assume the server is in Setting B and additionally the following holds:*

- *During $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ no leftmost extreme is served.*

- *During $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ the safe tour cannot be reached.*
- *We have $p_{t^{\mathrm{L}}} \geq 0$*
- *The server moves to the left during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$*
- *Denote by $t'' \in [t^{\mathrm{R}}, t^{\mathrm{L}}]$ the last time in this interval at which the server turns to the right after moving to the left before.*

Then the server only moves to the left during $[t^{\mathrm{R}}, t'']$ and to the right during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ and it holds that $\operatorname{sign}(p_{t^{\mathrm{R}}}) = \operatorname{sign}(p_{t''})$.

Proof. As we assumed that during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ the safe tour cannot be reached, Lemma 19 implies that

$$T^{\text{greedy}} = T_{\mathrm{RL}} \text{ whenever during } (t^{\mathrm{R}}, t^{\mathrm{L}}] \text{ a leftmost extreme is released.} \tag{22}$$

By assumption no leftmost request is served during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ and also no rightmost extreme is released during this interval. This in particular means, that at time t'' a new leftmost request is released. The release of $\sigma^{\mathrm{L}}(t'')$ has the effect that the server starts to move to the right, without having served the previous leftmost request. Note, that such a t'' always exists in the interval $(t^{\mathrm{R}}, t^{\mathrm{L}}]$ (if we assume that ALG is driving to the left at some point in this timeframe), because at time t^{L} the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. Also, by definition of t'', the algorithm only drives to the right during $[t'', t^{\mathrm{L}}]$. This proves the first assertion of the lemma.
Now, we have to take the tour of the server immediately before time t'' into account. By our assumption the server is not able to recover the safe tour in $[t^{\mathrm{R}}, t^{\mathrm{L}}]$. Thus, we can assume that immediately before time t'', the server was on a nonsafe tour first serving the leftmost extreme, because $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is not served during $[t^{\mathrm{R}}, t^{\mathrm{L}})$. Assume at time t''' the last request before time t'' is released. At time t''' the server starts a tour first moving to the left. This in particular means $|a^{\mathrm{L}}(t''')| > 0$. We make a case distinction regarding the position of the server at time t''.
Assume $p_{t''} \geq 0$. This implies $p_{t'''} \geq 0$ as the server by definition only moves to the left during $[t''', t'']$. In this case the server will only start a tour first moving to the left at time t''' if $\operatorname{sign}(p_t) \neq \operatorname{sign}(a_1(t'''))$, i.e., if $T^{\text{greedy}}(t''') = T_{\mathrm{LR}}(t''')$ (Case C). By (22) we know that $T^{\text{greedy}}(t''') = T_{\mathrm{RL}}(t''')$ if $t''' > t^{\mathrm{R}}$. Hence, we can assume that $t''' = t^{\mathrm{R}}$ which implies that during $[t^{\mathrm{R}}, t'']$ the server only moves to the left and it also holds $\operatorname{sign}(p_{t''}) = \operatorname{sign}(p_{t'''})$ which means that the assertions of the lemma are fulfilled.
Next assume $p_{t''} < 0$. This in particular means that $t'' \neq t^{\mathrm{L}}$, because $p_{t^{\mathrm{L}}} \geq 0$ by the assumptions of the lemma. Recall, that at time t'' a new leftmost request is released such that the algorithm starts a nonsafe tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. It is $t'' \in (t^{\mathrm{R}}, t^{\mathrm{L}}]$, hence we know by (22) that $T^{\text{greedy}}(t'') = T_{\mathrm{RL}}(t'')$. Thus, in the notation of the algorithm we have $\operatorname{sign}(p_{t''}) \neq \operatorname{sign}(a_1(t'')) = \operatorname{sign}(a^{\mathrm{R}}(t''))$. Hence, at time t'' the server only starts a tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ if it is in Case C. Thus, we must have,

$$\begin{aligned}\tau(t'', a^{\mathrm{L}}(t'')) &= t'' + 2|a^{\mathrm{L}}(t'')| - |p_{t''}| \\ &< \frac{\rho|a^{\mathrm{R}}(t^{\mathrm{L}})| - (2-\rho)t^{\mathrm{R}}}{2\rho - 3}.\end{aligned} \tag{23}$$

We have $p_{t''} < 0$ by assumption and it also holds $T^{\text{greedy}}(t'') = T_{\mathrm{RL}}(t'')$. Hence, by Lemma 20 we also know that the server can only start a tour first serving $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ at time t'' if the safe tour at time t'' serves $a^{\mathrm{R}}(t^{\mathrm{L}})$ first, i.e., $|a^{\mathrm{R}}(t^{\mathrm{L}})| \geq |a^{\mathrm{L}}(t'')|$. Using this and the fact that at time t''' the algorithm is going in the direction of a leftmost extreme, there are only two cases which can occur at time t'''.
Assume first that the greedy tour at time t''' serves $\sigma^{\mathrm{L}}(t''')$ first. This in particular implies that $t''' = t^{\mathrm{R}}$ by (22). At time t'' a new leftmost extreme is released and as no leftmost extremes are served during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ we have $|a^{\mathrm{L}}(t'')| > |a^{\mathrm{L}}(t^{\mathrm{R}})|$. Thus, it also holds $|a^{\mathrm{R}}(t^{\mathrm{L}})| > |a^{\mathrm{L}}(t^{\mathrm{R}})|$, i.e., the safe tour at time t^{R} serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first. Together with $T^{\text{greedy}}(t^{\mathrm{R}}) = T_{\mathrm{LR}}(t^{\mathrm{R}})$ Lemma 20 implies that the server can only start a tour first serving $\sigma^{\mathrm{L}}(t^{\mathrm{R}})$ if $p_{t^{\mathrm{R}}} < 0$. Thus it holds that $\operatorname{sign}(p_{t''}) = \operatorname{sign}(p_{t^{\mathrm{R}}})$ and by definition of $t''' = t^{\mathrm{R}}$ the server only moves to the left during $[t^{\mathrm{R}}, t'']$. Hence, again the assertions of the Lemma are fulfilled.
Assume now that the greedy tour at time t''' serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first. As $|a^{\mathrm{L}}(t''')| > 0$ by assumption, we can use Lemma 19 to deduce that $t''' > t^{\mathrm{R}}$, as otherwise the server could reach the safe tour at time t^{R}. We have in

the notation of the algorithm that $a_1(t''') = a^{\mathrm{R}}(t''')$. Thus the algorithm will only start a tour first serving $\sigma^{\mathrm{L}}(t''')$ if $\mathrm{sign}(p_{t'''}) \neq \mathrm{sign}(a_1(t''')) = \mathrm{sign}(a^{\mathrm{R}}(t'''))$, i.e., $p_{t'''} < 0$ and inequality (1) from Fact 1 is fulfilled when doing so (Case B2), i.e.,

$$\tau(t''', a^{\mathrm{L}}(t''')) = t''' + 2|a^{\mathrm{L}}(t''')| - |p_{t'''}| \geq \frac{\rho|a^{\mathrm{R}}(t^{\mathrm{L}})| - (2-\rho)t^{\mathrm{R}}}{2\rho - 3}. \tag{24}$$

We deduced above that $t''' > t^{\mathrm{R}}$. Hence, at time t''' a leftmost extreme with $|a^{\mathrm{L}}(t''')| < |a^{\mathrm{L}}(t'')|$ is released, because no leftmost extreme is served during $[t^{\mathrm{L}}, t^{\mathrm{R}}]$. Between t''' and t'' the server moves to the left without interruption. Thus, we have $(t'' - t''') = |p_{t'''} - p_{t''}|$. Since, $p_{t''} \leq p_{t'''} < 0$, it holds that $|p_{t'''} - p_{t''}| = |p_{t''}| - |p_{t'''}|$. Thus, we have $t'' + 2|a^{\mathrm{L}}(t''')| - |p_{t''}| = t''' + 2|a^{\mathrm{L}}(t''')| - |p_{t'''}|$. Using this and $|a^{\mathrm{L}}(t'')| > |a^{\mathrm{L}}(t''')|$, we get

$$\begin{aligned} t'' + 2|a^{\mathrm{L}}(t'')| - |p_{t''}| &> t''' + 2|a^{\mathrm{L}}(t''')| - |p_{t'''}| \\ &\overset{(24)}{\geq} \frac{\rho|a^{\mathrm{R}}(t^{\mathrm{L}})| - (2-\rho)t^{\mathrm{R}}}{2\rho - 3}, \end{aligned}$$

contradicting inequality (23).

□

Lemma 35. *Assume the algorithm is in Setting B and at time t' we have $T^{\mathrm{greedy}}(t') = T_{\mathrm{RL}}(t')$ while the safe tour $T^{\mathrm{safe}}(t')$ serves $\sigma^{\mathrm{L}}(t')$ first. Then the tour derived at time t is ρ-competitive.*

Proof. By Lemma 33 we can assume that $t' = t^{\mathrm{L}}$, $p_{t^{\mathrm{L}}} \geq 0$, that the safe tour cannot be reached during $[t^{\mathrm{L}}, t)$ and that

$$T^{\mathrm{greedy}}(\hat{t}) = T_{\mathrm{LR}}(\hat{t}), \text{ for all } \hat{t} \in (t^{\mathrm{L}}, t] \text{ at which a rightmost extreme appears.}$$

The next step is to derive an upper bound for $|T^{\mathrm{ALG}}|$. In the following we will show that

$$|T^{\mathrm{ALG}}| \leq t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}|. \tag{25}$$

Recall, that at time t^{L} the server starts a tour in the direction of $|a^{\mathrm{R}}(t^{\mathrm{L}})|$ and that during $(t^{\mathrm{L}}, t]$ only rightmost extremes are released. Our claim for $|T^{\mathrm{ALG}}|$ follows immediately, if between t^{L} and t no new requests are released or if all rightmost requests which are released after $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ was served, force the server onto a (nonsafe) tour in the direction of $\sigma^{\mathrm{L}}(t)$. We will in the following analyze which cases of the algorithm can occur at time t''.

Assume that at a time $t'' \leq t$ after ALG has served $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ a new rightmost request $\sigma^{\mathrm{R}}(t'')$ is released which forces the server onto a nonsafe tour first serving the rightmost request. Note, that we can assume that the server does not serve $\sigma^{\mathrm{R}}(t'')$ during $[t'', t]$ as in this case we can apply Lemma 32 which directly implies that the tour derived at time t is ρ-competitive. Also assume that t'' is the first time that after having served $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ the server moves towards a rightmost extreme. Note, that by Lemma 22 the server will not turn around again during $[t'', t]$ before $\sigma^{\mathrm{R}}(t'')$ was served. We know that at time t'' the greedy tour serves $\sigma^{\mathrm{L}}(t)$ first, i.e., $a_1(t'') = a^{\mathrm{L}}(t)$ in the notation of the algorithm. Thus, if $p_{t''} < 0$, we have $\mathrm{sign}(p_{t''}) = \mathrm{sign}(a_1(t''))$ and the server always starts a tour in the direction of $\sigma^{\mathrm{L}}(t)$ at time t'' (because we assume that the safe tour cannot be reached) (Case B1). This contradicts the definition of t''. Hence, we can assume that $p_{t''} \geq 0$. Thus, because between serving $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ and time t'' the server did not leave $[p_{t''}, a^{\mathrm{R}}(t^{\mathrm{L}})]$ and since $\sigma^{\mathrm{R}}(t'')$ is not served during $[t'', t]$, we have

$$\begin{aligned} |T^{\mathrm{ALG}}| &\leq t^{\mathrm{L}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| + |a^{\mathrm{R}}(t^{\mathrm{L}}) - p_{t''}| + |p_{t''}| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\ &\leq t^{\mathrm{L}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}|, \end{aligned}$$

as claimed. At time t^{L} we have $|a^{\mathrm{L}}(t)| \geq |a^{\mathrm{R}}(t^{\mathrm{L}})|$, because $T^{\mathrm{safe}}(t^{\mathrm{L}})$ serves $\sigma^{\mathrm{L}}(t^{\mathrm{L}})$ first by assumption. Also, recall that $|a^{\mathrm{R}}(t^{\mathrm{L}})| > 0$ and that at time t^{L} a leftmost request is released. Thus, at time t^{L} all conditions for Lemma 21 are fulfilled, which implies

$$2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}| < (\rho - 1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|). \tag{26}$$

We have to distinguish several cases. First assume that $|a^{\mathrm{R}}(t^{\mathrm{L}})| \leq |a^{\mathrm{R}}(t)|$. Inequalities (25) and (26) and $t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| \leq t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \leq |T^{\textsc{Opt}}|$ yield

$$\begin{aligned}
|T^{\textsc{Alg}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\
&\overset{(26)}{\leq} (\rho - 1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|) + t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\
&\leq \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) \\
&\leq \rho|T^{\textsc{Opt}}|.
\end{aligned}$$

Hence, in the following we can assume that

$$|a^{\mathrm{R}}(t^{\mathrm{L}})| > |a^{\mathrm{R}}(t)| \tag{27}$$

holds. Now we have to take the behaviour of the offline optimum into account.

- Assume that in the offline optimal solution $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is served after $\sigma^{\mathrm{L}}(t)$, i.e.,

$$|T^{\textsc{Opt}}| \geq t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|. \tag{28}$$

  Inequality (26) and $|a^{\mathrm{R}}(t^{\mathrm{L}})| > |a^{\mathrm{R}}(t)|$ imply $2|a^{\mathrm{R}}(t)| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}| < (\rho - 1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|)$. Thus, the following holds,

$$\begin{aligned}
|T^{\textsc{Alg}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\
&\leq (\rho - 1)(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})|) + t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| \\
&= \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|) \\
&\overset{(28)}{\leq} \rho|T^{\textsc{Opt}}|.
\end{aligned}$$

- Assume that in the offline optimal solution $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ and $\sigma^{\mathrm{R}}(t)$ are served before $\sigma^{\mathrm{L}}(t)$. Then

$$|T^{\textsc{Opt}}| \geq t + |a^{\mathrm{R}}(t)| + 2|a^{\mathrm{L}}(t)|. \tag{29}$$

  We assume that $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is served during $[t^{\mathrm{L}}, t]$. Thus, the server serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ before the release of $\sigma^{\mathrm{R}}(t)$ at time t, which implies

$$t \geq t^{\mathrm{L}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}|. \tag{30}$$

  Hence, again using (25), $|a^{\mathrm{R}}(t^{\mathrm{L}})| \leq |a^{\mathrm{L}}(t)|$ as the safe tour at time t^{L} serves $\sigma^{\mathrm{L}}(t^{\mathrm{L}})$ first by assumption, $p_t \in [a^{\mathrm{L}}(t), a^{\mathrm{R}}(t)]$ by Lemma 17, and our assumption regarding the offline optimum, we get that

$$\begin{aligned}
|T^{\textsc{Alg}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\
&\overset{(30)}{\leq} t + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\
&\overset{(29)}{\leq} |T^{\textsc{Opt}}| + |a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{R}}(t)| \\
&\leq |T^{\textsc{Opt}}| + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\
&\overset{\text{Lem.17}}{\leq} 3|T^{\textsc{Opt}}|/2.
\end{aligned}$$

- Assume that in the optimal offline solution $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is served before $\sigma^{\mathrm{L}}(t)$ while $\sigma^{\mathrm{R}}(t)$ is served after $\sigma^{\mathrm{L}}(t)$. Define $t^{\mathrm{R}} := t^{\mathrm{R}}(t^{\mathrm{L}})$ to be the release time of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. Our assumption on the offline optimum solution implies,

$$|T^{\textsc{Opt}}| \geq t^{\mathrm{R}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|. \tag{31}$$

  To analyze this case, we have to regard the behaviour of Alg before $\sigma^{\mathrm{L}}(t)$ is released, i.e., we are interested in the timeframe between t^{R} and t^{L}. By definition, no new rightmost request appears during $(t^{\mathrm{R}}, t^{\mathrm{L}}]$. We again distinguish several cases:

1. *During $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ no leftmost extreme is served*

   – *During $[t^{\mathrm{R}}, t^{\mathrm{L}})$ the server is able to reach the safe tour*
   Recall, that by assumption at time t^{L} the safe tour cannot be reached. Let $t^{\mathrm{safe}} \in [t^{\mathrm{R}}, t^{\mathrm{L}})$ be the last time at which the safe tour can be reached in this interval. Recall that at time t^{L} a new leftmost extreme is released, that the safe tour cannot be reached at time t^{L}, we have $|a^{\mathrm{R}}(t^{\mathrm{L}})| > 0$ and $T^{\mathrm{greedy}}(t^{\mathrm{L}}) = T_{\mathrm{RL}}(t)$ by assumption. Also during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ no leftmost extreme is served. Thus, the requirements for Corollary 26 are fulfilled and we have that $p_{t^{\mathrm{L}}} \geq 0$, the safe tour at time t^{safe} serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first and the server does not turn around after t^{safe} before $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is served. Using the fact that at time t^{safe} the safe tour first serving $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ can be reached and that the server only moves to the right during $[t^{\mathrm{safe}}, t^{\mathrm{L}}]$ we get with Definition 18

   $$t^{\mathrm{L}} \leq \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| - 2|a^{\mathrm{L}}(t^{\mathrm{L}})| - 2|a^{\mathrm{R}}(t^{\mathrm{safe}})| + |p_{t^{\mathrm{L}}}|. \tag{32}$$

   Also note that with $\hat{t} = t^{\mathrm{L}}$ and $\tilde{t} = t^{\mathrm{safe}}$ the requirements for Corollary 24 are fulfilled and it holds that

   $$\rho|T^{\mathrm{greedy}}(t)| = \rho|T_{\mathrm{RL}}(t^{\mathrm{L}})| > \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{L}}(t) - a^{\mathrm{L}}(t^{\mathrm{safe}})|. \tag{33}$$

   Recall, that at time t^{L} the greedy tour serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ first, which implies

   $$|T^{\mathrm{OPT}}| \overset{(31)}{\geq} t^{\mathrm{R}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| = T^{\mathrm{greedy}}(t^{\mathrm{L}}) + 2|a^{\mathrm{R}}(t)|. \tag{34}$$

   Thus, all in all we get

   $$\begin{aligned} |T^{\mathrm{ALG}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\ &\overset{(32)}{\leq} \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| - 2|a^{\mathrm{L}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\ &= \rho|T^{\mathrm{greedy}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{L}}(t) - a^{\mathrm{L}}(t^{\mathrm{safe}})| + 2|a^{\mathrm{R}}(t)| \\ &\overset{(33)}{<} \rho|T^{\mathrm{greedy}}(t^{\mathrm{L}})| + 2|a^{\mathrm{R}}(t)| \overset{(34)}{<} \rho|T^{\mathrm{OPT}}|. \end{aligned}$$

   – *During $[t^{\mathrm{R}}, t^{\mathrm{L}})$ the server cannot reach the safe tour*
   In this case Lemma 19 implies that

   $$T^{\mathrm{greedy}} = T_{\mathrm{RL}} \text{ whenever during } (t^{\mathrm{R}}, t^{\mathrm{L}}] \text{ a leftmost extreme is released.} \tag{35}$$

   In the following we will regard two cases. Either the server does not move to the left at all during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ or it moves to the left but is forced to turn around before reaching a leftmost request. Recall, that at time t^{L} the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. Hence, if the server moves to the left during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ it is forced to turn around at the latest at time t^{L}.

   Assume first that ALG does not move to the left at all between t^{R} and t^{L}. Then, since we assumed that no safe tour can be reached during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$, the server directly starts to move to the right without interruption at time t^{R} (Case B1,Case B2 or Case C). Thus, we have $t^{\mathrm{L}} = t^{\mathrm{R}} + |p_{t^{\mathrm{R}}} - p_{t^{\mathrm{L}}}|$, because $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ is still unserved at time t^{L}. Using (25), (31) and $|p_t - a^{\mathrm{R}}(t^{\mathrm{L}})| \leq |T^{\mathrm{OPT}}|/2$ by Lemma 17,

   $$\begin{aligned} |T^{\mathrm{ALG}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\ &\leq t^{\mathrm{R}} + |p_{t^{\mathrm{R}}} - p_{t^{\mathrm{L}}}| + |a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\ &\leq t^{\mathrm{R}} + |p_{t^{\mathrm{R}}} - a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\ &\overset{(31)}{\leq} 3|T^{\mathrm{OPT}}|/2. \end{aligned}$$

   Next, assume that the server does move to the left during $[t^{\mathrm{R}}, t^{\mathrm{L}}]$. Denote by $t'' \in (t^{\mathrm{R}}, t^{\mathrm{L}}]$ the last point in time at which the server turns around to drive to the right, after driving to the left before. As by assumption no leftmost request is served during this interval, this in particular means, that at time t'' a new leftmost request is released. The release of $\sigma^{\mathrm{L}}(t'')$ has the effect that the server starts to move to the right, without having served the previous leftmost request. Note, that such a

t'' always exists in the interval $(t^{\mathrm{R}}, t^{\mathrm{L}}]$ (if we assume that ALG is driving to the left at some point in this timeframe), because at time t^{L} the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. Now the requirements of Lemma 34 are fulfilled which implies that during $[t^{\mathrm{R}}, t'']$ the server only moves to the left and during $[t'', t^{\mathrm{L}}]$ the server only moves to the right and $\operatorname{sign}(p_{t^{\mathrm{R}}}) = \operatorname{sign}(p_{t''})$. In particular $|a^{\mathrm{L}}(t^{\mathrm{R}})| > 0$ as the server starts a tour in the direction of a rightmost extreme at time t^{R}. We now make a case distinction regarding the position of the server at time t''.
Assume $p_{t''} \geq 0$, thus we have $p_{t^{\mathrm{R}}} \geq 0$. During $[t^{\mathrm{R}}, t'']$ the server only moves to the left by assumption. Thus, we have $p_{t^{\mathrm{R}}} \geq p_{t''} \geq 0$ and $t'' = t^{\mathrm{R}} + |p_{t''} - p_{t^{\mathrm{R}}}| \leq t^{\mathrm{R}} + |p_{t^{\mathrm{R}}}| + |p_{t''}|$. During $[t'', t^{\mathrm{L}}]$ the server only moves to the right, thus $t^{\mathrm{L}} = t'' + |p_{t^{\mathrm{L}}} - p_{t''}| = t'' + |p_{t^{\mathrm{L}}}| - |p_{t''}|$, because $p_{t^{\mathrm{L}}} \geq p_{t''} \geq 0$ (it is $p_{t^{\mathrm{L}}} \geq 0$ by our assumption from the beginning of the proof). Note that we have $|p_{t^{\mathrm{R}}}| < |a^{\mathrm{R}}(t^{\mathrm{L}})| \leq |a^{\mathrm{L}}(t^{\mathrm{L}})|$ because $T^{\mathrm{safe}}(t^{\mathrm{L}})$ serves $\sigma^{\mathrm{L}}(t^{\mathrm{L}})$ first by assumption. Thus we have $|p_{t^{\mathrm{L}}}| + |a^{\mathrm{R}}(t^{\mathrm{L}})| \leq |a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| \leq |T^{\mathrm{OPT}}|/2$ This implies together with (25)

$$\begin{aligned} |T^{\mathrm{ALG}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\ &= t'' + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| + |p_{t''}| \\ &\leq t^{\mathrm{R}} + |p_{t^{\mathrm{R}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| \\ &\overset{(31)}{\leq} 3|T^{\mathrm{OPT}}|/2. \end{aligned}$$

Next assume $p_{t''} < 0$, i.e. also $p_{t^{\mathrm{R}}} < 0$. Between t^{R} and t'' the server moves to the left without interruption and between t'' and t^{L} it moves to the right without interruption. With $p_{t^{\mathrm{L}}} \geq 0$ by the assumptions from the beginning of the proof and $p_{t^{\mathrm{R}}} < 0$ this implies $t^{\mathrm{L}} \leq t^{\mathrm{R}} + |a^{\mathrm{L}}(t^{\mathrm{R}}) - p_{t^{\mathrm{R}}}| + |a^{\mathrm{L}}(t^{\mathrm{R}}) - p_{t^{\mathrm{L}}}| = t^{\mathrm{R}} + |a^{\mathrm{L}}(t^{\mathrm{R}})| - |p_{t^{\mathrm{R}}}| + |a^{\mathrm{L}}(t^{\mathrm{R}})| + |p_{t^{\mathrm{L}}}|$. Using inequality (25) this yields

$$\begin{aligned} |T^{\mathrm{ALG}}| &\overset{(25)}{\leq} t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\ &\leq t^{\mathrm{R}} + 2|a^{\mathrm{L}}(t^{\mathrm{R}})| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{R}}}|. \end{aligned}$$

As

$$|T^{\mathrm{OPT}}| \overset{(31)}{\geq} t^{\mathrm{R}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)|,$$

we are left to prove that

$$2|a^{\mathrm{L}}(t^{\mathrm{R}})| + |a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{R}}}| \leq (\rho - 1)|T^{\mathrm{OPT}}|.$$

At time t^{R} we have $|a^{\mathrm{L}}(t^{\mathrm{R}})| > 0$, $|a^{\mathrm{R}}(t^{\mathrm{R}})| = |a^{\mathrm{R}}(t^{\mathrm{L}})| \geq |a^{\mathrm{L}}(t^{\mathrm{R}})|$ by assumption (see (27)). Also the safe tour cannot be reached at time t^{R}. Hence, at time t^{R} all conditions for Lemma 21 are fulfilled which implies

$$\begin{aligned} 2|a^{\mathrm{L}}(t^{\mathrm{R}})| + |a^{\mathrm{R}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| &< (\rho - 1)(t^{\mathrm{R}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)|) \\ &\overset{(31)}{\leq} (\rho - 1)|T^{\mathrm{OPT}}|. \end{aligned}$$

2. *During* $[t^{\mathrm{R}}, t^{\mathrm{L}}]$ *a leftmost extreme is served*
Let $t''' \in [t^{\mathrm{R}}, t^{\mathrm{L}}]$ be the time at which the last request before ALG serves the leftmost extreme is released. In particular this means $|a^{\mathrm{L}}(t''')| > 0$. At time t''' the server starts a tour in the direction of the current leftmost extreme $\sigma^{\mathrm{L}}(t''')$ and does not turn around before serving it. Recall, that we have that at time t^{L} a new leftmost extreme is released and the safe tour cannot be reached. Also assume that at time t''' we have $T^{\mathrm{greedy}}(t''') = T_{\mathrm{RL}}(t''')$, or $T^{\mathrm{greedy}}(t''') = T_{\mathrm{LR}}(t''')$ and the safe tour $T^{\mathrm{safe}}(t''')$ serves $\sigma^{\mathrm{L}}(t''')$ first. Then the requirements for Corollary 31 are fulfilled and we have that

$$2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}| \leq (\rho - 1)|T^{\mathrm{OPT}}|.$$

Next assume that $T^{\mathrm{greedy}}(t''') = T_{\mathrm{LR}}(t''')$ and the safe tour at time t''' serves $\sigma^{\mathrm{R}}(t''')$, i.e. $a^{\mathrm{R}}(t''') \geq |a^{\mathrm{L}}(t''')|$. We make a case distinction regarding the position $p_{t'''}$.
Assume $p_{t'''} \geq 0$. Then together with $T^{\mathrm{greedy}}(t''') = T_{\mathrm{LR}}(t''')$ and $|a^{\mathrm{R}}(t''')| \geq |a^{\mathrm{L}}(t''')|$ the requirements for Lemma 20 are fulfilled which implies that the server starts a tour in the direction of $\sigma^{\mathrm{R}}(t''')$ at time t''', a contradiction.

Hence, we can assume that $p_{t'''} \leq 0$. As the safe tour at time t''' serves $\sigma^{\mathrm{R}}(t''')$ first while the server starts a tour in the direction of $\sigma^{\mathrm{L}}(t''')$, we know that the server does not reach the safe tour at time t'''. Assume $t''' > t^{\mathrm{R}}$. Then by definition of t''', a new leftmost extreme is released at time t'''. Together with $T^{\text{greedy}}(t''') = T_{\mathrm{LR}}(t''')$ and $|a^{\mathrm{R}}(t''')| = |a^{\mathrm{R}}(t^{\mathrm{L}})| > 0$ the requirements for Lemma 19 are fulfilled which implies that the server can reach the safe tour a time t''', a contradiction. Hence, we can assume that $t''' = t^{\mathrm{R}}$. Now, all requirements for Lemma 21 are fulfilled, with $|a^{\mathrm{L}}(t^{\mathrm{R}})| > 0$, $|a^{\mathrm{R}}(t^{\mathrm{L}})| \geq |a^{\mathrm{L}}(t^{\mathrm{R}})|$ and the fact that the safe tour cannot be reached at time t^{R}. This again implies using (31) that $2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t^{\mathrm{L}})| - |p_{t^{\mathrm{L}}}| \leq (\rho - 1)|T^{\textsc{Opt}}|$. Hence, in general it holds that

$$2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)| - |p_{t^{\mathrm{L}}}| \leq (\rho - 1)|T^{\textsc{Opt}}|.$$

Thus, we can deduce

$$\begin{aligned} |T^{\textsc{Alg}}| &\leq t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| - |p_{t^{\mathrm{L}}}| \\ &\leq t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + 2|a^{\mathrm{R}}(t)| + (\rho - 1)|T^{\textsc{Opt}}| \\ &\leq \rho |T^{\textsc{Opt}}|. \end{aligned}$$

□

B Proofs of Section 4

Lemma 7. *There exists $W \geq 0$ with*

$$\operatorname{delay}\left(|T_{\mathrm{LR}}| + \frac{W}{\rho' - 1}\right) = W. \tag{3}$$

Proof. First observe that Alg has served neither σ^{L} nor σ^{R} at time $|a_1^\star|$, by Definition 5 (ii). Since, by definition, σ^{L} is served after σ^{R}, and, hence, σ^{L} is still unserved at time $|a_1^\star| + |a^{\mathrm{L}}| + |a^{\mathrm{R}}| \geq |T_{\mathrm{LR}}| = 2|a^{\mathrm{L}}| + |a^{\mathrm{R}}|$. Therefore, $\operatorname{delay}(|T_{\mathrm{LR}}|)$ is defined. Now note that we have

$$t^\star = (2\rho' - 2) \cdot |a^{\mathrm{L}}| + (\rho' - 2) \cdot |a^{\mathrm{R}}| < 2.08 \cdot |a^{\mathrm{L}}| + 0.04 \cdot |a^{\mathrm{R}}| \leq 2|a^{\mathrm{L}}| + |a^{\mathrm{R}}| = |T_{\mathrm{LR}}|,$$

where the last inequality follows by Definition 5 (iv). Since Alg does not serve σ^{R} earlier than $t^\star$ by Definition 5 (iii), this implies that no tour $T \in \mathcal{T}$ (as in Fact 3) can arrive at a^{L} earlier than $t^{\star\star} = t^\star + |a^{\mathrm{L}}| + |a^{\mathrm{R}}|$, and thus

$$\operatorname{delay}(|T_{\mathrm{LR}}|) \geq 0 \tag{36}$$

by Fact 3 (ii).

If $\operatorname{delay}(|T_{\mathrm{LR}}|) = 0$, we have $W = 0$, and we are done. Otherwise, by Inequality (36), we must have $\operatorname{delay}(|T_{\mathrm{LR}}|) > 0$. Observe that, if no new requests appear, Alg must serve σ^{L} at some time in order to stay $(\rho - \varepsilon)$-competitive. Let $W^\star$ be chosen such that Alg serves σ^{L} at time $|T_{\mathrm{LR}}| + W^\star/(\rho' - 1)$, that is,

$$\operatorname{delay}\left(|T_{\mathrm{LR}}| + \frac{W^\star}{\rho' - 1} - \varepsilon'\right)$$

is defined for some sufficiently small $\varepsilon' \leq |a^{\mathrm{L}}|$. Thus, if

$$\operatorname{delay}\left(|T_{\mathrm{LR}}| + \frac{W^\star}{\rho' - 1} - \varepsilon'\right) \leq \frac{W^\star}{\rho' - 1} - \varepsilon',$$

we can find W within the interval $(0, W^\star - \varepsilon(\rho' - 1)]$ since $\operatorname{delay}(|T_{\mathrm{LR}}|) > 0$ and both sides of Equation (3) are continuous, and we are done. Otherwise

$$\operatorname{delay}\left(|T_{\mathrm{LR}}| + \frac{W^\star}{\rho' - 1} - \varepsilon'\right) > \frac{W^\star}{\rho' - 1} - \varepsilon'.$$

So, by Fact 3 (i), Alg has not served σ^{L} at time

$$\begin{aligned} t^{\star\star} + \frac{W^\star}{\rho'-1} - \varepsilon' &= (2\rho'-1)\cdot|a^{\mathrm{L}}| + (\rho'-1)\cdot|a^{\mathrm{R}}| + \frac{W^\star}{\rho'-1} - \varepsilon' \\ &> |T_{\mathrm{LR}}| + \frac{W^\star}{\rho'-1}, \end{aligned}$$

where we use the definition of $t^{\star\star}$ in the first step and $\rho' > 2$, $\varepsilon' \le |a^{\mathrm{L}}|$ in the second step. We obtain a contradiction to the definition of $W^\star$. □

Lemma 8. *With W as in Lemma 7,* Alg *serves σ^{R} no later than time $|T_{\mathrm{LR}}| + \frac{W}{\rho'-1}$.*

Proof. Fix $t = |T_{\mathrm{LR}}| + \frac{W}{\rho'-1}$ and let $\mathcal{T}$ be as in Fact 3. Assume that

$$|T_{\mathrm{LR}}| + \frac{W}{\rho'-1} \ge t^\star + W. \tag{37}$$

Then, by definition of W and by Fact 3 (ii), there must be $T \in \mathcal{T}$ that serves σ^{L} at time $t^{\star\star} + \mathrm{delay}(t) = t^{\star\star} + W$. Since $t \ge t^\star + W$ by assumption, this can only be the case if Alg serves σ^{R} no later than time $t^{\star\star} + W - |a^{\mathrm{L}}| - |a^{\mathrm{R}}| = t^\star + W \le t$ as claimed.

It remains to show (37). By Definition 5 (i) and since $|a_0^\star| \le |a_1^\star|$, the tour $\mathrm{move}(a_0^\star) \oplus \mathrm{move}(a_1^\star)$ is optimal. By Fact 3 (i) and for Alg to be $(\rho-\varepsilon)$-competitive, we thus have

$$t^{\star\star} + \mathrm{delay}(t) \le (\rho-\varepsilon)\cdot(2|a_0^\star| + |a_1^\star|)$$

and thus

$$\mathrm{delay}(t) \le (\rho-\varepsilon)\cdot(2|a_0^\star| + |a_1^\star|) - t^{\star\star} \le |a_0^\star| + |a_1^\star|. \tag{38}$$

We can now derive Inequality (37) (solved for W):

$$\begin{aligned} \frac{|T_{\mathrm{LR}}| - t^\star}{1 - 1/(\rho'-1)} &= \frac{(\rho'-1)(|T_{\mathrm{LR}}| - t^\star)}{\rho'-2} \\ &= \frac{(\rho'-1)(4-2\rho')}{(\rho'-2)}\cdot|a^{\mathrm{L}}| + \frac{(\rho'-1)(3-\rho')}{(\rho'-2)}\cdot|a^{\mathrm{R}}| \\ &> -2.08|a^{\mathrm{L}}| + 25|a^{\mathrm{R}}| \\ &> |a^{\mathrm{L}}| + |a^{\mathrm{R}}| && \text{(Definition 5 (iv))} \\ &= |a_0^\star| + |a_1^\star| \\ &\ge W. \end{aligned}$$

□

Lemma 6. *If there is a request sequence with two critical requests for* Alg*, we can release additional requests such that* Alg *is not $(\rho-\varepsilon)$-competitive on the resulting instance.*

Proof. We let W as in Lemma 7, present the request

$$\sigma_+^{\mathrm{R}} = (a_+^{\mathrm{R}}; t_+^{\mathrm{R}}) := \left(|a^{\mathrm{R}}| + \frac{W}{\rho'-1}; |T_{\mathrm{LR}}| + \frac{W}{\rho'-1}\right)$$

and distinguish two cases.

Case 1: At time t_+^{R}, Alg *is at least as close to a^{L} as to a_+^{R}, or it deliberately serves σ^{L} before σ_+^{R}.* We do not present additional requests. Since Alg has served σ^{R} at time t_+^{R} already (Lemma 8), it follows by

Fact 3 (i) and the definition of W that $\textsc{Alg}$ does not serve σ^{L} earlier than $t^{\star\star} + W$. Hence, the total cost of $\textsc{Alg}$ is

$$\begin{aligned}
|T^{\textsc{Alg}}| &\geq t^{\star\star} + W + |a^{\mathrm{L}}| + |a^{\mathrm{R}}| + \frac{W}{\rho' - 1} \\
&= \rho' \cdot \left(|T_{\mathrm{LR}}| + \frac{W}{\rho' - 1} \right) = \rho' \cdot t_{+}^{\mathrm{R}} \geq \rho' \cdot |T^{\textsc{Opt}}|,
\end{aligned}$$

which shows the claim.

Case 2: At time t_{+}^{R}, $\textsc{Alg}$ is closer to a_{+}^{R} than to a^{L}, and it serves σ_{+}^{R} first. The offline server first serving σ^{L} continues going to the right at time t_{+}^{R} (note that, in fact, it has been going to the right from time $|a^{\mathrm{L}}|$ on). For later times t, we denote by

$$M(t) \coloneqq \frac{-|a^{\mathrm{L}}| + (t - 2|a^{\mathrm{L}}|)}{2} = \frac{t - 3|a^{\mathrm{L}}|}{2}$$

the midpoint between the current position of the offline server and $|a^{\mathrm{L}}|$. We distinguish two cases.

Case 2a: $\textsc{Alg}$ does not serve σ_{+}^{R} until time $M^{-1}(a_{+}^{R})$, i.e., the midpoint reaches σ_{+}^{R} before $\textsc{Alg}$. We do not present additional requests and compute $\textsc{Alg}$'s cost. Because we are in Case 2, σ_{+}^{L} is also not served before $M^{-1}(a_{+}^{R})$. Using (37) from the proof of Lemma 8, we have $t_{+}^{\mathrm{R}} \geq t^{\star} + W$ and the resulting total cost for $\textsc{Alg}$ is

$$\begin{aligned}
|T^{\textsc{Alg}}| &\geq M^{-1}(a_{+}^{\mathrm{R}}) + |a_{+}^{\mathrm{R}}| + |a^{\mathrm{L}}| \\
&\overset{(37)}{\geq} t^{\star} + W + \left(M^{-1}(a_{+}^{\mathrm{R}}) - t_{+}^{\mathrm{R}}\right) + |a_{+}^{\mathrm{R}}| + |a^{\mathrm{L}}| \\
&= \frac{(\rho' + 1)W}{\rho' - 1} + \rho'|a^{\mathrm{R}}| + 2\rho'|a^{\mathrm{L}}| \\
&= \rho' \cdot \left(|T_{\mathrm{LR}}| + \frac{W}{\rho' - 1} \right) + \frac{W}{\rho' - 1} \geq \rho' \cdot |T^{\textsc{Opt}}|,
\end{aligned}$$

which shows the claim.

Case 2b: $\textsc{Alg}$ serves σ^{R} before time $M^{-1}(a_{+}^{R})$. By definition of W, the delay function is defined at for time t_{+}^{R}, hence $\textsc{Alg}$ cannot have served σ^{L} before this time. Since $\textsc{Alg}$ is to the right of the midpoint at time t_{+}^{R}, there must be a (first) time t_{mid} at which $M(t_{\mathrm{mid}}) = p_{t_{\mathrm{mid}}}$. We present a last request $\sigma_{++}^{\mathrm{R}} = (a_{++}^{\mathrm{R}}; t_{++}^{\mathrm{R}}) \coloneqq (t_{\mathrm{mid}} - 2|a^{\mathrm{L}}|, t_{\mathrm{mid}})$. Because at time t_{mid} $\textsc{Alg}$ is at the midpoint between a^{L} and a_{++}^{R}, at this time the tours $\mathrm{move}(\sigma_{++}^{\mathrm{R}}) \oplus \mathrm{move}(\sigma^{\mathrm{L}})$ and $\mathrm{move}(\sigma^{\mathrm{L}}) \oplus \mathrm{move}(\sigma_{++}^{\mathrm{R}})$ incur identical costs for $\textsc{Alg}$ and we have

$$|T^{\textsc{Alg}}| = t_{\mathrm{mid}} + 3\left(\frac{t_{\mathrm{mid}} - 3|a^{\mathrm{L}}|}{2} + |a^{\mathrm{L}}| \right) = \frac{5t_{\mathrm{mid}} - 3|a^{\mathrm{L}}|}{2}.$$

For $\textsc{Alg}$ to not be $(\rho - \varepsilon)$-competitive, we need

$$|T^{\textsc{Alg}}| \geq \rho' \cdot |T^{\textsc{Opt}}| = \rho' \cdot |a_{++}^{\mathrm{R}}| = \rho' \cdot t_{\mathrm{mid}} = \rho' \cdot p_{t_{\mathrm{mid}}},$$

which is equivalent to

$$(5 - 2\rho') \cdot t_{\mathrm{mid}} \geq 3|a^{\mathrm{L}}|. \tag{39}$$

Since the coefficient $5 - 2\rho'$ of t_{mid} is positive, we may assume that t_{mid} is minimal to show (39). By assumption, σ_{+}^{R} is already served at time t_{mid}. Hence, t_{mid} is minimized if, starting at time t_{+}^{R} at $p_{t_{+}^{\mathrm{R}}}$, $\textsc{Alg}$ serves σ_{+}^{R} and then goes to the left, both at full speed. We get

$$|a_{+}^{\mathrm{R}}| - \left(t'' - t_{+}^{\mathrm{R}} - \left| a_{+}^{\mathrm{R}} - p_{t_{+}^{\mathrm{R}}} \right| \right) = p_{t_{\mathrm{mid}}} = M(t_{\mathrm{mid}}) = \frac{t_{\mathrm{mid}} - 3|a^{\mathrm{L}}|}{2}. \tag{40}$$

By Lemma 8, σ^{R} is already served at time $|T_{\mathrm{LR}}| + \frac{W}{\rho'-1}$, and we can thus solve (3) for $p_{t_+^{\mathrm{R}}}$:

$$\begin{aligned}
p_{t_+^{\mathrm{R}}} &= \mathrm{delay}\left(|T_{\mathrm{LR}}| + \frac{W}{\rho'-1}\right) - |a^{\mathrm{L}}| - |T_{\mathrm{LR}}| - \frac{W}{\rho'-1} + t^{\star\star} \\
&= W - 2|a^{\mathrm{L}}| - \frac{W}{\rho'-1} + t^{\star} \\
&= (2\rho'-4)\cdot|a^{\mathrm{L}}| + (\rho'-2)\cdot|a^{\mathrm{R}}| + \frac{\rho'-2}{\rho'-1}\cdot W \qquad (41) \\
&< 0.08\cdot|a^{\mathrm{L}}| + 0.04\cdot|a^{\mathrm{R}}| + \frac{\rho'-2}{\rho'-1}\cdot W \\
&\le |a^{\mathrm{R}}| + \frac{W}{\rho'-1} = a_+^{\mathrm{R}}. \qquad \text{Definition 5 (iv)}
\end{aligned}$$

This implies that ALG is to the left of a_+^{R} at time t_+^{R}. Using this insight and plugging (41) into (40), we obtain

$$\begin{aligned}
t_{\mathrm{mid}} &= \frac{1}{3}\left(3|a^{\mathrm{L}}| + 4|a_+^{\mathrm{R}}| + 2t_+^{\mathrm{R}} - 2p_{t_+^{\mathrm{R}}}\right) \\
&= \frac{1}{3}\left(7|a^{\mathrm{L}}| + 6|a^{\mathrm{R}}| + \frac{6W}{\rho'-1} - 2p_{t_+^{\mathrm{R}}}\right) \\
&= \frac{1}{3}\left((15-4\rho')\cdot|a^{\mathrm{L}}| + (10-2\rho')\cdot|a^{\mathrm{R}}| + \frac{(10-2\rho')\cdot W}{\rho'-1}\right).
\end{aligned}$$

By substituting this into Inequality (39) and noting that it is hardest to satisfy when $W=0$, we get

$$\frac{|a^{\mathrm{L}}|}{|a^{\mathrm{R}}|} \le \underbrace{\frac{4\rho'^2 - 30\rho' + 50}{-8\rho'^2 + 50\rho' - 66}}_{>2.08},$$

which is true due to Definition 5 (iv).

This completes the proof. □

Lemma 9. *Let $i \ge 1$. At time t_i^{L}, ALG is to the right of ℓ_i^{L} and ℓ_i^{R} and crosses one of them after t_i^{L} and before it serves σ_i^{L}. We also have that $t_i^{\mathrm{R}} \le 2t_i^{\mathrm{L}}$.*

Proof. We use induction on i. For the induction base, note that $t_1^{\mathrm{L}} \in [2,\rho')$. Thus $\ell_1^{\mathrm{L}}(t_1^{\mathrm{L}}) = -\ell_1^{\mathrm{R}}(t_1^{\mathrm{L}}) \in [6\rho'-12, \rho'\cdot(3\rho'-6))$. In particular, $\ell_1^{\mathrm{L}}(t_1^{\mathrm{L}})$ and $\ell_1^{\mathrm{R}}(t_1^{\mathrm{L}})$ are both contained is contained in $(-t_1, p_{t_1^{\mathrm{L}}}) = (a_1^{\mathrm{L}}, 1)$. Since the coefficients of t in both $\ell_1^{\mathrm{L}}(t)$ and $\ell_1^{\mathrm{R}}(t)$ are positive, there is no way for ALG to serve σ_1^{L} before crossing the line.

We get an upper bound on t_1^{R} by neglecting the possible intersection of ALG with ℓ_1^{L}. Also note that t_1^{R} is maximized if ALG goes to the right at full speed throughout the interval (t_1, t_1^{R}). In this case,

$$\ell_1^{\mathrm{R}}(t_1^{\mathrm{R}}) = (4-\rho')\cdot t_1^{\mathrm{R}} - (2\rho'-2)\cdot t_1^{\mathrm{L}} = 1 + (t_1^{\mathrm{R}} - t_1^{\mathrm{L}}),$$

and, using $t_1^{\mathrm{L}} \ge 2$,

$$\frac{t_1^{\mathrm{R}}}{t_1^{\mathrm{L}}} = \frac{2\rho'-3}{3-\rho'} + \frac{1}{(3-\rho')t_1^{\mathrm{L}}} \le \frac{4\rho'-5}{6-2\rho'} < 1.63,$$

for $\rho' \in (2,\rho)$, in accordance with our claim.

For the induction step, consider some $i > 1$ and assume the statement was true for all smaller i. Observe that ALG served $\sigma_{i-1}^{\mathrm{R}}$ before $(2\rho'-2)\cdot t_{i-1}^{\mathrm{L}} + (\rho'-2)\cdot t_{i-1}^{\mathrm{R}}$ since otherwise we would have stopped the procedure. Using $t_{i-1}^{\mathrm{R}} \ge t_{i-1}^{\mathrm{L}}$, this implies

$$t_i^{\mathrm{L}} < (3\rho'-4)\cdot t_{i-1}^{\mathrm{R}} = (3\rho'-4)\cdot|a_{i-1}^{\mathrm{R}}| = (3\rho'-4)\cdot p_{t_i^{\mathrm{L}}}.$$

Hence, $\ell_i^{\mathrm{L}}(t_i^{\mathrm{L}}) = -\ell_i^{\mathrm{R}}(t_i^{\mathrm{L}}) = (3\rho' - 6) \cdot t_i^{\mathrm{L}} < (3\rho' - 6)(3\rho' - 4) \cdot p_{t_i^{\mathrm{L}}}$. Consequently $\ell_i^{\mathrm{L}}(t_i^{\mathrm{L}})$ and $\ell_i^{\mathrm{R}}(t_i^{\mathrm{L}})$ are both contained in $(-t_i^{\mathrm{L}}, p_{t_i^{\mathrm{L}}}) = (a_i^{\mathrm{L}}, p_{t_i^{\mathrm{L}}})$, which implies, as before, that ALG crosses ℓ_i^{L} or ℓ_i^{R} before serving σ_i^{L}.

We prove the upper bound on $t_i^{\mathrm{R}}/t_i^{\mathrm{L}}$. First, observe that in step $i-1$, ALG cannot have crossed ℓ_{i-1}^{R} at time t_{i-1}^{R}, because in this case σ_i^{R} can be served no earlier than

$$t_{i-1}^{\mathrm{R}} + a_{i-1}^{\mathrm{R}} - \ell_{i-1}^{\mathrm{R}}(t_{i-1}^{\mathrm{R}}) = (2\rho' - 2) \cdot t_{i-1}^{\mathrm{L}} + (\rho' - 2) \cdot t_{i-1}^{\mathrm{R}},$$

which is the condition of Case 2. Hence,

$$\ell_{i-1}^{\mathrm{R}}(t_{i-1}^{\mathrm{R}}) < \ell_{i-1}^{\mathrm{L}}(t_{i-1}^{\mathrm{R}}) \;\Leftrightarrow\; t_{i-1}^{\mathrm{L}} > \frac{7 - 3\rho'}{3\rho' - 5} \cdot t_{i-1}^{\mathrm{R}}$$

and, using $p_{t_{i-1}^{\mathrm{R}}} = \ell_{i-1}^{\mathrm{L}}(t_{i-1}^{\mathrm{R}})$,

$$\frac{p_{t_{i-1}^{\mathrm{R}}}}{t_{i-1}^{\mathrm{R}}} = (2\rho' - 3) - (3 - \rho')\frac{t_{i-1}^{\mathrm{L}}}{t_{i-1}^{\mathrm{R}}} < (2\rho' - 3) - \frac{(3 - \rho')(7 - 3\rho')}{3\rho' - 5}.$$

This implies

$$\frac{t_i^{\mathrm{L}}}{t_{i-1}^{\mathrm{R}}} \geq \frac{t_{i-1}^{\mathrm{R}} + a_{i-1}^{\mathrm{R}} - p_{t_{i-1}^{\mathrm{R}}}}{t_{i-1}^{\mathrm{R}}} > \frac{-3\rho'^2 + 9\rho' - 4}{3\rho' - 5}. \tag{42}$$

As $t_{i-1}^{\mathrm{R}} = |p_{t_i^{\mathrm{L}}}|$, Inequality (42) gives an upper bound to the quotient $|p_t|/t$ at the time t_i^{L} the request σ_i^{R} is served. This inequality can be interpreted as a bound on the speed of the server and it is essential for the design and proof of Algorithm 2.

To see the claimed upper bound on $t_i^{\mathrm{R}}/t_i^{\mathrm{L}}$, we again neglect the intersection with ℓ_i^{L}, and note that, t_i^{L}, t_i^{R} are maximized if ALG goes to the right at full speed throughout the interval $(t_i^{\mathrm{L}}, t_i^{\mathrm{R}})$. In this case,

$$\begin{aligned}\ell_i^{\mathrm{R}}(t_i^{\mathrm{R}}) &= (4 - \rho') \cdot t_i^{\mathrm{R}} - (2\rho' - 2) \cdot t_i^{\mathrm{L}} \\ &= p_{t_i^{\mathrm{L}}} + (t_i^{\mathrm{R}} - t_i^{\mathrm{L}}) = a_{i-1}^{\mathrm{R}} + (t_i^{\mathrm{R}} - t_i^{\mathrm{L}}) = t_{i-1}^{\mathrm{R}} + (t_i^{\mathrm{R}} - t_i^{\mathrm{L}}),\end{aligned}$$

and, using Inequality (42),

$$\frac{t_i^{\mathrm{R}}}{t_i^{\mathrm{L}}} \leq \frac{-6\rho'^3 + 27\rho'^2 - 32\rho' + 7}{3\rho'^3 - 18\rho'^2 + 31\rho' - 12} < 1.72,$$

for $\rho' \in (2, \rho)$, as claimed. □

Theorem 4. *Let $\rho \approx 2.04$ be the second-largest root (out of the four real roots) of $9\rho^4 - 18\rho^3 - 78\rho^2 + 210\rho - 107$. There is no $(\rho - \varepsilon)$-competitive algorithm for open TSP on the line for any $\varepsilon > 0$.*

Proof. The proof has two parts: We first show that we can indeed apply Lemma 6 in Cases 1 and 2, and then we argue that the procedure terminates.

Consider the step i in which the procedure terminates. In both cases, we apply Lemma 6, that is, we have to show that we have two critical requests. We set $\sigma_0^\star := \sigma_i^{\mathrm{L}}$ and $\sigma_1^\star := \sigma_i^{\mathrm{R}}$, which are obviously of the desired form $(a; |a|)$ for some a and fulfill $\operatorname{sign}(a_0^\star) \neq \operatorname{sign}(a_1^\star)$ as well as $0 < |a_0^\star| \leq |a_1^\star|$. Again let $\sigma_k^\star$ be the request served first by ALG in case no new requests appear. We argue that Properties (i)–(iv) of Definition 5 are fulfilled.

The tours $\operatorname{move}(a_0^\star) \oplus \operatorname{move}(a_1^\star)$ and $\operatorname{move}(a_1^\star) \oplus \operatorname{move}(a_0^\star)$ both serve all requests presented until time t_i^{R} as all these requests have the form $(a, |a|)$ for some a and there are no requests among them outside $[a_0^\star, a_1^\star]$. So Property (i) is fulfilled. Since ℓ_i^{L} has a positive slope and $\ell_i^{\mathrm{L}}(t_i^{\mathrm{L}}) > a_i^{\mathrm{L}}$, Lemma 9 implies that σ_i^{L} is still unserved at time t_i^{R}. From $p_1 \leq 0$ it follows that $p_t < t$ at all times, hence σ_i^{R} also remains unserved at time $t_i^{\mathrm{R}} = |a_i^{\mathrm{R}}|$. In particular $p_{t_i^{\mathrm{R}}} \in (a_i^{\mathrm{L}}, a_i^{\mathrm{R}})$, and Property (ii) is also fulfilled. Property (iv) follows by Lemma 9. To see Property (iii), we make a case distinction:

- Consider Case 1, that is, we have $\sigma_k^\star = \sigma_i^{\mathrm{L}}$. By definition of t_i^{R} we have

$$p_{t_i^{\mathrm{R}}} \geq \ell_i^{\mathrm{L}}(t_i^{\mathrm{R}}) = (2\rho' - 3) \cdot t_i^{\mathrm{R}} + (\rho' - 3) \cdot t_i^{\mathrm{L}} = (2\rho' - 3) \cdot |a_{1-k}^\star| + (\rho' - 3) \cdot |a_k^\star|.$$

With $p_{t_i^{\mathrm{R}}} > a_i^{\mathrm{L}}$, this implies that $\sigma_k^\star$ is served no earlier than

$$t^\star = t_i^{\mathrm{R}} + |a_i^{\mathrm{L}} - p_{t_i^{\mathrm{R}}}| = t_i^{\mathrm{R}} + |a_i^{\mathrm{L}}| + p_{t_i^{\mathrm{R}}} = (2\rho' - 2) \cdot |a_{1-k}^\star| + (\rho' - 2) \cdot |a_k^\star|.$$

- Consider Case 2. We can assume that Case 1 is not fulfilled, meaning that Alg serves $\sigma_i^{\mathrm{R}} = \sigma_k^\star$ first and, since we are in Case 2, not earlier than

$$t^\star = (2\rho' - 2) \cdot t_i^{\mathrm{L}} + (\rho' - 2) \cdot t_i^{\mathrm{R}} = (2\rho' - 2) \cdot |a_{1-k}^\star| + (\rho' - 2) \cdot |a_k^\star|.$$

It remains to show that the procedure terminates. We show this by contradiction. Assume Case 1 and Case 2 both never occur. First observe that, if ℓ_i^{R} is crossed at t_i^{R}, Case 2 is fulfilled as we have

$$t_i^{\mathrm{R}} + a_i^{\mathrm{R}} - \ell_i^{\mathrm{R}}(t_i^{\mathrm{R}}) = (2\rho' - 2) \cdot t_i^{\mathrm{L}} + (\rho' - 2) \cdot t_i^{\mathrm{R}}.$$

Thus we have $t_i^{\mathrm{R}} = \ell_i^{\mathrm{L}}(t_i^{\mathrm{R}})$.

We consider the difference between the release times $t_i^{\mathrm{L}}, t_{i-1}^{\mathrm{R}}$ of the two consecutive requests $\sigma_{i-1}^{\mathrm{R}}, \sigma_i^{\mathrm{L}}$. It is at least the time to move from $\ell_{i-1}^{\mathrm{L}}(t_{i-1}^{\mathrm{R}})$ to a_{i-1}^{R}:

$$t_i^{\mathrm{L}} - t_{i-1}^{\mathrm{R}} \geq a_{i-1}^{\mathrm{R}} - \ell_{i-1}^{\mathrm{L}}(t_{i-1}^{\mathrm{R}}) \Leftrightarrow t_i^{\mathrm{L}} \geq (5 - 2\rho') \cdot t_{i-1}^{\mathrm{R}} + (3 - \rho') \cdot t_{i-1}^{\mathrm{L}}. \tag{43}$$

Similarly, the difference between the release times $t_i^{\mathrm{R}}, t_i^{\mathrm{L}}$ is at least the time to move from a_{i-1}^{R} to $\ell_i^{\mathrm{L}}(t_i^{\mathrm{R}})$. If $\ell_i^{\mathrm{L}}(t_i^{\mathrm{R}}) < a_{i-1}^{\mathrm{R}}$, this is equivalent to the following inequality, which is true otherwise since $t_i^{\mathrm{R}} > t_i^{\mathrm{L}}$:

$$t_i^{\mathrm{R}} - t_i^{\mathrm{L}} \geq a_{i-1}^{\mathrm{R}} - \ell_i^{\mathrm{L}}(t_i^{\mathrm{R}}) \Leftrightarrow (2\rho' - 2) \cdot t_i^{\mathrm{R}} \geq t_{i-1}^{\mathrm{R}} + (4 - \rho') \cdot t_i^{\mathrm{L}}. \tag{44}$$

As Case 2 is not triggered when $\sigma_{i-1}^{\mathrm{R}}$ is served, we have

$$\begin{aligned}
& t_i^{\mathrm{L}} \overset{\neg(\text{Case 2})}{<} (2\rho' - 2) \cdot t_{i-1}^{\mathrm{L}} + (\rho' - 2) \cdot t_{i-1}^{\mathrm{R}} \\
& \quad \overset{(43)}{<} (2\rho' - 2) \cdot t_{i-1}^{\mathrm{L}} + \frac{\rho - 2}{5 - 2\rho'} \cdot (t_i^{\mathrm{L}} - (3 - \rho') \cdot t_{i-1}^{\mathrm{L}}) \\
\Leftrightarrow \quad & t_i^{\mathrm{L}} < \frac{-3\rho'^2 + 9\rho' - 4}{7 - 3\rho'} \cdot t_{i-1}^{\mathrm{L}} \\
\Rightarrow \quad & t_i^{\mathrm{L}} < \left(\frac{-3\rho'^2 + 9\rho' - 4}{7 - 3\rho'} \right)^{i-1} \cdot t_1^{\mathrm{L}}.
\end{aligned}$$

Combing Inequality (43) and the fact that Case 2 does not occur also yields

$$\begin{aligned}
(3 - \rho') \cdot t_i^{\mathrm{L}} + (5 - 2\rho') \cdot t_i^{\mathrm{R}} \overset{(43)}{<} t_{i+1}^{\mathrm{L}} & \overset{\neg(\text{Case 2})}{<} (2\rho' - 2) \cdot t_i^{\mathrm{L}} + (\rho' - 2) \cdot t_i^{\mathrm{R}} \\
\Leftrightarrow \quad t_i^{\mathrm{L}} & > \frac{7 - 3\rho'}{3\rho' - 5} \cdot t_i^{\mathrm{R}}.
\end{aligned} \tag{45}$$

We use this relation between the release time of right and left requests in Inequality (44)

$$\begin{aligned}
& (2\rho' - 2) \cdot t_i^{\mathrm{R}} > t_{i-1}^{\mathrm{R}} + (4 - \rho') \cdot \frac{7 - 3\rho'}{3\rho' - 5} \cdot t_i^{\mathrm{R}} \\
\Leftrightarrow \quad & t_i^{\mathrm{R}} > \frac{3\rho' - 5}{3\rho'^2 + 3\rho' - 18} \cdot t_{i-1}^{\mathrm{R}} \\
\Rightarrow \quad & t_i^{\mathrm{R}} > \left(\frac{3\rho' - 5}{3\rho'^2 + 3\rho' - 18} \right)^i \cdot t_0^{\mathrm{R}}.
\end{aligned}$$

Combining the bounds on t^{L} and t^{R} we get

$$\left(\frac{3\rho'-5}{3\rho'^2+3\rho'-18}\right)^i < t_i^{\mathrm{R}} \overset{(45)}{<} \frac{3\rho'-5}{7-3\rho'} \cdot t_i^{\mathrm{L}} > c \cdot \left(\frac{-3\rho'^2+9\rho'-4}{7-3\rho'}\right)^i.$$

for some constant $c > 1$. This can only be true for all i if

$$\begin{aligned} \frac{3\rho'-5}{3\rho'^2+3\rho'-18} &\leq \frac{-3\rho'^2+9\rho'-4}{7-3\rho'} \\ \Leftrightarrow \quad 9\rho'^4 - 18\rho'^3 - 78\rho'^2 + 210\rho' - 107 &\leq 0 \\ \Leftrightarrow \quad \rho' &\geq \rho, \end{aligned}$$

where we use $\rho' \geq 2$ in the last step. This contradicts our choice of ρ' and thus proves that the request procedure terminates. □

C Proofs of Section 5

Recall that we generally assume without loss of generality that $t \geq |a|$ holds for all request $\sigma = (a;t)$ because the server can not reach σ before time $|a|$ and it only helps the algorithm to know a request earlier. Note that Algorithm 2 is called only if the newly released request is an extreme. By the following lemma, we can assume for the analysis that every request σ is an extreme at its release time.

Lemma 36. *Any request that is not an extreme at its release is served by Algorithm 2 while serving the current extremes.*

Proof. Let $\sigma = (a;t)$ be a request that is not an extreme request at its release. Assume without loss of generality that $a \geq p_t$. As σ is not an extreme, there exists a rightmost extreme $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}(t); t^{\mathrm{R}}(t))$ with $a \leq a^{\mathrm{R}}(t)$. Thus, the server will serve σ while serving $\sigma^{\mathrm{R}}(t)$. □

Moreover, we will frequently use the following lower bounds for an optimal offline algorithm.

Lemma 37. *We have the following lower bounds for the the makespan of an optimal (offline) schedule $|T^{\text{OPT}}|$. If $\sigma_1 = (a_1; t_1)$ and $\sigma_2 = (a_2, t_2)$ are two requests such that $t_1 \leq t_2$, then*

$$|T^{\text{OPT}}| \geq t_1 + |a_1 - a_2| \geq |a_1 - a_2|.$$

Moreover, if $\sigma_1 = (a_1; t_1)$ is a request and p_t denotes the position of the server at an arbitrary time t in a tour computed by Algorithm 2, then

$$|T^{\text{OPT}}| \geq |p_t - a_1|.$$

Proof. The first inequality follows from the fact that OPT needs to serve both σ_1 and σ_2. For the second inequality, let $A^{\mathrm{R}}(t)$ be the rightmost point of a request seen until time t and $A^{\mathrm{L}}(t)$ be the leftmost point of a request seen until time t. The server following the tour computed by Algorithm 2 never moves left of $A^{\mathrm{L}}(t)$ or right of $A^{\mathrm{R}}(t)$, thus $p_t \in \left[A^{\mathrm{L}}(t), A^{\mathrm{R}}(t)\right]$ and $|T^{\text{OPT}}| \geq |A^{\mathrm{L}}(t) - A^{\mathrm{R}}(t)| \geq |p_t - a_1|$. □

First, we proof one of the central properties of Algorithm 2.

Lemma 11. *The position of the server in a tour computed by Algorithm 2 satisfies*

$$\frac{|p_t|}{t} \leq \frac{3\rho - 5}{-3\rho^2 + 9\rho - 4} \approx 0.58$$

for all times $t \geq 0$.

Proof. Observe that the server in a tour computed by Algorithm 2 always either moves with unit speed or it waits at the origin. When the server moves with unit speed towards the origin, then the ratio $|p_t|/t$ is decreasing and when it moves away from the origin with unit speed towards a request, then the ratio $|p_t|/t$ is increasing (unless it is 1). The ratio $|p_t|/t$ is therefore maximized at a time t when a request is served and the server moves back to the origin afterwards. In order to show the claim it is thus sufficient to consider only times when an extreme request is served. We distinguish between the cases that occurred when the tour serving the extreme request was computed by Algorithm 2.

1. *Cases (P1) and (E1)*
   The server follows the preferred or enforced tour serving a single extreme request $\sigma = (a_1; t_1)$. Let t be the time this request is served. In Case (P1) the server arrives at a_1 exactly at time ρt_1 and in Case (E1) it arrives after time ρt_1. Thus, we have $t \geq \rho t_1$ and using $\rho > 2$ and $a_1 \leq t_1$, we obtain
   $$\frac{p_t}{t} \leq \frac{a_1}{\rho t_1} \leq \frac{a_1}{\rho a_1} = \frac{1}{\rho} < \frac{1}{2} < 0.583.$$
2. *Case (O)*
   The server follows the tour defined according to Case (O). As long as the server moves towards the request $\sigma_1 = (a_1; t_1)$ that is closer to the origin the ratio $|p_t|/t$ decreases. The request $\sigma_2 = (a_2, t_2)$ is not served until time ρt_2. We again apply $\rho > 2$ and $a_2 \leq t_2$ and get
   $$\frac{p_{\rho t_2}}{\rho t_2} = \frac{a_2}{\rho t_2} \leq \frac{a_2}{\rho a_2} = \frac{1}{\rho} < \frac{1}{2} < 0.583.$$
3. *Cases (P2) and (E2)*
   The algorithm follows the preferred or enforced tour serving first $\sigma_1 = (a_1; t_1)$ and then $\sigma = (a_2; t_2)$. Let t be the time when the server arrives at a_1. In Case (P2) we have $t = L^{\sigma_1,\sigma_2}$ and in Case (E2) we have $t > L^{\sigma_1,\sigma_2}$. Using $|a_1| \leq t_1 \leq t_2$, we obtain
   $$\begin{aligned} t &\geq L^{\sigma_1,\sigma_2} \\ &= \min\{\rho t_1 + (\rho-1)|a_1| + (\rho-1)|a_2|,\ \rho t_2 + (\rho-2)|a_1| + (\rho-2)|a_2|\}, \\ &\geq (2\rho-2)|a_1|. \end{aligned}$$
   Thus, $|p_t|/t \leq 1/(2\rho-2) < 1/2 < 0.583$ holds, which implies the claim for σ_1.
   After serving σ_1, the server moves to a_2 to serve σ_2. Let t' be the time the server arrives at a_2. We have $|a_2| \leq t_2$ and therefore
   $$\begin{aligned} t' &\geq L^{\sigma_1,\sigma_1} + |a_1| + |a_2| \\ &= \min\{\rho t_1 + \rho|a_1| + \rho|a_2|, \rho t_2 + (\rho-1)|a_1| + (\rho-1)|a_2|\} \\ &\geq \rho|a_2|. \end{aligned}$$
   We get $|p_{t'}|/t' \leq 1/\rho < 1/2 < 0.583$ as claimed.
4. *Case (A2)*
   In the anticipating tour, Algorithm 2 serves $\sigma_2 = (a_2; t_2)$ before $\sigma_1 = (a_1; t_1)$. Moreover, as the conditions for the anticipating tour are satisfied, we also have
   $$|a_2| \leq \frac{3\rho-5}{(2\rho-2)(7-3\rho)} \left(\rho t_1 + (\rho-2)|a_1|\right). \tag{46}$$
   Let t be the time when the server arrives at a_2. By using Inequality (46) and $t_1 \leq t_2$, we obtain
   $$\begin{aligned} t &= L^{\sigma_2,\sigma_1} \\ &\geq \min\{\rho t_2 + (\rho-1)|a_2| + (\rho-1)|a_1|, \rho t_1 + (\rho-2)|a_2| + (\rho-2)|a_1|\} \\ &= \rho t_1 + (\rho-2)|a_2| + (\rho-2)|a_1| \\ &\geq |a_2| \left((\rho-2) + \frac{(2\rho-2)(7-3\rho)}{3\rho-5}\right) \\ &= |a_2| \frac{-3\rho^2+9\rho-4}{3\rho-5}. \end{aligned}$$

Thus, the inequality in our claim is satisfied with equality.

Now, let t' be the time the server arrives at a_1 after serving σ_2. Using $t_1 \geq a_1$, we obtain

$$\begin{aligned} t' &= L^{\sigma_2,\sigma_1} + |a_2| + |a_1| \\ &= \rho t_1 + (\rho-1)|a_2| + (\rho-1)|a_1| \\ &\geq (2\rho-1)|a_1|. \end{aligned}$$

Therefore, $|p_{t'}|/t' \leq 1/(2\rho-2) < 1/2 < 0.583$ holds, as claimed.

□

The main result, the ρ-competitiveness of Algorithm 2, is shown in the following theorem.

Theorem 10. *Algorithm 2 is ρ-competitive with $\rho \approx 2.04$ being the second-largest root of the polynomial $9\rho^4 - 18\rho^3 - 78\rho^2 + 210\rho - 107$.*

Proof. At time t a new request is released and Algorithm 2 computes a new tour T^{ALG}. We show that this new tour is ρ-competitive, i.e., $|T^{\text{ALG}}| \leq \rho|T^{\text{OPT}}|$. In the proof we first distinguish between the cases that can occur at time t and also in previous calls of Algorithm 2.

Without loss of generality, we assume that the new extreme request released at time t is a rightmost request denoted by $\sigma^{\text{R}}(t) = (a^{\text{R}}(t); t)$. If at the call of Algorithm 2 at time t Case (P1), (P2) or (A2) occurs, i.e., the newly computed tour T^{ALG} is preferred or anticipating, we know that T^{ALG} is ρ-competitive. This is because the request served last will be served at time ρt at the latest and obviously $|T^{\text{OPT}}| \geq t$. Thus, we analyze the other cases that can occur at time t. If at time t a leftmost extreme $\sigma^{\text{L}}(t)$ exists, we denote its release time by $t^{\text{L}} := t^{\text{L}}(t)$.

1. *Case (E1) at time t*
   In this case, we have one extreme request $\sigma^{\text{R}}(t) = (a^{\text{R}}(t); t)$ and the server immediately moves towards $\sigma^{\text{R}}(t)$ as $t + |p_t - a^{\text{R}}(t)| > \rho t$. It holds that $|T^{\text{ALG}}| = t + |p_t - a^{\text{R}}(t)|$. By Lemma 37, we have $|T^{\text{OPT}}| \geq t$ as well as $|T^{\text{OPT}}| \geq |p_t - a^{\text{R}}(t)|$. Thus, $|T^{\text{ALG}}| = t + |p_t - a^{\text{R}}(t)| \leq 2|T^{\text{OPT}}| < \rho|T^{\text{OPT}}|$ holds, as desired.

2. *Case (O) at time t*
   This means that two extremes $\sigma^{\text{L}}(t)$ and $\sigma^{\text{R}}(t)$ are present such that $0 \notin [a^{\text{L}}(t), a^{\text{R}}(t)]$. Recall that the position of the server is strictly between both extremes because $a^{\text{R}}(t)$ is right of the server and $a^{\text{L}}(t)$ left of the server by definition and both requests are unserved. If there is no unserved request on one side of the server, we would be in Case (P1) or (E1).

   If after serving σ_1 the server follows the tour T_0 for some time, then it serves σ_2 by time ρt_2 and the tour computed by Algorithm 2 is ρ-competitive. So from now on we assume that the server moves towards σ_1 and then σ_2 without interruption. We distinguish the two cases $|a^{\text{R}}(t)| < |a^{\text{L}}(t)|$ and $|a^{\text{R}}(t)| > |a^{\text{L}}(t)|$.

   We first consider the case where $|a^{\text{R}}(t)| < |a^{\text{L}}(t)|$. As we are in the Case (O), this implies that $a^{\text{R}}(t), a^{\text{L}}(t) < 0$. The tour in this case chooses to serve $\sigma^{\text{R}}(t)$ first. Using $|a^{\text{L}}(t) - a^{\text{R}}(t)| \leq |a^{\text{L}}(t)| \leq t^{\text{L}}$, $p_t \in (a^{\text{L}}(t), a^{\text{R}}(t))$ and the bounds $|T^{\text{OPT}}| \geq t$ and $|a^{\text{L}}(t) - a^{\text{R}}(t)| + t^{\text{L}} \leq |T^{\text{OPT}}|$ by Lemma 37 ($t \geq t^{\text{L}}$), we obtain

$$\begin{aligned} |T^{\text{ALG}}| &= t + |p_t - a^{\text{R}}(t)| + |a^{\text{R}}(t) - a^{\text{L}}(t)| \\ &\leq t + |p_t - a^{\text{R}}(t)| + t^{\text{L}} \\ &\leq t + |a^{\text{L}}(t) - a^{\text{R}}(t)| + t^{\text{L}} \\ &\leq 2|T^{\text{OPT}}|. \end{aligned}$$

   Now assume $|a^{\text{R}}(t)| > |a^{\text{L}}(t)|$ and hence $a^{\text{L}}(t), a^{\text{R}}(t) > 0$. In this case $\sigma^{\text{L}}(t)$ is served first and we obtain

$$|T^{\text{ALG}}| = t + |p_t - a^{\text{L}}(t)| + |a^{\text{L}}(t) - a^{\text{R}}(t)|.$$

   Recall that $p_t \in [a^{\text{L}}(t), a^{\text{R}}(t)]$. If $p_t = a^{\text{L}}_t$ this implies that $|T^{\text{ALG}}| = t + |a^{\text{L}}(t) - a^{\text{R}}(t)| \leq 2|T^{\text{OPT}}|$. If $p_t > a^{\text{L}}(t)$, we have to take the situation during the time interval $[t^{\text{L}}, t]$ into account. As $\sigma^{\text{L}}(t)$

was released at t^{L}, only rightmost extremes are released during $(t^{\mathrm{L}}, t]$, which means that Case (O) occurs whenever a new request is released during this time frame since $a^{\mathrm{L}}(t) > 0$. At time t^{L} only the Cases (P1), (E1) or (O) can occur, again because $a^{\mathrm{L}}(t) > 0$. In all these cases the server will immediately move towards a^{L} or 0, i.e., to the left, and $p_t > a^{\mathrm{L}}(t)$ means that the server does not reach $a^{\mathrm{L}}(t)$ by time t. Thus, the server moves uninterrupted to the left during $[t^{\mathrm{L}}, t]$. Applying Lemma 37, we get

$$|T^{\mathrm{ALG}}| = t^{\mathrm{L}} + |p_{t^{\mathrm{L}}} - a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)| \leq 2|T^{\mathrm{OPT}}|$$

3. *Case (E2) at time t*
   As Case (E2) occurred, we have $a^{\mathrm{L}}(t) \leq 0$ and $a^{\mathrm{R}}(t) \geq 0$. Hence, from time t on the tour computed by Algorithm 2 serves $\sigma^{\mathrm{L}}(t)$ and $\sigma^{\mathrm{R}}(t)$ without interruption in this order provided no new requests are released after time t. If $p_t \leq 0$, we can apply Lemma 37 and obtain

$$|T^{\mathrm{ALG}}| \leq t + |a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)| \leq t + (t^{\mathrm{L}} + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)|) \leq 2|T^{\mathrm{OPT}}|.$$

   So from now on we assume that $p_t > 0$. In this case we get

$$|T^{\mathrm{ALG}}| = t + |p_t| + |a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)|.$$

   It is not possible to directly show ρ-competitiveness in this case. Instead we need to take the behavior of the server during times $[t^{\mathrm{L}}, t]$ into account. We distinguish two cases. Either the server moves to the left towards $\sigma^{\mathrm{L}}(t)$ without interruption during this interval or the server gets delayed because it waits at the origin or moves to the right for some time. Let $A^{\mathrm{R}}(t)$ be the rightmost point of a request seen until time t. In the first case, we know that $0 < |p_t| \leq |p_{t^{\mathrm{L}}}| \leq A^{\mathrm{R}}(t)$ as the server moves left during $[t^{\mathrm{L}}, t]$. Applying Lemma 37 and and using $|p_{t^{\mathrm{L}}}| \leq |A^{\mathrm{R}}(t)|$, we obtain

$$\begin{aligned} |T^{\mathrm{ALG}}| &\leq t^{\mathrm{L}} + |p_{t^{\mathrm{L}}}| + |a^{\mathrm{L}}(t)| + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)| \\ &\leq (|A^{\mathrm{R}}(t)| + |a^{\mathrm{L}}(t)|) + (t^{\mathrm{L}} + |a^{\mathrm{L}}(t) - a^{\mathrm{R}}(t)|) \leq 2|T^{\mathrm{OPT}}|. \end{aligned}$$

   We now consider the case where the server does not move towards $\sigma^{\mathrm{L}}(t)$ without interruption during $[t^{\mathrm{L}}, t]$. For this to happen, there has to be a call of the algorithm Algorithm 2 at a time $\hat{t} \in [t^{\mathrm{L}}, t)$ such that the tour followed by the server in $[\hat{t}, t]$ waits at the origin or moves to the right for some time. We choose $\hat{t}$ maximal with this property. By Lemma 38, only the Cases (A2), (P2) and (E2) can occur at time $\hat{t}$. The Cases (A2) and (P2) are both treated in Lemma 40. Finally, Case (E2) is treated in Lemma 41. This concludes the proof. ☐

We still have to show ρ-competitiveness of Algorithm 2 in the following setting.

Setting C. Consider the following setting of requests of an instance for Algorithm 2.

- At time t the request $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}(t), t)$ appears with $a^{\mathrm{R}}(t) \geq 0$ and Case (E2) occurs.
- At time t there is a leftmost extreme $\sigma^{\mathrm{L}}(t) = (a^{\mathrm{L}}(t), t^{\mathrm{L}})$ present with $a^{\mathrm{L}}(t) \leq 0$.
- The position of the server satisfies $p_t > 0$.
- Time $\hat{t} \in [t^{\mathrm{L}}, t)$ is the last time Algorithm 2 is called such that the tour followed by the server in $[\hat{t}, t]$ waits at the origin or moves to the right for some time, i.e., is not identical to "move($a^{\mathrm{L}}(t)$)".

In the next lemma, we limit the cases that can occur at time $\hat{t}$ in the setting above.

Lemma 38. *In Setting C, the cases that can occur when Algorithm 2 is called at time $\hat{t}$ are Case (A2) when $\hat{t} > t^{\mathrm{L}}$, and Cases (P2) and (E2) when $\hat{t} = t^{\mathrm{L}}$.*

Proof. We first establish the following two claims:

- *The tour computed at $\hat{t}$ does not serve the extreme request $\sigma^{\mathrm{L}}(t)$ first.*
  Assume, on the contrary, that $\sigma^{\mathrm{L}}(t)$ is served first. Note that by the choice of $\hat{t}$ any tour computed after $\hat{t}$ immediately moves left. If $p_{\hat{t}} \leq 0$, then also $p_t \leq 0$ holds because the computed tour does not move right of the origin before serving $\sigma^{\mathrm{L}}(t)$ in this case. But this contradicts $p_t > 0$. If $p_{\hat{t}} > 0$, then the server will move towards the origin first. If it waits there, we have $p_t \leq 0$, which is again a contradiction. If it does not wait at the origin, the server moves without interruption towards $a^{\mathrm{L}}(t)$ during $[\hat{t}, t]$, contradicting our assumptions about $\hat{t}$ in Setting C.

- *Case (O) cannot occur.*
  As $a^{\mathrm{L}}(t) \leq 0$ holds, Case (O) would imply that $p_{\hat{t}} \leq \sigma^{\mathrm{R}}(\hat{t}) < 0$. But this contradicts $p_t > 0$ because the tour computed at $\hat{t}$ only moves to $\sigma^{\mathrm{R}}(\hat{t})$ and then left and any tour computed after $\hat{t}$ immediately moves to the left by the choice of $\hat{t}$.

Now, consider the case $t^{\mathrm{L}} < \hat{t} < t$. This means that ALG is forced to divert from the direct tour towards $\sigma^{\mathrm{L}}(t)$ caused by the release of a rightmost extreme at time $\hat{t}$. The release of a leftmost extreme would contradict that $\sigma(t)$ is still the leftmost extreme at time t. The Cases (P1) and (E1) cannot occur at time $\hat{t}$, because two different extremes exists at this point in time. By the first claim above, (P2) and (E2) cannot occur and by the second claim, Case (O) cannot occur. This leaves Case (A2) as the only possibility.

Next, let us consider the case $\hat{t} = t^{\mathrm{L}}$. This means that Algorithm 2 computes a tour at the release of $\sigma^{\mathrm{L}}(t)$ at time t^{L} that waits at the origin or moves to the right for some time. By the first claim above, Case (P1), (E1) and (A2) cannot occur. Moreover, Case (O) cannot occur by the second claim. Thus, only Case (P2) and Case (E2) are possible for $\hat{t} = t^{\mathrm{L}}$. □

Later, we will also consider the following analogous version of the setting above, where the role of the leftmost and rightmost extreme are exchanged, the server is on the other side of the origin and we again consider the time interval between the release times of both extremes.

Setting D. Consider the following setting of requests of an instance for Algorithm 2.

- At time t^{L} the request $\sigma^{\mathrm{L}}(t) = (a^{\mathrm{L}}(t), t^{\mathrm{L}})$ appears with $a^{\mathrm{L}}(t) \leq 0$ and Case (E2) occurs.
- At time t^{L} there is a rightmost extreme $\sigma^{\mathrm{R}}(t^{\mathrm{L}}) = (a^{\mathrm{R}}(t^{\mathrm{L}}), t^{\mathrm{R}})$ present with $a^{\mathrm{R}}(t^{\mathrm{L}}) \geq 0$.
- The position of the server satisfies $p_{t^{\mathrm{L}}} < 0$.
- Time $\tilde{t} \in [t^{\mathrm{R}}, t^{\mathrm{L}})$ is the last time Algorithm 2 is called such that the tour followed by the server in $[\tilde{t}, t^{\mathrm{L}}]$ waits at the origin or moves to the left for some time, i.e., is not identical to "move($a^{\mathrm{R}}(t^{\mathrm{L}})$)".

With a completely analogous proof, we obtain the same result as in Lemma 38 for the setting above.

Corollary 39. *In Setting D, the cases that can occur when Algorithm 2 is called at time $\tilde{t}$ are Case (A2) when $\tilde{t} > t^{\mathrm{R}}$, and Cases (P2) and (E2) when $\tilde{t} = t^{\mathrm{R}}$.*

The following lemma now treats the Cases (P2) and (A2) at $\hat{t}$.

Lemma 40. *In Setting C, if Case (P2) or Case (A2) occur at time $\hat{t}$, no additional requests appear after $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}(t), t)$ and the tour computed by Algorithm 2 at $\hat{t}$ serves $\sigma^{\mathrm{R}}(\hat{t})$ before $\sigma^{\mathrm{L}}(t)$, then $|T^{\mathrm{ALG}}| \leq \rho |T^{OPT}|$.*

Proof. By the choice of $\hat{t}$, we know that the server moves immediately to $a^{\mathrm{L}}(t)$ without waiting if any new extreme arrives in $(\hat{t}, t]$. From then on it will serve $\sigma^{\mathrm{L}}(t)$ and afterwards $\sigma^{\mathrm{R}}(t)$ without interruption because Case (E2) occurs at time t. Thus, the latest point in time after $\hat{t}$ at which the server starts moving to the left is when it reaches $a^{\mathrm{R}}(\hat{t})$. Hence, we have $p_t \leq a^{\mathrm{R}}(\hat{t})$ and additionally using $p_t > 0$ we obtain

$$\begin{aligned} |T^{\mathrm{ALG}}| &\leq t + |p_t| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\ &\leq t + |a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|. \end{aligned} \tag{47}$$

By the definition of Cases (P2) and (A2), the server reaches $a^{\mathrm{R}}(\hat{t})$ at the latest at time $L^{\sigma^{\mathrm{R}}(\hat{t}),\sigma^{\mathrm{L}}(t)}$. From then on it will serve $\sigma^{\mathrm{L}}(t)$ and afterwards $\sigma^{\mathrm{R}}(t)$ without interruption. Using $L^{\sigma^{\mathrm{R}}(\hat{t}),\sigma^{\mathrm{L}}(t)} \leq \rho t^{\mathrm{L}} + (\rho-2)a^{\mathrm{R}}(\hat{t}) + (\rho-2)a^{\mathrm{L}}(t)$, we also have the following bound:

$$\begin{aligned}|T^{\mathrm{ALG}}| &\leq L^{\sigma^{\mathrm{R}}(\hat{t}),\sigma^{\mathrm{L}}(t)} + |a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\ &\leq \rho t^{\mathrm{L}} + (\rho-1)|a^{\mathrm{R}}(\hat{t})| + \rho|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|.\end{aligned}$$

For $|a^{\mathrm{R}}(\hat{t})| \leq |a^{\mathrm{R}}(t)|$, we immediately obtain $|T^{\mathrm{ALG}}| \leq \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|) \leq \rho|T^{\mathrm{OPT}}|$. Otherwise, we have $|a^{\mathrm{R}}(t)| < |a^{\mathrm{R}}(\hat{t})|$ and get

$$|T^{\mathrm{ALG}}| \leq \rho(t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(\hat{t})|). \tag{48}$$

In order to show ρ-competitiveness in this case we consider the six possible orders in which the requests $\sigma^{\mathrm{L}}(t)$, $\sigma^{\mathrm{R}}(\hat{t})$ and $\sigma^{\mathrm{R}}(t)$ can be served in an optimal tour.

- In the three cases where the optimal tour serves $\sigma^{\mathrm{L}}(t)$ before $\sigma^{\mathrm{R}}(\hat{t})$, we have $|T^{\mathrm{OPT}}| \geq t^{\mathrm{L}} + |a^{\mathrm{R}}(\hat{t})| + |a^{\mathrm{L}}(t)|$ and Inequality (48) guarantees ρ-competitiveness.
- If the optimal tour serves $\sigma^{\mathrm{R}}(t)$ before $\sigma^{\mathrm{L}}(t)$, we obtain $|T^{\mathrm{OPT}}| \geq t + |a^{\mathrm{R}}(t)| + |a^{\mathrm{L}}(t)|$. By the bound in (47), we obtain

$$\begin{aligned}|T^{\mathrm{ALG}}| &\leq t + |a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\ &= (t + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|) + (|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(\hat{t})|) \leq 2|T^{\mathrm{OPT}}|.\end{aligned}$$

- In the remaining case where the optimal server uses the order $\sigma^{\mathrm{R}}(\hat{t}), \sigma^{\mathrm{L}}(t), \sigma^{\mathrm{R}}(t)$ we have that $|T^{\mathrm{OPT}}| \geq 2|a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|$. Using (47) together with $|T^{\mathrm{OPT}}| \geq t$, we can conclude that

$$\begin{aligned}|T^{\mathrm{ALG}}| &\leq t + |a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\ &\leq t + (2|a^{\mathrm{R}}(\hat{t})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|) \leq 2|T^{\mathrm{OPT}}|.\end{aligned}$$

□

Finally, the following lemma treats the remaining Case (E2) at time $\hat{t}$.

Lemma 41. *In Setting C, if Case (E2) occurs at time $\hat{t} = t^{\mathrm{L}}$ and no additional requests appear after $\sigma^{\mathrm{R}}(t) = (a^{\mathrm{R}}(t), t)$, then $|T^{\mathrm{ALG}}| \leq \rho|T^{OPT}|$.*

Proof. By assumption Case (E2) occurs at time t^{L} and the server immediately moves towards $a^{\mathrm{R}}(t^{\mathrm{L}})$. It either serves $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ and then moves left to serve $\sigma^{\mathrm{L}}(t)$ or it moves to $\sigma^{\mathrm{L}}(t)$ earlier when a new extreme arrives in $(\hat{t}, t]$. By time t the server has not reached $a^{\mathrm{L}}(t)$ as $\sigma^{\mathrm{L}}(t)$ is still an extreme request. Now Case (E2) occurs and the server serves $\sigma^{\mathrm{L}}(t)$ and afterwards $\sigma^{\mathrm{R}}(t)$ without interruption. We denote the release time of $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ by t^{R}. We first consider the case where $p_{t^{\mathrm{L}}} \geq 0$. Using $0 \leq p_{t^{\mathrm{L}}} \leq |a^{\mathrm{R}}(t^{\mathrm{L}})| \leq t^{\mathrm{R}}$ and applying Lemma 37 twice, we obtain

$$\begin{aligned}|T^{\mathrm{ALG}}| &\leq t^{\mathrm{L}} + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \\ &\leq (t^{\mathrm{R}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|) + (t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|) \leq 2|T^{\mathrm{OPT}}|.\end{aligned}$$

So from now on we assume $p_{t^{\mathrm{L}}} < 0$. In this case we have

$$\begin{aligned}|T^{\mathrm{ALG}}| &\leq t^{\mathrm{L}} + |p_{t^{\mathrm{L}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)| \qquad (49)\\ &= (t^{\mathrm{L}} + |a^{\mathrm{L}}(t)| + |a^{\mathrm{R}}(t)|) + (|p_{t^{\mathrm{L}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|) \\ &\leq |T^{\mathrm{OPT}}| + (|p_{t^{\mathrm{L}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|). \qquad (50)\end{aligned}$$

The remaining proof of the lemma proceeds along the following key claims. Note that Claim 2 and Claim 3 together show that Algorithm 2 is ρ-competitive because either Inequality (52) or Inequality (53) is satisfied.

Claim 1: We can assume the following lower bound for OPT (otherwise Algorithm 2 is ρ-competitive)

$$|T^{\text{OPT}}| \geq t^{\text{L}} + |a^{\text{L}}(t)| + \max(|a^{\text{R}}(t^{\text{L}})|, |a^{\text{R}}(t)|). \tag{51}$$

Claim 2: We can show $|T^{\text{ALG}}| \leq \rho |T^{\text{OPT}}|$ provided that

$$|a^{\text{L}}(t)| \leq \frac{3\rho - 5}{(2\rho - 2)(7 - 3\rho)}(\rho t^{\text{R}} + (\rho - 2)|a^{\text{R}}(t^{\text{L}})|). \tag{52}$$

Claim 3: We can show $|T^{\text{ALG}}| \leq \rho |T^{\text{OPT}}|$ provided that

$$|a^{\text{L}}(t)| > \frac{3\rho - 5}{(2\rho - 2)(7 - 3\rho)}(\rho t^{\text{R}} + (\rho - 2)|a^{\text{R}}(t^{\text{L}})|). \tag{53}$$

Claim 1. Assume at first that $|a^{\text{R}}(t^{\text{L}})| > |a^{\text{R}}(t)|$. In this case we know that by time t Algorithm 2 has served $a^{\text{R}}(t^{\text{L}})$ since at that time $\sigma^{\text{R}}(t)$ is the extreme rightmost of p_t rather than $\sigma^{\text{R}}(t^{\text{L}})$. As Case (E2) occurs at time t, we have $p_t \in [a^{\text{L}}(t), a^{\text{R}}(t)]$ and therefore $t \geq t^{\text{L}} + |p_{t^{\text{L}}}| + |a^{\text{R}}(t^{\text{L}})| + (|a^{\text{R}}(t^{\text{L}})| - |a^{\text{R}}(t)|)$. Using this inequality and $p_{t^{\text{L}}} < 0$, we obtain

$$\begin{aligned}
|T^{\text{ALG}}| &\overset{(49)}{\leq} t^{\text{L}} + |p_{t^{\text{L}}}| + 2|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \\
&= t^{\text{L}} + |p_{t^{\text{L}}}| + |a^{\text{R}}(t^{\text{L}})| + (|a^{\text{R}}(t^{\text{L}})| - |a^{\text{R}}(t)|) + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \\
&\leq t + 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| \\
&= (t + |a^{\text{R}}(t)| + |a^{\text{L}}(t)|) + (|a^{\text{L}}(t)| + |a^{\text{R}}(t)|) \\
&\leq (t + |a^{\text{R}}(t)| + |a^{\text{L}}(t)|) + |T^{\text{OPT}}|.
\end{aligned}$$

If the optimal tour serves $\sigma^{\text{R}}(t)$ before $\sigma^{\text{L}}(t)$, we have $|T^{\text{OPT}}| \geq t + |a^{\text{L}}(t)| + |a^{\text{R}}(t)|$ which implies that $|T^{\text{ALG}}| \leq 2|T^{\text{OPT}}|$. If on the other hand the optimal tour serves the requests in the order $\sigma^{\text{R}}(t^{\text{L}}), \sigma^{\text{L}}(t), \sigma^{\text{R}}(t)$, we know that $|T^{\text{OPT}}| \geq 2|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \geq 2|a^{\text{L}}(t)| + 2|a^{\text{R}}(t)| + |a^{\text{R}}(t^{\text{L}})|$. This implies that

$$|T^{\text{ALG}}| \leq t + (2|a^{\text{R}}(t)| + 2|a^{\text{L}}(t)|) \leq 2|T^{\text{OPT}}|.$$

Thus, we assume from now on that if $|a^{\text{R}}(t^{\text{L}})| > |a^{\text{R}}(t)|$ holds, then in the optimal tour $\sigma^{\text{L}}(t)$ is served before $\sigma^{\text{R}}(t^{\text{L}})$. Of course it might still be the case that $|a^{\text{R}}(t^{\text{L}})| \leq |a^{\text{R}}(t)|$, however in both cases we obtain the claimed lower bound $|T^{\text{OPT}}| \geq t^{\text{L}} + |a^{\text{L}}(t)| + \max(|a^{\text{R}}(t^{\text{L}})|, |a^{\text{R}}(t)|)$.

Claim 2. Using $t^{\text{R}} \geq a^{\text{R}}(t^{\text{L}})$, the assumption implies

$$|a^{\text{L}}(t)| \leq \frac{3\rho - 5}{(2\rho - 2)(7 - 3\rho)}(\rho t^{\text{R}} + (\rho - 2)|a^{\text{R}}(t^{\text{L}})|) \leq \frac{3\rho - 5}{7 - 3\rho} t^{\text{R}}. \tag{54}$$

Note that the assumed inequality is exactly the second condition for Case (A2) at time t^{L}. Because Case (E2) occurs at time t^{L} by assumption, we know that the first condition of Case (A2) does not hold, i.e., $t^{\text{L}} + |p_{t^{\text{L}}} - a^{\text{L}}(t)| > L^{\sigma^{\text{L}}(t), \sigma^{\text{R}}(t^{\text{L}})}$. We have $t^{\text{R}} \leq t^{\text{L}}$ and thus $L^{\sigma^{\text{L}}(t), \sigma^{\text{R}}(t^{\text{L}})} = \rho t^{\text{R}} + (\rho - 2)|a^{\text{L}}(t)| + (\rho - 2)|a^{\text{R}}(t^{\text{L}})|$. Moreover, $|p_{t^{\text{L}}} - a^{\text{L}}(t)| = |a^{\text{L}}(t)| - |p_{t^{\text{L}}}|$ holds because $p_{t^{\text{L}}} < 0$. Hence, we obtain

$$\begin{aligned}
|p_{t^{\text{L}}}| &< t^{\text{L}} + |a^{\text{L}}(t)| - L^{\sigma^{\text{L}}(t), \sigma^{\text{R}}(t^{\text{L}})} \\
&= t^{\text{L}} + (3 - \rho)|a^{\text{L}}(t)| + (2 - \rho)|a^{\text{R}}(t^{\text{L}})| - \rho t^{\text{R}}
\end{aligned} \tag{55}$$

Using Inequality (55), $2t^{\text{L}} \leq \rho t^{\text{L}} + (2 - \rho)|a^{\text{L}}(t)|$ and Inequality (54), we obtain

$$\begin{aligned}
|T^{\text{ALG}}| &\overset{(49)}{\leq} t^{\text{L}} + |p_{t^{\text{L}}}| + 2|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \\
&\overset{(55)}{\leq} 2t^{\text{L}} + (5 - \rho)|a^{\text{L}}(t)| + (4 - \rho)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| - \rho t^{\text{R}} \\
&\leq \rho t^{\text{L}} + (7 - 3\rho)|a^{\text{L}}(t)| + \rho|a^{\text{L}}(t)| + (4 - \rho)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| - \rho t^{\text{R}} \\
&\overset{(54)}{\leq} \rho t^{\text{L}} + (3\rho - 5)t^{\text{R}} + \rho|a^{\text{L}}(t)| + (4 - \rho)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| - \rho t^{\text{R}} \\
&= \rho t^{\text{L}} - (5 - 2\rho)t^{\text{R}} + \rho|a^{\text{L}}(t)| + (4 - \rho)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| \\
&\leq \rho t^{\text{L}} - (5 - 2\rho)|a^{\text{R}}(t^{\text{L}})| + \rho|a^{\text{L}}(t)| + (4 - \rho)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| \\
&\leq \rho t^{\text{L}} + (\rho - 1)|a^{\text{R}}(t^{\text{L}})| + \rho|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \\
&\leq \rho\left(t^{\text{L}} + |a^{\text{L}}(t)| + \max(|a^{\text{R}}(t^{\text{L}})|, |a^{\text{R}}(t)|)\right) \overset{(51)}{\leq} \rho|T^{\text{OPT}}|.
\end{aligned}$$

Claim 3. In order to show $|T^{\text{ALG}}| \leq \rho|T^{\text{OPT}}|$ in this case, we have to take the situation during $[t^{\text{R}}, t^{\text{L}}]$ into account. Let us first assume that during this interval Algorithm 2 moves towards $\sigma^{\text{R}}(t^{\text{L}})$ without interruption. In particular, this implies that $p_{t^{\text{R}}} < p_{t^{\text{L}}} < 0$. Let us additionally assume that the following inequality holds

$$|p_{t^{\text{R}}}| \leq (2\rho - 2)|a^{\text{L}}(t)| - (3 - \rho)|a^{\text{R}}(t^{\text{L}})| - t^{\text{R}}. \tag{56}$$

Recall that at time t^{L} Case (E2) occurs, $\sigma^{\text{R}}(t^{\text{L}})$ is still the rightmost extreme and the server continues moving towards $\sigma^{\text{R}}(t^{\text{L}})$. The server then moves towards $\sigma^{\text{L}}(t)$ at the latest by the time it reaches $\sigma^{\text{R}}(t^{\text{L}})$. From then on it serves $\sigma^{\text{L}}(t)$ and then $\sigma^{\text{R}}(t)$ without interruption. Using Inequality (56) and Claim 1, we therefore obtain

$$\begin{aligned}
|T^{\text{ALG}}| &\leq t^{\text{R}} + |p_{t^{\text{R}}}| + 2|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \\
&\overset{(56)}{\leq} (2\rho - 2)|a^{\text{L}}(t)| - (3 - \rho)|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{R}}(t^{\text{L}})| + 2|a^{\text{L}}(t)| + |a^{\text{R}}(t)| \\
&= 2\rho|a^{\text{L}}(t)| + (\rho - 1)|a^{\text{R}}(t^{\text{L}})| + |a^{\text{R}}(t)| \\
&\leq \rho|t^{\text{L}}| + \rho|a^{\text{L}}(t)| + \rho\max(|a^{\text{R}}(t)|, |a^{\text{R}}(t^{\text{L}})|) \\
&\overset{(51)}{\leq} \rho|T^{\text{OPT}}|.
\end{aligned}$$

Thus, it is sufficient to show that Inequality (56) is satisfied. Using Inequality (53), Lemma 11, $|a^{\text{R}}(t^{\text{L}})| \leq t^{\text{R}}$ and the definition of ρ, we get

$$\begin{aligned}
(2\rho - 2)|a^{\text{L}}(t)| - (3 - \rho)|a^{\text{R}}(t^{\text{L}})| - t^{\text{R}} &\overset{(53)}{>} \frac{3\rho - 5}{7 - 3\rho}(\rho t^{\text{R}} + (\rho - 2)|a^{\text{R}}(t^{\text{L}})|) - (3 - \rho)|a^{\text{R}}(t^{\text{L}})| - t^{\text{R}} \\
&= \frac{3\rho^2 - 2\rho - 7}{7 - 3\rho} t^{\text{R}} - \frac{11 - 5\rho}{7 - 3\rho}|a^{\text{R}}(t^{\text{L}})| \\
&\geq \frac{3\rho^2 - 2\rho - 7}{7 - 3\rho} t^{\text{R}} - \frac{11 - 5\rho}{7 - 3\rho} t^{\text{R}} \\
&= \frac{3\rho^2 + 3\rho - 18}{7 - 3\rho} t^{\text{R}} \\
&\overset{\textit{Lem. } 11}{\geq} \frac{3\rho^2 + 3\rho - 18}{7 - 3\rho} \cdot \frac{-3\rho^2 + 9\rho - 4}{3\rho - 5}|p_{t^{\text{R}}}|.
\end{aligned}$$

Hence, it suffices to show

$$\frac{3\rho^2 + 3\rho - 18}{7 - 3\rho} \cdot \frac{-3\rho^2 + 9\rho - 4}{3\rho - 5} \geq 1,$$

which is equivalent to $-9\rho^4 + 18\rho^3 + 78\rho^2 - 210\rho + 107 \geq 0$. The left-hand side is zero by the definition of ρ. This completes the case that the server moves constantly towards $a^{\text{R}}(t^{\text{L}})$ during $[t^{\text{R}}, t^{\text{L}}]$.

Next, we consider the case that the server does not constantly move towards $a^{\text{R}}(t^{\text{L}})$ during $[t^{\text{R}}, t^{\text{L}}]$. For this to happen, there has to be a call of Algorithm 2 at a time $\tilde{t} \in [t^{\text{R}}, t^{\text{L}})$ such that the tour followed by the server in $[t^{\text{R}}, t^{\text{L}})$ waits at the origin or goes to the left for some time. We choose $\tilde{t}$ maximal with this property. This setting is stated in Setting D and analogous to the setting Setting C we considered before. By Corollary 39 we know that only Case (A2) can occur when $\tilde{t} > t^{\text{R}}$ and Cases (P2) and (E2) can occur at time $\tilde{t} = t^{\text{R}}$.

We first establish that in all three cases the following inequality holds:

$$t^{\text{L}} + |p_{t^{\text{L}}} - a^{\text{L}}(\tilde{t})| \geq L^{\sigma^{\text{L}}(\tilde{t}), \sigma^{\text{R}}(t^{\text{L}})}. \tag{57}$$

Let $f(\tau) := \tau + |p_\tau - a^{\text{L}}(\tilde{t})|$, then we need to show $f(t^{\text{L}}) \geq L^{\sigma^{\text{L}}(\tilde{t}), \sigma^{\text{R}}(t^{\text{L}})}$. Observe that the function f is monotonously increasing in τ because the server moves with at most unit speed. Thus, if $f(t') \geq L^{\sigma^{\text{L}}(\tilde{t}), \sigma^{\text{R}}(t^{\text{L}})}$ holds for some $t' < t^{\text{L}}$, then Inequality (57) follows. If Case (E2) occurred at time $\tilde{t} = t^{\text{R}}$, then $f(t^{\text{R}}) > L^{\sigma^{\text{L}}(\tilde{t}), \sigma^{\text{R}}(t^{\text{L}})}$ because the condition for Case (P2) is not satisfied. Thus we get Inequality (57) in Case (E2) by monotonicity. In the Cases (P2) and (A2), we know that the tour computed at $\tilde{t}$ by Algorithm 2 serves $\sigma^{\mathrm{L}}(\tilde{t})$ first and then $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$. More precisely, the computed tour is

$$T_0.\operatorname{until}(\tau + |p_\tau - a^{\mathrm{L}}(\tilde{t})| = L^{\sigma^{\mathrm{L}}(\tilde{t}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})}) \oplus \operatorname{move}(a^{\mathrm{L}}(\tilde{t})) \oplus \operatorname{move}(a^{\mathrm{R}}(t^{\mathrm{L}})).$$

If the until-condition becomes satisfied at a time $t' \leq t^{\mathrm{L}}$, i.e., $f(t') = L^{\sigma^{\mathrm{L}}(\tilde{t}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})}$, then Inequality (57) holds by the monotonicity of f. Otherwise, there is a request appearing before the until-condition is satisfied. Let t' be the time the first request appearing after $\tilde{t}$. By the choice of $\tilde{t}$, we know that the server will immediately move towards $a^{\mathrm{R}}(t^{\mathrm{L}})$ at t'. Thus, if the server is waiting at the origin at t', i.e., $p_{t'} = 0$, then we have $p_{t^{\mathrm{L}}} \geq 0$ contradicting that $p_{t^{\mathrm{L}}} < 0$ by assumption. Next assume, that the server is still moving towards the origin at t'. If additionally $p_{t'} < 0$ holds, then the server moves continuously to the right in $[\tilde{t}, t^{\mathrm{L}}]$ contradicting the choice of $\tilde{t}$. If $p_{t'} > 0$ holds, we also get $p_{t^{\mathrm{L}}} > 0$ because the server moves to the right from t' onwards and does not reach $\sigma^{\mathrm{R}}(t^{\mathrm{L}})$ as this request is still the rightmost extreme at t^{L}. Thus, the until-condition must become satisfied as all other cases were contradictions.

We have shown that Inequality (57) holds in any of the three Cases (A2), (P2) and (E2). Now, we analyze the cases $\tilde{t} > t^{\mathrm{R}}$ (Case (A2)) and $\tilde{t} = t^{\mathrm{R}}$ (Cases (P2) and (E2)) separately.

1. *Case (A2) at time $\tilde{t} > t^{\mathrm{R}}$.*
   Note that we cannot use the same argumentation as in Lemma 40 because there we rely on the fact that after t (here it would be after t^{L}) no new requests appear.

   We have $t^{\mathrm{R}} < \tilde{t}$ and $t^{\mathrm{R}} \geq |a^{\mathrm{R}}(t^{\mathrm{L}})|$ and therefore

$$\begin{aligned} L^{\sigma^{\mathrm{L}}(\tilde{t}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})} &= \rho t^{\mathrm{R}} + (\rho-2)|a^{\mathrm{L}}(\tilde{t})| + (\rho-2)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &\geq (\rho-2)|a^{\mathrm{L}}(\tilde{t})| + (2\rho-2)|a^{\mathrm{R}}(t^{\mathrm{L}})| \end{aligned} \tag{58}$$

   Using $|p_{t^{\mathrm{L}}} - a^{\mathrm{L}}(\tilde{t})| = |a^{\mathrm{L}}(\tilde{t})| - |p_{t^{\mathrm{L}}}|$ since $a^{\mathrm{L}}(\tilde{t}) \leq p_{t^{\mathrm{L}}} < 0$, we can simplify Inequality (57) as follows:

$$|p_{t^{\mathrm{L}}}| \leq t^{\mathrm{L}} + |a^{\mathrm{L}}(\tilde{t})| - L^{\sigma^{\mathrm{L}}(\tilde{t}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})} \overset{(58)}{\leq} t^{\mathrm{L}} + (3-\rho)|a^{\mathrm{L}}(\tilde{t})| + (2-2\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})|. \tag{59}$$

   As Case (A2) occurred at time $\tilde{t}$, we know that

$$|a^{\mathrm{L}}(\tilde{t})| \leq \frac{3\rho-5}{(2\rho-2)(7-3\rho)}(\rho t^{\mathrm{R}} + (\rho-2)|a^{\mathrm{R}}(t^{\mathrm{L}})|) \leq \frac{3\rho-5}{7-3\rho} t^{\mathrm{R}}. \tag{60}$$

   By Inequality (53), this implies that $|a^{\mathrm{L}}(\tilde{t})| < |a^{\mathrm{L}}(t)|$. Using the last two statements together with Equation (59) and Claim 1, we obtain

$$\begin{aligned} |p_{t^{\mathrm{L}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| &\overset{(59)}{\leq} t^{\mathrm{L}} - \rho t^{\mathrm{R}} + (3-\rho)|a^{\mathrm{L}}(\tilde{t})| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &= t^{\mathrm{L}} - \rho t^{\mathrm{R}} + (7-3\rho)|a^{\mathrm{L}}(\tilde{t})| + (2\rho-4)|a^{\mathrm{L}}(\tilde{t})| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &\overset{(60)}{\leq} t^{\mathrm{L}} - \rho t^{\mathrm{R}} + (7-3\rho)\frac{3\rho-5}{7-3\rho} t^{\mathrm{R}} + (2\rho-4)|a^{\mathrm{L}}(\tilde{t})| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &\leq t^{\mathrm{L}} - \rho t^{\mathrm{R}} + (7-3\rho)\frac{3\rho-5}{7-3\rho} t^{\mathrm{R}} + (2\rho-4)|a^{\mathrm{L}}(t)| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &= t^{\mathrm{L}} - (5-2\rho)t^{\mathrm{R}} + (2\rho-4)|a^{\mathrm{L}}(t)| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &\leq t^{\mathrm{L}} - (5-2\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| + (2\rho-4)|a^{\mathrm{L}}(t)| + (4-\rho)|a^{\mathrm{R}}(t^{\mathrm{L}})| \\ &= t^{\mathrm{L}} + (\rho-2)|a^{\mathrm{L}}(t)| + (\rho-1)|a^{\mathrm{L}}(t)| + (\rho-1)|a^{\mathrm{R}}(t^{\mathrm{L}})| - |a^{\mathrm{L}}(t)| \\ &\leq ((\rho-1)t^{\mathrm{L}} + (\rho-1)|a^{\mathrm{L}}(t)| + (\rho-1)|a^{\mathrm{R}}(t^{\mathrm{L}})|) - |a^{\mathrm{L}}(t)| \\ &\overset{(51)}{\leq} (\rho-1)|T^{\mathrm{OPT}}| - |a^{\mathrm{L}}(t)|. \end{aligned}$$

   This implies $|T^{\mathrm{ALG}}| \leq \rho|T^{\mathrm{OPT}}|$ by Inequality (50).

2. *Cases (P2) or (E2) at time $t^{\mathrm{R}} = \tilde{t}$*
   Using $t^{\mathrm{R}} \geq |a^{\mathrm{L}}(t^{\mathrm{R}})|$, $t^{\mathrm{R}} \geq |a^{\mathrm{R}}(t^{\mathrm{L}})|$ and $\rho \geq 2$, we can bound $L^{\sigma^{\mathrm{L}}(t^{\mathrm{R}}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})}$ as follows

$$\begin{aligned}
L^{\sigma^{\mathrm{L}}(t^{\mathrm{R}}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})} &= \min\{\rho t^{\mathrm{L}}(t^{\mathrm{R}}) + (\rho-1)|a^{\mathrm{L}}(t^{\mathrm{R}})| + (\rho-1)|a^{\mathrm{R}}(t^{\mathrm{L}})|,\\
&\qquad \rho t^{\mathrm{R}} + (\rho-2)|a^{\mathrm{L}}(t^{\mathrm{R}})| + (\rho-2)|a^{\mathrm{R}}(t^{\mathrm{L}})|\}\\
&\geq \min\{|a^{\mathrm{L}}(t^{\mathrm{R}})| + |a^{\mathrm{R}}(t^{\mathrm{L}})|, \rho t^{\mathrm{R}}\}\\
&\geq |a^{\mathrm{L}}(t^{\mathrm{R}})| + |a^{\mathrm{R}}(t^{\mathrm{L}})|. \qquad (61)
\end{aligned}$$

   Recall that $a^{\mathrm{L}}(t^{\mathrm{R}}) \leq p_{t^{\mathrm{L}}} < 0$ and hence $|p_{t^{\mathrm{L}}} - a^{\mathrm{L}}(t^{\mathrm{R}})| = |a^{\mathrm{L}}(t^{\mathrm{R}})| - |p_{t^{\mathrm{L}}}|$. This means that Inequality (57) for $\tilde{t} = t^{\mathrm{R}}$ can be simplified to $t^{\mathrm{L}} + |a^{\mathrm{L}}(t^{\mathrm{R}})| - |p_{t^{\mathrm{L}}}| \geq L^{\sigma^{\mathrm{L}}(t^{\mathrm{R}}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})}$. Applying Inequality (61) above yields

$$\begin{aligned}
|p_{t^{\mathrm{L}}}| &\leq t^{\mathrm{L}} + |a^{\mathrm{L}}(t^{\mathrm{R}})| - L^{\sigma^{\mathrm{L}}(t^{\mathrm{R}}),\sigma^{\mathrm{R}}(t^{\mathrm{L}})}\\
&\overset{(61)}{\leq} t^{\mathrm{L}} + |a^{\mathrm{L}}(t^{\mathrm{R}})| - |a^{\mathrm{L}}(t^{\mathrm{R}})| - |a^{\mathrm{R}}(t^{\mathrm{L}})| = t^{\mathrm{L}} - |a^{\mathrm{R}}(t^{\mathrm{L}})|. \qquad (62)
\end{aligned}$$

   Finally, using Inequality (50) and Lemma 37, we obtain

$$\begin{aligned}
|T^{\textsc{Alg}}| &\overset{(50)}{\leq} |T^{\textsc{Opt}}| + |p_{t^{\mathrm{L}}}| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|\\
&\overset{(62)}{\leq} |T^{\textsc{Opt}}| + t^{\mathrm{L}} - |a^{\mathrm{R}}(t^{\mathrm{L}})| + 2|a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|\\
&= |T^{\textsc{Opt}}| + t^{\mathrm{L}} + |a^{\mathrm{R}}(t^{\mathrm{L}})| + |a^{\mathrm{L}}(t)|\\
&\leq 2|T^{\textsc{Opt}}|.
\end{aligned}$$

   This concludes the proof. ☐

D Proofs of Section 6

Lemma 42. *Without loss of generality, the following holds:*

1. *If $T^{\textsc{Opt}}$ visits position $p \neq p_0$, then there is a request with source or destination q, such that $(q - p_0)/(p - p_0) \geq 1$.*
2. *Let R_i^0 be the requests in R_i with their original sources and destinations (before being moved). Then, $|T^{\textsc{Opt}}_{R_i}| \leq |T^{\textsc{Opt}}_{R_i^0}|$.*

Proof. Let $p_{\min}$ and $p_{\max}$ be the minimum and maximum values, respectively, among p_0 and all source and destinations of requests. We can replace all subtours of $T^{\textsc{Opt}}$ beyond $[p_{\min}, p_{\max}]$ by waiting at $p_{\min}$ or $p_{\max}$, without increasing the overall length of the tour. Hence, the first property holds.

For the second property, it is sufficient to show that our algorithm never moves requests away from their destination, since then $T^{\textsc{Opt}}_{R_i^0}$ is a valid tour also for R_i, and hence $|T^{\textsc{Opt}}_{R_i}| \leq |T^{\textsc{Opt}}_{R_i^0}|$. To see this, consider $T^{\textsc{Opt}}_S$ for an arbitrary set of requests S. Consider the path that request $\sigma = (a, b; t) \in S$ is taking in $T^{\textsc{Opt}}_S$ and replace all subpaths that increase σ's distance to b by unloading the request and leaving it at its current position. This does not increase the length of the tour, and we obtain that, without loss of generality, we can assume that $T^{\textsc{Opt}}_S$ never moves requests away from their destination. Since our algorithm consists in the repeated application of tours of the form $T^{\textsc{Opt}}_{R_i}$, the same holds for our algorithm. ☐

In the following, we denote the i-th request by $\sigma_i = (a_i, b_i; t_i)$ and let $R_i := R_{t_i}$.

Theorem 12. *Algorithm 3 is $(1 + \sqrt{2}) \approx 2.41$-competitive for the preemptive open online* Dial-A-Ride *problem with capacity $c \geq 1$.*

Proof. Let $|T^{\textsc{Opt}}|$ be the length of an optimal tour and $|T^{\textsc{Alg}}|$ be the length of the tour produced by our algorithm. Assume that the server can always return to p_0 before time $\sqrt{2} \cdot |T^{\textsc{Opt}}_{R_i}|$ upon receiving a new request σ_i. Then, after the last request σ_n is received, the server stays in p_0 until time $\sqrt{2} \cdot |T^{\textsc{Opt}}_{R_n}|$ and then finishes all requests within time $|T^{\text{OPT}}_{R_n}|$. Thus $|T^{\text{ALG}}| \leq \sqrt{2} \cdot |T^{\text{OPT}}_{R_n}| + |T^{\text{OPT}}_{R_n}|$. Using the second property of Lemma 2, we have $|T^{\text{OPT}}_{R_n}| \leq |T^{\text{OPT}}_{R^0_n}| \leq |T^{\text{OPT}}|$, and the algorithm is $(1+\sqrt{2})$-competitive, as claimed.

It remains to show that the server can always return to p_0 before time $\sqrt{2} \cdot |T^{\text{OPT}}_{R_i}|$ when receiving a new request σ_i at time t_i. We prove this by induction over the number i of released requests. Clearly, the statement is true if $t_i = 0$ or $i = 1$, as the server does not leave p_0 before time t_i. Now consider the request σ_i with $i > 1, t_i > 0$. As an optimal tour $T^{\text{OPT}}_{R_i}$ over the set of requests R_i cannot finish before all requests are released, we have $|T^{\text{OPT}}_{R_i}| \geq t_i$. For the sake of contradiction, suppose that the server cannot return to p_0 before time $\sqrt{2} \cdot |T^{\text{OPT}}_{R_i}| \geq \sqrt{2} \cdot t_i$. This implies that the distance from p_i to p_0 is at least $|p_i - p_0| > \sqrt{2} \cdot t_i - t_i$. In order to reach position p_i at time t_i, the server cannot have been at p_0 after time $t_i - (\sqrt{2} \cdot t_i - t_i) = 2t_i - \sqrt{2} \cdot t_i$. Consider the last request σ_{i-1} that was released before σ_i at time t_{i-1}. The tour that the server was following until time t_i was thus the tour serving R_{i-1}. By induction, the server was at p_0 at time $\sqrt{2} \cdot |T^{\text{OPT}}_{R_{i-1}}|$, and from before we know that it cannot have been at p_0 after time $2t_i - \sqrt{2} \cdot t_i$. We thus have $\sqrt{2}|T^{\text{OPT}}_{R_{i-1}}| \leq 2t_i - \sqrt{2}t_i$, or

$$|T^{\text{OPT}}_{R_{i-1}}| \leq \sqrt{2}t_i - t_i. \tag{63}$$

Using the first property of Lemma 2, there must have been a request in R_{i-1} with source or destination at distance at least $|p_i - p_0|$ from p_0. With $|p_i - p_0| > \sqrt{2} \cdot t_i - t_i$, we thus have $|T^{\text{OPT}}_{R_{i-1}}| \geq |p_i - p_0| > \sqrt{2} \cdot t_i - t_i$, a contradiction with (63), since $t_i > 0$. ☐

Theorem 13. *No algorithm for the non-preemptive closed* DIAL-A-RIDE *problem on the line with fixed capacity $c \geq 1$ has competitive ratio lower than $\rho = 1.75$.*

Proof. Consider any ρ-competitive online algorithm ALG. We define an instance with three types of requests $\sigma^{(0)} = (0, 0; 1)$, $\sigma^{(1)}_i = (t, 0; t)$, $\sigma^{(2)}_i = (-t, 0; t)$, where $i = 1, \ldots, c$ and $t \geq 1$ is the time when ALG serves σ_0. First, assume that ALG serves the c requests $\sigma^{(1)}_1, \ldots, \sigma^{(1)}_c$ not together in one tour from t to 0. As ALG can pick up any $\sigma^{(1)}_i$ at the earliest at time $2t$ and ALG has to return to t to pick up the remaining requests $\sigma^{(1)}_i$ it did not bring to 0 in the first trip, we have $|T^{\text{ALG}}| \geq 7t$ in this case. We have $|T^{\text{OPT}}| = 4t$ and hence obtain $|T^{\text{ALG}}|/|T^{\text{OPT}}| \geq 7/4$ in this case. An analogous argument shows that ALG has to take the c requests $\sigma^{(2)}_1, \ldots, \sigma^{(2)}_c$ together to the origin or we have $|T^{\text{ALG}}|/|T^{\text{OPT}}| \geq 7/4$.

Thus, from now on we assume that ALG picks up all c requests $\sigma^{(j)}_1, \ldots, \sigma^{(j)}_c$ before going to the origin for $j = 1, 2$. Let $t' \geq 2t$ be the first time when ALG picks up all c requests $\sigma^{(1)}_i$ or all c requests $\sigma^{(2)}_i$. We have $|T^{\text{OPT}}| = 4t$ and $|T^{\text{ALG}}| \geq t' + 3t$. If $t' \geq 4t$, then $\rho \geq |T^{\text{ALG}}|/|T^{\text{OPT}}| \geq 7/4$ as claimed.

Otherwise, without loss of generality, we assume that the c requests $\sigma^{(1)}_i$ are picked up before the c requests $\sigma^{(2)}_i$. We introduce a new request $\sigma_3 = (t, t; t' + 1/7)$. For the new instance, we have $|T^{\text{ALG}}| \geq t' + 5t$, since at time $t' + 1/7$ ALG still needs to return to the origin to finish serving the c requests it is currently transporting (no preemption), then serve the remaining requests $\sigma^{(2)}_i$, σ_3, and finally return to the origin. If $t' \leq 3t - 1/7$, we have $|T^{\text{OPT}}| = 4t$ and, using $t' \geq 2t$, we get

$$\rho \geq \frac{|T^{\text{ALG}}|}{|T^{\text{OPT}}|} \geq \frac{t' + 5t}{4t} \geq \frac{7}{4},$$

as claimed. Now if $t' > 3t - 1/7$, then $|T^{\text{OPT}}| = t' + 1/7 + t$ and $\rho \geq (t' + 5t)/(t' + 1/7 + t)$, which is monotonically decreasing in t' for $t, t' \geq 1$. Since we have $t' < 4t$ from above, we get $\rho > 9/(5 + 1/7) = 1.75$. ☐

E Proofs of section 7

Theorem 14. *For every $k \in \mathbb{N}$, there is an instance of closed* TSP *on the line such that any optimal solution turns around at least $2k$ times.*

Proof. We analyze the following instance consisting of $2k + 1$ requests with $M := 2k(k + 1)$.

$$\begin{aligned} \sigma_i = (a_i; t_i) \text{ with } & a_i = 0, 1, -1, 2, -2, 3, -3, \ldots, k, -k \\ \text{and } & t_i = M, M-1, M-3, M-6, M-10, M-15, \ldots, k \end{aligned}$$

We show this instance has a unique optimal server tour with $2k$ turnarounds. Observe first, that the difference between two consecutive release times $t_i - t_{i-1}$ is exactly the travel time of the server between the two requested positions $|a_i - a_{i-1}|$. Now consider the server tour T serving each of the requests exactly at its release time, which is obviously feasible and also optimal, as it ends at time $M = 2k(k+1)$, which is the release time of the request to position 0. By construction, the server alternatively serves requests left and right of position 0 and thus turns around $2k$ times.

Now assume there is another optimal server tour T'. Then T' must serve the request at position 0 at time M, as this is its release time as well as the makespan of T'. As we observed $t_i - t_{i-1} = |a_i - a_{i-1}|$, the tour T' also serves the previous request exactly at its release time. Iteratively the same must hold for all other positions. This shows any request is served exactly at its release time and hence, the tour is exactly the one described above. □

To describe the dynamic for the closed offline TSP problem on the line, we make the following non-restrictive assumptions on the requests:

- Each position is requested at most once, as several requests to the same position can all be fulfilled when the one with the largest release time is served.
- There is a request $\sigma_0 = (0; 0)$.
- The requests $\sigma_i = (a_i; t_i)$ are labeled by increasing position, i.e. $a_i > a_{i-1}$, and requests with $a_i < 0$ have negative indices, while the other ones have a positive index. This yields a sequence $a_{-\ell} < \cdots < a_{-1} < a_0 = 0 < a_1 < \cdots < a_r$.

Our dynamic program, Algorithm 4 relies on the fact that an optimal server tour has a 'zig-zag shape' with decreasing amplitude. We compute for each index pair $-\ell - 1 \leq i < j \leq r + 1$ the completion time of two tours. We use $C^+_{i,j}$ for the best tour serving all requests σ_k with $k \leq i$ or $k \geq j$ and ending at position a_j. This means the tour serves all requests to position a_j and larger positions as well as all requests to position a_i and smaller positions. The time $C^-_{i,j}$ is the completion time of the best tour that serves the same set of requests and ends at position a_i.

We start by considering the two request sets $\{\sigma_{-\ell}\}$ and $\{\sigma_r\}$. By assumption, the release time t_i exceeds the travel time between start position 0 and requested position a_i for each request. Thus the initial values are $C^+_{-\ell-1,r} = t_r$ and $C^-_{-\ell,r+1} = t_{-\ell}$. We then compute by recursion the completion time of the tours $C^+_{i,j}$ computing $C^+_{-\ell-1,r}$ and $C^-_{-\ell,r+1}$ for the two request sets $\{\sigma_r\}$ and $\{\sigma_{-\ell}\}$. Then we recursively compute $C^+_{i,j}$ and $C^-_{i,j}$, starting with large difference $d = j - i$ and iteratively decreasing it. We output $C^+_{0,0}$ as the minimum completion time of a feasible server tour.

Algorithm 6: Dynamic Program for Closed Offline TSP on the line.

Input: A set of requests $\sigma_i = (a_i; t_i)$ with $t_i \geq |a_i|$ for $-\ell \leq i \leq r$ and $a_i < a_{i+1}$ for $-\ell \leq i < r$.
Output: The minimum completion time $C_{\max}$ of a tour.

$C^+_{-\ell-1,r} \leftarrow t_r$.
$C^-_{-\ell,r+1} \leftarrow t_{-\ell}$.
for $i = -\ell, ..., r$ **do**
  $C^-_{i,r+1} \leftarrow \max\{t_i, C^-_{i-1,r+1} + a_i - a_{i-1}\}$.
for $j = r, .., \ell$ **do**
  $C^+_{-\ell-1,j} \leftarrow \max\{t_j, C^+_{\ell-1,j+1} + a_{j+1} - a_j\}$.
for $d = r + \ell, ..., 0$ **do**
  for $i = -\ell, .., r - d$ **do**
    $j \leftarrow i + d$.
    $C^-_{i,j} \leftarrow \max\{t_i, \min\{C^+_{i-1,j} + a_j - a_i, C^-_{i-1,j} + a_i - a_{i-1}\}\}$.
    $C^+_{i,j} \leftarrow \max\{t_j, \min\{C^+_{i,j+1} + a_{j+1} - a_j, C^-_{i,j+1} + a_j - a_i\}\}$.
return $C_{\max} = C^+_{0,0}$.

To prove the correctness of Algorithm 4 we use two lemmas. We first prove a structural result about feasible server tours and then show that the recursion we use in the dynamic program is correct. Contrary to before, we assume without loss of generality that each request is served when its requested position is visited for the last time in the tour.

Lemma 43. *At any time in a feasible server tour ending at position p, the set of served requests is the union of two disjoint sets $S_1 = \{\sigma_\ell, ..., \sigma_i\}, a_i \leq p$ and $S_2 = \{\sigma_j, ..., \sigma_r\}, p \leq a_j$, both of which are contiguous.*

Proof. Assume there is a time t, at which the set of requests does not have the claimed structure. Then there is a request σ_1 served until time t and a request σ_2 served after t with either $p \leq a_1 < a_2$ or $a_1 < a_2 \leq p$. Without loss of generality assume the first to be true. At time $t' > t$, when request σ_2 is served, the server is at position a_2. The server tour ends at position p and thus passes position a_1 at some time after t. This contradicts that request σ_1 has been served until time t, which was our assumption. ☐

Lemma 44. *Given an instance of* TSP *on the line with requests $\sigma_\ell, ..., \sigma_r$ to positions $a_\ell < a_{\ell+1} < ... < a_r$ and completion times $C^+_{i,j+1}, C^-_{i,j+1}, C^+_{i-1,j}$ and $C^-_{i-1,j}$ for some indices $\ell \leq i \leq j \leq r$; then the minimal completion times $C^+_{i,j}$ and $C^-_{i,j}$ are given by the following recursion:*

$$C^+_{i,j} = \max\{t_j, \min\{C^+_{i,j+1} + a_{j+1} - a_j, C^-_{i,j+1} + a_j - a_i\}\},$$
$$C^-_{i,j} = \max\{t_i, \min\{C^+_{i-1,j} + a_j - a_i, C^-_{i-1,j} + a_i - a_{i-1}\}\}.$$

Proof. The completion time $C^+_{i,j}$ we give is feasible, as it can be achieved by executing the tour attaining $C^+_{i,j+1}$ or the tour attaining $C^-_{i,j+1}$ and then moving to position a_j, waiting there until the release time t_j has passed.

By Lemma 43, the tour serving requests $\sigma_\ell, ..., \sigma_i$ and $\sigma_j, ..., \sigma_r$ serves request σ_j last and request σ_{j+1} or σ_i immediately before. Thus, the completion time exceeds the minimum of $C^+_{i,j+1} + a_{j+1} - a_j$ and $C^-_{i,j+1} + a_j - a_i$. Furthermore, we cannot serve request σ_j before its release time t_j. Symmetrically the recursion for $C^-_{i,j}$ is constructed. ☐

These two lemmas allow us to prove the correctness of Algorithm 4.

Theorem 15. *Algorithm 4 computes the minimum completion time of a server tour for offline TSP on the line in time $O(n^2)$.*

Proof. By Lemma 43, the optimal server tour has the structure that it has served the union of two disjoint contiguous sets at any fixed point in time. We start computing the best completion time of a tour for pairs of request sets containing one request each and increase the number of contained requests $r - \ell - d$ in each iteration by decreasing the parameter d. We end with $d = 0$ and compute the earliest completion time of tours serving the complete request set in this iteration. The recursion we use is correct by Lemma 44 and the computation order is feasible, as the recursion formula only contains request set pairs of strictly smaller size.

The algorithm uses quadratic time for the two nested loops and all other steps take linear or constant time. Hence the algorithm runs in time $O(n^2)$. ☐

Observation. The algorithm of [23] for open TSP on the line can be obtained from Algorithm 4 by changing the computation order of the completion times and returning the minimum completion time of all possible end positions $\min_{-\ell \leq i \leq r} C^+_{i,i}$.

For our hardness proofs for the non-preemptive offline Dial-A-Ride problem on the line, we reduce from the NP-complete problem Circular Arc Coloring [12].

Definition 45 (Circular Arc Coloring)**.** Let $\mathcal{I}$ be a family of intervals on a circle, and let $k \in \mathbb{Z}_{>0}$ be a fixed parameter. Decide, if a coloring of all intervals $I \in \mathcal{I}$ with k colors exists, such that no two intervals of the same color overlap.

Theorem 46. *The non-preemptive closed offline* Dial-A-Ride *problem on the line with capacity $c = 1$ is NP-complete.*

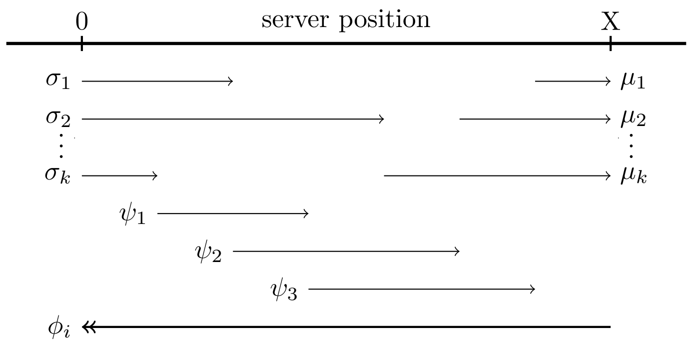

Figure 1: Construction of a Dial-A-Ride instance for Theorem 46

Proof. The problem is in NP, as we can decide whether or not a server tour completes before a fixed deadline in polynomial time.

Let an instance of the circular arc coloring problem be given by a circle with circumference X, a set $\mathcal{I}$ of intervals $I = [\ell, r), \ell, r \in [0, X)$, on the circle and coloring number k. Without loss of generality we can assume that there are exactly k intervals $I_1^0, ..., I_k^0$ overlapping 0. If there are fewer, we can add intervals $[0, \varepsilon)$ to the instance for sufficiently small ε. If there are more, the instance is trivial. We define the following set of requests for the Dial-A-Ride problem: For each interval $I_j^0 = [\ell_j, r_j), j \leq k$, we create two requests $\sigma_j = (a_j, b_j; t_j) = (0, r_j; 2(j-1)X)$ and $\mu_j = (\ell_j, X; 2(j-1)X)$. For any other interval $I = [\ell, r)$ define a request $\psi = (\ell, r; 0)$ and furthermore define k requests $\phi_i = (0, X; 0)$ for $i = \{1, ..., k\}$. The construction of requests with their start and end position is displayed in Figure 1. Note that the server position increases from left to right and that the release times are not displayed. The bold arrow with a double tip represents the k identical requests ϕ_i.

Any feasible server tour for this instance with start position 0 travels at least k times from 0 to X and back to serve the k requests ϕ_i. Hence, any feasible tour has length at least $2kX$. A tour of exactly that length is closed and partitions the requests into k sets, each served during one trip from 0 to X and back. Observe that requests σ_k and μ_k must be served on the last trip because they are released at time $2(k-1)X$. Then requests σ_{k-1} and μ_{k-1} must be served on the previous trip. They cannot be served before because of their release time and they cannot be served on the last trip as the server has capacity 1 and serves requests σ_k whose interval overlaps that of σ_{k-1} and also serves request μ_k whose interval overlaps that of μ_{k-1}. Iteratively we deduce that requests σ_j and μ_j, $j \leq k$, must be served on the same trip. This shows, that if we use the k-partitioning induced by the tour to color the intervals on the circle, we get a feasible circular arc coloring. Conversely, we can transform any circular arc coloring into a feasible server tour of length $2kX$ by scheduling the requests corresponding to one color on the same trip from O to X. ☐

Remark. Note that we can extend this proof to the case $c \geq 1$ if we add $c-1$ requests from 0 to X for each release time $0, 2X, 4X, \ldots, 2(k-1)X$. Then any feasible tour finishing at time $2kX$ serves exactly $c-1$ of these additional requests on each trip from 0 to X and thus still partitions the other requests in the desired fashion. However, this result is subsumed by the next theorem.

Theorem 47. *The non-preemptive closed offline* Dial-A-Ride *problem on the line with capacity $c \geq 2$ is NP-complete, even when all release times are* 0*.*

Proof. The problem is in NP, as we can decide whether or not a server tour completes before a fixed deadline in polynomial time.

Let an instance of the circular arc coloring problem be given by a circle with circumference X, a set $\mathcal{I}$ of intervals $I = [\ell, r), \ell, r \in [0, X)$, on the circle and coloring number k. Without loss of generality we can assume that there are exactly k intervals overlapping any point p on the circle. If there are less, we can add intervals $[p, p + \varepsilon \bmod X)$ to the instance. If there are more, then the instance is trivial. Let the k intervals that overlap the point 0 be $I_1^0, \ldots, I_k^0$. We define the following set of requests for the Dial-A-Ride problem: For each interval $I_j^0 = [\ell_j, r_j), j \in \{1, \ldots, k\}$, we create two requests $\sigma_j = (-j, r_j; 0)$ and $\mu_j = (\ell_j, X + j; 0)$. For any other interval $I = [\ell, r)$ define a request $\psi = (\ell, r; 0)$. Furthermore define for each $j \in \{1, \ldots, k\}$ exactly $c-1$

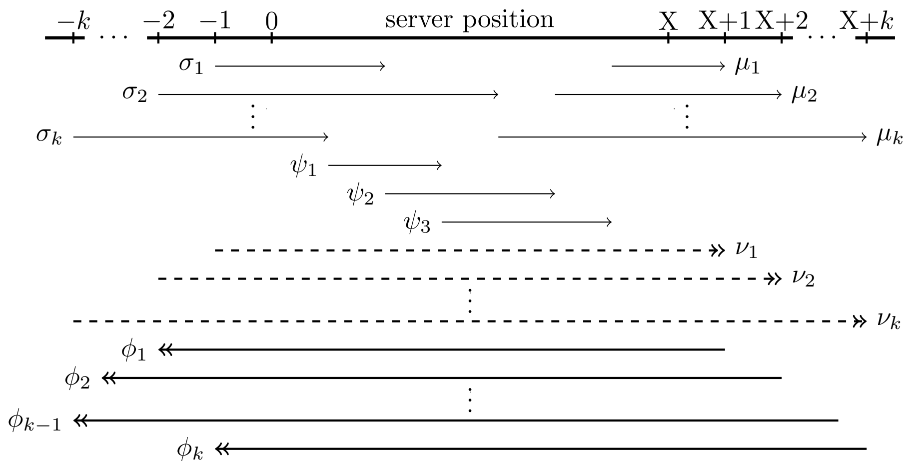

Figure 2: Construction of a Dial-A-Ride instance for Theorem 47

identical requests $\nu_j = (-j, X + j; 0)$, for $j \in \{1, \ldots, k-1\}$ exactly c identical requests $\phi_j = (X + j, -j-1; 0)$ and another c identical requests $\phi_k = (X + k, -1; 0)$. Contrary to before, let the start position of the server be -1 for this instance. The set of requests with their start and end position is displayed in Figure 2. In the Figure, the server position increases from left to right and release times are not displayed. We depict a set of $c-1$ identical requests by a dashed thick arrow with a double tip, and a set of c identical requests by a thick arrow with a double tip.

Let d be given by $d = 2kX + 2\sum_{j=1}^{k} j$. This is $2c$ times the total length of requests in positive direction because exactly k intervals overlap each point of the circle and there is one request starting at $-j$ and one ending at $X + j$ for all $j \leq k$. Furthermore for each $j \in \{1, ..., k\}$ there are $c-1$ requests from $-j$ to $X + j$, so it is also $2c$ times the total length of requests in negative direction. Thus, d is the minimum length of any feasible server tour. A server tour meeting this length bound must travel with full capacity at any time. As the server capacity c is at least 2, the tour must start with serving the $c-1$ requests ν_1, continue with serving requests ϕ_1 and then serve ν_j and ϕ_j for increasing indices $j = 2, \ldots, k$. This tour leaves capacity 1 on each trip from $-j$ to $X + j$ for requests from the set $\bigcup_{j=1}^{k}\{\sigma_j, \mu_j\}$ and the requests of type ψ. Hence, any tour of length d yields a partition of the requests σ, μ and ψ into k sets of non-overlapping requests. Moreover, requests σ_j and μ_j occur in the same set of the partition. Thus assigning one color to each set of the partition yields a feasible coloring of the circular arc instance. Conversely, we can use any circular arc coloring to design a server tour of length d. We serve requests ν and ϕ as described above and use for the $(j+1)$-st trip in positive direction all requests with the same color as request σ_j in the arc coloring. $\square$

Theorem 16. *The non-preemptive open and closed offline* Dial-A-Ride *problem on the line are NP-complete. For capacity $c \geq 2$ this even holds when all release times are* 0.

Proof. Theorems 46 and 47 show the statement for the closed problem variant. For the open case, we show in the proof of Theorem 46, that deciding if there is a tour of length $2kX$ is hard. Also for the open problem variant, any feasible tour has at least this length as a feasible server tour starts at position 0 and travels at least k times from X to 0 to serve the k requests ϕ_i. This also shows that any tour attaining this bound is a closed tour.

In the proof of Theorem 47 we consider closed tours of full capacity. Again this shows, that there cannot be a smaller open tour and that all minimal open tours are closed. $\square$

Remark. It is particularly interesting, that in the hardness reduction the server's start position is at one of the two extreme positions occurring in the tour. For the same problem allowing preemption, Karp showed that it is solvable in polynomial time [16], while the hardness of the problem with an arbitrary start position is not known.

Remark. Non-preemptive offline DIAL-A-RIDE on the line is known to be easy without release dates [9] and becomes hard if each request comes with a deadline in addition to its release time [24]. Thus the only complexity question remaining open is the case with unbounded capacity.