New Bounds on the Performance of SBP for the Dial-a-Ride Problem with Revenues

Barbara M. Anthony
Southwestern University, Georgetown, TX, USA

Christine Chung
Connecticut College, New London, CT, USA

Ananya Das
Middlebury College, VT, USA

David Yuen
Independent Researcher, Kapolei, HI, USA

Abstract

We revisit the Segmented Best Path (SBP) algorithm for online DARP in an offline setting with revenues and a time limit. The goal is to find a subset of the inputted ride requests that can be served within the time limit while maximizing the total revenue earned. SBP divides the time into segments and greedily chooses the highest-revenue path of requests to serve within each time segment. We show that SBP's performance has an upper bound of 5. Further, while SBP is a tight 4-approximation in the uniform-revenue case, we find that with non-uniform revenues, the approximation ratio of SBP has a lower bound strictly greater than 4; in particular, we provide a lower bound of $(\sqrt{e}+1)/(\sqrt{e}-1) \approx 4.08299$, which we show can be generalized to instances with ratio greater than 4.278.

2012 ACM Subject Classification Theory of computation → Routing and network design problems; Applied computing → Transportation

Keywords and phrases Dial-a-Ride problems, Lower bounds, Vehicle routing

Digital Object Identifier 10.4230/OASIcs.ATMOS.2024.8

Category Short Paper

1 Introduction

We study the Dial-a-Ride Problem in an offline setting with revenues and a time limit T. The goal is to find a subset of the inputted ride requests that can be served within the time limit while maximizing the total revenue earned. We consider the Segmented Best Path (SBP) algorithm, originally proposed in [6] for an online variant of DARP. It was later adapted by [1] for the offline setting where revenues are uniform and the goal is to maximize the number of requests served. We present SBP in a form that applies to our offline non-uniform revenue setting. This modified SBP algorithm starts by partitioning the total time limit into *time windows*, where each window (except possibly the last) is split into two equal *time segments*. The algorithm uses the first segment of each window to determine a maximum revenue set of requests that can be served within a segment, moving (if needed) to this set. It then uses the second segment of each window to serve the requests in this set.

For a literature review of some of the numerous DARP variants, see a recent survey by Ho et al. [8]. DARP problems are generalizations of the Traveling Salesperson Problem (TSP), so we mention TSP work that is most closely related to the time-limited variant of DARP that we study in this paper. Balas [2] first introduced the Prize Collecting Traveling Salesperson Problem (PCTSP), in which the server earns a prize (similar to our revenues) for each location visited, with the goal of collecting a prescribed amount of prize money

24th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2024).
Editors: Paul C. Bouman and Spyros C. Kontogiannis; Article No. 8; pp. 8:1–8:6
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

while minimizing travel costs and penalties. Bienstock et al. [3] gave the first approximation algorithm for PCTSP with ratio 2.5. Recently, Blauth and Nägele [4] achieved a significant improvement, obtaining an approximation guarantee of 1.774. Blum et al. [5] provide a constant-factor approximation algorithm for the Orienteering Problem (OP), another special case of the DARP problem we study. The goal of OP is, given a weighted graph with rewards on the nodes, to maximize the total reward collected on a path of a predefined maximum length. Our problem generalizes OP (and TSP), since in DARP we must visit pairs of points, rather than single points. The limit on the path length for OP is analogous to the time limit T in our DARP setting.

In this work, we highlight how SBP's performance changes when the revenues switch from uniform to non-uniform. Previously, we showed that an adapted version of the SBP algorithm which enforced an even number of time segments gave an approximation ratio of 4 in the uniform-revenue setting [1]. In this work we show that when revenues are non-uniform, SBP approximates the optimal revenue that can be earned within the time limit to within a factor of 5. We then show that when the number of time segments is odd, the approximation ratio of SBP is no better than 5, before showing that when the number of time segments is even, the ratio is strictly greater than 4.

2 New Upper and Lower Bounds

We formally define RDARP, the *Revenue-Dial-a-Ride-Problem*, as follows. The input to RDARP is a complete weighted graph, a set of requests given as source-destination node-pairs where each request has an associated revenue, and a time limit $T > 0$. We note that any simple, connected, weighted graph is allowed as input, with the simple preprocessing step of adding an edge wherever one is not present whose weight is the length of the shortest path between its two endpoints. We further note that the input can be regarded as a metric space if the graph is undirected and the edge weights satisfy triangle inequality. We treat the edge weights as travel-times, but for expository convenience may also refer to them as distances. Let t_{max} denote the maximum length of an edge in the graph.

Algorithm 1 SEGMENTED BEST PATH (SBP) Algorithm as adapted from [6]. Input: time limit $T > 0$, a complete graph with $T \geq 2t_{max}$, and a set of requests with associated revenues.

Let $t_1, t_2, \dots t_f$ denote time segments of length $X = T/f$ ending at times $T/f, 2T/f, \dots, T$, respectively, where $f = \lfloor T/t_{max} \rfloor$.
Let $i = 1$.
while $i < f$ and there are still unserved requests **do**
  At the start of t_i, find the *max-revenue-sequence*, $\mathcal{R}$.
  Move to the source location of the first request in $\mathcal{R}$.
  At the start of t_{i+1}, serve the requests in $\mathcal{R}$.
  Let $i = i + 2$.
end while

We begin by adapting the SBP algorithm for online DARP to the offline RDARP setting. In the online setting, SBP was shown in [6] to have competitive ratio 6, which was then improved to 5 and shown to be tight [7]. At the beginning of SBP (see Algorithm 1), set f, the number of time segments, to $\lfloor T/t_{max} \rfloor$. Let $X = T/f$ be the length of a *time segment*; note $X \geq t_{max}$. Let every pair of consecutive time segments, starting from the first time segment t_1, form a *time window*. A *max-revenue-sequence*, $\mathcal{R}$, is a sequence of requests of maximum total revenue that can be served within one time segment of length T/f.

We use the term *drive* to refer to any move of the server from one point to another, whether or not there is a request being served during the move. We use OPT to denote an optimal solution: a sequence of requests that maximizes total revenue that can be served within the time limit T. We let $|\text{OPT}(I)|$ denote the total revenue earned by the optimal solution on an instance I of RDARP.

2.1 Upper bound

In what follows, we allow OPT to choose its desired position at time 0. We will show that even with this extra flexibility, SBP is a 5-approximation.

▶ **Lemma 1.** *Let r denote the revenue of all requests that* OPT begins *serving by the end of the first time window, and s denote the revenue earned by* SBP *within the first time window. Then $s \geq r/4$.*

Proof. Consider the initial subpath of OPT that contains the first time window and any requests OPT begins serving by the end of the first time window. This subpath thus has revenue r. We subdivide this subpath into four further subpaths:

1. The subpath that is entirely contained in the time interval $[0, X]$.
2. The subpath that is entirely contained in the time interval $[X, 2X]$.
3. The drive (not necessarily a request), if any, between (1) and (2).
4. The drive (not necessarily a request), if any, that comes after (2) and overlaps time $2X$.

Each of these four subpaths has total length no greater than X. Because their collective revenue is r, at least one of the four subpaths must have revenue at least $r/4$. Since SBP could have greedily chosen any of these four subpaths, SBP must thus earn revenue $s \geq r/4$. ◀

▶ **Theorem 2.** *For the offline general metric with nonuniform revenues and time limit $T \geq 2t_{max}$, we have $|\text{OPT}| \leq 5|\text{SBP}|$.*

Proof. We will show by induction on the number of time windows that $|\text{SBP}| \geq |\text{OPT}|/5$.

Base case 1: There is only one time window, so $T = 2X$. ($T \geq 2X$ since $T \geq 2t_{max}$ by assumption, so $f \geq 2$.) Using Lemma 1, $|\text{OPT}| = r \leq 4s = 4|\text{SBP}| \leq 5|\text{SBP}|$.

Base case 2: $T = 3X$. For this case, consider the above proof of Lemma 1, but add one more subdivision so that there are 5 subpaths instead of 4. The fifth subpath is the subpath of OPT that is entirely contained within the time interval $[2X, 3X]$. We now redefine r to be the revenue earned by OPT by time $3X$. The rest of the proof remains the same, except we divide into five subpaths instead of four, so $|\text{OPT}| = r \leq 5s = 5|\text{SBP}|$.

For the inductive step, we may now assume $T \geq 4X$, with the theorem holding for any smaller time limit than T. Let I refer to the original input instance. We want to show that $|\text{SBP}(I)| \geq |\text{OPT}(I)|/5$. We consider the smaller instance after SBP has completed its first time window. This smaller instance, I', has time limit $T - 2X$, and the requests served by SBP in the first time window are removed.

Consider the path in I' formed by taking the OPT path in I and removing the initial part that earned revenue r; this path has length at most $T - 2X$, and revenue at least $|\text{OPT}(I)| - r - s$. In essence, it is the OPT path with the first portion removed and potentially some 'holes' from requests that are not present in I'. By the inductive hypothesis, $\text{SBP}(I')$ would have revenue at least $|\text{OPT}(I')|/5 \geq (|\text{OPT}(I)| - r - s)/5$.

Thus, using the inductive hypothesis and Lemma 1, we have $|\text{SBP}(I)| = s + |\text{SBP}(I')| \geq s + (|\text{OPT}(I)| - r - s)/5 = (|\text{OPT}(I)| - r + 4s)/5 \geq (|\text{OPT}(I)| - r + r)/5 = |\text{OPT}(I)|/5$, completing the induction. ◀

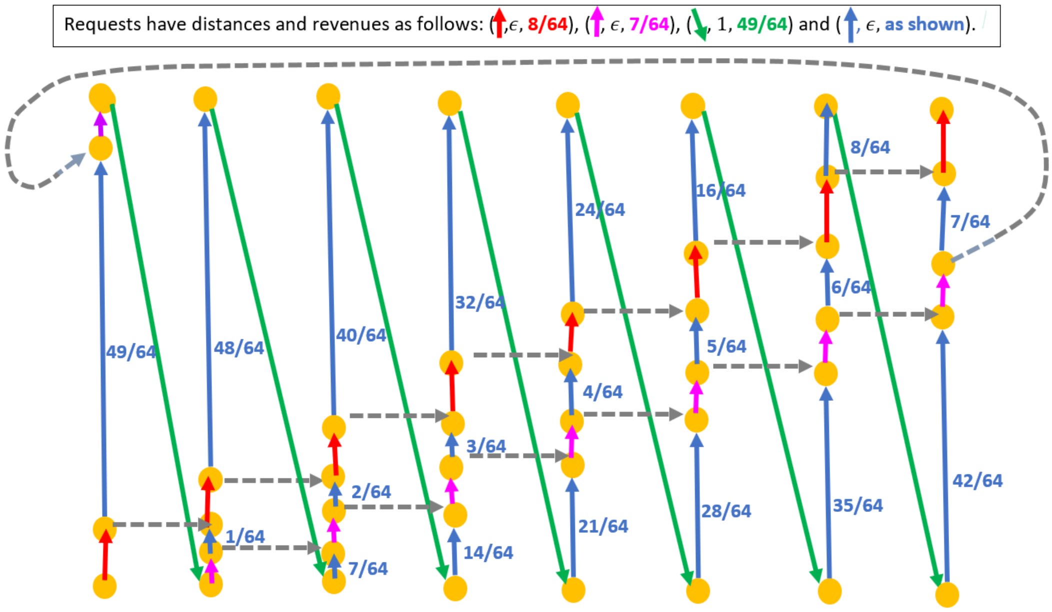

Figure 1 A small instance with $a = 4$ and $b = 2$ that illustrates the overall structure of the lower bound (but does not yield a ratio greater than 4). Requests are shown in color with distances and revenues as indicated. Gray dashed edges are empty drives of distance ϵ. All edges not shown have distance 1, including the reverse of existing directed edges.

2.2 Tight lower bound when the number of time segments is odd

Consider an instance with $T = 6$, $t_{max} = 2$, and five (or more) requests of unit revenue and length $1 + \delta$, for some small $\delta > 0$, so $f = \lfloor T/t_{max} \rfloor = 3$. These requests are such that OPT is able to serve five consecutively until the time limit, earning total revenue of 5. By contrast, SBP serves only one request: two consecutive requests take time $2 + 2\delta > 2 = T/f$, so only one request can be served within a single time segment of length T/f, and the first segment of each window is used, by definition of SBP, to move. Thus, we achieve a ratio of 5 which matches base case 2 of Theorem 2.

In [1] we proposed a version of SBP that ensured f, the number of time segments, was even. There we showed that version of SBP earned a (tight) approximation ratio of 4; however, in what follows we show that even if the number of time segments is even in Algorithm 1, we cannot guarantee an upper bound of 4 on the ratio.

2.3 Lower bound when the number of time segments is even

We describe how to construct instances with an even number of time segments that have a lower bound greater than 4. We note that these instances are not metric spaces, as they lack symmetry. Let a, b be positive integers with $b \leq a$. Let $t_{max} = 1$ and $T = 2a$, resulting in $f = 2a$ and $X = 1$. See Figure 1 for a representative example with $a = 4$ and $b = 2$.

Let the OPT path consist of $P_1, E_1, P_2, E_2, ..., P_{2a-1}, E_{2a-1}, P_{2a}$ where each P_i (depicted in Fig. 1 as vertical paths oriented upward) is a path of up to $2b + 1$ requests with total distance less than $1/(2a)$ and total revenue 1, and each E_i (depicted as a green downward diagonal edge in the figure) is a request of distance 1 and revenue c that we will specify later. These are the only requests in the input. Hence, the OPT path has total distance at most $2a/(2a) + 2a - 1 = 2a$, and the optimal solution can be completed in time T.

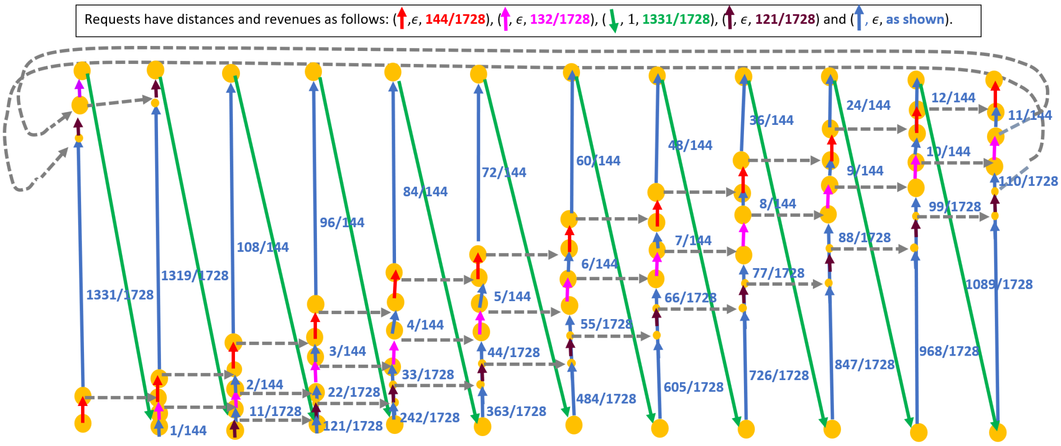

Figure 2 Instance with $a = 6, b = 3$. Requests are shown in color with distances and revenues as indicated. Gray dashed edges are empty drives of distance ϵ. All edges not shown have distance 1, including the reverse of existing directed edges.

All other distances between nodes will be 1 except that specific nodes of each P_i are very close to nodes in P_{i+1} so that there are paths $Q_1, Q_2, \ldots, Q_b$ defined as follows. Q_1 (depicted in Figures 1 and 2 as the red and gray staircase pattern) starts with a request in P_1 of revenue $1/(2a)$ and then cuts through $P_2, P_3, \ldots P_{2a}$, serving a request of revenue $1/(2a)$ from each P_i, followed by an empty drive (gray dashed edges) of a sufficiently small distance $\epsilon > 0$ after each request, accruing a total revenue of 1. (Note: ϵ must be small enough so that Q_1 can be served within one time segment.)

Setting $c \leq 1$, no paths of total distance 1 or less have revenue larger than 1, so we can assume that SBP will move to and then drive along Q_1 during the first two time segments. At time $t = 2$, each path P_i now has remaining revenue $(1 - \frac{1}{2a})$, which we denote by ρ.

Now suppose we have path Q_2 (depicted as the magenta and gray staircase in the figures) that similarly cuts through $P_2, P_3, \ldots, P_{2a}, P_1$, serving a request of revenue $\rho/(2a)$ each, so that Q_2 has total revenue ρ. Again, as long as $c \leq \rho$, no path of length 1 or less has revenue more than 1 at time $t = 0$, and no path of length 1 or less has revenue more than ρ at time $t = 2$. So we can assume that SBP moves to and serves Q_2 from time $t = 2$ to time $t = 4$.

In general, define Q_i for $i = 1 \ldots b$ as a path that cuts through $P_i, P_{i+1} \ldots P_{2a}, P_1 \ldots P_{i-1}$, serving a request of revenue $\rho^{i-1}/(2a)$ each, so that Q_i has a total revenue of ρ^{i-1}. To ensure that SBP chooses Q_i in time segment $[2i - 2, 2i]$, we need $c \leq \rho^{i-1}$, for all $i = 1 \ldots b$. Note that each Q_i, $1 \leq i \leq b$, consists of $4a - 1$ drives: $2a$ requests with an empty drive of distance ϵ between each pair for a total of $2a - 1$ empty drives.

After time $t = 2b$, the remnants of each path P_i have revenue ρ^b. (After SBP serves Q_i for $i = 1 \ldots b$, the remaining revenue of each P_i shrinks by a factor of ρ.) If $c \geq \rho^b$, SBP may serve paths of revenue c for the remainder of time. We choose $c = \rho^b$, so SBP serves $a - b$ paths of revenue ρ^b for the remainder of time. Summarizing,

$$|\text{OPT}| = 2a + (2a - 1)\rho^b = 2a + (2a - 1)\left(1 - \frac{1}{2a}\right)^b$$

and since $1 - \rho = 1/(2a)$,

$$|\text{SBP}| = 1 + \rho + \rho^2 + \ldots \rho^{b-1} + (a - b)\rho^b = (1 - \rho^b)/(1 - \rho) + (a - b)\rho^b$$

$$= 2a(1 - \rho^b) + (a - b)\rho^b = 2a - \rho^b(a + b) = 2a - \left(1 - \frac{1}{2a}\right)^b (a + b).$$

Table 1 Some sample instance parameters and their corresponding ratios.

b	a	$\lvert\text{OPT}\rvert/\lvert\text{SBP}\rvert$
2	6	4.025
3	6	4.03985
3	10	4.09867
1000000	1877946	4.27805

If $a = b$, we can take the limit as $a \to \infty$ to get $\rho^b = (1 - \frac{1}{2a})^a$ is $1/\sqrt{e}$. Then $|\text{OPT}|/|\text{SBP}|$ has a limit of $(\sqrt{e}+1)/(\sqrt{e}-1) \approx 4.08299$.

Table 1 shows some sample instance parameters and their corresponding ratios. The instance shown in Figure 2 is reflected in the second row. Thus far, preliminary testing suggests that a ratio much greater than 4.27805 (in the final row of the table) is unachievable.

Since our upper bound is tight only when f is odd, we continue to investigate a version of the SBP algorithm that enforces an even number of time segments. An open question is if the upper bound of 5 is no longer tight for this adjusted algorithm, and whether the true upper bound matches the above family of instances, or can be shown to be strictly below 5.

References

1 Barbara M. Anthony, Ananya D. Christman, Christine Chung, and David Yuen. Serving Rides of Equal Importance for Time-Limited Dial-a-Ride. In Panos Pardalos, Michael Khachay, and Alexander Kazakov, editors, *Mathematical Optimization Theory and Operations Research*, pages 35–50. Springer International Publishing, 2021. doi:10.1007/978-3-030-77876-7_3.

2 Egon Balas. The prize collecting traveling salesman problem. *Networks*, 19(6):621–636, 1989. doi:10.1002/net.3230190602.

3 Daniel Bienstock, Michel X. Goemans, David Simchi-Levi, and David Williamson. A note on the prize collecting traveling salesman problem. *Mathematical programming*, 59(1-3):413–420, 1993. doi:10.1007/BF01581256.

4 Jannis Blauth and Martin Nägele. An improved approximation guarantee for prize-collecting TSP. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, pages 1848–1861, New York, NY, USA, 2023. doi:10.1145/3564246.3585159.

5 Avrim Blum, Shuchi Chawla, David R. Karger, Terran Lane, Adam Meyerson, and Maria Minkoff. Approximation algorithms for orienteering and discounted-reward TSP. *SIAM Journal on Computing*, 37(2):653–670, 2007. doi:10.1137/050645464.

6 Ananya Christman, Christine Chung, Nicholas Jaczko, Marina Milan, Anna Vasilchenko, and Scott Westvold. Revenue Maximization in Online Dial-A-Ride. In *17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017)*, volume 59, pages 1:1–1:15, Dagstuhl, Germany, 2017. doi:10.4230/OASIcs.ATMOS.2017.1.

7 Ananya D. Christman, Christine Chung, Nicholas Jaczko, Tianzhi Li, Scott Westvold, Xinyue Xu, and David Yuen. Improved Bounds for Revenue Maximization in Time-Limited Online Dial-a-Ride. *SN Operations Research Forum*, 2(3):1–38, September 2021. doi:10.1007/S43069-021-00076-X.

8 Sin C. Ho, W.Y. Szeto, Yong-Hong Kuo, Janny M.Y. Leung, Matthew Petering, and Terence W.H. Tou. A survey of dial-a-ride problems: Literature review and recent developments. *Transportation Research Part B: Methodological*, 111:395–421, 2018.