Optimal Planning on a Single-Route Transit System with Modular Buses

Karan Jagdale, Zaid Saeed Khan, Mónica Menéndez, and Melkior Ornik

***Abstract*— This paper considers optimal planning of modular buses on a single route to minimize the average passenger travel time. Modular vehicles consist of individual units that can attach and detach from other modules at any stop, allowing for increased flexibility. As the passenger demand and the travel time between stops are stochastic variables, we use a Markov decision process to model the system. The problem of minimizing passenger travel time is converted to optimally planning the actions of all the vehicles, which include stopping at a stop, skipping the stop, and attaching or detaching from other modules. We develop a cost formulation to capture the impact of these actions on passenger travel time, and propose a control policy that selects an action with the minimum cost. To evaluate our proposed policy, we compare it in simulation with two other policies: when vehicles serve every stop without attaching or detaching, and a previously proposed rule-based attaching/detaching policy. Our proposed policy significantly outperforms both of these strategies.**

I. Introduction

Bus transit in urban regions involves a multitude of stochastic features, including passenger loads and the variable travel times caused by traffic conditions such as congestion and traffic signals. Planning under this stochasticity often leads to a phenomenon of *bus bunching* [1]. Namely, a bus whose distance from the preceding bus is shorter than intended will on average board fewer passengers. The shorter boarding time further decreases its headway, resulting in increasingly uneven passenger distribution among buses. Buses with shorter headways will be thus underutilized, while those with longer headways will be overutilized, eventually leading to some passengers being unable to board the bus due to capacity constraints. Additionally, some passengers may face extremely long wait times at the bus stop.

Previous studies and practices have proposed various strategies to optimize passenger travel time and prevent bunching. These strategies include *bus holding* [2], where some buses are held longer at different stops to increase short headways, *stop skipping* [3], where certain stops are skipped to decrease long headways, and *bus insertion* [4] or *substitution* [5] where standby buses are inserted at specific stops to split long headways. While these strategies are shown to mitigate bus bunching, they all suffer from drawbacks: a bus holding strategy only slows down an early bus but cannot speed up a bus which is late, while stop skipping cannot slow down an early bus. Bus insertion requires a larger fleet size, resulting in higher costs.

In addition to classical routing strategies, recent work introduced policies based on the emergent framework of *autonomous modular vehicles (AMVs)* [6]. AMVs can quickly join with each other while on the road, forming a larger vehicle with higher capacity. Once joined, passengers can move from one module to another, and the modules can split when beneficial. Utilizing this modularity, [7] and [8] investigated the joint optimization of an AMV system's dispatch headway and vehicle capacity. Similarly, [9] used the concept of "in-motion transfer" to evaluate the benefits of the modular bus system in large-scale urban networks. Most importantly, [10] proposed a rule-based bus-splitting policy to decrease the average passenger travel time using the modular bus system, and [11] combined it with bus holding to improve its performance.

Taking the motivation from previous work on AMVs, this paper proposes a planning strategy for modular vehicles that takes into account the relative significance of different stops in terms of passenger arrival rates. We use the notion of a *Markov decision process* (MDP) to model the underlying system, with states that include the locations, loads, and other salient properties of all modules. The problem of minimizing passenger travel time is then reduced to that of optimal planning on an MDP. The cost of different actions in the given state is formulated to capture the effects of those actions on average travel time. Depending on the state, available actions include stopping at the upcoming stop, skipping the stop, splitting into two modules, and joining with other modules. The costs are computed using the current state of the system and parameters specific to bus stops such as passenger arrival and alighting rates. At each time step, the strategy commands the modules to implement the actions which correspond to the minimum cost. The proposed policy's performance is compared extensively against — and shown to outperform — the rule-based policy proposed by [10] and a policy where all buses serve all stops as in the traditional bus transit system.

II. Problem Formulation

We consider a single-route passenger transportation system composed of n modular vehicles. In the rest of the paper, a *module* refers to a vehicle that cannot split further, whereas *bus* refers to a vehicle that is formed by joining multiple modules. We use *vehicle* to refer generally to both modules and buses.

This work was partially supported by the NYUAD Center for Interacting Urban Networks (CITIES), funded by Tamkeen under the NYUAD Research Institute Award CG001, as well as by a Discovery Partners Institute Science Team award.

K. Jagdale and M. Ornik are with the Department of Aerospace Engineering and the Coordinated Science Laboratory, University of Illinois Urbana-Champaign, Urbana, USA. karansj2@illinois.edu, mornik@illinois.edu

Z. S. Khan and M. Menéndez are with the Division of Engineering, New York University Abu Dhabi, Abu Dhabi, United Arab Emirates. zaid.khan@nyu.edu, monica.menendez@nyu.edu

While autonomous modular vehicles can actually join and detach while moving [6], in this paper we consider a conservative scenario that allows this maneuver only at the bus stops. When the modules are joined, they create an open common area, allowing passengers to walk from one module to another. We assume that no more than two modules can be joined at any given time and no more than one vehicle can serve a stop at a time. We recognize that the technical assumption of only two modules joining is limiting and does not allow us to fully reap the benefits of a modular framework. Yet, as we will show, strategies that make use of modularity significantly outperform classical plans even under such a constraint.

In addition to attaching and detaching at bus stops, vehicles can also choose to skip a stop when it is beneficial to the overall system performance, even if some passengers wish to board or disembark there.

The vehicles are assumed to operate on a single route in the form of a loop with m stops. We assume that, as in many high-frequency urban transportation systems [12], the bus schedule is not publicly available, and the number of passengers arriving at the bus stops does not depend on the actions taken by the vehicles. Under this setting, we consider the problem of finding a strategy for module joining, splitting and stop skipping that minimizes the *average travel cost* of passengers. The travel time of a passenger consists of three elements — waiting time at their origin stop, time spent traveling inside the vehicle (in-vehicle time), and the time spent walking back to the destination stop if the vehicle has skipped their destination stop. Thus, we define the average travel cost as a weighted sum of average waiting, walking, and in-vehicle time with weights $w_{wait} > 1, w_{walk} > 1$ and 1, respectively.

While the definition of travel cost through weighted travel element times is motivated by the work of [13] and used in recent work including [10], [11], we recognize that an attempt to capture the greater valuation of walking and waiting time is simplistic, particularly for passengers with limited mobility. A possible way to partially accommodate this issue is to announce a stop skipping action slightly ahead of time, allowing passengers to get off or change buses at an earlier stop. However, for simplicity we proceed without such an model in this paper.

We provide a more rigorous definition of the above problem after formally describing the system as a Markov decision process in Section III.

III. System Model

We begin with modeling the bus system using a discrete-time Markov decision process [14]. While we recognize that the underlying system is naturally continuous in time, we use discrete time steps in order to simplify modeling and develop a policy suitable for implementation.

A. Markov Decision Process Environment

A Markov decision process (MDP) consists of three primary components: *state space*, i.e., the set of all possible system states, *action set*, i.e., the set of actions available to the system at every time, and *transition function*, i.e., a function determining the transition probabilities from one state to another, given a system action. We describe the three components below.

If there are N vehicles present at a given time step — noting that the number of vehicles changes as they split and join — we denote the state of vehicle i by $(s_i, l_i, q_i, j_i, t_i)$, where s_i is the location of the vehicle, defined as the last stop which the vehicle has visited, l_i is the vehicle's current passenger load, q_i and j_i are binary variable indicating whether the vehicle is currently at a stop and whether it is formed by attaching two modules, respectively, and t_i is a variable indicating the minimum remaining time to reach the next stop — if the vehicle is not at a stop — or the minimum remaining time left before it can leave the stop, either due to serving passengers or due to waiting for the next vehicle to join with it. The overall state of the system is composed of the states of each vehicle, as well as the auxiliary variables such as the number of passengers waiting at each stop and headway, described below.

We now elaborate on the actions available to the vehicles. Before a vehicle reaches the next stop, the available actions are $stop$ — the vehicle stops at the stop and lets the passengers board and alight, $skip$ — the vehicle does not stop at the stop and continues onwards, and $split$ — the vehicle splits at the stop; the front module $skips$ the stop and the rear module $stops$ at the stop. Naturally, the $split$ action is available only if the vehicle is a bus with two modules. If a vehicle is at a stop and all passengers have boarded, its available actions are $next$ — the vehicle leaves for the next stop, and $join$ — the vehicle waits for the following module and attaches to it when that module reaches stop s. The $join$ action is available only if the vehicle leaving the stop and the vehicle following it are both individual modules. For simplicity, we do not allow a vehicle to remain at a stop after all passengers have boarded unless it is waiting to join another module.

To describe the system's transition function, i.e., dynamics, we first introduce some notation. The passenger capacity of each of n identical modules is denoted by K. The per-passenger boarding and alighting times are denoted by β and α. We assume that passengers board and alight concurrently. A fixed amount of time E is lost at a stop due to acceleration, deceleration, opening, and closing of the doors. Largely for notational simplicity, we assume that vehicles are prevented from overtaking each other. Let the vehicle behind and in front of vehicle i be denoted by i^+ and i^-, respectively. Thus, $i^+ = i+1$ when $i \neq n$, if $i = n$, $i^+ = 1$. We define i^- analogously. Similarly, s^+ and s^- denote the stops following and preceding stop s, respectively.

We define $la_{i,s}$ as the number of *leftover alighting passengers* in vehicle i when it reaches stop s, i.e., the number of the passengers who could not alight at s^- due to the vehicle skipping it. Similarly, we define $lo_{i,s}$ as the passengers waiting at stop s at that time who could not board vehicle i^-, either due to its limited capacity or because it had skipped stop s. Suppose $ta_{s,i}$ and $td_{i,s}$ are the arrival and departure times of vehicle i to/from stop s respectively. We define the *arriving* and *departing headway* of vehicle i by $ha_{i,s} = ta_{i,s} - ta_{i^-,s}$ and $hd_{i,s} = td_{i,s} - td_{i^-,s}$, respectively. We are now ready to describe the probability distributions that drive the system dynamics.

B. Uncertainty

We recognize two sources of uncertainty: in the arrival of passengers to a stop and in vehicle travel times.

Passenger load uncertainty. We use the same specifications described in [10]. We assume that passengers arrive at stop s with a fixed rate λ_s per second following a Poisson process [15]. Thus, the number of passengers arriving in time interval of length T at stop s is distributed as $\mathbf{Poisson}(T\lambda_s)$. Hence, when vehicle i arrives at stop s, the total number of passengers waiting at the bus stop is distributed as

$$pa_{i,s} \sim \mathbf{Poisson}(\lambda_s ha_{i,s}) + lo_{i,s}. \tag{1}$$

We assume that stop s is the destination of each passenger on the bus with probability p_s. Thus, the number of passengers alighting from the vehicle i as stop s is given by

$$pd_{i,s} \sim \mathbf{Binomial}(l_{i,s} - la_{i,s}, p_s) + la_{i,s}. \tag{2}$$

After the passengers have alighted, the load of the vehicle, excluding passengers boarding the vehicle at stop s, is given by $l'_{i,s} = l_{i,s} - pd_{i,s}$, and the number of passengers boarding the vehicle is given by $pb_{i,s} = \min(pa_{i,s}, 2^{j_i}K - l'_{i,s})$.

Travel time model. Let V_{veh} be the average cruising speed of the vehicle, which is assumed to be the same whether it is a bus or an individual module. If D_s is the distance between stop s and s^+, we assume that the travel time to go from stop s to stop s^+ is

$$c_s = D_s/V_{veh} + \epsilon_s, \tag{3}$$

where ϵ_s is a Gamma-distributed zero-mean error $\epsilon_s \sim \mathbf{Gamma}(\kappa, \theta) - \kappa\theta$, and κ, θ are the standard parameters of the Gamma distribution, chosen to model a realistic degree of stochasticity.

We now proceed with defining the action-driven state transitions.

C. State Dynamics

When vehicle i reaches a stop, s_i is updated to s_i^+. Variable q_i is set to 1 when vehicle i reaches a stop and is set to 0 when vehicle i leaves a stop. Similarly, we set j_i to 0 if the vehicle i splits and set it to 1 if the vehicle i joins to another vehicle. We now elaborate on the transitions of the other state variables.

Action *stop***.** After completing passenger boarding and alighting, the load of the vehicle is updated as $l_{i,s^+} = l_{i,s} - pd_{i,s} + pb_{i,s}$. and the leftover passengers at stop s are given by

$$lo_{i,s} = pa_{i,s} - pb_{i,s}. \tag{4}$$

As vehicle i is allowing passengers to alight at stop s, the number of leftover alighting passengers for the next stop s^+ is given by $la_{i,s^+} = 0$.

We define $tp_{i,s}$ as the delay faced by vehicle i when it reaches stop s due to the previous vehicle i^- still being at the stop. If the previous vehicle departed stop s by the time vehicle i reaches the stop, then $tp_{i,s} = 0$. The time duration after which vehicle i is able to leave stop s is given by

$$t_{i,s} = tp_{i,s} + \max(\alpha pd_{i,s}, \beta pb_{i,s}) + E. \tag{5}$$

Adding $tp_{i,s}$ in equation (5) ensures that vehicle i does not overtake the vehicle in front and it is consistent with our assumptions that only one vehicle can serve a stop at a time.

Action *skip***.** In this case, vehicle i does not serve the stop s. Hence, the passenger load of the vehicle remains the same, i.e., $l_{i,s^+} = l_{i,s}$, and the leftover alighting passengers for the next stop s^+ are given by $la_{i,s^+} = pd_{i,s}$. The leftover passengers at the stop s for vehicle i^+ reaching stop s are given by $lo_{i^+,s} = pa_{i,s}$. The time duration after which the vehicle i can leave stop s is given by $t_{i,s} = tp_{i,s}$, ensuring that vehicle i does not overtake the vehicle in front.

Action *split***.** The passengers are redistributed between the two joined modules before splitting the bus, with passengers who intend to alight at stop s moving — capacity permitting — to the trailing module, which will stop at s, and others moving — again capacity permitting — to the front module, which will skip s. Namely, let $lf_{i,s}$ and $lb_{i,s}$ denote loads of the leading and trailing modules after redistributing the passengers. Let $\mathbb{1}$ denote the indicator function, i.e., $\mathbb{1}_{(\cdot)} = 1$ if condition $(\cdot)$ is true and $\mathbb{1}_{(\cdot)} = 0$ otherwise. Then, $lb_{i,s} = \mathbb{1}_{pd_{i,s} > K} K + \mathbb{1}_{pd_{i,s} \le K} pd_{i,s}$ and $lf_{i,s} = \mathbb{1}_{l_{i,s} - lb_{i,s} > K} K + \mathbb{1}_{l_{i,s} - lb_{i,s} \le K} (l_{i,s} - lb_{i,s})$.

The leftover alighting passengers that cannot fit in the trailing module and are left in the leading module are denoted by $laf_{i,s}$ and are given as $laf_{i,s} = \mathbb{1}_{pd_{i,s} > K}(pd_{i,s} - K)$. The time taken at stop s for the leading and trailing modules of the vehicle i is denoted by $tf_{i,s}$ and $tb_{i,s}$, respectively, and is given by $tf_{i,s} = tp_{i,s}$ and $tb_{i,s} = tp_{i,s} + max(\alpha pd_{i,s}, \beta pb_{i,s}) + E$.

After implementing the *split* action the leading module becomes vehicle i, trailing module becomes vehicle i^+ and all the remaining vehicles behind the trailing module shift by a unit. Similar to the case of the *stop* action, the number of leftover passengers at stop s is given by equation (4).

Action *join***.** The bus formed by joining module i and module i^+ is the new vehicle i, and the indices of all trailing vehicles are decremented by 1. We add the loads and the leftover alighting passengers of the modules and associate them with index i. Thus, $l_{i,s} := l_{i,s} + l_{i^+,s}$, and $la_{i,s} := la_{i,s} + la_{i^+,s}$. Let $td'_{i,s}$ be the time at which vehicle i performed the *join* action. Note that $td'_{i,s}$ is different from $td_{i,s}$, which is the time when vehicle i will leave stop s. Then, the time spent by vehicle i at stop s is equal to the waiting time for the trailing module, given by $t_{i,s} = ta_{i^+,s} + t_{i^+,s} - td'_{i,s}$, where $t_{i^+,s}$ is the time duration for which module i^+ stops at stop s.

Action $next$**.** If the vehicle uses the $next$ action, it leaves for the next bus stop. Noting that we do not allow vehicles to overtake each other, the time duration for which vehicle i travels between stops s and s^+ is given by $tc_{i,s} = max(c_s, cp_{i,s})$, where c_s is given by (3) and $cp_{i,s}$ is the remaining travel time for the downstream vehicle to reach stop s^+ if it has not already done so.

Having defined the relevant Markov decision process, we can now describe the problem of optimal policy design.

D. Optimal Planning Problem

Consider a Markov decision process M described above, noting that we restrict the set of available actions to only those that are possible at a given state — e.g., *split* is not possible if the modules are not joined. A *deterministic, stationary policy* π of an MDP assigns an action to be performed when the system is at given state. We denote the set of all possible deterministic, stationary policies by $\Pi(M)$. Let $T^\pi_{wait}, T^\pi_{walk}, T^\pi_{veh}$ denote the average of the waiting, walking, and in-vehicle time, respectively, across all passengers given a particular policy π and some relevant finite time interval. Let w_{wait}, w_{walk} be the waiting time and walking time weights as described in the Section II. The problem considered by this paper is thus to find a policy π^* such that $\pi^* = \arg\min_{\pi\in\Pi(M)} T^\pi_{veh} + w_{wait}T^\pi_{wait} + w_{walk}T^\pi_{walk}$.

With a formal presentation of the MDP setting and the problem statement, we proceed to describe the proposed approach to approximately solve the stated problem.

IV. Proposed Approach

Owing to the complicated and stochastic nature of the relationship between a policy and passenger waiting, walking, and in-vehicle times, the problem of minimizing $T^\pi_{veh} + w_{wait}T^\pi_{wait} + w_{walk}T^\pi_{walk}$ is difficult to solve directly. Instead, following a standard approach for optimal planning in MDPs [14], we will encode the approximate impact of each vehicle action on passenger travel times by a *cost function*, and determine a policy which minimizes the overall incurred cost. We start by formulating the cost function.

A. Cost Design

We formulate the cost of an action in a way that captures the expected increase in passengers' *ideal travel time* due to the action. The ideal travel time of a passenger is defined as the time spent in the vehicle from the origin bus stop to the destination bus stop assuming that the vehicle does not stop at intermediate stops and the passenger does not have to wait for vehicle at the origin bus stop.

In the cost formulation, the increase in the travel time resulting from an action is weighted by the number of affected passengers. We thus begin by computing the expected numbers of passengers boarding and alighting at all stops.

Say that vehicle i is serving stop s. Let $pd^e_{i,s}$ denote the *expected* number of alighting passengers. Then $pd^e_{i,s}$ can be calculated using equation (2), i.e., $pd^e_{i,s} = \mathbb{E}[pd_{i,s}] = (l_{i,s} - la_{i,s})p_s + la_{i,s}$. Similarly, we compute the approximate value of expected number of passengers arriving at stop s, $pa^e_{i,s}$, by $pa^e_{i,s} = \mathbb{E}(pa_{i,s}) = \lambda_s hd_{i,s^-} + lo_{i,s}$. Since the arriving headway becomes known only after reaching the stop, we have taken the expectation of $pa_{s,i}$ using equation (1) and used the departing headway instead of the arriving headway.

The number of passengers boarding vehicle i of total passenger capacity $K' = 2^{j_i}K$ with average passenger arrival rates is given by $pb^e_{i,s} = \min(\lambda_s hd_{i,s^-} + lo_{i,s}, K' - l_{i,s} - pd^e_{i,s})$. We approximate the time that vehicle i spends serving stop s by $t^e_{i,s} = tp_{i,s} + \max(\alpha pd^e_{i,s}, \beta pb^e_{i,s}) + E$. Assuming that the headway of vehicle i remains constant between stops s_i and s_{i^+} and that vehicle i^- has stopped at all the stops from s_i to s_{i^+}, the expected number of passengers that intend to board vehicle i at stops s to s_{i^+} is

$$p_{s,i,s_{i^+}} = ha_{i,s} \sum_{s'=s^+}^{s_{i^+}} \lambda_{s'}. \tag{6}$$

For computational reasons, we simplistically assume that an action implemented on vehicle i at stop s only affects the passengers inside that vehicle and the passengers waiting to board at stops s_i to s_{i^+}. We can now move to determining the cost of each action.

Action $stop$**.** If the vehicle stops at a bus stop, the ideal travel time of the passengers that are not alighting at stop s increases by $t^e_{i,s}$. Moreover, the waiting time of passengers at the stops between s and s_{i^+} increases by $t^e_{i,s}$. Additionally, the travel time is increased for some of the passengers who arrive at stops s_i to s_{i^+} after vehicle i departs from s_i. However, computing the expected increase in their travel time is difficult, as some of those passengers would have missed vehicle i if it did *not* stop as s. For these reasons, we use the parameter p_{st} as a tuning parameter to the cost in the increase of travel time for passengers at stops s_i to s_{i^+}, and choose p_{st} by evaluating the proposed policy for different values of p_{st}.

Combining the above considerations, if $c(s, i, a)$ gives the cost of action a applied to vehicle i in state s, then

$$\begin{aligned} c(s, i, stop) = {} & w_{wait}\, t^e_{i,s}\, (l_{i,s} - pd^e_{i,s}) \\ & + w_{wait}\, t^e_{i,s}\, p_{st}\, p_{s,i,s_{i^+}}. \end{aligned} \tag{7}$$

Action $skip$**.** If the bus skips the stop, passengers who intended to board the bus will have to wait for the next one, increasing their waiting time. Additionally, passengers who intended to alight at the stop will have to walk from the next stop to the current stop, increasing their walking time. The expected walking time for passengers from stop s^+ to s is given by $W_s = D_s/v_{pass}$, where v_{pass} is the average passenger walking speed. Thus, the cost of $skip$ action is given by $c(s, i, skip) = pb^e_{i,s} w_{wait} ha_{i^-,s} + pd^e_{i,s} W_s w_{walk}$.

Actions $next$ **and** $split$**.** The cost of $next$ action is zero, as it does not increase the passenger travel time. In other words, $c(s, i, next) = 0$. If there were no capacity limitations, the same would be true for the $split$ as it distributes the passengers based on their desired alighting stop. Due to the limited capacity of the modules, the $split$ action might increase the travel time of certain passengers in cases where the module capacity is insufficient to serve all the passengers who wish to alight or board at a certain stop. However, this event is unlikely in practice unless a particular stop is an extremely busy hub, in which case the split action can be disallowed at that stop. For simplicity, we assume $c(s, i, split) = 0$.

Action $join$**.** The $join$ action is only available when a separated module is about to leave the stop, with the only other possible action being $next$. As previously stated, the cost of the $next$ action is zero. However, the cost of the $join$ action is positive because the passengers in the leading module have to wait for the trailing module to complete passenger boarding/alighting at stop s. As a result, cost minimization with an immediate decision horizon would never result in a join action. On the other hand, longer decision horizons lead to computational issues, as well as an increasingly poor correlation between the cost model and actual increases in travel time. Thus, instead of considering the cost of $join$ directly, we decide whether to pursue $join$ based on a rule-based policy described below.

B. Proposed Policy

When a bus is about to reach a stop, we choose the action among $stop$, if available $split$, and $skip$ which gives the lowest immediate cost described above. When a separated module is about to leave the stop, we consider the additional dwell time required for the trailing module to be ready to depart from the stop. If this time is shorter than a tuned threshold τ, we join the modules.

We now describe the numerical results obtained with the proposed policy.

V. Simulation Experiments

In order to validate the policy and compare it to the state of the art, we use the exact system specifications as described in [10]. We draw the benchmark policy from [10] as well, referring to it as *split policy*, and begin the section by describing this strategy.

A. Benchmark Policy

In the split policy [10], a bus splits at stop s if its headway exceeds a particular threshold described in [10]; otherwise, it stops at the stop. After splitting, the trailing module serves stop s, and the leading module leaves for the next stop s^+. The leading module waits for the trailing module at stop s^+. After serving stop s, the trailing module alights the passengers at stop s^+, where the leading module is waiting, and then gets attached to the leading module. It was shown in [10] that such a simple policy already outperforms an alternative state-of-the-art stop-skipping policy. We now proceed to describe our simulation setting.

B. Simulation Scenario

We follow the setting of [10]. The parameters p_s, λ_s, D_s are drawn from a normal distribution with mean $\bar{p}, \bar{\lambda}, \bar{D}$ and the standard deviation equal to 10% of the mean for the respective parameter. We assume uniform origin-destination demand, with the passengers traveling $\frac{s}{2}$ stops on average resulting in $\bar{p} = \frac{2}{s}$. Upon performing multiple simulations with different values of p_{st} — the parameter identified in (7) used to tune the cost of $stop$ action — the optimal value is obtained to be 0.5. Input parameters used in the numerical simulations are all taken directly from [10]. Namely, we consider a scenario with 20 stops, at an average distance of 400 meters from each other, and 24 modules, each with a capacity of 40 passengers. We set the wait time weight w_{wait} to 2.1 and the walk time weight w_{walk} to 2.2, based on [13]. We invite the reader to consult [10] for a full description of the simulation parameters. Preprint [11] posted roughly in parallel with this paper's initial version proposes an alternative policy, but uses different parameters in its numerical work, so for the purpose of our simulations we focus on the peer-reviewed results of [10].

The vehicles initially start from stop 1 at time 0, with start times staggered to dispatch them with the *ideal headway*, which is computed using the method given in [10]. To ensure the system reaches its usual operating conditions before policy evaluation, we define a *evaluation period*, which begins after all the vehicles complete two rounds of the bus route and lasts for an hour. We present all the results by assessing the policy performance in the evaluation period. All modules are initially separated.

C. Numerical Results

To examine the relative performance of the proposed policy, we simulate three different policies, namely, (i) a *no-control policy* — a policy in which each module stops at each stop, without joining or splitting, (ii) the split policy from [10], and (iii) the proposed policy, with threshold τ equal to 31 seconds and stop cost parameter p_{st} set to 0.5. We plot the number of in-vehicle passengers and the waiting passengers as a function of time in Fig. 1 for one illustrative simulation. The number of in-vehicle passengers for the proposed policy is smaller than in the other two policies, making it more capable of handling higher passenger demand. The number of waiting passengers for the proposed policy is significantly smaller than the no-control and the benchmark policy. This feature is reflected in Table I as the average waiting time for the proposed policy is significantly smaller for the proposed policy. A small number of walking passengers, shown in Fig. 1, are observed for the proposed policy as it involves performing $skip$ actions.

Table I, based on 100 simulations, extensively compares the proposed policy with the benchmark policy and no-control case. The values of metrics for the policy used in [10] are taken directly from [10]. The proposed policy performs significantly better than the benchmark policy and decreases the average travel cost by around 13% compared to the latter. In addition, using the proposed policy, the passengers experience the least waiting time, and the average load for the proposed policy is the lowest among all policies, making it more capable of handling higher passenger demand.

To analyze the policies' impact on bus bunching, we plot time-space diagrams of the vehicle trajectories in Fig. 2,

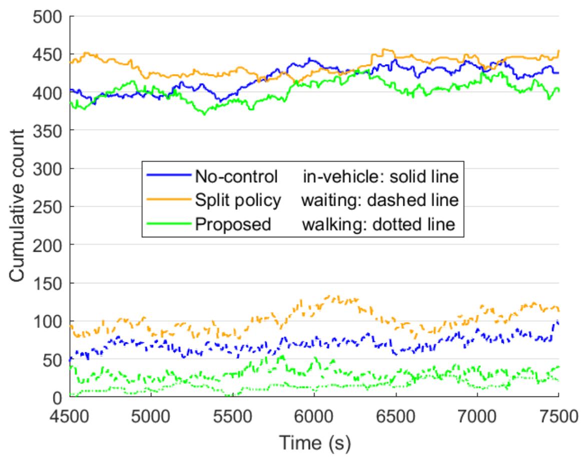

Fig. 1. Number of in-vehicle, waiting, and walking passengers obtained with the no-control policy, split policy from [10], and the proposed policy.

TABLE I
PERFORMANCE COMPARISON

Metric	Units	No Control	Benchmark Policy [10]	Proposed Policy
Average wait time	min	2.6	2.2	1.46
Average walk time	min	-	-	0.82
Average in-vehicle time	min	19.5	20.2	16.98
Weighted travel cost Q	min	24.96	24.90	21.87
Average cycle length	min	36.1	40.2	32.7
Average load per module	pax	18.31	19.95	17.53
Fraction of full vehicles	-	10%	0	0

plotting only trajectories of buses under the no-control policy and when using the proposed policy. The figure shows that, in the former case, buses are bunched together in large groups. On the other hand, the proposed policy results in nearly uniform bus headways with insignificant bus bunching, once again validating the strength of the proposed policy.

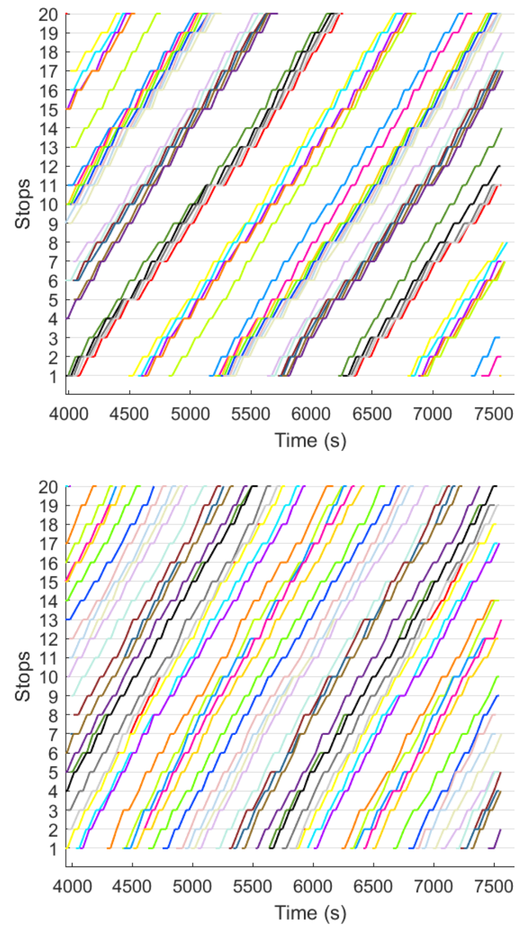

Fig. 2. Time-space plots of all the buses during the evaluation period, comparing two strategies: (top) when every module stops at every stop and all the modules are separated, (bottom) policy proposed in this paper. Different colors corresponds to different modules.

References

[1] Carlos F. Daganzo. A headway-based approach to eliminate bus bunching: Systematic analysis and comparisons. *Transportation Research Part B: Methodological*, 43(10):913–921, 2009.

[2] Simon J. Berrebi, Etienne Hans, Nicolas Chiabaut, Jorge A Laval, Ludovic Leclercq, and Kari E Watkins. Comparing bus holding methods with and without real-time predictions. *Transportation Research Part C: Emerging Technologies*, 87:197–211, 2018.

[3] Zhiyuan Liu, Yadan Yan, Xiaobo Qu, and Yong Zhang. Bus stop-skipping scheme with random travel time. *Transportation Research Part C: Emerging Technologies*, 35:46–56, 2013.

[4] Diego Morales, Juan Carlos Muñoz, and Pedro Gazmuri. A stochastic model for bus injection in a public transport service. *Transportation Research Procedia*, 38:688–708, 2019.

[5] Antoine Petit, Yanfeng Ouyang, and Chao Lei. Dynamic bus substitution strategy for bunching intervention. *Transportation Research Part B: Methodological*, 115:1–16, 2018.

[6] Next Future Mobility. https://www.next-future-mobility.com/.

[7] Zhiwei Chen, Xiaopeng Li, and Xuesong Zhou. Operational design for shuttle systems with modular vehicles under oversaturated traffic: Discrete modeling method. *Transportation Research Part B: Methodological*, 122:1–19, 2019.

[8] Igor Dakic, Kaidi Yang, Mónica Menéndez, and Joseph Y. J. Chow. On the design of an optimal flexible bus dispatching system with modular bus units: Using the three-dimensional macroscopic fundamental diagram. *Transportation Research Part B: Methodological*, 148:38–59, 2021.

[9] Jiaming Wu, Balázs Kulcsár, Selpi, and Xiaobo Qu. A modular, adaptive, and autonomous transit system (MAATS): An in-motion transfer strategy and performance evaluation in urban grid transit networks. *Transportation Research Part A: Policy and Practice*, 151:81–98, 2021.

[10] Zaid Saeed Khan, Weili He, and Mónica Menéndez. Application of modular vehicle technology to mitigate bus bunching. *Transportation Research Part C: Emerging Technologies*, 146, 2023.

[11] Zaid Saeed Khan and Mónica Menéndez. Bus splitting + bus holding: A new strategy using autonomous modular buses for preventing bus bunching. *Preprint available at SSRN 4356986*, 2023.

[12] Jaime Soza-Parra, Oded Cats, Yvonne Carney, and Catherine Vanderwaart. Lessons and evaluation of a headway control experiment in Washington, DC. *Transportation Research Record*, 2673(8):430–438, 2019.

[13] Mark Wardman. Public transport values of time. *Transport Policy*, 11(4):363–377, 2004.

[14] Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.

[15] Guy Desaulniers and Mark D. Hickman. Public transit. In Cynthia Barnhart and Gilbert Laporte, editors, *Transportation*, volume 14 of *Handbooks in Operations Research and Management Science*, pages 69–127. Elsevier, 2007.