Computing Headway Bounds under Worst-Case Bunching in Fixed-Line Transit Systems

Michael Yuhas[1,2], George Gunter[1], Jose Paolo Talusan[1], Aron Laszka[3], Dan Freudberg[4], Abhishek Dubey[1]
[1]*Institute for Software Integrated Systems*, *Vanderbilt University*, Nashville, TN, USA
[2]*School of Science, Engineering and Technology*, *St. Mary's University*, San Antonio, TX, USA
[3]*College of Information Sciences and Technology*, *Pennsylvania State University*, University Park, PA, USA
[4]*WeGo Public Transit*, Nashville, TN, USA

***Abstract*—Vehicle bunching is a major problem for transit operators. When vehicles bunch together, the lead vehicle will service the majority of passenger demand, leaving the following vehicles to operate below capacity, wasting fuel and money. Furthermore, after the last vehicle in the bunch passes, the time before the next vehicle's arrival (headway) will be large. Transit operators can combat bunching by holding buses at stops along a route, trading riding time for even headway times. While prior work has focused on developing holding policies to minimize average case bunching, no work has focused on analyzing the longest and shortest possible headway times under a broad group of such policies. We assume that dwell times at stops and travel times between stops are bounded and develop a dynamic program that computes the maximum and minimum headway times for a single bus route with an arbitrary number of control points, vehicles, and holding policies. These bounds are tight in the sense that it is always possible to identify the specific sequence of events that lead to their occurrence. We use these bounds to investigate the effects of different holding policies, stop placement, and number of vehicles on route headways and worst-case bunching. Finally, we apply these analysis techniques to a real-world transit system in Nashville, TN and show their utility for transit planning.**

I. Introduction

In a fixed-line transit system, vehicles travel along predetermined routes and serve a set of predetermined stops, e.g., a bus line in an urban area. Fig. 1 provides a visualization of the stop order and a snapshot of the vehicle locations along such a route at a fixed point in time in an example system. While vehicles often depart from the terminal according to a fixed schedule, as they travel along their route, random variations in travel time and number of passengers lead to bunching: when many vehicles clump together and service stops in rapid succession. A common pattern is when the lead vehicle encounters a large number of passengers boarding or alighting at a single stop: this delays the vehicle longer than expected, giving its follower a chance to catch up. Because less time has passed between two vehicles servicing the same stop, there are fewer passengers waiting to board at the subsequent stop and the follower spends less time waiting, causing the gap to progressively narrow [1]. This leads to vehicles clustering in *bunches*, with the gap between bunches increasing beyond what is scheduled [2]. This is bad from a transit operator's perspective (the follower buses in a bunch are now operating below capacity) and from a rider's perspective (if the rider arrives at a stop after the last vehicle in a bunch has passed, they must now wait longer to begin their journey).

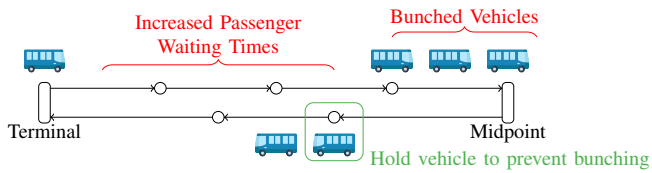

Fig. 1. As vehicles bunch, the headway for the vehicle behind the bunch increases. Operators can hold vehicles at certain stops to regulate headway but cannot make vehicles speed up to close the gap.

In fixed-line transit planning, the interval of time between vehicles at a stop is referred to as *headway* [3]. In general, smaller headways are good for passengers, ensuring frequency of service, however, bunching leads to a scenario where some headway times are extremely small, leading to inefficiencies, and other headway times are extremely large, disgruntling riders. The goal then is, given a fixed number of vehicles serving a route, to even out headway times, such that a rider's waiting time is bounded. Transit operators can employ a variety of strategies to reduce bunching and manage headway. One strategy is to operate buses on a timetable at intermediate stops along the route, introducing delays for buses that arrive early [4]. With real-time vehicle location tracking, operators can also dynamically issue hold commands to vehicles getting too far ahead of their followers [5]. The hold times can be determined by a human operator, rules-based logic, or data driven model [6]. While holding a vehicle leads to an increase in in-vehicle travel time for passengers already aboard a bus, studies have shown that passengers perceive wait time at stops as worse than time spent traveling after boarding [7]. Furthermore, in areas with extreme weather (heat in the summer and cold in the winter), it can be dangerous for riders to wait for long periods at unsheltered stops. Thus, many transit agencies choose to pursue such policies to improve rider experience and improve operational efficiency.

Prior analysis of headway times has focused heavily on statistical models: travel times between stops and dwell times at each stop are drawn from distributions fit to real-world data with the objective of optimizing the average case [8]. While this is important, human perception of a transit system

Published in the Proceedings of the 32nd IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2026)

is not necessarily shaped by the average case behavior [9] and transit operators cannot make operational decisions based only on average case behavior. Thus, in addition to providing an analysis of the average case performance of various holding policies, it is necessary to provide a bound (both upper and lower) on headway times at each stop. While these times are potentially unbounded (e.g. due to vehicle breakdown), during normal operation it is safe to assume bounded dwell times at stops [10] and travel times between stops [11]. Furthermore, if dwell or travel times from the long tail of these distributions occur, they are easy to detect and flag as violations.

Contributions of this paper – We propose an analytical bound on the maximum and minimum headway times for a single transit route with an arbitrary number of stops, each governed by an arbitrary holding policy and serviced by an arbitrary number of vehicles. This work applies real-time system analysis to a cyber-physical system: automated headway control policies gather data through remote sensing and effect a change in a physical system by holding vehicles, which in turn affects the timing performance of system. To the best of our knowledge this is the first work of its kind: to place analytical bounds on headway times under broad classes of holding policies. This work allows transit operators to evaluate their current holding policies and guide the selection of new proposals to ensure a safer and more enjoyable ride for all passengers as well as more economical operation. We use our bounds to provide an analysis of the tradeoffs involved in selecting a policy and include a case study using data from a real bus route in Nashville, TN.

II. Background

The terminology and metrics used in this paper have been defined differently by different authors in slightly different transit contexts. *Bunching* itself is seldom formally defined (definitions on how many and how far apart vehicles must be to constitute a bunch vary) but all authors agree that it can be measured indirectly through the unevenness of *headway* times [12]. We use the Federal Transit Administration's definition of headway: "the time interval between vehicles moving in the same direction on a particular route" [3]. Other authors use this working definition, but the expected headway at stops is the only location where headway directly impacts system performance [13]. In literature, *tailway* is also used to describe the time between a lead vehicle and its follower [14], but this terminology is not as popular.

Many works focus on the problem of managing headway times. Infrastructure enhancements, such as traffic signal priority for buses can give vehicles falling behind schedule a chance to regulate headway [15], however we do not consider them in this work as the required infrastructure is not readily available in every locality. Other works consider using a headway control tool to give commands to speed up or slow down to drivers at any point along route [16], but the commands may not always be carried out depending on driver skills and traffic conditions. The bulk of prior research on mitigating bunching focuses on holding buses at specific stops along a route called *control points* [6]. We focus on analyzing these techniques since they are relatively easy for operators to implement without costly infrastructure upgrades and whether or not a driver holds is not subject to external factors like traffic. In general, these fall into two categories: schedule-based holding and headway-based holding. In schedule-based holding policies, operators focus on creating a schedule for buses that provides enough slack to mitigate headway variations, while keeping vehicle frequency high and travel times low [4]. The other approach is to use information about the headway between vehicles to dynamically adjust holding time based on real-world conditions. Daganzo provided a stochastic analysis of holding based only on headway between two vehicles on a fixed line transit route [17]. Cats et al. developed a simulator to track per stop waiting time distributions and considered two dynamic headway-based policies: always maintaining a minimum headway at a stop and holding so as to even headway times across the route [18]. Zhou et al. considered holding based on the ratio of a vehicle's headway to tailway and included a minimum and maximum holding time for operational convenience [19]. Fonzone et al. extend these works by considering the case when travel and dwell times are dependent on non-uniform passenger arrival rates [20]. Recently, works have also considered using reinforcement learning to determine holding time based on headway [21]. In all these studies, however, the priority is always optimizing the average case behavior, not the headway bounds.

III. System Model

We now present a mathematical model of a fixed-line transit route. While this work focuses on bus transit systems, the same model generalizes to other fixed-line transit modalities. A table of notation is provided in Table I. We use bold letters to represent a vector of values and italic text with subscripts to represent individual elements, e.g., x_i is the ith value in $\mathbf{x}$. For compactness, we use the notation $[\cdot]_{lb}^{ub}$ to represent a quantity clamped between ub and lb, i.e., $[x]_{lb}^{ub} \triangleq \max\{lb, \min\{ub, x\}\}$. We also indicate the upper bound of a variable with an overline and a lower bound with and underline (e.g. $\underline{x} \leq x \leq \overline{x}$).

A. Transit System

We consider a route as a sequence of N stops $s^{(0)} \dots s^{(N)}$. Each stop has an associated dwell time, which is the time a vehicle spends at the stop to load and unload passengers. This is typically assumed to be drawn from a distribution [22], which we assume has bounded support $[\underline{\delta}^{(i)}, \overline{\delta}^{(i)}]$ at stop $s^{(i)}$ where $\underline{\delta}^{(i)}$ is the minimum possible dwell time and $\overline{\delta}^{(i)}$ is the maximum possible dwell time. In the real world, dwell time is always lower bounded by at least $\underline{\delta}^{(i)} = 0$ (corresponding to the case where a driver skips a stop) and an unbounded dwell time ($\overline{\delta}^{(i)} = \infty$) implies a bus break down or some other extreme event that prevents a bus from completing its route. In practice, we care about the upper bound on dwell time that corresponds to some percentile of the underlying distribution. This allows us to consider how well a holding policy performs

TABLE I

TABLE OF NOTATION. FOR CONSISTENCY, A SUPERSCRIPT ALWAYS DENOTES STOP INDEX AND SUBSCRIPT ALWAYS DENOTES VEHICLE INDEX.

Symbol	Definition
N	Number of stops in a route
$s^{(i)}$	The ith stop along a route
M	Number of vehicles serving a route
$r_j^{(0)}$	The time the jth vehicle begins service at stop $s^{(0)}$
$\underline{\tau}^{(i)}/\overline{\tau}^{(i)}$	The minimum/maximum travel time for the route segment between stops $s^{(i)}$ and $s^{(i+1)}$
$\tau_j^{(i)}$	The travel time for the jth vehicle to depart from $s^{(i)}$ on the route segment between stops $s^{(i)}$ and $s^{(i+1)}$
$\underline{\delta}^{(i)}/\overline{\delta}^{(i)}$	The minimum/maximum dwell time at stop $s^{(i)}$
$\delta_j^{(i)}$	The dwell time for the jth vehicle to arrive at $s^{(i)}$
$\pi^{(i)}(\cdot)$	The policy in effect at stop $s^{(i)}$
$\underline{\phi}_j^{(i)}/\overline{\phi}_j^{(i)}$	The minimum/maximum policy holding time for all buses at stop $s^{(i)}$
$\phi_j^{(i)}$	The delay imposed by holding policy $\pi^{(i)}$ for the jth vehicle to arrive at stop $s^{(i)}$
$\underline{a}_j^{(i)}/a_j^{(i)}/\overline{a}_j^{(i)}$	The earliest/actual/latest arrival time of the jth vehicle at $s^{(i)}$
$\hat{a}_j^{(i)}$	The estimated time of the jth arrival at $s^{(i)}$
$\underline{d}_j^{(i)}/d_j^{(i)}/\overline{d}_j^{(i)}$	The earliest/actual/latest departure time of the jth vehicle from $s^{(i)}$
$h_j^{(i)}$	The jth vehicle's headway at $s^{(i)}$
$\overline{HT}^{(i)}$	The upper bound on headway time at stop $s^{(i)}$
$\underline{HT}^{(i)}$	The lower bound on headway at $s^{(i)}$

given that dwell times conform to a higher density region of the underlying distribution and assign a probability that our analytical bound will not be exceeded during operation. Likewise, we assume the travel time between stops $s^{(i)}$ and $s^{(i+1)}$ are drawn from the interval $[\underline{\tau}^{(i)}, \overline{\tau}^{(i)}]$. By the laws of physics (and most countries) a lower bound must exist and $\underline{\tau}^{(i)} > 0$. In the real world $\overline{\tau}^{(i)}$ could be unbounded due to an extreme event like a breakdown, but given a past distribution of travel times between stops, it is possible to select $\overline{\tau}^{(i)}$ to cover any arbitrary percentile of travel times. Prior studies have considered bounded travel and dwell times [10], [11].

We consider a fleet of M vehicles serving stops along the route. Each vehicle must serve stops $s^{(0)} \dots s^{(N)}$ in order. We consider the case where vehicles can overtake each other either during travel or at stops, which is possible and common in bus transit systems [23]. We also assume no vehicles are added or removed while serving the route. While operators generally add vehicles for peak hours, the following analysis can simply be redone with the new number of vehicles. We assume each vehicle is released into service at stop $s^{(0)}$ at time $r_j^{(0)}$ where j is the order in which the vehicle enters service (i.e., $j < j' \implies r_j^{(0)} < r_{j'}^{(0)}$). The arrival time of the jth vehicle at stop $s^{(i)}$ is given by $a_j^{(i)}$. Note that j is the index of arrival since the start of service, so j counts up from $0 \dots \infty$. Each stop is governed by a holding policy, which instructs the driver of the jth vehicle to wait at the stop for at least $\phi_j^{(i)}$. The holding policy at the ith stop is a function $\pi^{(i)}(\cdot) : X \mapsto \mathbb{R}_+$ where X is the system state. We assume that $\phi_j^{(i)}$ is computed upon a vehicle's arrival $a_j^{(i)}$ and cannot be subsequently updated.

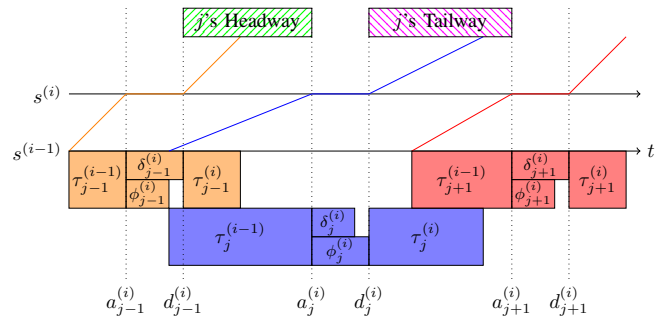

Fig. 2. Illustration of headway and tailway calculation for the jth arrival at stop $s^{(i)}$. Vehicle $j-1$ is the leader and $j+1$ is the follower making vehicle j's headway $h_j^{(i)} = a_j^{(i)} - d_{j-1}^{(i)}$ and the tailway $h_{j+1}^{(i)} = a_{j+1}^{(i)} - d_j^{(i)}$.

This assumption is based on real-world constraints where a bus driver can only safely check for a policy induced hold when stopped and a hold request cannot be rescinded once issued. Policy induced holding occurs concurrently with dwell time, thus the total time the jth arrival spends stationary at $s^{(i)}$ is $\max\{\delta_j^{(i)}, \phi_j^{(i)}\}$ where $\delta_j^{(i)} \in [\underline{\delta}^{(i)}, \overline{\delta}^{(i)}]$. We let $d_j^{(i)}$ represent the jth vehicle departure at $s^{(i)}$. Letting $\tau_j^{(i)} \in [\underline{\tau}^{(i)}, \overline{\tau}^{(i)}]$ be the travel time of the jth departure from $s^{(i)}$ to $s^{(i+1)}$, the progression of a single vehicle from its arrival at $s^{(i)}$ to its arrival at the subsequent stop is then given by Eqs. (1) and (2). Note that the jth arrival is not necessarily the jth departure, so Eq. (1) is written in terms of the j'th departure – likewise the j''th arrival in Eq. (2) – where $j, j', j'' \in 0 \dots \infty$.

$$d_{j'}^{(i)} = a_j^{(i)} + \max\{\delta_j^{(i)}, \phi_j^{(i)}\} \tag{1}$$

$$a_{j''}^{(i+1)} = d_{j'}^{(i)} + \tau_j^{(i)} \tag{2}$$

Fig. 2 provides an illustration of the time for three consecutive vehicles to travel between $s^{(i-1)}$ and $s^{(i)}$.

B. Metrics

Headway time refers to the spacing between vehicles, i.e., the duration between the departure of one vehicle and the arrival of the subsequent vehicle at the same station [3]. For the purposes of this paper, headway time is only defined at stops, not every point along a route; riders are only impacted by headway as measured at a station and the economic impact of uneven headway times is only realized based on vehicle occupancy, which only changes at the stops.

Definition 1. *The headway $h_j^{(i)}$ of the jth arrival at $s^{(i)}$ is given by Eq. (3).*

$$h_j^{(i)} = \max\{0, a_j^{(i)} - d_{j-1}^{(i)}\} \tag{3}$$

Note 1. *$h^{(i)}$ is undefined at each $s^{(i)}$ until the second vehicle arrives ($t \geq a_1^{(i)}$). Furthermore, at $s^{(0)}$ headway remains undefined until the release of the last vehicle into service ($t \geq r_M^{(0)}$).*

Fig. 2 illustrates the concept of headway calculation. For convenience, we also introduce the term *tailway*, which refers

to the headway of the subsequent arrival, i.e., $h_{j+1}^{(i)}$. While many studies focus on computing the average headway time, we focus on computing headway time bounds.

Definition 2. *$\overline{HT}^{(i)}$ is the upper bound on headway time that can occur at $s^{(i)}$ during the service period and $\underline{HT}^{(i)}$ is the lower bound, i.e.:*

$$\overline{HT}^{(i)} = \max_{0<j<\infty} \{h_j^{(i)}\} \tag{4}$$

$$\underline{HT}^{(i)} = \min_{0<j<\infty} \{h_j^{(i)}\} \tag{5}$$

Given a fixed number of vehicles serving a route, bunching causes $\overline{HT}^{(i)}$ to increase and $\underline{HT}^{(i)}$ to decrease. In areas where waiting for the bus means standing unsheltered in extreme cold or heat, a bound on maximum time spent waiting ($\overline{HT}^{(i)}$) may be more important than the average value. Likewise, bounding minimum headway can ensure good bus utilization and efficient use of resources.

C. Bus Holding Policies

A holding policy guarantees that a vehicle remains stationary at a stop for a minimum amount of time regardless of $\delta_j^{(i)}$. This interval is determined by the holding policy at the stop $\pi^{(i)}(\cdot)$, which computes $\phi_j^{(i)}$ for every vehicle at the moment it reaches the stop. This reflects real world conditions where issuing policy updates when a driver is approaching the end of their hold period can result in the updates being ignored due to the notification not arriving on time or the driver being preoccupied with other tasks. On the other hand, most control systems seek to apply actions as late as possible so as to minimize the affect of random variations in system state that happen between when a hold is requested and when it is executed. Since the vehicle's dwell time at a stop is unknown and drawn from a distribution, arrival is the latest possible moment when this should be calculated. Policies are applied on a per stop basis, not to individual vehicles and not across an entire route. The reason for this is practical: it is not feasible for operators to idle vehicles at every stop, or the time may be limited due to infrastructure constraints. Consider, for example, stops along a two-lane road with no turnouts: in these cases, a vehicle can still stop and dwell for passengers to alight and board, but idling for too long interferes with other traffic. Furthermore, some stops along a route may require a vehicle to synchronize with other transit times that cannot be adjusted (e.g. commuter rail). The space of possible policies is infinite, so we focus on three sets that encompass the policies most commonly deployed by transit operators. In every case, a policy is a function that maps some set of state measurements to a positive real number (hold time), so machine learning-based policies fit into our analysis framework as long as they are a member of one of the following sets.

1) Null Policy (πNIL): If $\pi^{(i)} \in \pi\texttt{NIL}$ vehicles are never held, i.e., for stop $s^{(i)}$, $\phi_j^{(i)} = 0 \ \forall j \in 0 \ldots \infty$. This is the default policy for most stops in real world transit systems.

2) Static Schedule-Based Policies (πSCH): Under this set of policies vehicles depart stops at predetermined times. We let $\mathbf{T}$ be a list of scheduled departure times. Only one vehicle is permitted to depart from the stop at each of the entries in $\mathbf{T}$. If multiple vehicles are held at the stop waiting for the next scheduled departure, we assume they depart in the order of their arrival, i.e., a stop governed by this policy behaves as a FIFO queue. If no vehicle departs at a time listed in $\mathbf{T}$, the next vehicle to arrive at the stop receives no holding time, corresponding to the case where vehicles are running behind schedule. Mathematically, we model this as the stateful function shown in Eq. (6) where l is the index of the last used departure time in $\mathbf{T}$ and t is the current time. Note that if dwell time $\delta_j^{(i)}$ exceeds the holding time $\phi_j^{(i)}$, it is still possible for a vehicle to depart after its scheduled release.

$$\pi^{(i)}(t; l, \mathbf{T}) = \max\{0, \mathbf{T}_l - t\} \tag{6}$$

3) Headway-Based Control (πDYN): Under this set of policies the time a vehicle idles at a given stop is a function of the locations of other vehicles along the route. The goal is to hold a vehicle so that the gap between headway is roughly equal across all arrival and departure pairs. We assume that the holding calculated by headway-based control is only a function of the current (and past) positions of vehicles along a route, i.e., $\pi^{(i)}(X) : \mathbb{R}_+^{2N\times M} \mapsto \mathbb{R}_+$. where X refers to the system state – vectors of the last arrival times $\mathbf{a}^{(i')}$ and departure times $\mathbf{d}^{(i')}$ for each of M vehicles at stop i', for $i' \in [0 \ldots N]$. Note that it is not required for a dynamic policy to use all of this information and only a subset may be required. For our analysis to work, we make the following assumptions about $\pi^{(i)} \in \pi\texttt{DYN}$:

1) The output is bounded: $0 \leq \phi_j^{(i)} < \infty$.
2) The policy is 1-left-Lipschitz [24] with respect to headway, i.e., when headway is computable from the system state, the following inequality must hold: $\pi^{(i)}(h_j^{(i)}) - \pi^{(i)}(h_{j'}^{(i)}) \leq h_{j'}^{(i)} - h_j^{(i)} \forall h_j^{(i)} \leq h_{j'}^{(i)}$.

Assumption (1) is a requirement for any practical policy as unbounded $\phi_j^{(i)}$ would mean a bus is stuck at a stop forever. Assumption (2) requires that a vehicle arriving with more tailway is not held longer than if it had arrived with less tailway. Prior work has shown it is possible to train Lipschitz bounded [25] neural networks, so πDYN still represents a broad class of data driven policies. Many headway-based holding policies [17], [19] already satisfy these properties.

Problem Statement – The goal of this paper is to identify the upper and lower bounds on headway times for all stops in any bus route modeled by the system described above. This problem is challenging because headway bounds are dependent not just on the maximum and minimum travel and dwell times between two stops, but also all the possible travel and dwell times between any two stops taken by every vehicle servicing the route, at every time since the beginning of service, as these affect policy holding for stops governed by πSCH and πDYN, which in turn affect $\overline{HT}^{(i)}$ and $\underline{HT}^{(i)}$.

IV. Headway Time Bounds

We first show that in a route with all stops governed by $\pi^{(i)} \in \pi\texttt{NIL}$ (and in some cases $\pi\texttt{SCH}$) headway is upper bounded by the maximum travel time and lower bounded by 0. Finally, we use a dynamic program to establish tight upper and lower headway bounds for any route. In order to accomplish this, we must formally define a *bunch*.

Definition 3. *A bunch is a sequence of v vehicles with leader arriving at $s^{(i)}$ at $a_j^{(i)}$ and the last bus in the bunch departing the same $s^{(i)}$ at $d_{j+v-1}^{(i)}$. The time for a bunch to pass the stop is given as $d_{j+v-1}^{(i)} - a_j^{(i)}$.*

A. Insufficiency of Static Policies

Theorem 1. *In a route with all stops governed by $\pi\texttt{NIL}$, if $\overline{\tau}^{(i)} > \underline{\tau}^{(i)}$ or $\overline{\delta}^{(i)} > \underline{\delta}^{(i)}$ for at least one stop $s^{(i)}$, then $\underline{HT}^{(k)} = 0$ and $\overline{HT}^{(k)} = \sum_{i'=0\ldots N}(\overline{\tau}^{(i')} + \overline{\delta}^{(i')}) - \overline{\delta}^{(k)} \forall k \in 0\ldots N$ (which is equivalent to the maximum time required to travel the entire route, including dwell time).*

Proof. The maximum and minimum headway time occurs when all M vehicles are bunched together such that $\exists j | a_j^{(k)} = a_{j+m}^{(k)} \forall m \in 0\ldots M$, i.e., all buses arrive and depart $s^{(k)}$ simultaneously ($\underline{HT}^{(k)} = 0$). Given this precondition, it is always possible for all M vehicles to take the maximum travel time and dwell times for all stops back to $s^{(k)}$, i.e., all vehicles arrive at and depart from the remaining stops simultaneously. This is functionally equivalent to a single vehicle servicing the route, in which case maximum headway is the time between the vehicle departing $s^{(k)}$ and arriving after traversing the entire route: $\overline{HT}^{(k)} = \sum_{i'=0\ldots N}(\overline{\tau}^{(i')} + \overline{\delta}^{(i')}) - \overline{\delta}^{(k)}$ (recall that under $\pi\texttt{NIL}$, $\phi_j^{(i)} = 0\ \forall i, j$). To show that $a_j^{(k)} = a_{j+m}^{(k)} \forall m \in 0\ldots M$ is always possible, assume either $\overline{\tau}^{(i)} > \underline{\tau}^{(i)}$ or $\overline{\delta}^{(i)} > \underline{\delta}^{(i)}$ for at least one stop $s^{(i)}$ along the route. Consider two vehicles at $s^{(i)}$: a leader $j-1$ and a follower j. Under the system model they can take the same transit and dwell times for all other stops excluding $s^{(i)}$. Let $h_j^{(i)} = a_j^{(i)} - d_{j-1}^{(i)}$ denote the initial headway time between leader and follower. Suppose that for the next $\lfloor h_j^{(i)} / (\overline{\tau}^{(i)} - \underline{\tau}^{(i)}) \rfloor$ arrivals at $s^{(i)}$, the leader takes $\overline{\tau}^{(i)}$ to complete $s^{(i)}$'s travel and the follower takes $\underline{\tau}^{(i)}$. On the very next iteration, if the leader still takes $\overline{\tau}^{(i)}$ to complete the leg, and the follower takes $\underline{\tau}^{(i)} + h_j^{(i)} \bmod (\overline{\tau}^{(i)} - \underline{\tau}^{(i)})$ the headway between them will be zero. In the worst case, both vehicles now take the same time to complete each route segment. The same reasoning can be applied for $\overline{\delta}^{(i)} > \underline{\delta}^{(i)}$. □

Note 2. *It follows from Theorem 1 that for a route with stop $s^{(k)}$ controlled by $\pi^{(k)} \in \pi\texttt{SCH}$, as long as $\sum_{i=k'\ldots k}(\overline{\tau}^{(i)} + \overline{\delta}^{(i)} - \underline{\tau}^{(i)} - \underline{\delta}^{(i)}) > \mathbf{T}_l - \mathbf{T}_{l-1} \forall l = 0\ldots\infty$ where k' is the stop index after the last schedule controlled stop and all other stops are controlled by $\pi\texttt{NIL}$, then the maximum headway time is equivalent to that given in Theorem 1.*

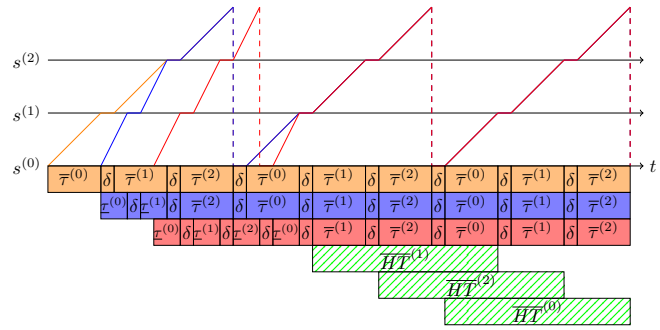

Fig. 3. Illustration of the sequence of events that leads to the worst case upper and lower bounds in Theorem 1 for a route with 3 stops and 3 vehicles.

Note 2 implies that even if transit operators use schedule-driven control at the terminal, it will not be sufficient to guarantee that headway is less than the route completion time if it is possible for vehicles to arrive behind schedule. To build intuition, Fig. 3 illustrates how bunching can lead to the bounds established in Theorem 1 and Note 2 in a route with 3 stops and 3 buses.

B. Headway Bounds for Routes with Any Set of Policies

Further analysis is required to make a meaningful statement about headway bounds in routes with stops under $\pi\texttt{DYN}$. By our definition $\pi\texttt{NIL} \subset \pi\texttt{DYN}$ and strict monotonicity with respect to headway is not required, so there is no guarantee that bunch times increase on the output of a stop governed by $\pi\texttt{DYN}$. Instead, we first focus on identifying the condition that leads to maximum headway (when vehicles are as tightly bunched as possible) and then show the conditions that must occur for bunch time to be minimized.

Lemma 2. *Given that the $j-1$st departure from $s^{(i-1)}$ occurs at fixed time $d_{j-1}^{(i-1)}$, a later arrival of the jth vehicle at $s^{(i-1)}$ can always lead to a later arrival of the jth vehicle at $s^{(i)}$, i.e.:*

$$a_j^{(i-1)} \leq a_j'^{(i-1)} \implies \overline{a}_j^{(i)} \leq \overline{a}_j'^{(i)} \tag{7}$$

Proof. When $\overline{\delta}^{(i-1)} \geq \max\{\phi_j'^{(i-1)}, \phi_j^{(i-1)}\}$, the maximum arrival time at $s^{(i)}$ is given by $\overline{a}_j^{(i)} = a_j^{(i-1)} + \overline{\delta}^{(i-1)} + \overline{\tau}^{(i-1)}$ or $\overline{a}_j'^{(i)} = a_j'^{(i-1)} + \overline{\delta}^{(i-1)} + \overline{\tau}^{(i-1)}$. Algebraic manipulation reveals that $\overline{a}_j^{(i)} \leq \overline{a}_j'^{(i)}$. When $\overline{\delta}^{(i-1)} < \max\{\phi_j'^{(i-1)}, \phi_j^{(i-1)}\}$, we need to show that the departure of the later arrival is always later, i.e., $d_j^{(i-1)} \leq d_j'^{(i-1)}$. While this is not true for any arbitrary policy, we show that it is true for the three policy classes considered in our formulation:

- $\pi^{(i-1)} \in \pi\texttt{NIL}$: $\phi_j^{(i-1)} = \phi_j'^{(i-1)}$. This means departure time equals arrival time for both $a_j^{(i-1)}$ and $a_j'^{(i-1)}$, thus $d_j^{(i-1)} \leq d_j'^{(i-1)}$.
- $\pi^{(i-1)} \in \pi\texttt{SCH}$: Consider the next scheduled release time $\mathbf{T}_l$. For the jth arrival: $\phi_j^{(i-1)} = \max\{0, \mathbf{T}_l - a_j^{(i-1)}\}$. For $a_j'^{(i-1)} \geq a_j^{(i-1)}$, j will either have the same scheduled

release time or a later one: $\phi_j'^{(i-1)} = \max\{0, \mathbf{T}_{l+n} - a_j'^{(i-1)}\} | n \in 0 \dots \infty$.

- $\pi^{(i-1)} \in \pi\texttt{DYN}$: Recall that the dynamic policy is 1-left-Lipschitz with respect to headway. Since the lead vehicle $j-1$ departs at $d_{j-1}^{(i-1)}$ regardless if j arrives at $a_j^{(i-1)}$ or $a_j'^{(i-1)}$, headway is only impacted by arrival time. Ignoring dwell time, substituting $a_j^{(i-1)} - d_{j-1}^{(i-1)}$ and $a_j'^{(i-1)} - d_j^{(i-1)}$ for headway into the definition of a 1-left-Lipschitz function gives the result $a_j^{(i-1)} + \phi_j^{(i-1)} \leq a_j'^{(i-1)} + \phi_j'^{(i-1)}$. Thus even if the dynamic policy holds a vehicle arriving later longer ($\phi_j'^{(i-1)} > \phi_j^{(i-1)}$), at the very latest, the vehicle will still depart $s^{(i-1)}$ at the same time as if it had arrived earlier, i.e., $d_j^{(i-1)} \leq d_j'^{(i-1)}$.

Even if a vehicle departing at $d_j^{(i-1)}$ takes $\overline{\tau}_j^{(i-1)}$, a vehicle departing at $d_j'^{(i-1)}$ could do the same, meaning $\overline{a}_j^{(i)} \leq \overline{a}_j'^{(i)}$. □

Theorem 3. *The minimum time for a bunch of M vehicles to pass a stop $s^{(i)}$ is a necessary condition for the maximum headway between the $j+M-1$st departure and the $j+M$th arrival for any $j \in 0 \dots \infty$ at that same stop.*

Proof. The headway between the last vehicle in the bunch and the next arrival is given by $h_{j+M}^{(i)} = a_{j+M}^{(i)} - d_{j+M-1}^{(i)}$. If we fix the $j+M-1$st vehicle's departure at $d_{j+M-1}^{(i)}$, then by Lemma 2 the latest possible arrival $\overline{a}_{j+M}^{(i-1)}$ is a necessary condition for $\overline{a}_{j+M}^{(i)}$, which maximizes $h_{j+M}^{(i)}$. Likewise, $\overline{a}_{j+M}^{(i-n)}$ is a necessary condition for $\overline{a}_{j+M}^{(i-n+1)}$ all the way back to stop $s^{(0)}$. Assuming all M vehicles were released from $s^{(0)}$ (which is a requirement of our system model), then $\overline{a}_{j+M}^{(0)}$ corresponds to $\overline{a}_j^{(N)}$. Again, by Lemma 2, $\overline{a}_j^{(N-n)}$ is a necessary condition for $\overline{a}_j^{(N-n-1)}$ all the way back to stop $s^{(i)}$. Thus, $\overline{a}_j^{(i)}$ is a necessary condition for the maximum headway to occur. Under Definition 3, the time taken for a bunch of M vehicles to pass the stop is given by $d_{j+M-1}^{(i)} - a_j^{(i)}$. This value is minimized when $a_j^{(i)} = \overline{a}_j^{(j)}$. □

From Theorem 3, it is clear that in addition to the condition that the time for a bunch of M vehicles to pass a stop is minimized, another condition exists: the lead vehicle in the bunch must always achieve the maximum arrival time ($\overline{a}_j^{(n)} | i < n \leq N$ and $\overline{a}_{j+M}^{(n)} | 0 \leq n \leq i$) for all stops in the route until it comes back to service $s^{(i)}$. Together these two conditions are sufficient to guarantee $\overline{HT}^{(i)}$. We now must show what conditions are required to cause the minimum bunch time to occur.

Lemma 4. *Given a bunch of v vehicles with leader arriving at $s^{(i)}$ at time $a_j^{(i)}$ and last vehicle departing $s^{(i)}$ at $d_{j+v-1}^{(i)}$, a smaller bunch time at $s^{(i)}$ always leads to a smaller possible bunch time at $s^{(i+1)}$, i.e.:*

$$d_{j+v-1}^{(i)} - a_j^{(i)} \leq d_{j+v-1}'^{(i)} - a_j^{(i)} \implies d_{j+v-1}^{(i+1)} - a_{j-1}^{(i+1)} \leq d_{j+v-1}'^{(i+1)} - a_{j-1}^{(i+1)} \quad (8)$$

Proof. We fix $a_j^{(i)}$ and compare the bunch time at $s^{(i+1)}$ for two bunches with different final vehicle departure times $d_{j+v-1}^{(i)} \leq d_{j+v-1}'^{(i)}$. To minimize bunch time at $s^{(i+1)}$ we need to maximize the leader's arrival ($a_j^{(i+1)} = \overline{a}_j^{(i+1)}$) and minimize the last vehicle's departure ($d_{j+v-1}^{(i+1)} = \underline{d}_{j+v-1}^{(i+1)}$ and $d_{j+v-1}'^{(i+1)} = \underline{d}_{j+v-1}'^{(i+1)}$). $a_j^{(i+1)}$ is maximized when dwell time at $s^{(i)}$ and travel to $s^{(i+1)}$ are maximized: $\delta_j^{(i)} = \overline{\delta}^{(i)}$ and $\tau_j^{(i)} = \overline{\tau}_j^{(i)}$. Note that the policy at $s^{(i)}$ may hold the vehicle for longer, but the effect is the same since we assumed the same lead vehicle arrival for both cases and holding is not affected by future arrivals. To minimize $d_{j+v-1}^{(i+1)}$ and $d_{j+v-1}'^{(i+1)}$ we need vehicles $j+1 \dots j+v-1$ to travel to $s^{(i+1)}$ as fast as possible without overtaking the lead vehicle. This is $a_j^{(i+1)} - d_{j+k}^{(i)}$ and $a_j^{(i+1)} - d_{j+k}'^{(i)}$ for $k = 1 \dots v-1$, which is bounded from above and below by $\overline{\tau}^{(i)}$ and $\underline{\tau}^{(i)}$. At this point we can see that $a_{j+v-1}^{(i+1)} \leq a_{j+v-1}'^{(i+1)}$ regardless of the positions of the other vehicles in the bunch. Likewise, while waiting at $s^{(i+1)}$, the gap can be narrowed if $\delta_{j+k}^{(i+1)} = d_j^{(i+1)} - a_{j+k}^{(i+1)}$ and $\delta_{j+k}'^{(i+1)} = d_j^{(i+1)} - a_{j+k}'^{(i+1)}$, which are of course bounded above by $\overline{\delta}^{(i+1)}$ and below by $\underline{\delta}^{(i+1)}$. At this point, we see that the only way for $d_{j+v-1}^{(i+1)} \geq d_{j+v-1}'^{(i+1)}$ is if $\phi_{j+v-1}^{(i+1)} \geq \phi_{j+v-1}'^{(i+1)} + d_{j+v-1}'^{(i)} - d_{j+v-1}^{(i)}$. Considering the same three cases as in Lemma 2, however, shows us that this is impossible: under all three policy classes, a later arrival at the stop necessitates a later departure. □

In Lemma 4 we not only showed that a smaller bunch time at a preceding stop is a necessary condition for a smaller bunch time at a following stop, we showed the travel times and dwell times required to achieve it for each vehicle in the bunch. We can now describe how to calculate the smallest bunch time at any given stop for a bunch of M vehicles.

Theorem 5. *For a bunch of M vehicles, bunch time is minimized at $s^{(i+1)}$ when, starting at $s^{(0)}$, the lead vehicle in the bunch j dwells with $\delta_j^{(i)} = \overline{\delta}^{(i)}$ and travels with $\tau_j^{(i)} = \overline{\tau}^{(i)}$ for all stops and the subsequent $M-1$ vehicles always travel with parameters $\tau_{j+m}^{(i)} = [a_j^{(i+1)} - d_{j+m}^{(i)}]_{\underline{\tau}^{(i)}}^{\overline{\tau}^{(i)}}$ and $\delta_{j+m}^{(i+1)} = [d_j^{(i+1)} - a_{j+m}^{(i+1)}]_{\underline{\delta}^{(i+1)}}^{\overline{\delta}^{(i+1)}} \forall m \in 1 \dots M; j \in 0, M, 2M, \dots \infty; i \in 0 \dots N$.*

Proof. We start with the base case where vehicles are released from $s^{(0)}$ at $r_0^{(0)} \dots r_{M-1}^{(0)}$. Since there is only one possible departure time for each vehicle, this case only has one possible bunch time, and is therefore minimized. We now show that for a bunch of M vehicles with the first vehicle arriving at $s^{(i)}$ at $a_j^{(i)}$ and the last vehicle departing at $d_{j+M-1}^{(i)}$, minimum bunch size at $s^{(i+1)}$ can only occur if $d_{j+M-1}^{(i)} - a_j^{(i)}$ is already minimized. By Lemma 4, we know this is true for the same

bunch of M vehicles, however, we still need to show that this is true for all future bunches of M vehicles. We prove this by contradiction: assume that for a given $a_j^{(i)}$, there is some $d_{j+M-1+x}^{(i+1)} - a_{j+x}^{(i+1)} < d_{j+M-1}^{(i+1)} - a_j^{(i+1)}$ for $x > 0$. If this was the case, then by Lemma 4 it must also hold true that $d_{j+M-1+x}^{(i)} - a_{j+x}^{(i)} < d_{j+M-1}^{(i)} - a_j^{(i)}$, however, this would violate the assumption that $d_{j+M-1}^{(i)} - a_j^{(i)}$ is minimized: a contradiction. Since the bunch size at the start is minimized in the base case, we can use the travel and dwell parameters computed in Lemma 4 to recursively calculate the minimum bunch size at $s^{(1)} \dots s^{(N)}$. After $s^{(N)}$, the vehicles travel back to $s^{(0)}$ and begin traveling the route again. Now that all the stops have been visited by a vehicle at least once, headway is computable. Minimum bunch time is now calculated for the bunch where the lead vehicle is the Mth arrival and the last vehicle is the $2M-1$st arrival at each stop. By our induction hypothesis, this continues for all bunches of M vehicles for all N stops with the lead vehicle arriving at $M, 2M, \dots \infty$. □

Using Theorem 5, we can now compute the headway upper bound. We note that the travel and dwell times taken by the vehicles in Theorem 5 are sufficient to cause the maximum headway to occur, but not always necessary: the same result can sometimes be reached with different travel and dwell parameters. While it is not true in general that a smaller bunch time of M vehicles implies a lower minimum headway time across all M vehicles in a bunch (e.g., all vehicles in the bunch are evenly spaced instead of two successive vehicles having zero headway), we now show that by following the travel and dwell times prescribed by Theorem 5, the lower bound on headway time must occur.

Theorem 6. *The headway lower bound, $\underline{HT}^{(i)}$ occurs at $s^{(i)}$, when, starting at $s^{(0)}$, the lead vehicle j in a bunch of M vehicles dwells with $\delta_j^{(i)} = \overline{\delta}^{(i)}$ and travels with $\tau_j^{(i)} = \overline{\tau}^{(i)}$ for all stops and the subsequent $M-1$ vehicles always travel with parameters $\tau_{j+m}^{(i)} = [a_j^{(i+1)} - d_{j+m}^{(i)}]_{\underline{\tau}^{(i)}}^{\overline{\tau}^{(i)}}$ and $\delta_{j+m}^{(i+1)} = [d_j^{(i+1)} - a_{j+m}^{(i+1)}]_{\underline{\delta}^{(i+1)}}^{\overline{\delta}^{(i+1)}} \forall m \in 1 \dots M; j \in 0, M, 2M, \dots \infty; i \in 0 \dots N$.*

Proof. To minimize $\underline{HT}^{(i)}$, we want $d_{j-1}^{(i)} = \overline{d}_{j-1}^{(i)}$ and $a_j^{(i)} = \overline{a}_j^{(i)}$. The base case is the departure of the 0th vehicle and the arrival of the 1st vehicle at $s^{(1)}$. Under the travel conditions, listed above, it is clear that $d_0^{(1)}$ is maximized and $a_1^{(1)}$ is minimized since $\tau_0^{(0)} = \overline{\tau}^{(0)}$ and $\tau_1^{(0)} = [a_0^{(1)} - d_1^{(0)}]_{\underline{\tau}^{(0)}}^{\overline{\tau}^{(0)}}$ will clamp the difference between these values as close to 0 as possible. The induction hypothesis is that the smallest possible headway for the jth arrival at $s^{(i)}$ leads to the smallest possible headway for the jth vehicle at $s^{(i+1)}$ under the given travel conditions. It is clear that $a_{j-1}^{(i+1)} = d_{j-1}^{(i)} + \overline{\tau}^{(i)}$ is the latest possible arrival given the starting condition and any other travel time cannot lead to a later departure $d_{j-1}^{(i+1)}$ under Lemma 2. Meanwhile, the jth arrival at $s^{(i)}$ dwells for $\delta_j^{(i)} = [d_{j-1}^{(i)} - a_{j]}^{(i)}]_{\underline{\delta}^{(i)}}^{\overline{\delta}^{(i)}}$, which results in its $d_j^{(i)}$ being as close to $d_{j-1}^{(i)}$ as possible without overtaking. The travel condition $\tau_j^{(i)} = [a_{j-1}^{(i+1)} - d_j^{(i)}]_{\underline{\tau}^{(i)}}^{\overline{\tau}^{(i)}}$ then minimizes $a_j^{(i+1)}$ given that $a_{j-1}^{(i+1)}$ is fixed (if an overtake were to occur, the jth departure would become the $j-1$st arrival, so traveling at $\underline{\tau}^{(i)}$, would not improve $a_j^{(i+1)}$). Since $d_{j-1}^{(i+1)}$ is maximized and $a_j^{(i+1)}$ is minimized, $h_j^{(i+1)}$ will be minimized given the initial condition. From the base case, this is minimized for $h_1^{(1)}$. □

In Theorem 6, we see that the lower bound on headway can be calculated using a similar process to the upper bound but instead focusing on arrival of the $1, M+1, 2M+1, \dots$ vehicle at every stop. This means $\overline{HT}^{(i)}$ and $\underline{HT}^{(i)}$ can be calculated simultaneously.

C. Computing Headway Bounds

Algorithm 1 uses Theorem 5 and Theorem 6 to calculate the headway upper and lower bounds. It requires the stop parameters $(\underline{\tau}^{(i)}, \overline{\tau}^{(i)}, \underline{\delta}^{(i)}, \overline{\delta}^{(i)}, \pi^{(i)})$ for all N stops in a route. To satisfy the base case in Theorem 5 and Theorem 6, the initial release times $r_j^{(0)}$ of all M vehicles are required as well. i is used as an index to track the current stop for which headway time is being calculated (initialized in Line 1 and updated in Line 14). Initially, $\overline{HT}^{(i)}$ is set to $-\infty$ for all stops, since it has not yet been computed (Line 2). A negative value indicates no valid headway while $\overline{HT}^{(i)} \geq 0$ indicates the worst possible headway time at $s^{(i)}$ up to the time of the latest considered departure. Likewise, $\underline{HT}^{(i)}$ is initially set to ∞ for all stops (Line 3) and decreases as the loop progresses. We use $\omega^{(i)}$ to track the departure time of the last bus at each stop $s^{(i)}$. This is initialized to undefined for each stop since at the start of the route, no buses have departed (Line 4). Two vectors $\mathbf{a}$ and $\mathbf{d}$ track the latest arrivals and departures for all M buses at the stop i under consideration. We know the initial departures for all buses (corresponding to the base case in Theorem 5), so $\mathbf{d}$ is initialized to release times $r_0^{(0)} \dots r_M^{(0)}$ (Line 5). Following Theorem 5, we loop through the successive stops calculating arrivals and departures given that buses travel with $\tau^{(i)}$ and $\delta^{(i)}$ determined by Lemma 4. We use the variable `converged` (initially set to false in Line 6) to track the convergence on headway bounds. Inside the loop, the first element of $\mathbf{a}$ (corresponding to the first arrival at $s^{(i)}$) receives the earliest previously tracked departure plus the maximum travel time (Line 9). For the remaining vehicles the travel time is computed to be as fast as possible, without over taking the lead (Lines 10–12). Next, we clear the previous departure values (Line 13) and update the stop index (Line 14). The lead vehicle's dwell time is set as long as possible (Line 15), and the remaining vehicles are set to dwell such that the times between their departures and the lead vehicle's are minimized (Line 17). Note that calculation of $\phi^{(i)}$ may require additional variables not tracked here, for example, a stop controlled by π`SCH` will require the current time since start of service and a stop controlled by π`DYN` will require an estimate of the next vehicle's arrival time $\hat{a}_{j+1}^{(i)}$. This estimate needs to be for the earliest possible arrival of the subsequent

vehicle so that the bunch time will shrink and the algorithm captures the worst-case behavior. Once departure times have been calculated, we can update our headway bounds: if a stop has been previously visited ($\omega^{(i)}$ exists – Line 19) and the $\overline{HT}^{(i)}$ is larger than previously observed (Line 20), we record it (Line 21). Convergence of $\overline{HT}^{(i)}$ depends only on bunch time: the same bunch time of M vehicles always leads to the same next minimum bunch time, so if there is no change, it is possible that our algorithm has converged (Line 23). We also update $\underline{HT}^{(i)}$ if this is lower than previously observed (Line 24). It is possible that $\underline{HT}^{(i)}$ converges slower than $\overline{HT}^{(i)}$ since the positions of vehicles in a bunch can continue changing even if bunch time remains unchanged, thus an update to $\underline{HT}^{(i)}$ causes us to unset the convergence flag (Line 26). Before considering the next stop, we set $\omega^{(i)}$ to the departure of the last vehicle, which will be used to calculate $\overline{HT}^{(i)}$ when the bunch visits $s^{(i)}$ again (Line 27).

We note that Algorithm 1 is not guaranteed to converge in a finite amount of time. Consider, for example, a stop controlled by $\pi^{(i)} \in \pi\texttt{DYN}$, where holding times are calculated based on headway such that $\phi_j^{(i)} = h_j^{(i)}/t$ (where t is time since service start) and $\underline{\delta}^{(i)} = \overline{\delta}^{(i)} = 0$. This policy still respects all the requirements of $\pi\texttt{DYN}$, but the bunch time at $s^{(i)}$ can only approach 0: it will never reach it in a finite time window. In practice however, a time cutoff can be introduced since most routes do not serve stops continuously and pause at night or on the weekends. Since convergence is not provable, we consider the time complexity to update the bounds for one stop, which is $O(NM)$ because arrivals and departures must be calculated for M vehicles at each of the N stops leading to the update. Memory complexity is linear with respect to number of stops and vehicles ($O(NM)$). The vectors of arrivals and departures ($\mathbf{a}$ and $\mathbf{d}$) never exceed length M and are freed and reallocated on each iteration of the loop. The last departure for each stop, ω, is also required for all N stops.

Algorithm 1: Dynamic program to calculate $\overline{HT}^{(i)}$ and $\underline{HT}^{(i)}$.

Data: Stop params $\underline{\tau}^{(i)}, \overline{\tau}^{(i)}, \underline{\delta}^{(i)}, \overline{\delta}^{(i)}, \pi^{(i)} \forall i \in 0 \dots N$
Vehicle release times $r_j^{(0)} \forall j \in 0 \dots M$
Result: $\overline{HT}^{(i)}, \underline{HT}^{(i)} \forall i \in 0 \dots N$

1 $i \leftarrow 1$;
2 $\overline{HT}^{(i)} \leftarrow -\infty \forall i \in 0 \dots N$;
3 $\underline{HT}^{(i)} \leftarrow \infty \forall i \in 0 \dots N$;
4 $\omega^{(i)} \leftarrow \texttt{undef}. \forall i \in 0 \dots N$;
5 $d_j \leftarrow r_j^{(0)} \forall j \in 0 \dots M$;
6 $\texttt{converged} \leftarrow \perp$;
7 **while** $\neg\texttt{converged}$ **do**
8 $\quad \mathbf{a} \leftarrow []$;
9 $\quad$ Push $d_0 + \overline{\tau}^{(i)}$ to $\mathbf{a}$;
10 $\quad$ **for** $j \in 1 \dots M$ **do**
11 $\quad\quad$ Push $d_j + [a_{j-1} - d_j]_{\underline{\tau}^{(i)}}^{\overline{\tau}^{(i)}}$ to $\mathbf{a}$;
12 $\quad$ **end**
13 $\quad \mathbf{d} \leftarrow []$;
14 $\quad i \leftarrow (i+1) \mod N$;
15 $\quad$ Push $a_0 + \max\{\overline{\delta}^{(i)}, \phi^{(i)}\}$ to $\mathbf{d}$;
16 $\quad$ **for** $j \in 1 \dots M$ **do**
17 $\quad\quad$ Push $a_j + \max\{[d_{j-1} - a_j]_{\underline{\delta}^{(i)}}^{\overline{\delta}^{(i)}}, \phi^{(i)}\}$ to $\mathbf{d}$;
18 $\quad$ **end**
19 $\quad$ **if** $\exists \omega^{(i)}$ **then**
20 $\quad\quad$ **if** $a_0 - \omega^{(i)} > \overline{HT}^{(i)}$ **then**
21 $\quad\quad\quad \overline{HT}^{(i)} \leftarrow \max\{0, a_0 - \omega^{(i)}\}$;
22 $\quad\quad$ **else**
23 $\quad\quad\quad \texttt{converged} \leftarrow \top$;
24 $\quad$ **if** $\max\{0, a_1 - d_0\} < \underline{HT}^{(i)}$ **then**
25 $\quad\quad \underline{HT}^{(i)} \leftarrow \max\{0, a_1 - d_0\}$;
26 $\quad\quad \texttt{converged} \leftarrow \perp$;
27 $\quad \omega^{(i)} \leftarrow d_M$;
28 **end**

V. Experiments

To build intuition on the effects of different holding policies on headway bounds, we first demonstrate our algorithm on an example bus route with 5 stops. All code for the following sections is publicly available on GitHub[1]. The travel and dwell time parameters for each stop are $\tau^{(i)} \in [4, 4.5]$ and $\delta^{(i)} \in [0, 0.5]$, with units given in minutes. This yields a minimum route completion time of 20 min. and a maximum route completion time of 25 min. in the uncontrolled case, making it easier to link the computed headway times to other critical timing parameters. We compare the following policies:

- $\pi\texttt{NIL}$ – $\phi_j^{(i)} = 0 \forall i \in 0 \dots N, j \in 0 \dots \infty$.
- $\pi\texttt{SCH}$ – We set scheduled departure times at intervals of the desired headway (10 min.) and add an offset for the maximum time it takes the lead bus to reach a stop after beginning the route. For example, the schedule list at $s^{(1)}$ would be $\mathbf{T} = \langle 5, 15, 25, \dots \rangle$, since the maximum time to reach $s^{(1)}$ is 5 minutes ($\overline{\tau}^{(i)} + \overline{\delta}^{(i)}$). This offset is chosen to prevent the worst-case scenario for a schedule-driven policy outlined in Note 2.
- $\pi\texttt{DYN}$ – We use the same headway-based dynamic holding policy as Zhou et al. [19], adapting their notation to match our system model in Eq. (9).

$$\phi_j^{(i)} = \begin{cases} \left[\frac{1}{2}(\hat{a}_{j+1}^{(i)} + d_{j-1}^{(i)} - 2a_j^{(i)}) \right]_0^{\Phi} & \text{if} \quad \frac{h_j^{(i)}}{h_{j+1}^{(i)}} \leq R \\ 0 & \text{else} \end{cases} \tag{9}$$

In addition to the arrivals and departures of the surrounding vehicles, the policy is parameterized by Φ, the maximum amount of holding that can be applied, and R, the headway to tailway ratio needed to activate the policy. These parameters are desirable for transit operators as

[1] https://github.com/smarttransit-ai/bus-transit-headway-bounds

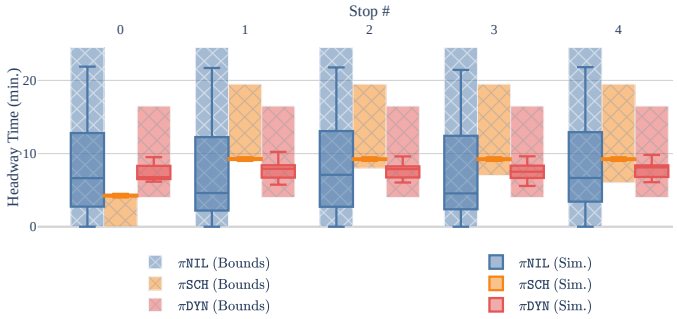

Fig. 4. Simulated headway measurements (boxes) and headway bounds (hatched bars) for a 5 stop route with all stops controled by πNIL (blue), πSCH (orange), or πDYN with $R = 0.75$, $\Phi = 30$ (red).

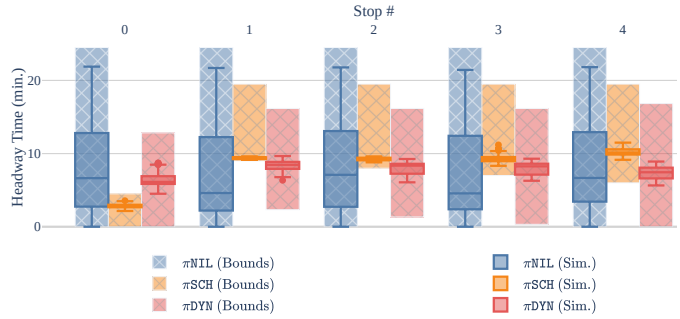

Fig. 5. Simulated headway measurements (boxes) and headway bounds (hatched bars) for a 5 stop route with all stops governed by πNIL (blue) except for $s^{(0)}$, which can be governed by πSCH (orange) or πDYN with $R = 0.75$, $\Phi = 30$ (red).

they bound holding times from above and below.

Holding policies at every stop – We begin by examining a route where all 5 stops have an identical holding policy applied, i.e., all stops governed by either πNIL, πSCH, or πDYN, in Fig. 4. The route is served by 3 buses released from $s^{(0)}$ at times 0, 10, and 20 min. Since the maximum completion time of the uncontrolled route is 30 min., three buses should be sufficient to maintain a headway of 10 min. at each stop. In Fig. 4, the hatched bars represent the upper and lower headway time bounds and the box plots represent the headways actually observed in a probabilistic simulation of 10,000 trips along the route with travel and dwell times drawn uniformly at each stop. As predicted in Theorem 1, when there is no control policy, worst case bunching occurs and $\overline{HT}^{(i)}$ approaches the maximum route completion time while $\underline{HT}^{(i)}$ approaches zero. Furthermore, we observe that the statistical simulation also approaches this bound. When all stops are scheduled so that the driver has to hold at a stop if they arrive early, we see that headway times are maintained around 10 minutes with tight variance, except for $s^{(0)}$. This is because the schedule release times at $s^{(4)}$ are $\langle 20, 30, 40 \ldots \rangle$ while at $s^{(0)}$ they are at $\langle 30, 40, 50, \ldots \rangle$. In the average case, buses will reach $s^{(0)}$ around 5 minutes after their predecessor left. We observe that despite the relatively tight statistical distribution of headway, maintaining a schedule can incur up to 20 min. of headway ($\overline{HT}^{(i)}$) at all stops other than $s^{(0)}$ due to follower buses in a bunch becoming unacceptably delayed. For πDYN we set $R = 0.75$ and $\Phi = 30$. In this case, the mean headway time is less than that of the schedule driven route although the average variance at each stop is higher. For all stops excluding $s^{(0)}$, the headway upper and lower bounds are less than πSCH meaning that the πDYN is able keep buses evenly spaced allowing headway to approach the average travel time divided by number of buses, which is lower than the πSCH's strictly enforced 10 min.

Holding policies at one stop – Some works argue that a single control point is sufficient to regulate headway [26], which we analyze in Fig. 5 for holding only applied at $s^{(0)}$ using the same πSCH and πDYN parameters as Fig. 4. We see that for πSCH, the bounds are equivalent to when all stops

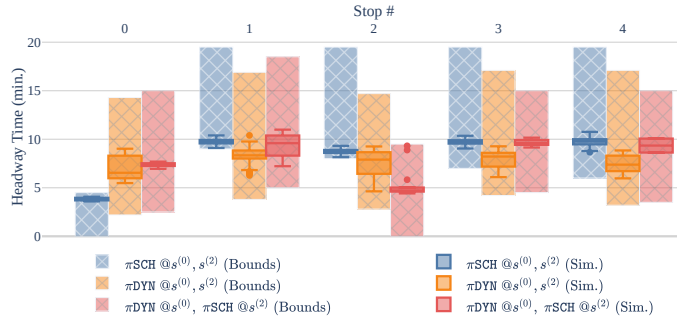

Fig. 6. Simulated headway measurements (boxes) and headway bounds (hatched bars) for a 5 stop route with holding polices applied at $s^{(0)}$ and $s^{(2)}$. πDYN is applied with parameters $R = 0.75$, $\Phi = 30$.

are controlled by πSCH, but that the variance of the simulated results is now greater. However, for the πDYN route, $\underline{HT}^{(i)}$ is lower for all stops, meaning 2 of the 3 vehicles are approaching each other. Despite this, $\overline{HT}^{(i)}$ remains unchanged meaning one stop is sufficient to maintain separation.

Combinations of holding policies – Other authors argue that multiple control points lead to better performance [27]. We investigate this claim and also compare with the case when different strategies are combined on the same route. In Fig. 6, $s^{(0)}$ and $s^{(2)}$ are control points and we consider (1) both stops controlled by πSCH, (2) both stops controlled by πDYN, (3) one stop controlled by πSCH ($s^{(2)}$) and one stop controlled by πDYN ($s^{(0)}$). The policy parameters are the same as Fig. 4. Case (1) maintains control similar to its counterparts in Figs. 4 and 5 indicating that more scheduled stops are not necessary to increase performance in this case. In case (2) we see that by adding a control point midroute, the $\underline{HT}^{(i)}$ improves for the stops, though not to the extent of Fig. 4. In case (3) we see that the bounds are tighter than (1) and (2) immediately following the πSCH control point showing that πSCH's ability to regulate headway is somewhat enhanced by the πDYN at $s^{(0)}$, despite the marginally worse headway on the output of that stop. This shows that policies can be combined along a route to shape headway bounds, but that they may inadvertently negatively impact bounds at other locations as well.

Hyperparameters of πDYN – We now consider the route

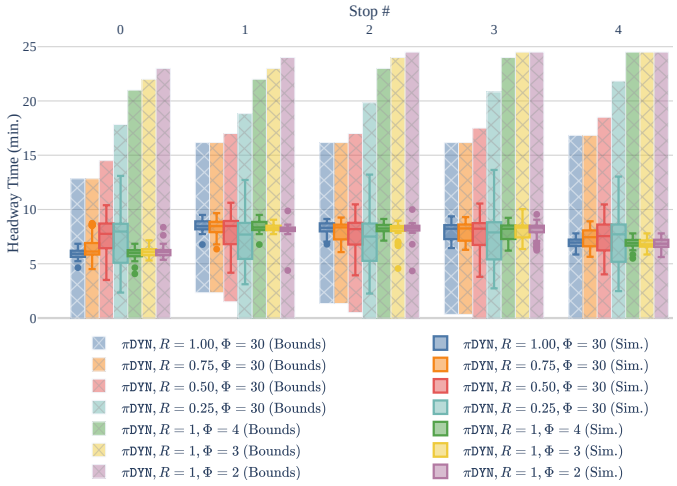

Fig. 7. Simulated headway measurements (boxes) and headway bounds (hatched bars) for a 5 stop route with πDYN applied at $s^{(0)}$ while sweeping parameters R and Φ.

where πDYN is applied only at $s^{(0)}$ and adjust the maximum holding (Φ) and the headway/tailway ratio at which the policy activates (R). For transit operators, issuing holds of less than a minute can be difficult for drivers to execute (necessitating a lower R), and too much holding can delay passengers already onboard a vehicle (necessitating a bound on Φ). The results are shown in Fig. 7. We observe that the best performance is when $R = 1$, indicating that even minor deviations in headway will trigger holding, but that there is no difference in bounds with $R = 0.75$, despite more variance in the simulated results. Below $R = 0.75$ the upper bound increases and the lower bound decreases indicating reduced performance, which is matched in the simulated results. We also observe that decreasing Φ (reducing the amount of hold time the policy can issue), has little effect on the simulated performance, but causes the upper headway bound to increase. As $\Phi \to 0$, $\overline{HT}^{(i)}$ will approach the performance of πNIL.

VI. Case Study

We now demonstrate the utility of our analytical bound when selecting holding policies for a real bus route in Nashville, TN. Route 55 runs from Murfreesboro to downtown Nashville and consists of about 80 stops round-trip with 16 buses in service at a time, which leads to a high probability of bunching during peak commute times. While we cannot arbitrarily test different holding policies for drivers and passengers, we have real world data for dwell times and travel times across all hours of the day. These are fit to log-normal distributions as is common practice in transit system modeling [28] of which we take the 5th to 95th percentile to generate our $\underline{\tau}^{(i)}$, $\overline{\tau}^{(i)}$, $\underline{\delta}^{(i)}$, and $\overline{\delta}^{(i)}$ for each stop. In the existing route, all buses depart the Hickory Hollow terminal ($s^{(0)}$) and City Center terminal ($s^{(40)}$) according to a static schedule. This is our baseline case in Fig. 8, which shows the headway bounds, and case (A) in Fig. 9, which shows the cumulative distributions of departure-to-departure times (a

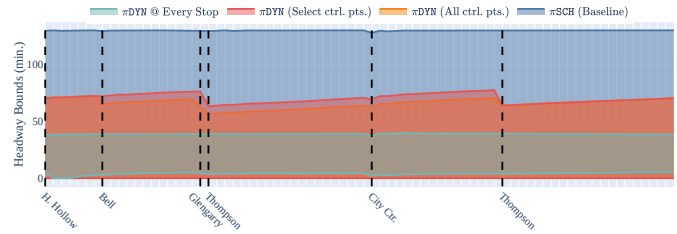

Fig. 8. Headway bounds for WeGo route 55 in Nashville, TN under a schedule driven baseline (A), πDYN with $R = 1$ and unbounded Φ applied at every stop (B), holding at 2 control points with $\Phi = 15$ min. (C), and holding at 4 control points with $\Phi = 7.5$ min. (D).

proxy for headway times) for all stops in a transit simulator internal to WeGo that comprehensively models passenger arrivals at stops, bus crowding, and vehicle overtakes. We observe that in this case, buses can arrive behind schedule and the situation described in Note 2 can occur leading to a headway upper bound of 120 min. The simulated results bear this out with a relatively flat distribution of waiting times with some as long as an hour. In (B), we also include headway times for the ideal case, where we could apply πDYN at every stop with $R = 1$ and no limit on Φ.

Our goal is to replace the schedule with dynamic holding at key points along the route to regulate headway. Due to physical constraints, this can only occur at Bell, Glengarry, and Thompson stations, whose positions along the route are marked in Fig. 8. In addition, we want to constrain total possible waiting time to 30 min. to mitigate the probability of delayed releases from the terminals. Thus, we consider two cases: (C) applying πDYN at Thompson station (inbound and outbound) with $\Phi = 15$ min. in each direction, and (D) applying πDYN at all four potential control points but with $\Phi = 7.5$ min. (C) is operationally easier to implement, but we see that even with the limitation on Φ, by spreading the holding over more control points (D) is able to outperform (C) in headway bounds along the middle of the route, which corresponds to the downtown core. The simulated results bear a similar story with (D) marginally outperforming (C) in the long tail. While these results show that our analysis can quickly identify headway bounds, it is important to remember that the specific result for this route – that spreading holding over many control points outperforms two major control points – may be different for different routes with different numbers of buses and scheduled departure times. Our analysis technique is meant to facilitate exploration of such cases.

VII. Conclusion

We have shown how to calculate upper and lower bounds on headway times under the assumption that travel and dwell times between stops in a transit route are bounded. Through experiments on an example bus route, we have built some intuition about the affects of various policies and holding point locations on the route performance. Finally, we have demonstrated the utility of this analysis by applying it to a real transit system and showing how the results can be applied

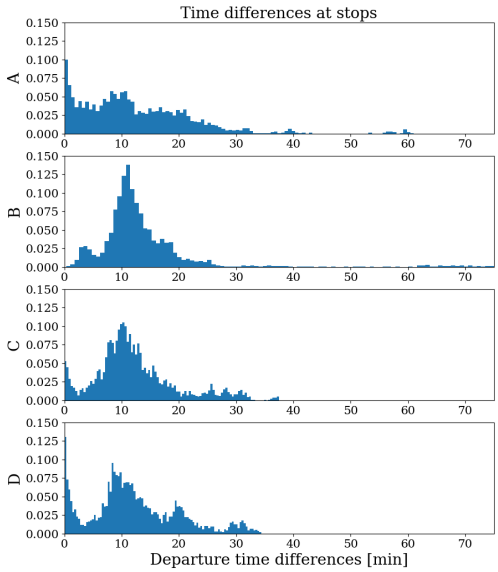

Fig. 9. Simulated departure-to-departure times aggregated across all stops for the schedule-driven baseline (A), πDYN applied at every stop (B), holding at 2 control points with $\Phi = 15$ min. (C), and holding at 4 control points with $\Phi = 7.5$ min. (D).

in transit planning. In the future, we hope to extend this work to include an upper bound on travel time on routes with arbitrary holding policies. We would also like to develop a formal design strategy based on the results of this analysis that efficiently identifies optimal holding policies and stop locations for arbitrary routes. We hope this work serves as a starting point for future worst-case analysis in transit planning.

Acknowledgment

This work was supported in part by the U.S. Department of Transportation through the SMART Grants Program Phase 1 award in partnership with WeGo Public Transit and the Metropolitan Government of Nashville and Davidson County. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. Department of Transportation.

References

[1] F. Enayatollahi, A. O. Idris, and M. A. A. Atashgah, "Modelling bus bunching under variable transit demand using cellular automata," *Public Transport*, vol. 11, no. 2, pp. 269–298, 2019.

[2] L. Pan, Y. Zhou, Q. Meng, and Y. Wang, "Impact analysis of traffic factors on urban bus bunching," *IEEE Intelligent Transportation Systems Magazine*, vol. 15, no. 5, pp. 6–24, 2023.

[3] "National transit database (NTD) glossary," Federal Transit Administration, Washington, DC, USA, 2025. [Online]. Available: https://www.transit.dot.gov/ntd/national-transit-database-ntd-glossary

[4] W. Kim, B. Son, J.-H. Chung, and E. Kim, "Development of real-time optimal bus scheduling and headway control models," *Transportation Research Record*, vol. 2111, no. 1, pp. 33–41, 2009.

[5] C. Pangilinan, N. Wilson, and A. Moore, "Bus supervision deployment strategies and use of real-time automatic vehicle location for improved bus service reliability," *Transportation Research Record*, vol. 2063, no. 1, pp. 28–33, 2008.

[6] M. Rezazada, N. Nassir, and E. Tanin, "Public transport bunching: A critical review with focus on methods and findings for implications for policy and future research," in *Australasian Transport Research Forum (ATRF)*, 2022.

[7] P. He, G. Jiang, S.-K. Lam, and D. Tang, "Travel-time prediction of bus journey with multiple bus trips," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 11, pp. 4192–4205, 2019.

[8] S. Liu, X. Luo, and P. J. Jin, "Improving bus operations through integrated dynamic holding control and schedule optimization," *Journal of Advanced Transportation*, vol. 2018, no. 1, p. 9714046, 2018.

[9] A. Carrel, A. Halvorsen, and J. L. Walker, "Passengers' perception of and behavioral adaptation to unreliability in public transportation," *Transportation Research Record*, vol. 2351, no. 1, pp. 153–162, 2013.

[10] K. J. Dueker, T. J. Kimpel, J. G. Strathman, and S. Callas, "Determinants of bus dwell time," *Journal of Public Transportation*, vol. 7, no. 1, pp. 21–40, 2004.

[11] M. M. Rahman, S. Wirasinghe, and L. Kattan, "Analysis of bus travel time distributions for varying horizons and real-time applications," *Transportation Research Part C: Emerging Technologies*, vol. 86, pp. 453–466, 2018.

[12] Y.-J. Byon, C. E. Cortés, Y.-S. Jeong, F. J. Martínez, M. A. Munizaga, and M. Zúñiga, "Bunching and headway adherence approach to public transport with GPS," *International Journal of Civil Engineering*, vol. 16, no. 6, pp. 647–658, 2018.

[13] M. A. Esfeh, S. C. Wirasinghe, S. Saidi, and L. Kattan, "Waiting time and headway modelling for urban transit systems – a critical review and proposed approach," *Transport Reviews*, vol. 41, no. 2, pp. 141–163, 2021.

[14] H. A. Mohammed, D. S. Hurwitz, D. A. Noyce, and X. Wang, "Influence of following vehicle's tailway and classification on subject driver's response to the circular yellow indication," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 77, pp. 73–86, 2021.

[15] K. Long, J. Wei, J. Gu, and X. Yang, "Headway-based multi-route transit signal priority at isolated intersection," *IEEE Access*, vol. 8, pp. 187 824–187 831, 2020.

[16] Y. Martínez-Estupiñan, F. Delgado, J. C. Muñoz, and K. E. Watkins, "Understanding what elements influence a bus driver to use headway regularity tools: case study of Santiago public transit system," *Transportmetrica A: Transport Science*, vol. 19, no. 2, p. 2025950, 2023.

[17] C. F. Daganzo, "A headway-based approach to eliminate bus bunching: Systematic analysis and comparisons," *Transportation Research Part B: Methodological*, vol. 43, no. 10, pp. 913–921, 2009.

[18] O. Cats, W. Burghout, T. Toledo, and H. N. Koutsopoulos, "Evaluation of real-time holding strategies for improved bus service reliability," in *13th International IEEE Conference on Intelligent Transportation Systems*, 2010, pp. 718–723.

[19] C. Zhou, Q. Tian, and D. Z. Wang, "A novel control strategy in mitigating bus bunching: Utilizing real-time information," *Transport Policy*, vol. 123, pp. 1–13, 2022.

[20] A. Fonzone, J.-D. Schmöcker, and R. Liu, "A model of bus bunching under reliability-based passenger arrival patterns," *Transportation Research Procedia*, vol. 7, pp. 276–299, 2015.

[21] W. Chen, K. Zhou, and C. Chen, "Real-time bus holding control on a transit corridor based on multi-agent reinforcement learning," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, 2016, pp. 100–106.

[22] S. Rashidi, S. Ataeian, and P. Ranjitkar, "Estimating bus dwell time: A review of the literature," *Transport Reviews*, vol. 43, no. 1, pp. 32–61, 2023.

[23] V.-L. Saw, N. N. Chung, W. L. Quek, Y. E. I. Pang, and L. Y. Chew, "Bus bunching as a synchronisation phenomenon," *Scientific Reports*, vol. 9, no. 1, p. 6887, May 2019.

[24] P. Duetting, G. Guruganesh, J. Schneider, and J. R. Wang, "Optimal no-regret learning for one-sided Lipschitz functions," in *Proceedings of the 40th International Conference on Machine Learning*, vol. 202, 2023, pp. 8836–8850.

[25] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of Lipschitz constants for deep neural networks," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[26] X. J. Eberlein, N. H. M. Wilson, and D. Bernstein, "The holding problem with real–time information available," *Transportation Science*, vol. 35, no. 1, pp. 1–18, 2001.

[27] Q. Chen, E. Adida, and J. Lin, "Implementation of an iterative headway-based bus holding strategy with real-time information," *Public Transport*, vol. 4, no. 3, pp. 165–186, 2013.

[28] R. Rajbhandari, S. I. Chien, and J. R. Daniel, "Estimation of bus dwell times with automatic passenger counter information," *Transportation Research Record*, vol. 1841, no. 1, pp. 120–127, 2003.