Coordinated Demand-side Management and Traffic Control for Tight Areas

March 2023

A Research Report from the Pacific Southwest Region University Transportation Center

Ketan Savla, University of Southern California

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
PSR-21-25	N/A	N/A

4. Title and Subtitle	5. Report Date
Coordinated Demand-side Management and Traffic Control for Tight Areas	March 2023
	6. Performing Organization Code N/A

7. Author(s)	8. Performing Organization Report No.
Ketan Savla, 0000-0002-1668-6380	PSR-21-25

9. Performing Organization Name and Address	10. Work Unit No.
METRANS Transportation Center University of Southern California University Park Campus, RGL 216 Los Angeles, CA 90089-0626	N/A
	11. Contract or Grant No. USDOT Grant 69A3551747109

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE, Washington, DC 20590	Final report (August 15, 2021-August 16, 2022)
	14. Sponsoring Agency Code USDOT OST-R

15. Supplementary Notes
https://www.metrans.org/research/coordinated-demand-side-management-and-traffic-control-for-tight-areas

16. Abstract
We develop traffic management strategies for urban areas characterized by tight areas by control of demand and supply. Specifically, on the supply side, we design boundary control, in the form of ramp metering. On the demand side, we design time-dependent incentives/payment schemes to spread the demand during periods of peak demand. We systematically analyze the social optimum under a user equilibrium as well as budget and maximum toll constraints for a single bottleneck. We cast the congestion pricing problem as a bilevel optimization problem and provide several analytical and numerical results. Specifically, we show that the bilevel optimization problem can be converted into a convex optimization problem under some weak assumptions of the schedule delay cost under an inelastic demand setting. The methodological contributions are supplemented with illustrative simulation results. We consider ramp metering (RM) at the microscopic level subject to vehicle following safety constraints for a single freeway with arbitrary number of on- and off-ramps. The arrival times of vehicles to the on-ramps, as well as their destinations are modeled by exogenous stochastic processes. Once a vehicle is released from an on-ramp, it accelerates towards the free flow speed if it is not obstructed by another vehicle; once it gets close to another vehicle, it adopts a safe gap vehicle following behavior. The vehicle exits the freeway once it reaches its destination off-ramp. We design traffic-responsive RM policies that maximize the freeway *throughput*. All the proposed policies are *reactive*, meaning that they only require real-time traffic measurements without the need for demand prediction. The throughput of these policies is proven to be maximized when the merging speed of all the on-ramps equals the free flow speed. Simulations are provided to illustrate the performance of our policies and compare with a well-known RM policy from the literature.

17. Key Words	18. Distribution Statement
Traffic, congestion pricing, ramp metering	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	44	N/A

About the Pacific Southwest Region University Transportation Center

The Pacific Southwest Region University Transportation Center (UTC) is the Region 9 University Transportation Center funded under the US Department of Transportation's University Transportation Centers Program. Established in 2016, the Pacific Southwest Region UTC (PSR) is led by the University of Southern California and includes seven partners: Long Beach State University; University of California, Davis; University of California, Irvine; University of California, Los Angeles; University of Hawaii; Northern Arizona University; Pima Community College.

The Pacific Southwest Region UTC conducts an integrated, multidisciplinary program of research, education and technology transfer aimed at *improving the mobility of people and goods throughout the region.* Our program is organized around four themes: 1) technology to address transportation problems and improve mobility; 2) improving mobility for vulnerable populations; 3) Improving resilience and protecting the environment; and 4) managing mobility in high growth areas.

U.S. Department of Transportation (USDOT) Disclaimer

Disclosure

Ketan Savla conducted this research titled, "Coordinated Demand-side Management and Traffic Control for Tight Areas" at the Viterbi School of Engineering at the University of Southern California. The research took place from August 15, 2021, to August 16, 2022, and was funded by a grant from the U.S. Department of Transportation in the amount of $100,000. The research was conducted as part of the Pacific Southwest Region University Transportation Center research program.

Disclaimer Statement

The contents of this report reflect the views of the authors, who are responsible for the accuracy of the data and information presented herein. This document is disseminated under the sponsorship of the Department of Transportation, University Transportation Centers Program, the California Department of Transportation and the METRANS Transportation Center in the interest of information exchange. The U.S. Government, the California Department of Transportation, and the University of Southern California assume no liability for the contents or use thereof. The contents do not necessarily reflect the official views or policies of the State of California, USC, or the Department of Transportation. This report does not constitute a standard, specification, or regulation.

Abstract

We develop traffic management strategies for urban areas characterized by tight areas by control of demand and supply. Specifically, on the supply side, we design boundary control, in the form of ramp metering. On the demand side, we design time-dependent incentives/payment schemes to spread the demand during periods of peak demand.

We systematically analyze the social optimum under a user equilibrium as well as budget and maximum toll constraints for a single bottleneck. We cast the congestion pricing problem as a bilevel optimization problem and provide several analytical and numerical results. Specifically, we show that the bilevel optimization problem can be converted into a convex optimization problem under some weak assumptions of the schedule delay cost under an inelastic demand setting. The methodological contributions are supplemented with illustrative simulation results.

We consider ramp metering (RM) at the microscopic level subject to vehicle following safety constraints for a single freeway with arbitrary number of on- and off-ramps. The arrival times of vehicles to the on-ramps, as well as their destinations are modeled by exogenous stochastic processes. Once a vehicle is released from an on-ramp, it accelerates towards the free flow speed if it is not obstructed by another vehicle; once it gets close to another vehicle, it adopts a safe gap vehicle following behavior. The vehicle exits the freeway once it reaches its destination off-ramp. We design traffic-responsive RM policies that maximize the freeway *throughput*. All the proposed policies are *reactive*, meaning that they only require real-time traffic measurements without the need for demand prediction. The throughput of these policies is proven to be maximized when the merging speed of all the on-ramps equals the free flow speed. Simulations are provided to illustrate the performance of our policies and compare with a well-known RM policy from the literature.

Contents

1 Disclosure **6**

2 Acknowledgements **6**

3 Introduction **7**

4 Task 1: Demand Control of Bottlenecks through Incentive Design **7**
4.1 MODEL . . . 8
4.1.1 Equilibrium Formulation For A Single Bottleneck With Inelastic Demand . . 9
4.2 EQUIVALENT MATHEMATICAL PROGRAMS . . . 10
4.2.1 A Single Bottleneck With Inelastic Demand . . . 10
4.3 SOCIAL OPTIMUM WITH PRICING CONSTRAINTS . . . 11
4.3.1 Social Optimum For A Single Bottleneck . . . 12
4.3.2 Equilibrium Formulation For A Single Bottleneck With Elastic Demand . . . 14
4.3.3 A Single Bottleneck With Elastic Demand . . . 14
4.3.4 Social Optimum For A Single Bottleneck With Elastic Demand . . . 16
4.4 SIMULATION . . . 16
4.4.1 Example 1 . . . 17
4.4.2 Example 2 . . . 17

5 Task 2: Supply Control through Microscopic Ramp Metering **19**
5.1 Problem Formulation . . . 21
5.1.1 Vehicle-Level Objectives . . . 22
5.1.2 Probabilistic Demand Model and Throughput . . . 24
5.1.3 Ramp Metering . . . 26
5.2 Ramp Metering Policies and Performance Analysis . . . 27
5.2.1 Renewal Policy . . . 27
5.2.2 Dynamic Release Rate Policy . . . 29
5.2.3 Distributed Dynamic Release Rate Policy . . . 32
5.2.4 Dynamic Space Gap Policy . . . 33
5.2.5 Local and Greedy Policies . . . 34
5.2.6 An Outer Estimate . . . 35
5.2.7 Discussion of the Straight Road Geometry . . . 36
5.3 Simulations . . . 36
5.3.1 Greedy Policy for Low Merging Speed . . . 37
5.3.2 Effect of Cycle Length on Queue Size . . . 37
5.3.3 Relaxing the V2X Requirements . . . 38
5.3.4 Comparing the Total Travel Time . . . 39

6 Conclusion **40**

7 Implementation **41**

8 References **41**

9 Data Management Plan **43**

List of Tables

1 Summary of the RM policies studied in this project. 27
2 Performance of the policies. 40

1 Disclosure

K. Savla has financial interest in Xtelligent, Inc.

2 Acknowledgements

This research was performed in collaboration with M. Pooladsanj, G. Rostomyan, and P. A. Ioannou at the University of Southern California.

3 Introduction

Traffic systems in several spots of a large city are characterized by tight spaces. Therefore, during peak period, when traffic demand exceeds supply, congestion could spill to adjoining areas, which could then lead to cascading effects. We develop traffic management strategies for such systems by control of demand and supply. Specifically, on the supply side, we design boundary control, in the form of ramp metering. On the demand side, we design time-dependent incentives/payment schemes to spread the demand during periods of peak demand. In contrast to the current literature, we explicitly include budget constraint of a city as well as cap on payment for an individual driver when designing incentives. Unlike existing literature, our ramp metering algorithms operate at the microscopic level and hence are able to explicitly handle vehicle safety constraints.

4 Task 1: Demand Control of Bottlenecks through Incentive Design

Bottleneck related congestion involves traffic demand exceeding a bottleneck capacity in a limited time interval. A typical example is a rush hour traffic between residential areas and a central business district or a popular sporting/cultural event. A traffic incident could also lead to congestion. If left uncontrolled, the congestion could lead to long and persistent queues, which not only result in a wasteful loss of time for travelers, but can also spill over into adjacent areas.

A natural approach for eliminating congestion is the demand-side decentralized control mechanism which spreads the arrivals to the congested bottleneck over time to achieve an optimal trade-off between the traveled time in congestion and unpunctuality, (i.e., early or late arrival at the destination). However, prior approaches did not consider a strategy of using a combination of incentives and tolls under budget and maximum toll constraints. The budget and maximum toll constraints are motivated by the fact that relying solely on incentives may put a large strain on the budget of the traffic management authorities and social fairness, necessitating a maximum toll for the use of the congested bottleneck.

The congestion problem is commonly studied using the single bottleneck(See, for example, [1]). Previous works typically focus only on pricing the congestion (See, for example [2, 3]). Comprehensive reviews of the congestion pricing using marginal cost-based pricing can be found in [4–6]. Even though using tolls to price congestion may eliminate queuing, [7] shows that tolls can make some travelers worse off because they do not consider travelers' income levels, occupations, or the costs of travel times. In addition, there may be resistance from elected officials towards introducing a price for accessing public roads that have been freely available for a long time. In fact, the recent survey paper [2] suggests that schemes which do not involve money transfer from travelers to the traffic management authorities can significantly improve public engagement in the congestion management schemes.

Studies on the effects of incentives on congestion management are relatively few. For example, [8], using [9] marginal-cost based congestion pricing, states that optimal price can be a toll, a reward, or a combination of the two for a perfect inelastic demand setting. However, the marginal-cost based congestion pricing mentioned in [8] is restrictive and does not provide a generalized framework for extending the bottleneck model to optimizing the total social cost given constraints on the offered pricing.

We consider the possibility of incentives for some travelers and tolls for others. We design a range of incentives and tolls to respect the budget constraints as well as the maximum toll while minimizing the social cost and ensuring a user equilibrium. We start with the departure-time based

formulation of the user equilibrium, where the travelers choose their departure times given the exogenous pricing profile. We then propose a systematic approach for determining a combination of incentives and tolls that achieves the social optimum. We further determine the optimal pricing for a single bottleneck under an elastic demand setting. Our proposed approach relies on the following: (1) given an exogenous pricing profile, a corresponding unique equilibrium solution can be obtained by solving equivalent mathematical programs and (2) in the inelastic demand setting and under some weak assumptions on the schedule delay cost, the bilevel optimization problem for obtaining the social optimum can be converted into a single level optimization problem.

Proofs of all the technical results of this section can be found in [10].

4.1 MODEL

The model studied herein assumes that each traveler is travelling on a road segment that connects for example a residential area ("origin") to a central business district ("destination"). It is further assumed, the road segment has only one bottleneck. At the bottleneck, the traffic congestion can be modeled as a point queue having a First-In-First-Out (FIFO) discipline. In addition, we assume that at the beginning of congestion period every traveler that intends to travel via the road segment can choose their departure times from the bottleneck by knowing the equilibrium delays at the bottleneck, and hence select their arrival times by subtracting the equilibrium delays from the bottleneck departure times.

Moreover, we assume each traveler is aware of the pricing profile offered by the system planner and considers his/her private travel cost which includes the schedule delay cost at the destination, travel delay cost at the bottleneck and offered pricing when selecting a departure time from the bottleneck. In one implementation, the system planner displays the pricing profile and the a suggested departure time, resulting in the minimum cost to each traveler. For example, the system planner may display information to the travelers by means of a mobile application or communicate the information directly to the autonomous vehicles that intend to use the road segment containing the bottleneck.

Under the above assumptions, in equilibrium, no traveler can reduce his/her private travel cost by changing his/her departure/arrival time unilaterally. There are two equivalent formulations for the equilibrium: (1) arrival-time based, where the time variable is given in terms of the arrival at the bottleneck and (2) departure-time based, where the time variable is given in terms of the departure from the bottleneck. In this study, we adopt the departure-time based formulation and subsequently show that the equilibrium in the departure time formulation is equivalent to a mathematical program.

The following additional assumptions are made in our model formulation (See [2, 11]):

1. travelers are homogeneous. In other words, the travelers are assumed to have identical values of the times associated with queuing and schedule delay costs.

2. The physical length of the bottleneck does not affect the queuing delays. In other words, the bottleneck is modeled as a point FIFO queue.

3. There is a latent demand for using the bottleneck over the time horizon $\mathcal{T}$, which covers the period within which travelers' departure time may be adjusted and in which all travelers arrive at their destinations.

4. All travelers have the same desired arrival time at the destination; that is travelers arriving at the desired arrival time incur no scheduling cost.

Notation

- t denotes the arrival time at the bottleneck
- s denotes the departure time from the bottleneck
- $l(s)$ denotes the exit rate from the bottleneck at the departure time s
- μ denotes the capacity of the bottleneck e.g., number of travelers which pass within each time interval Δs
- $\psi(s)$ denotes the time-varying incentive/toll faced by a traveler departing from the bottleneck at time s
- $\phi(s)$ denotes the departure time based schedule delay cost faced by a traveler at the departure time s
- $d(t)$ denotes the travel time delay through the bottleneck for a traveler arriving at time t to the bottleneck
- $b(s)$ denotes the travel time through the bottleneck for a traveler departing at time s from the the bottleneck
- τ desired arrival time at the destination

4.1.1 Equilibrium Formulation For A Single Bottleneck With Inelastic Demand

In this study, we adopt a discrete time formulation of the equilibrium conditions. As such, the departure time from the bottleneck is discretized by $s = i\Delta s$, where i is an integer from $\mathcal{I} = [1, ..., M]$ and Δs is the time interval between each discrete time. We first consider the equilibrium conditions for a single bottleneck with inelastic demand. As such, the first equilibrium condition; expressed in terms of the time cost are:
$\forall i \in \mathcal{I}$

$$\begin{cases} C_{dep} = b(i) + \phi(i) + \alpha^{-1}\psi(i), \text{ if } l(i) > 0 \\ C_{dep} \leq b(i) + \phi(i) + \alpha^{-1}\psi(i), \text{if } l(i) = 0, \end{cases} \tag{1}$$

where C_{dep} is the minimum (equilibrium) trip cost for travelers, $b(i)$, $\phi(i)$, and $\psi(i)$ are the sampled version of $b(s)$, $\phi(s)$, and $\psi(s)$, and α is the value of time for each traveler. Note, the equilibrium condition in (1) is given in terms of the departure times from the bottleneck and represents the no-arbitrage condition. The schedule delay cost is typically modeled as a piecewise linear function, [9]:

$$\phi(i) = \begin{cases} \frac{\beta}{\alpha}\left(\tau - i\Delta s\right), & \text{if } i\Delta s < \tau \\ \frac{\gamma}{\alpha}\left(i\Delta s - \tau\right), & \text{if } i\Delta s \geq \tau, \end{cases} \tag{2}$$

where β is the cost of early arrival per unit time and γ is the cost of late arrival per unit time. We will assume $\gamma > \alpha > \beta > 0$, which is consistent with the experimental data (e.g., [12]). The second equilibrium condition models the queuing dynamics of the bottleneck given by $\forall i \in \mathcal{I}$

$$\begin{cases} l(i) = \mu, \text{if } b(i) > 0 \\ l(i) \leq \mu, \text{if } b(i) = 0. \end{cases} \tag{3}$$

The third equilibrium condition is the total demand condition, which requires

$$\sum_{i\in\mathcal{I}} l(i) = N, \tag{4}$$

where $l(i)$ is the number of travelers exiting the bottleneck in the time interval $[i, i+\Delta s]$. Under the FIFO service discipline, the first traveler arriving at the bottleneck earlier than the second traveler cannot depart from it later than the second traveler. This proposition can be expressed as

$$\Delta s i_1 - b(i_1) \le \Delta s i_2 - b(i_2) \Rightarrow i_1 \le i_2. \tag{5}$$

The proposition (5) is equivalent to

$$i_2 - i_1 \ge 0 \Rightarrow b(i_2) - b(i_1) \le \Delta s(i_2 - i_1). \tag{6}$$

Remark 1. *In discrete time, the condition on the schedule delay cost $\phi(i)$ and the incentive $\psi(i)$ under which the FIFO condition* (6) *holds are*

$$\begin{aligned} &\forall i_1, i_2 \in \mathcal{I} \text{ and } i_2 \ge i_1, \\ &\phi(i_1) - \phi(i_2) + \alpha^{-1}(\psi(i_1) - \psi(i_2)) \le \Delta s(i_2 - i_1). \end{aligned} \tag{7}$$

or equivalently $\phi(i) + \alpha^{-1}\psi(i) - \phi(i+1) - \alpha^{-1}\psi(i+1) \le \Delta s$.

The condition (7) follows from applying the equilibrium conditions (1) and (3) in the FIFO condition (6).

In the next section, we show the equivalence between the equilibrium conditions introduced above and a finite dimensional mathematical program. We will show that the departure time based equilibrium formulation in the discrete time posseses a simple mathematical structure for expressing its solution, which is absent in the arrival time based equilibrium formulation.

4.2 EQUIVALENT MATHEMATICAL PROGRAMS

4.2.1 A Single Bottleneck With Inelastic Demand

Consider the following finite dimensional linear program (LP):

$$\begin{aligned} &\min_{l(i)} \sum_{i\in\mathcal{I}} \big[\phi(i) + \alpha^{-1}\psi(i)\big] l(i) \\ &\text{subject to:} \\ &\quad l(i) \ge 0 \\ &\quad l(i) \le \mu \\ &\quad \sum_{i\in\mathcal{I}} l(i) = N. \end{aligned} \tag{LP}$$

Theorem 1. *The optimal solution of the problem (LP) satisfies the departure time based equilibrium conditions* (1), (3), *and* (4) *for a single bottleneck model with an inelastic demand if the Lagrange multipliers associated with the inequality constraints $l(i) \le \mu$ are considered as the queuing delay $b(i)$.*

Remark 2. *For this interpretation of the Lagrange multiplier $u(i)$ to be valid, the FIFO conditions* (6) *or* (7) *should hold.*

The associated dual of (LP), which can be used to calculate the bottleneck delay $b(i)$ in the discrete time is given by

$$\begin{aligned} &\max_{b(i),\nu} -\mu \sum_{i\in\mathcal{I}} b(i) - N\nu \\ &\text{subject to:} \\ &\quad \phi(i) + \alpha^{-1}\psi(i) + b(i) \geq -\nu,\ \forall i \in \mathcal{I} \\ &\quad b(i) \geq 0,\ \forall i \in \mathcal{I}. \end{aligned} \tag{D-LP}$$

Note that the existence of the solution to the dual problem (D-LP) does not guarantee the FIFO condition (6). However, if the parameters of the dual program are known (i.e., $\phi(\cdot)$ and $\psi(\cdot)$), one can determine whether these parameters satisfy the condition (7) or if the candidate equilibrium solution $b^*(i)$ satisfies the FIFO condition (6). If the solution obtained by the (D-LP) does not satisfy the FIFO condition (6), then there is no equilibrium. Clearly, the departure time based formulation has a benefit of reducing to a linear program, whose solution can be easily obtained.

In summary, the above equivalent mathematical formulation of the user equilibrium allows us to consider a more general problem of the pricing design that achieves the social optimum under the budget and maximum price constraints, which is the topic of the next section.

4.3 SOCIAL OPTIMUM WITH PRICING CONSTRAINTS

In the case of an inelastic demand, the maximization of the social surplus is equivalent to minimizing the total social cost. The total social cost with respect to the departure time is the sum of queuing delay and schedule delay costs for all travelers during the commute period $\mathcal{T}$. Thus, the social optimum is determined by minimizing the total social cost subject to the user equilibrium and constraints on the incentive $\psi(\cdot)$. The constraints restrict the maximum toll to a predetermined ψ_{max} and do not allow the total incentives paid by the system planner to exceed a preset budget B_{max}. Moreover, we consider incentives that satisfy the FIFO condition, thus ensuring that only physically feasible arrival flows at the bottleneck are allowed.

4.3.1 Social Optimum For A Single Bottleneck

Now consider the following finite dimensional bilevel optimization problem:

$$\min_{\psi(i)} \sum_{i\in\mathcal{I}} l(i)\big[b(i)+\phi(i)\big] \tag{IP}$$

subject to:

$$\psi(i) \le \psi_{max},\ \forall i \in \mathcal{I}$$

$$-\sum_{i\in\mathcal{I}} \psi(i)l(i) \le B_{max}$$

$$\alpha^{-1}\psi(i) - \alpha^{-1}\psi(i+1) \le \Delta s - \phi(i) + \phi(i+1)$$

$$(b(i),\nu) \in \arg\max \Big\{ -\mu \sum_{i\in\mathcal{I}} b(i) + N\nu \Big\}$$

subject to:

$$b(i) + \phi(i) + \alpha^{-1}\psi(i) - \nu \ge 0 \tag{IPa}$$

$$b(i) \ge 0$$

$$l(i) \in \arg\min \sum_{i\in\mathcal{I}} \big[\phi(i) + \alpha^{-1}\psi(i)\big] l(i)$$

subject to:

$$l(i) \ge 0,\ \forall i \in \mathcal{I}$$

$$l(i) \le \mu,\ \forall i \in \mathcal{I} \tag{IPb}$$

$$\sum_{i\in\mathcal{I}} l(i) = N.$$

In (P1), the cost function expresses the total social cost and the inner problems (IPa) and (IPb) ensure that the feasible solution is also a candidate equilibrium solution. Clearly, by the rearrangement inequality, the optimal solution to the inner problem (IPb), given any feasible $\psi(i)$, simply assigns the maximum allowed exit rate $l^*(i) = \mu$ at time slots i which have $\frac{N}{\mu}$ as the smallest values of $\big[\phi(i) + \alpha^{-1}\psi(i)\big]$ and assigns $l^*(i) = 0$, otherwise. For example, in Figure 1, three time slots with smallest values $\big[\phi(i) + \alpha^{-1}\psi(i)\big]$ are i_2, i_3, and i_4. If only three time slots are required to satisfy the inelastic demand N, then at time slots i_2, i_3, and i_4 there will be a non-zero exit rate. Note, the assumption that $\frac{N}{\mu}$ is an integer is not restrictive since the residual demand will not cause queuing and can arrive before or after $\frac{N}{\mu}$ as the smallest $\big[\phi(i) + \alpha^{-1}\psi(i)\big]$ time slots.

In (I1), the optimal $\nu^* = \min_{i\in\mathcal{I}} \big\{b(i) + \phi(i) + \alpha^{-1}\psi(i)\big\}$ since $N \ge 0$. Substituting for ν^* into the cost function of the (I1), we can convert the (I1), for a feasible $\psi(i)$, into a concave optimization problem given by

$$\max_{b(i)\ge 0} -\mu \sum_{i\in\mathcal{I}} b(i) + N \min_{i\in\mathcal{I}} \big\{b(i) + \phi(i) + \alpha^{-1}\psi(i)\big\}. \tag{8}$$

Theorem 2. *The optimal solution to the concave optimization problem 8 is given by*

$$b^*(i) = \begin{cases} \min\limits_{k\in\mathcal{I}\setminus\mathcal{A}(\psi)} \big(\phi(k) + \alpha^{-1}\psi(k)\big) - \\ \qquad - \big(\phi(i) + \alpha^{-1}\psi(i)\big), i \in \mathcal{A}(\psi) \\ 0, i \in \mathcal{I} \setminus \mathcal{A}(\psi), \end{cases}$$

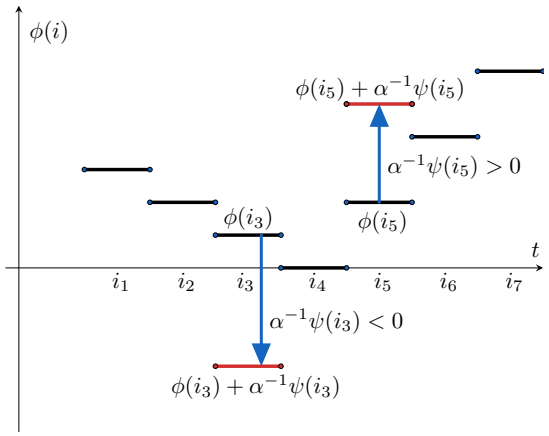

Figure 1: Schedule delay cost with toll applied at time i_5 and incentive applied at time i_3.

where $\mathcal{A}(\psi)$ is the set of time slots i with $\frac{N}{\mu}$ as the smallest $\phi(i)+\alpha^{-1}\psi(i)$ values.

Note that if there are two time slots i_1, i_2 with the same $\phi(i_1)+\alpha^{-1}\psi(i_1) = \phi(i_2)+\alpha^{-1}\psi(i_2)$, then ambiguity may be resolved by selecting the time slot i from i_1, i_2 that is closest to the desired arrival time $\frac{\tau}{\Delta s}$.

Taking into account that the exit rate $l(i) \neq 0$ on the $\frac{N}{\mu}$ number of time slots and the analytical solutions of the inner problems (IPa) and (IPb) enable us to further simplify (IP) as follows

$$
\begin{aligned}
&\min_{\psi(i)} \sum_{i\in\mathcal{A}(\psi)} \mu\big[b(i)+\phi(i)\big] \\
&\quad \text{subject to:} \\
&\qquad \psi(i) \le \psi_{max},\ \forall i \in \mathcal{A}(\psi) \\
&\qquad -\mu \sum_{i\in\mathcal{A}(\psi)} \psi(i) \le B_{max},\ \forall i \in \mathcal{A}(\psi) \\
&\alpha^{-1}\psi(i) - \alpha^{-1}\psi(i+1) \le \Delta s - \phi(i) + \phi(i+1) \\
&\qquad b(i) = \begin{cases} \min\limits_{j\in\mathcal{I}\setminus\mathcal{A}(\psi)} \big(\phi(k)+\alpha^{-1}\psi(k)\big) - \\ \qquad\qquad -\big(\phi(i)+\alpha^{-1}\psi(i)\big), i \in \mathcal{A}(\psi) \\ 0, i \in \mathcal{I}\setminus\mathcal{A}(\psi). \end{cases}
\end{aligned}
\tag{SIP}
$$

Assumption 1. *Assume the time based schedule delay cost $\phi(\cdot)$ is strictly convex, $\phi(\frac{\tau}{\Delta s}) = 0$, and $\phi(i) > 0$ for all $i \neq \frac{\tau}{\Delta s}$.*

Assumption 2. *Assume the time based schedule delay cost $\phi(\cdot)$ satisfies the following condition for $\forall i \in \mathcal{I}$*

$$\frac{\phi(i+1)-\phi(i)}{\Delta s} > -1.$$

Theorem 3. *If Assumptions 1 and 2 hold, then for the optimal $\psi^*(i)$ in the simplified Problem (SIP), $A(\psi^*) = S$, where S is the set of time slots with $\frac{N}{\mu}$ as the smallest $\phi(i)$.*

Remark 3. *The piecewise linear schedule delay given in* (2) *satisfies both Assumptions 1 and 2.*

Note, given the Assumption 1, the set $\mathcal{S}$ is uniquely defined. Moreover, the objective function in the problem (SIP) is also convex. Therefore, problem (SIP) under the assumptions of Theorem 3 is convex.

Remark 4. *Given the optimal pricing profile $\psi^*(i)$ in the departure times, we can recover the approximate time delay $d(\cdot)$ in terms of the arrival time at the bottleneck from $d(i - b(i)) = b(i)$. Similarly, the incentive profile at the arrival time $i - b(i)$ can be obtained from the departure time pricing $\psi(i)$.*

Since the bottleneck pricing has a long-term effect on the demand, the price elasticity of the travel demand should be considered. Next, we extend the approach developed above to the design of the optimal incentives under the elastic demand setting.

4.3.2 Equilibrium Formulation For A Single Bottleneck With Elastic Demand

Suppose the travelers demand for the use of the single bottleneck depends of the generalized cost C_{dep}, which is the equilibrium private cost faced by each traveler. Let $N = D(C_{dep})$ denote the demand for the equilibrium cost C_{dep}. We further assume that the demand function $D(C_{dep})$ is a strictly monotone, invertible, and decreasing function of the equilibrium cost C_{dep}. Let $D^{-1}(N)$ be the inverse or benefit function. The inverse function $D^{-1}(N)$ can be regarded as an amount that each traveler is willing to pay for his or her travel through the bottleneck, or a benefit that he or she can obtain from this travel. The equilibrium conditions in the departure time based formulation for the elastic demand are,
$\forall i \in \mathcal{I}$

$$\begin{cases} C_{dep} = b(i) + \phi(i) + \alpha^{-1}\psi(i), \text{ if } l(i) > 0 \\ C_{dep} \leq b(i) + \phi(i) + \alpha^{-1}\psi(i), \text{if } l(i) = 0, \end{cases} \tag{9}$$

$$\begin{cases} l(i) = \mu, \text{ if } b(i) > 0 \\ l(i) \leq \mu, \text{ if } b(i) = 0, \end{cases} \tag{10}$$

and

$$\sum_{i \in \mathcal{I}} l(i) = D(C_{dep}). \tag{11}$$

The FIFO condition in elastic demand setting is identical to (6).

4.3.3 A Single Bottleneck With Elastic Demand

Consider the following finite dimensional mathematical program,

$$\begin{aligned} &\min_{l(i),N} -\int_0^N D^{-1}(w)\,dw + \sum_{i \in \mathcal{I}} \big[\phi(i) + \alpha^{-1}\psi(i)\big] l(i) \\ &\text{subject to:} \\ &\quad l(i) \geq 0 \\ &\quad l(i) \leq \mu \\ &\quad \sum_{i \in \mathcal{I}} l(i) = N. \end{aligned} \tag{CP}$$

The mathematical problem (CP) is strictly convex and hence has a unique solution in terms of the exit rate $l(i)$ and the demand N since the first term in the objective function is strictly convex (the benefit function $D^{-1}(w)$ is strictly monotonically decreasing in demand), the second term is linear, and the constraints are convex.

Theorem 4. *The optimal solution of the convex problem (CP) satisfies the equilibrium conditions* (9), (10), *and* (11) *if the Lagrange multipliers associated with the inequality constraints* $l(i) \leq \mu$ *are considered as the queuing delay* $b(i)$, *and* ν *is considered as the equilibrium cost.*

Since the Slater's condition is satisfied for (CP), then the strong duality holds. The dual of (CP) is

$$\begin{aligned} &\max_{b(i),\nu} -\int_0^{D(-\nu)} D^{-1}(w)\,dw - \mu \sum_{i\in\mathcal{I}} b(i) - N\nu \\ &\text{subject to:} \qquad \text{(D-CP)} \\ &\quad \phi(i) + \alpha^{-1}\psi(i) + b(i) \geq -\nu,\ \forall i \in \mathcal{I} \\ &\quad b(i) \geq 0,\ \forall i \in \mathcal{I}. \end{aligned}$$

The dual problem (D-CP) can be used to calculate the bottleneck delay $b(i)$ in the elastic demand setting.

4.3.4 Social Optimum For A Single Bottleneck With Elastic Demand

We consider the following finite dimensional bilevel optimization problem for the single bottleneck with an elastic demand setting,

$$\min_{\psi(i)} \sum_{i\in\mathcal{I}} l(i)\big[b(i)+\phi(i)\big] \tag{EP}$$

subject to:

$$\psi(i) \le \psi_{max},\ \forall i \in \mathcal{I}$$

$$-\sum_{i\in\mathcal{I}} \psi(i) l(i) \le B_{max},$$

$$\alpha^{-1}\psi(i) - \alpha^{-1}\psi(i+1) \le \Delta s - \phi(i) + \phi(i+1)$$

$$(b(i),\nu) \in \arg\max \Bigg\{ -\int_0^{D(\nu)} D^{-1}(w)\,dw -$$

$$-\mu \sum_{i\in\mathcal{I}} b(i) + D(\nu)\nu \Bigg\}$$

subject to:

$$b(i) + \phi(i) + \alpha^{-1}\psi(i) - \nu \ge 0 \tag{EPa}$$

$$b(i) \ge 0$$

$$(l(i), N) \in \arg\min - \int_0^{N} D^{-1}(w)\,dw +$$

$$+\sum_{i\in\mathcal{I}} \big[\phi(i) + \alpha^{-1}\psi(i)\big] l(i)$$

subject to:

$$l(i) \ge 0,\ \forall i \in \mathcal{I}$$

$$l(i) \le \mu,\ \forall i \in \mathcal{I} \tag{EPb}$$

$$\sum_{i\in\mathcal{I}} l(i) = N.$$

Problems (EPba) and (EPb) (parameterized by $\psi(\cdot)$) are strictly convex because $D^{-1}(w)$ is strictly monotonically decreasing in demand w. Since the lower level problems (EPa) and (EPb) are convex and sufficiently regular (i.e., Slater's condition holds), it is possible to replace problems (EPa) and (EPb) with their Karush-Kuhn-Tucker (KKT) conditions. Using the KKT conditions, we can reformulate the bilevel optimization (EP) into a single-level mathematical program with equilibrium constraints (MPEC).

4.4 SIMULATION

In this section, two numerical examples are presented. The first example is designed to illustrate the findings in Section 4.1 for a single bottleneck, and the second example illustrates the results of Section 4.2 for a single bottleneck with elastic demand.

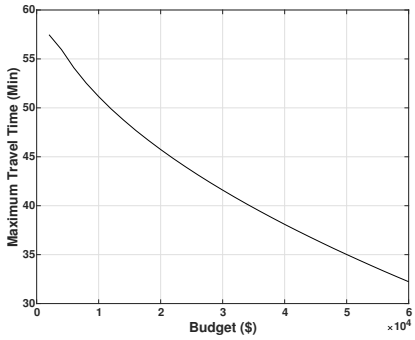

Figure 2: Maximum travel time $b(s)$ for the maximum price $\psi_{max} = \$0$.

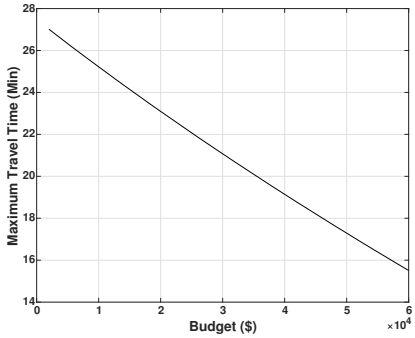

Figure 3: Maximum travel time $b(s)$ for the maximum price $\psi_{max} = \$1$.

4.4.1 Example 1

The basic input parameters used in the first numerical example are: $\tau = 300min, s = 600\frac{cars}{min}, \alpha = \$0.0952min, \beta = \$0.0582/min, \gamma = \$0.2263/min$ $N = 78600$ $cars,$ $\Delta s = 1$ min

We selected the demand that may be present at the Los Angeles Memorial Coliseum on a typical game day. Note that the desired time τ is referenced from beginning of the morning/evening commute start time at which there is no congestion present at the bottleneck. In this example, we consider the effect of the price cap and the budget constraint on the maximum travel time $b(s)$ through the bottleneck. Figure 2 shows the maximum travel time for the price cap $\psi_{max} = \$0$. Our simulations show that the maximum travel time monotonically decreases as the system planner increases the allocated budget. This behavior is consistent with the price cap $\psi_{max} = \$1$ shown in figure 3; however, for the same budget constraint the maximum travel time is more than 2 times smaller than in the case of price cap $\psi_{max} = \$0$. Figure 4 shows the maximum travel time dependence on the price caps for $B_{max} = \$0$ and Figure 5 shows the maximum travel time dependence on the maximum price for the fixed budget $B_{max} = \$5,000$.

4.4.2 Example 2

We assumed the elastic demand function to be $D(\nu) = \hat{D}exp(-\omega\nu)$, where $\hat{D}$ represents the potential level of demand over the entire time horizon, ν is the equilibrium generalized cost for each traveler, and ω is the cost sensitivity parameter. The input data are: $\hat{D} = 5000$ $\tau = 30$ $min.$ We solved the single-level MPEC by relaxing the complementary constraints with a positive relaxation

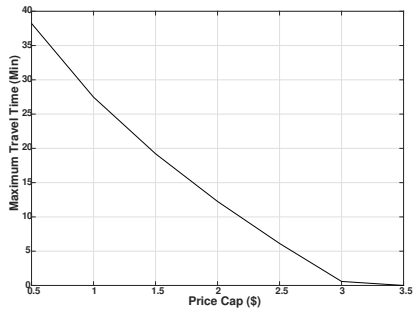

Figure 4: Maximum travel time $b(s)$ for the budget constraint $B_{max} = \$0$.

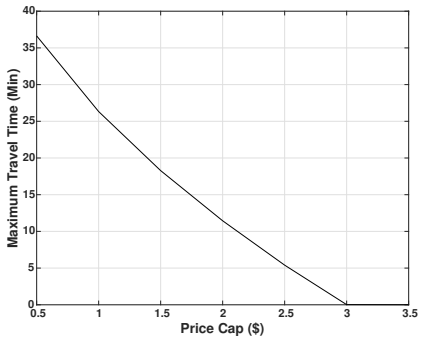

Figure 5: Maximum travel time $b(s)$ for the budget constraint $B_{max} = \$5K$.

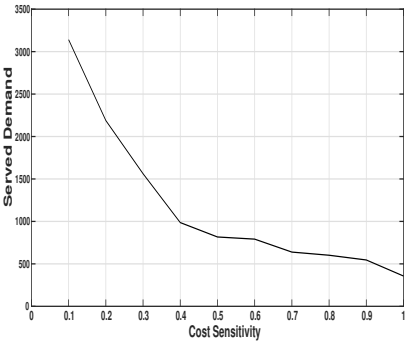

Figure 6: Served demand for $B_{max} = \$500$ and $\psi_{max} = \$5$.

parameter and solved the relaxed nonlinear program with the fmincon solver. We simulated fmincon solver with randomized initial conditions for 10 repeatedly runs. Figure 6 shows served demand for various cost sensitivity levels. As the cost sensitivity decreases, demand is less sensitive to the the equilibrium generalized cost. Note, in 6 as the cost sensitivity increase (i.e., demand becomes more elastic) less traveler are served, which is due to the fact that the cost function in the (EP) solely minimizes the system cost.

5 Task 2: Supply Control through Microscopic Ramp Metering

Freeway congestion is caused by high demand competing to use the limited supply of the freeway system. One of the most effective tools to combat this congestion is Ramp Metering (RM), where the inflows to a freeway are regulated in order to balance the supply-demand and ultimately improve some measure of performance ([13–15]). The problem of designing a RM policy is often studied by using macroscopic traffic flow models. In spite of their generality, these models do not have the resolution to distinguish between safety protocols under different connectivity and automation scenarios. On the other hand, the use of microscopic-level approach to study the interplay between safety, connectivity and freeway performance in the context of RM is limited to heuristics or simulations. Our objective is to address this gap.

There is an overwhelming body of literature on designing RM policies using macroscopic traffic flow models. We review them here only briefly; interested readers are referred to [13] for a comprehensive review. RM policies can be generally classified as fixed-time or traffic-responsive ([13]). Fixed-time policies such as [16] are fine-tuned offline and operate based on historical traffic data. Due to the uncertainty in the traffic demand and the absence of real-time measurements, these policies would either lead to congestion or under-utilization of the capacity of the freeway ([17]). Traffic-responsive policies, on the other hand, use real-time measurements. These policies can be further sub-classified as local or coordinated depending on whether the on-ramps make use of the measurements obtained from their vicinity (local) or other regions of the freeway (coordinated) ([13, 18]). A well-known example of a local policy is ALINEA ([19]) which has been shown, both analytically and in practice, to yield a (locally) good performance. A caveat in employing local policies is that there is no guarantee that they can improve the overall performance of the freeway while providing a fair access to the freeway from different on-ramps ([17]). This motivates the study of coordinated policies such as the ones considered in [17, 20, 21].

The aforementioned studies adopt macroscopic traffic flow models for design purposes. An alternative is the microscopic-level approach, which is an appropriate choice in the context of Connected and Automated Vehicles (CAVs) ([14]). CAVs come with promising capabilities: they can compensate for human errors by means of automation, and can provide more accurate traffic measurements by means of Vehicle-to-Vehicle (V2V), or Vehicle-to-Infrastructure (V2I) communications ([14, 22, 23]). On the other hand, they are required to obey certain constraints such as safety rules. It is natural to investigate how these capabilities and constraints can be incorporated in the design of traffic control techniques such as RM in order to optimize their performance. In the literature, the RM design problem at the microscopic level is often stated in terms of determining the *merging sequence* for an on-ramp. That is, determining the order in which vehicles on the mainline and at the on-ramp can cross the merging point. The objective is to choose a merging sequence that optimizes some measure of performance such as fuel consumption subject to vehicle constraints such as speed, safety, and comfort ([14, 24]). Most, if not all, previous work considers an isolated on-ramp for this purpose. Relatively little attention has been given to analyzing the implications of such merging-sequence design on the overall performance of the freeway. For example, do they optimize certain system-level performance? Indeed it is possible that a *greedy* policy, i.e., a policy where each on-ramp acts independent of the others, limits entry from the downstream on-ramps and thereby creating long queues or congestion. Motivated by increasing vehicle connectivity and automation, our primary objective is to systematically design and analyze the system-level performance of traffic-responsive RM policies subject to vehicle following safety constraints.

We consider a freeway with arbitrary number of on- and off-ramps. The freeway geometry can be modeled either as a *ring* road, in which case every entry point is controlled, or a *straight* road,

in which case the upstream entry point is not controlled. Previous studies on a ring road with no on/off-ramps have suggested that this setup has some theoretical advantages over the straight road network ([25–29]). For example, the creation and dissipation of stop-and-go waves can be captured using a ring road ([25, 26]). For the sake of completeness, we consider both geometries in this work. However, our results are not specific to the choice of road geometry. Vehicles in the network are assumed to have the same length, same acceleration and braking capabilities, and are equipped with V2V and V2I communication systems. Each vehicle follows the standard rules for safety and speed: it accelerates and maintains the free flow speed when it is sufficiently far away from the leading vehicle, or maintains a safe gap if it gets close. We do not specify the exact transient behavior for maintaining a safe gap, nor require vehicles to adopt the same behavior during the transient. However, at the steady state free flow speed, we assume that each vehicle keeps a safe constant time headway plus an additional constant gap from its leading vehicle ([30]). The entry points to the freeway are the on-ramps and, in the case of the straight road geometry, the upstream (uncontrolled) entry point. Vehicle arrivals to each on-ramp is modeled by a Bernoulli process that is independent across different on-ramps. The destination of each vehicle is one of the off-ramps, and, in the case of the straight road geometry, also the downstream exit point. The destination of each vehicle is sampled independently from a routing matrix. It should be emphasized that our main results do not depend on this specific demand model. Once a vehicle enters the freeway, it follows the aforementioned safety and speed rules until it reaches its destination, at which point it exits the network. We design traffic-responsive RM policies that maximize the freeway throughput subject to vehicle following safety constraints. For a given routing matrix, the throughput of a RM policy is characterized by the set of on-ramp arrival rates for which the queue sizes at all the on-ramps remain bounded in expectation. Roughly, the throughput of a RM policy is the highest traffic demand that it can handle without creating long queues at the on-ramps.

The proposed RM policies work in synchronous *cycles* during which an on-ramp does not release more vehicles than the number of vehicles waiting in its queue, i.e., its queue size, at the beginning of the cycle. Furthermore, all of the proposed policies operate under vehicle following safety constraints, where the on-ramps release new vehicles only if there is enough gap on the mainline at the moment of release. We provide three policies under which each on-ramp: (i) pauses release for a time-interval at the end of a cycle, or (ii) modulates the time between successive releases during a cycle, or (iii) adopts a conservative safe gap criterion for release during a cycle, all based on the traffic state. None of the policies however require information about the on-ramp arrival rates or the routing matrix, i.e., they are reactive. The throughput of these policies is characterized by studying stochastic stability of the induced Markov chains, and is proven to be maximized when the merging speed of all the on-ramps equals the free flow speed.

In summary, the three main contributions are as follows:

- Formalizing the notion of throughput as a performance index for freeways with arbitrary number of on- and off-ramps. The throughput of a RM policy measures the highest traffic demand that it can handle without creating long on-ramp queues. In case of a single on-ramp, the throughput of any policy cannot exceed the on-ramp's *flow capacity*.

- Design of traffic-responsive RM policies at the microscopic level subject to vehicle following safety constraints. This allows to understand the impact of different safety, connectivity, and automation protocols on the freeway throughput. The proposed policies require V2X communication but do not require vehicle autonomy. Hence, they are suitable to implement in mixed-autonomy scenarios.

- Design of policies that are reactive yet maximize the freeway throughput. In other words, the proposed policies only require real-time traffic measurements, and do not need any information about the arrival rates or the routing matrix (thus, they can adapt to time-varying demand); yet, their throughput is at least as good as any other policy, including the policies that know the demand in advance.

The following standard notations are used throughout this section. Let $\mathbb{N}$, $\mathbb{N}_0$, and $\mathbb{R}$ respectively denote the set of positive integers, non-negative integers and real numbers. For $m \in \mathbb{N}$, $[m]$ denotes the set $\{1, \ldots, m\}$. Proofs of the technical results can be found in [31].

5.1 Problem Formulation

Consider a simple model of a freeway, where the mainline is abstracted either as a straight road or a ring road of length P, as illustrated in Figure 7. We use the ring road geometry to formulate the problem and state the main results; the discussion of the straight road geometry can be found in Section 5.2.7. The freeway has m on- and off- ramps; they are placed alternately, and are numbered in an increasing order along the direction of travel, such that, for all $i \in [m]$, off-ramp i comes after on-ramp i [1]. The section of the mainline between the i-th on- and off- ramps is referred to as *link* i.

Vehicles arrive at the on-ramps from outside the freeway and join the on-ramp queues. We assume a point queue model for vehicles waiting at the on-ramps, with the queue on an on-ramp co-located with its ramp meter. The on-ramp vehicles are released into the mainline by the ramp meters installed at each on-ramp. Upon release, each vehicle follows the standard speed and safety rules until it reaches its destination off-ramp at which point it leaves the freeway without creating any obstruction for the upstream vehicles. Vehicles are equipped with V2V and V2I communication systems. They use these systems to communicate their state, e.g., speed, to nearby vehicles and the on-ramp control units. The exact information communicated by a vehicle will be specified in later sections.

The objective is to design RM policies that perform well under vehicle following safety constraints. The performance of a RM policy is evaluated in terms of its *throughput* defined as follows: let λ_i be the *arrival rate* to on-ramp i, $i \in [m]$, and $\lambda := [\lambda_i]$ be the vector of arrival rates. Let $R = [R_{ij}]$ be the *routing matrix*, where R_{ij} specifies the fraction of arrivals to on-ramp i that want to exit from off-ramp j. For a given routing matrix R, the *under-saturation region* of a RM policy is defined as the set of vector of arrival rates λ for which the queue sizes at all the on-ramps remain bounded in expectation [2]. The boundary of the under-saturation region is called the throughput. We are interested in finding RM policies that "maximize" the throughput for any given R. We will formalize this in Section 5.1.2.

The remainder of this section is organized as follows: in Section 5.1.1, we discuss the vehicle-level rules. We specify the demand model and formalize the notion of throughput in Section 5.1.2. We summarize the RM policies in Section 5.1.3.

[1] By this numbering scheme, we do not mean to imply that the location of an on-ramp must be close to the next off-ramp as in practice, the opposite is usually the case; see Figure 7. In fact, our setup is quite flexible and can also deal with cases where the number of on- and off-ramps are not the same. For simplicity of notation, we have used the same number of on- and off-ramps.

[2] Note that λ and R are macroscopic quantities. In order to specify vehicle arrivals and their destination at the microscopic level, a more detailed (probabilistic) demand model is required; see Section 5.1.2. The expected queue size is defined with respect to the probabilistic demand model.

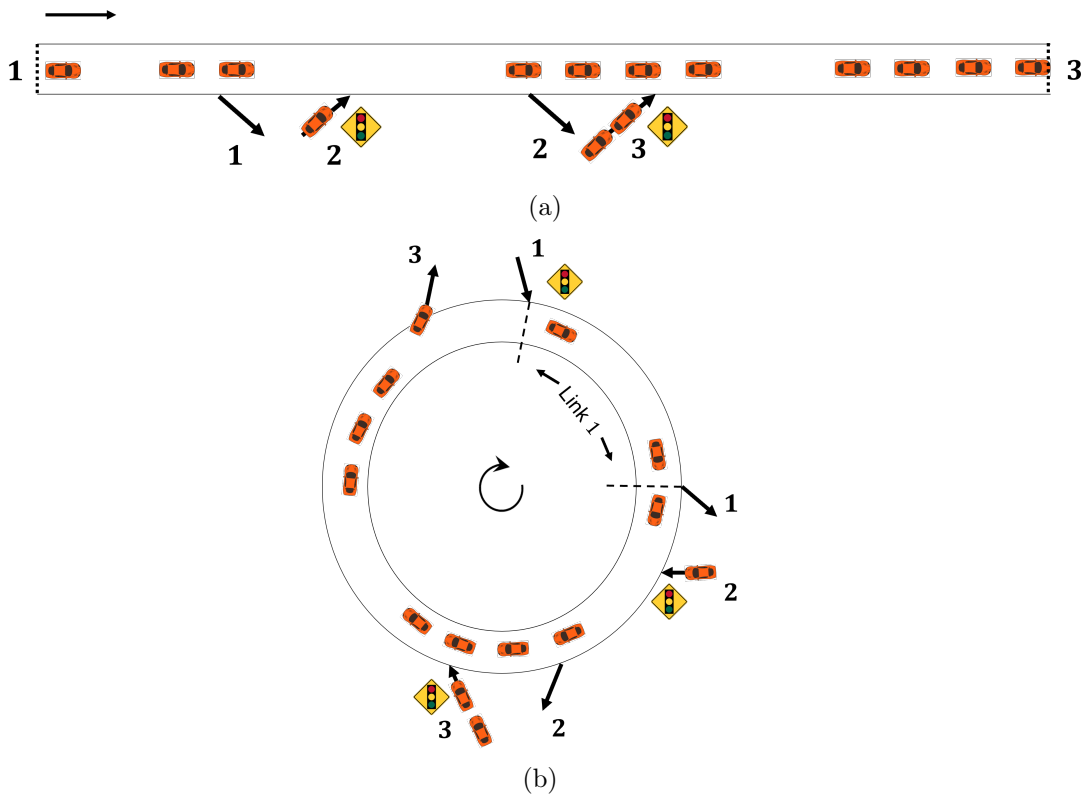

Figure 7: Example of a (a) straight road, (b) ring road.

5.1.1 Vehicle-Level Objectives

We consider vehicles of length L that have the same acceleration and braking capabilities, and are equipped with V2V and V2I communication systems. We use the term *ego vehicle* to refer to a specific vehicle under consideration, and denote it by e. Consider a vehicle following scenario and let v_e (resp. v_l) be the speed of the ego vehicle (resp. its leading vehicle), and S_e be the *safety distance* between the two vehicles required to avoid collision. We assume that S_e satisfies

$$S_e = hv_e + S_0 + \frac{v_e^2 - v_l^2}{2|a_{min}|}, \tag{12}$$

which is calculated based on an emergency stopping scenario with details given in [30]. Here, $h > 0$ is a safe time headway constant, $S_0 > 0$ is an additional constant gap, and $a_{min} < 0$ is the minimum possible deceleration of the leading vehicle. For simplicity, we assume a third-order vehicle dynamics. We consider two general modes of operation for each vehicle: the *speed tracking* mode and the *safety* mode. The main objective in the speed tracking mode is to adjust the speed to the *free flow speed* V_f when the ego vehicle is far from any leading vehicle. The main objective in the safety mode is to avoid collision when the ego vehicle gets close to a leading vehicle.

We define the *acceleration lane* of an on-ramp as the section of the network starting immediately downstream of the ramp meter and ending on the mainline, such that if the ego vehicle is in the speed tracking mode throughout the entire section, it achieves the speed V_f at the end of it. With this definition, the acceleration lane may or may not overlap with the mainline, depending on the distance from the ramp meter to the point where the merging occurs on the mainline, i.e., the

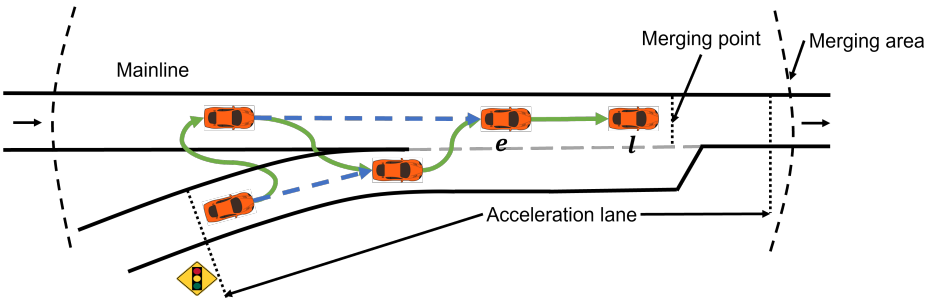

Figure 8: A merging scenario. For each vehicle, the dotted arrow indicates its leading vehicle while the solid arrow indicates its virtual leading vehicle. Note that vehicle l is both the leading and the virtual leading vehicle of vehicle e.

merging point; see Figure 8. We assume that all the vehicles coming from an on-ramp merge at a fixed merging point. We define the *merging speed* of an on-ramp as the speed of the ego vehicle at the merging point, if the ego vehicle is in the speed tracking mode between release and reaching the merging point. Note that the merging speed of an on-ramp is at most V_f.

Consider a merging scenario as shown in Figure 8. An ego vehicle entering the merging area is assigned a *virtual* leading vehicle, which is the vehicle that is predicted to be in front of the ego vehicle once the ego vehicle has crossed the merging point. The virtual leading vehicle is determined (and continuously updated) as follows: the ego vehicle predicts (and continuously updates) its speed and the time it will cross the merging point, and communicates this information to all other vehicles in the merging area using V2V communication. Other vehicles communicate the same information to the ego vehicle. We assume that the ego vehicle uses the following speed prediction rule: if in the speed tracking mode, the ego vehicle assumes that it remains in this mode in the future; if in the safety mode, it assumes a constant speed trajectory. At each time t before crossing the merging point, let $t_m|t$ be the time the ego vehicle predicts to cross the merging point based on the information available at time t. According to the prediction received from other vehicles at time t, the virtual leading vehicle is the vehicle that is to be in front of the ego vehicle at time $t_m|t$. We assume that vehicles obey the following rules:

(VC1) the ego vehicle maintains the constant speed V_f at time t if:

- (a) it is at a safe gap with respect to its leading vehicle, i.e., $y_e(t) \geq S_e(t)$, where y_e is the gap between the two vehicles. Note that if the leading vehicle is also at the constant speed V_f, then $y_e \geq S_e = hV_f + S_0$. This gap is equivalent to a time headway of at least $\tau := h + (S_0 + L)/V_f$ between the front bumpers of the two vehicles. This rule is believed to be widely adopted by human drivers as well as standard adaptive cruise control systems ([30]).
- (b) it predicts to be at a safe gap with respect to its virtual leading vehicle in merging areas, i.e., $\hat{y}_e(t_m|t) \geq \hat{S}_e(t_m|t) := h\hat{v}_e(t_m|t) + S_0 + \frac{\hat{v}_e^2(t_m|t) - \hat{v}_{\hat{l}}^2(t_m|t)}{2|a_{min}|}$, where $\hat{y}_e(t_m|t)$ is the predicted gap between the two vehicles at the moment of merging based on the information available at time $t \leq t_m$. Similarly, $\hat{v}_e$ and $\hat{v}_{\hat{l}}$ are the predicted speeds of the ego vehicle and its virtual leading vehicle, respectively.

(VC2) the ego vehicle is initialized to be in the speed tracking mode upon release. It changes mode at time t if $y_e(t) < S_e(t)$ or $\hat{y}_e(t_m|t) < \hat{S}_e(t_m|t)$. Also, if two vehicles are released from the

same on-ramp at least τ seconds apart, and neither change mode because of other vehicles, then the following vehicle remains in the speed tracking mode.

(VC3) there exists T_{empty} such that for any initial condition, vehicles reach the *free flow state* after at most T_{empty} time if no other vehicle is released from the on-ramps. The free flow state refers to a state where: (i) if a vehicle is in the safety mode, then it moves at the constant speed V_f and maintains this speed in the future, and (ii) if a vehicle is in the speed tracking mode, then it remains in this mode in the future.

Remark 1. *Note that we intentionally do not specify the total number of submodes within the safety mode, the exact control logic within each submode, or the exact logic for switching back to the speed tracking mode. Such details will be introduced only if and when needed for performance analysis of RM policies in the paper. Also, note that the control logic in the safety mode during the transient is allowed to be different for different vehicles.*

5.1.2 Probabilistic Demand Model and Throughput

It will be convenient for performance analysis later on to adopt a discrete time setting. Let the duration of each time step be τ, representing the minimum safe time headway between the front bumpers of two consecutive vehicles that are moving at the constant speed V_f; see (VC1) in Section 5.1.1. Let vehicles arrive to on-ramp $i \in [m]$ according to an i.i.d. Bernoulli process with parameter $\lambda_i \in [0, 1]$ independent of the other on-ramps. That is, in any given time step, the probability that a vehicle arrives at the i-th on-ramp is λ_i independent of everything else. Note that λ_i specifies the arrival rate to on-ramp i in terms of the number of vehicles per τ seconds. Let $\lambda := [\lambda_i]$ be the vector of arrival rates. The destination off-ramp for individual arriving vehicles is i.i.d. and is given by the routing matrix $R = [R_{ij}]$, where $0 \leq R_{ij} \leq 1$ is the probability that an arrival to on-ramp i wants to exit from off-ramp j. Note that R_{ij} specifies the (long-run) fraction of arrivals at on-ramp i that want to exit from off-ramp j. Naturally, for every on-ramp i we have $\sum_j R_{ij} = 1$. Finally, we let $\tilde{R} = [\tilde{R}_{ij}]$ be the *cumulative* routing matrix, where $\tilde{R}_{ij}$ is the fraction of arrivals at on-ramp i that need to cross link j in order to reach their destination. Note that $\lambda_i \tilde{R}_{ij}$ is the rate of arrivals at on-ramp i that need to use link j in order to reach their destination, i.e., the *load* induced on link j by on-ramp i. Let $\rho_j := \sum_i \lambda_i \tilde{R}_{ij}$ be the total load induced on link j by all the on-ramps, and let $\rho := \max_{j \in [m]} \rho_j$ be the *maximum load*.

Remark 2. *The current demand model, i.e., Bernoulli arrivals and Bernoulli routing, is chosen to simplify the technical details in the proofs. We believe that our results are far more general and hold for more practical demand models used in the literature, e.g., see [32] for an example of arrival models.*

Example 1. *Let the routing matrix for a 3-ramp freeway, i.e., $m = 3$ as in Figure 7b, be given by*

$$R = \begin{pmatrix} R_{11} & R_{12} & R_{13} \\ R_{21} & R_{22} & R_{23} \\ R_{31} & R_{32} & R_{33} \end{pmatrix}.$$

Then, the corresponding cumulative routing matrix is

$$\tilde{R} = \begin{pmatrix} 1 & 1 - R_{11} & 1 - (R_{11} + R_{12}) \\ 1 - (R_{22} + R_{23}) & 1 & 1 - R_{22} \\ 1 - R_{33} & 1 - (R_{31} + R_{33}) & 1 \end{pmatrix} = \begin{pmatrix} 1 & R_{12} + R_{13} & R_{13} \\ R_{21} & 1 & R_{21} + R_{23} \\ R_{31} + R_{32} & R_{32} & 1 \end{pmatrix}.$$

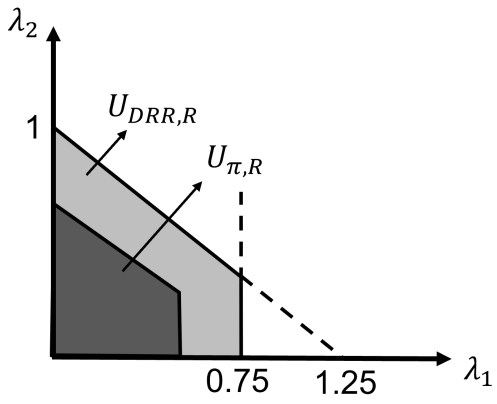

Figure 9: An illustration of the under-saturation region of some policy π (dark grey area) and the DRR policy (dark + light grey areas) from Example 2

We now formalize the notion of "throughput" which is the key performance metric in this work. For $i \in [m]$, let $Q_i(t)$ be the vector of destination off-ramps of the vehicles waiting at on-ramp i, arranged in the order of their arrival, at t. We use $|Q_i(t)|$ to denote the queue size at on-ramp i at time t. Let $|Q(t)| = [|Q_i(t)|]$ be the vector of queue sizes at all the on-ramps at time t. For a given routing matrix R, the under-saturation region of a RM policy π is defined as follows:

$$U_{\pi,R} = \{\lambda : \limsup_{t\to\infty} \mathbb{E}|Q_i(t)| < \infty \ \ \forall i \in [m] \text{ under policy } \pi\}.$$

This is the set of λ's for which the queue sizes at all the on-ramps remain bounded in expectation. The boundary of this set is called the throughput of the policy π. We are interested in finding a RM policy π such that for every R, $U_{\pi',R} \subseteq U_{\pi,R}$ for all policies π', including those that have information about λ and R. In other words, for every R, if the freeway remains under-saturated using some policy π', then it also remains under-saturated using the policy π. In that case, we say that policy π maximizes the freeway throughput. One of our main results is to introduce policies that maximize the saturation limit for all practical purposes but do not require any information about λ or R, i.e., they are reactive.

Remark 3. *A rigorous definition of the throughput should also include its dependence on the initial condition of the vehicles on the freeway and the initial queue sizes. We have removed this dependence for simplicity as the performance results given for our proposed policies do not depend on the initial condition.*

Example 2. *Consider a* 3*-ramp freeway with*

$$R = \begin{pmatrix} 0.2 & 0.7 & 0.1 \\ 0 & 0.8 & 0.2 \\ 0.5 & 0 & 0.5 \end{pmatrix}, \ \lambda_3 = 0.5 \ [veh/time\ step],$$

and suppose that the merging speed at all the on-ramps is V_f. Let us consider one of the policies introduced in Section 5.2, called the Dynamic Release Rate (DRR) policy. According to Theorem 2, the under-saturation region of this policy is given by

$$U_{DRR,R} = \{(\lambda_1, \lambda_2) : \lambda_1 < 0.75, \ 0.8\lambda_1 + \lambda_2 < 1\}.$$

This set is illustrated in Figure 9. Moreover, according to Theorem 4, for any policy π we have $U_{\pi,R} \subseteq U_{DRR,R}$, except maybe at the boundary of $U_{DRR,R}$; see Figure 9. Also, note that the boundary of $U_{DRR,R}$ has zero volume, which implies that the vector (λ_1, λ_2) lies either inside or outside $U_{DRR,R}$ in practice. More generally, Theorem 2 and Theorem 4 state that the previous conclusions hold for any other choice of the routing matrix. Therefore, the DRR policy maximizes the throughput for all practical purposes.

5.1.3 Ramp Metering

To conveniently track vehicle locations in discrete time, we introduce the notion of *slot*. A slot is associated with a particular point on the mainline or acceleration lanes at a particular time. We first define the mainline slots. Let n_c be the maximum number of distinct points that can be placed on the mainline, such that the space gap between adjacent points is $hV_f + S_0 + L$. This gap is governed by the safety distance S_e as explained in Section 5.1.1. In particular, n_c is the maximum number of vehicles that can safely travel at the constant free flow speed V_f on the mainline. Consider a configuration of these n_c points at $t = 0$. Each point represents a slot on the mainline that moves at the free flow speed; without loss of generality, we let the length of the mainline P be such that each slot replaces the next slot at the end of each time step τ.

We next define the acceleration lane slots. Suppose that the ego vehicle is released from the i-th on-ramp at $t = 0$ such that it remains in the speed tracking mode in the future. Consider the ego vehicle's location at the end of each time step τ until it exits the acceleration lane. Each of these location points represents a slot for the i-th acceleration lane at $t = 0$. For example, if the ego vehicle exits the acceleration after 2.5τ seconds upon release, there are three slots corresponding to its location at times τ, 2τ, and 3τ. Let n_i be the number of i-th acceleration lane slots, and $n_a = \sum_i n_i$. Note that by definition, the last acceleration lane slot of every on-ramp is on the mainline. Without loss of generality, we consider a configuration of slots at $t = 0$, such that all the last acceleration lane slots coincide with a mainline slot. The details to justify the no loss in generality are given as follows: for a given configuration of mainline slots at $t = 0$, there exists $t_i \in [0, \tau)$ such that the last acceleration lane slot of on-ramp i coincides with a mainline slot at time t_i. Thereafter, the last acceleration lane slot coincides with a mainline slot at the end of every time step τ, i.e., at times $k\tau + t_i$ for all $k \in \mathbb{N}_0$. The times $k\tau + t_i$, $k \in \mathbb{N}_0$, are the release times of on-ramp i in the proposed RM policies. Therefore, the assumption that all the last acceleration lane slots initially coincide with a mainline slot (which corresponds to $t_i = 0$ for all $i \in [m]$) only means a shifted sequence of release times, which justifies the no loss in generality.

Consider an initial condition of the vehicles on the freeway, where the vehicles are in the free flow state such that the location of each vehicle coincides with a slot for all times in the future. For this initial condition and under the proposed RM policies, the following sequence of events occurs during each time step: (i) the mainline slots rotate one position in the clockwise direction and replace the next slot; similarly, the acceleration lane slots of each on-ramp replace the next slot with the last slot replacing the first slot; the numbering of the slots is reset with the new first mainline slot after on-ramp 1 numbered 1, and the rest of the mainline slots numbered in an increasing order; the acceleration lane slots are numbered similarly; (ii) vehicles that reach their destination off-ramp exit the freeway without obstructing the upstream vehicles; (iii) if permitted by the RM policy, a new vehicle is released; for the given initial condition and under the proposed RM policies, the location of the newly released vehicle coincides with a slot for all times in the future.

The information available to the proposed policies will be a combination of $|Q|$ and the state

(or part of the state) of all the vehicles $X := (x_e)_{e\in[n]}$, where n is the number of vehicles on the mainline and acceleration lanes, and $x_e = (p_e, v_e, a_e, I_e)$ is the state of the ego vehicle, where p_e is its location, v_e is the speed, a_e is the acceleration, and I_e is a binary variable which is equal to one if the ego vehicle is in the safety mode, and zero otherwise. These information are communicated to the on-ramps using V2I communication systems. Table 1 provides a summary of the communication cost of the RM policies. In Table 1, n_m is the total number of slots in all the merging areas, T_{per} is a design update period, and $\overline{Q}_i$, $i = 1, 2$, is the contribution of the queue size to the communication cost.

Table 1: Summary of the RM policies studied in this project.

Ramp Metering Policy	Worst-case Communication Cost
Renewal	$(n_c + n_a)m + \overline{Q}_1$
Dynamic Release Rate (DRR)	$m(n_c + n_a)/T_{\text{per}} + n_m + \overline{Q}_2$
Distributed Dynamic Release Rate (DisDRR)	$(n_c + n_a)/T_{\text{per}} + n_m + \overline{Q}_2$
Dynamic Space Gap (DSG)	$(n_c + n_a)m + \overline{Q}_2$
Greedy	n_m

5.2 Ramp Metering Policies and Performance Analysis

In this section, we provide traffic-responsive RM policies subject to vehicle following safety constraints and analyze their performance. This is done in Sections 5.2.1-5.2.6. In Section 5.2.7, we discuss the extension of the results to the straight road geometry.

For easier navigation, we briefly review the proposed policies here. The policies in Sections 5.2.1, 5.2.2, and 5.2.4 are coordinated, the one in Section 5.2.3 is a distributed version of the coordinated policy in Section 5.2.2, and the one in Section 5.2.5 is a local policy. In all of our policies, the on-ramps work in synchronous *cycles* during which an on-ramp does not release more vehicles than its queue size at the beginning of the cycle. The synchronization of cycles is done in real time in Section 5.2.1, whereas it can be done once offline in Sections 5.2.2-5.2.5. Furthermore, all of our policies operate under vehicle following safety constraints, where the on-ramps release new vehicles only if there is enough gap on the mainline at the moment of release (cf. (VC1)). The policies differ in using the vehicle information to either pause release for a time interval at the end of a cycle (Section 5.2.1), to modulate the time between successive releases during a cycle (Sections 5.2.2-5.2.3), or to adopt a conservative dynamic safe gap criterion for release during a cycle (Section 5.2.4). The actions of all the three policies in Sections 5.2.2-5.2.4 at the free flow state is equivalent to that of the local policy in Section 5.2.5. All the policies are reactive, meaning that they do not require any information about the vector of arrival rates or the routing matrix. An inner estimate to the under-saturation region is provided for each policy, and is then compared to an outer estimate in Section 5.2.6.

5.2.1 Renewal Policy

The first policy is inspired by the queuing theory literature in the context of communication networks, e.g., see [33, 34]. Once an on-ramp releases all the vehicles waiting at the beginning of a cycle, it pauses release until all other on-ramps have done so, and these vehicles exit the freeway, i.e., until the mainline and acceleration lanes are empty – hence we refer to it as the *Renewal* policy.

Definition 1. ***(Renewal ramp metering policy)*** *No vehicle is released until all the initial vehicles exit the freeway, i.e., until the mainline and acceleration lanes are empty, say at time t_1. Thereafter, the policy works in cycles of* variable *length. At the beginning of the k-th cycle at time t_k, each on-ramp allocates itself a "quota" equal to the queue size at that on-ramp at t_k. At time t during the cycle, an on-ramp releases the ego vehicle if:*

(M1) $t = k\tau$ *for some* $k \in \mathbb{N}_0$.

(M2) the on-ramp has not reached its quota.

(M3) $y_e(t) \geq S_e(t)$*, i.e., there is a safe gap in front of the ego vehicle (cf. (VC2)).*

(M4) it predicts that the ego vehicle will be at a safe gap with respect to its virtual leading and following vehicles between merging and exiting the acceleration lane.

Once an on-ramp reaches its quota, it does not release a vehicle during the rest of the cycle. The next cycle begins when the mainline and acceleration lanes are empty.

Remark 4. *A simpler form of this policy, called the* Quota *policy, is analyzed in [33]. Direct adaptation of the Quota policy to the current transportation setting requires additional analysis because of the vehicle dynamics.*

We need an additional notation for future results. Consider a situation where on-ramp i releases the ego vehicle under (M1)-(M4), its virtual leading and following vehicles are at the constant speed V_f, and their location coincides with a mainline slot. We let τ_i be the minimum time headway between the front bumpers of the virtual leading and following vehicles such that they maintain a safe gap with respect to the ego vehicle after merging. Note that τ_i is an exact multiple of τ and $\tau_i \geq 2\tau$, where the equality holds if and only if the merging speed of on-ramp i is V_f.

Theorem 1. *For any initial condition, the Renewal policy keeps the freeway under-saturated if* $(\frac{\tau_i}{\tau} - 1)\rho_i - (\frac{\tau_i}{\tau} - 2)\lambda_i < 1$ *for all* $i \in [m]$.

V2I communication requirements: the Renewal policy uses information about $|Q|$ and X. Its worst-case communication cost is calculated as follows: at each time step during a cycle, any vehicle that is on the mainline or an acceleration lane must communicate its state to all on-ramps. After a finite time, the number of these vehicles is no more than $n_c + n_a$. Furthermore, at the beginning of every cycle, all the vehicles in an on-ramp queue must communicate their presence in the queue to that on-ramp. The contribution of the queue size to the communication cost is

$$\overline{Q}_1 := \limsup_{K\to\infty} \frac{1}{K} \sum_{t_k \leq K\tau} \sum_{i\in[m]} |Q_i(t_k)|,$$

where t_k is the beginning of the k-th cycle in the Renewal policy. Hence, the communication cost C is upper-bounded by $(n_c + n_a)m + \overline{Q}$.

Remark 5. *Under the constant time headway safety rule in Section 5.1.1, the flow capacity of the mainline is* 1 *vehicle per τ seconds. Theorem 1 provides an inner estimate of the under-saturation region in terms of the induced loads ρ_i, arrival rates λ_i, and the mainline capacity. In particular, if for every $i \in [m]$, $(\tau_i/\tau - 1)\rho_i - (\tau_i/\tau - 2)\lambda_i$ is less than the capacity, then the Renewal policy keeps the freeway under-saturated.*

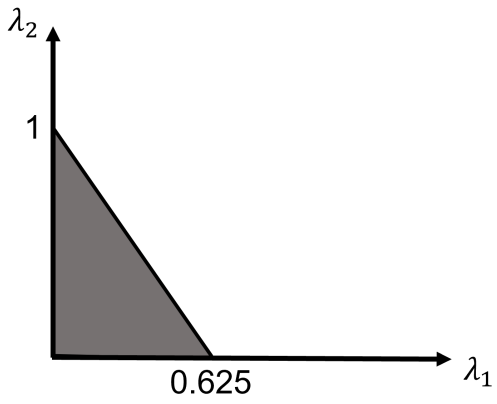

Figure 10: An inner estimate of the under-saturation region of the Renewal policy (grey area) from Example 3.

Remark 6. *The interplay between safety during merging and throughput is captured using the parameter τ_i, $i \in [m]$. In particular, as the merging speed of an on-ramp decreases, the required safety distance S_e in* (12) *increases. This puts a limit on the rate at which the on-ramp can release new vehicles under the vehicle following safety constraint, which in turn decreases the throughput.*

Remark 7. *Under the Renewal policy, vehicles that arrive during a cycle cannot enter the freeway until the next cycle begins. This increases the chance that these vehicles enter the freeway in platoons, rather than individually, once the next cycle begins. The platoon formation increases the release rate of other on-ramps as compared to the individual formation, when the merging speeds are less than V_f. As a consequence, the inner estimate of the under-saturation region of the Renewal policy is larger than the ones of other policies described in the following sections.*

Example 3. *Consider a 3-ramp freeway with $\tau_1 = \tau_3 = 2\tau$, $\tau_2 = 3\tau$, i.e., the merging speed of on-ramps 1 and 3 are V_f and is lower for on-ramp 2. Suppose that R and λ_3 are the same as Example 2. Then, the inner estimate of the under-saturation region given by Theorem 1 is illustrated in Figure 10.*

5.2.2 Dynamic Release Rate Policy

This policy imposes dynamic minimum time gap criterion, in addition to (M1), between release of successive vehicles from the same on-ramp. Changing the time gap between release of successive vehicles by an on-ramp is akin to changing its release rate, and hence the name of the policy.

Definition 2. ***(Dynamic Release Rate (DRR) ramp metering policy)*** *The policy works in cycles of* fixed *length $T\tau$, where $T \in \mathbb{N}$. At the beginning of the k-th cycle at $t_k = (k-1)T\tau$, each on-ramp allocates itself a "quota" equal to the queue size at that on-ramp at t_k. At time $t \in [t_k, t_{k+1}]$ during the k-th cycle, on-ramp i releases the ego vehicle if (M1)-(M4), and the following condition is satisfied:*

(M5) at least $g(t)$ time has passed since the release of the last vehicle from on-ramp i, where $g(t)$ is a piecewise constant minimum time gap, updated periodically at $t = T_{per}, 2T_{per}, \ldots$ as described in Algorithm 1.

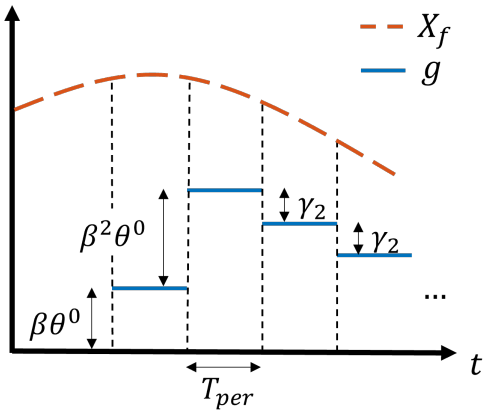

Figure 11: An illustration of the update rule for the minimum time gap g in Algorithm 1.

Once an on-ramp reaches its quota, it does not release a vehicle during the rest of the cycle.

In Algorithm 1, $X_f(t) := X_{f_1}(t)+X_{f_2}(t)$, where X_{f_1} and X_{f_2} are defined as follows: let $X_{f_1}(t) := \sum_{e\in[n]}(|v_e(t)-V_f|+|a_e(t)|)I_e(t)$, where n is the number of vehicles on the mainline and acceleration lanes. Furthermore, for $t \geq T_{per}$, let $X_{f_2}(t) := \sum(\delta_e(t)+\hat{\delta}_e(t))$, where the sum is over all the vehicles that either: (i) violated the safety distance S_e at some time in $[t-T_{per}, t]$, or (ii) predicted at some time in $[t-T_{per}, t]$ that they would violate the safety distance once they reach a merging point. The terms δ_e and $\hat{\delta}_e$ are, respectively, the maximum error and predicted error in the relative spacing, and are given by

$$\delta_e(t) = \max_{t'\in[t-T_{per},t]} |y_e(t') - S_e(t')|\mathbb{1}_{\{y_e(t')<S_e(t')\}},$$
$$\hat{\delta}_e(t) = \max_{t'\in[t-T_{per},t]} |\hat{y}_e(t_m|t') - \hat{S}_e(t_m|t')|\mathbb{1}_{\{\hat{y}_e(t_m|t')<\hat{S}_e(t_m|t')\}},$$

where $\mathbb{1}$ is the indicator function, and $t_m|t'$ is the time the ego vehicle predicts to cross a merging point based on the information available at time $t' \leq t_m$. If $\delta_e + \hat{\delta}_e$ is non-zero, then the ego vehicle communicates it either right before leaving the freeway, or at the update times $T_{per}, 2T_{per}, \ldots$, whichever comes earlier. Otherwise, it is not communicated by the ego vehicle.

Algorithm 1 Update rule for the minimum time gap between release of vehicles under the DRR policy

Require: **design constants:** $T_{\text{per}} > 0, \alpha_> 0, \gamma_2 > 0, \theta^\circ > 0, \beta > 1$
initial condition: $g(0) = 0,\ \theta(0) = \theta^\circ,\ X_f(0) = X_{f_1}(0)$
for $t = T_{\text{per}}, 2T_{\text{per}}, \cdots$ **do**
  if $X_f(t) \leq \max\{X_f(t - T_{\text{per}}) - \alpha, 0\}$ **then**
    $\theta(t) \leftarrow \theta(t - T_{\text{per}})$
    $g(t) \leftarrow \max\{g(t - T_{\text{per}}) - \gamma_2, 0\}$
  else
    $\theta(t) \leftarrow \beta\theta(t - T_{\text{per}})$
    $g(t) \leftarrow g(t - T_{\text{per}}) + \theta(t)$
  end if
end for

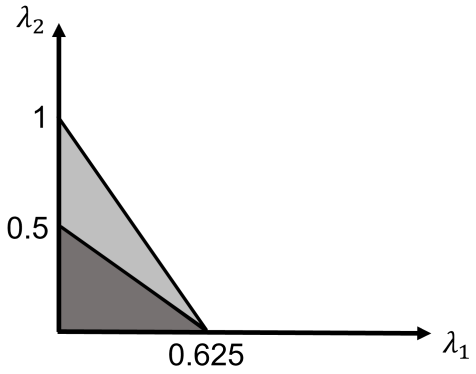

Figure 12: An inner estimate of the under-saturation region (dark grey area) of the DRR policy from Example 4. The light grey area indicates the additional under-saturation region if we use the Renewal policy.

Remark 8. *Recall the safety distance S_e in* (12)*. Since the speed of the vehicles are bounded, S_e is also bounded. Hence, δ_e (and similarly $\hat{\delta}_e$) in X_{f_2} is bounded. We use this in the following theorem.*

Theorem 2. *For any initial condition, $T \in \mathbb{N}$, and design constants in Algorithm 1, the DRR policy keeps the freeway under-saturated if $(\frac{\tau_i}{\tau} - 1)\rho_i < 1$ for all $i \in [m]$.*

V2I communication requirements: This policy uses information about $|Q|$ (for $T > 1$) and X. Its worst-case communication cost is calculated as follows: after a finite time, X is communicated to all on-ramps only at the end of each update period T_{per}. After such finite time, the number of vehicles that constitute X is no more than $n_c + n_a$. Furthermore, at each time step during a cycle, the vehicles in the merging area of an on-ramp communicate their state to that on-ramp for safe gap evaluation (cf. (M3)-(M4)). The number of vehicles in all the merging areas is at most n_m. Finally, if $T > 1$, then all the vehicles in an on-ramp queue must communicate their presence in the queue to that on-ramp at the beginning of every cycle. The contribution of the queue size to the communication cost is

$$\overline{Q}_2 := \left(\limsup_{K\to\infty} \frac{1}{K} \sum_{k \le K/T+1} \sum_{i\in[m]} |Q_i((k-1)T\tau)| \right) \mathbb{1}_{\{T>1\}}.$$

Hence, C is upper bounded by $m(n_c + n_a)/T_{\text{per}} + n_m + \overline{Q}_2$.

Remark 9. *Similar to Theorem 1, Theorem 2 provides an inner estimate of the under-saturation region of the DRR policy in terms of the induced loads ρ_i and the mainline capacity. The region specified by this estimate is the same for all $T \in \mathbb{N}$ and is contained in the one given for the Renewal policy in Theorem 1. However, this does not necessarily mean that the throughput of the DRR policy is the same for different cycle lengths, or the Renewal policy gives a better throughput as the inner estimates may not be exact. A simulation comparison of the throughput of the DRR policy for different cycle lengths is given in Section 5.3.2.*

Example 4. *Let the freeway parameters be as in Example 3. Then, the inner estimate of the under-saturation region given by Theorem 2 is illustrated in Figure 12, and is compared with the one in Figure 10. Also, comparing with Example 2 where the merging speed of on-ramp* 2 *was also V_f, one can see that the region specified in this example is smaller.*

5.2.3 Distributed Dynamic Release Rate Policy

This policy imposes dynamic minimum time gap criterion, just like its coordinated counterpart. However, each on-ramp only receives information from the vehicles in its vicinity and the downstream on-ramps.

Definition 3. ***(Distributed Dynamic Release Rate (DisDRR) ramp metering policy)*** *The policy works in cycles of fixed length $T\tau$, where $T \in \mathbb{N}$. At the beginning of the k-th cycle at $t_k = (k-1)T\tau$, each on-ramp allocates itself a "quota" equal to the queue size at that on-ramp at t_k. At time $t \in [t_k, t_{k+1}]$ during the k-th cycle, on-ramp i releases the ego vehicle only if (M1)-(M4), and the following condition is satisfied:*

(M5) at least $g_i(t)$ time has passed since the release of the last vehicle from on-ramp i, where $g_i(t)$ is a piecewise constant minimum time gap, updated periodically at $t = T_{per}, 2T_{per}, \ldots$ according to Algorithm 2.

Once an on-ramp reaches its quota, it does not release a vehicle during the rest of the cycle.

For $i \in [m]$, let X_f^i be the part of X_f associated with all the vehicles located between the i-th and $(i+1)$-th on-ramps. Thus, $X_f = \sum_{i\in[m]} X_f^i$. We assume that X_f^i is available to on-ramp i. Furthermore, if $g_{i+1}(t) > T_{max}$ for some design constant T_{max}, then all the on-ramps j downstream of on-ramp i communicate X_f^j to on-ramp i. In that case, $\sum_{j>i} X_f^j$ is available to on-ramp i, where the notation "$j > i$" means that on-ramp j is downstream of on-ramp i. Note that for the ring road geometry, all the on-ramps downstream of on-ramp i is equivalent to all the on-ramps. Hence, $\sum_{j>i} X_f^j = X_f$ for the ring road geometry. Informally, when $g_{i+1} \leq T_{max}$, g_i depends only on the traffic condition in the vicinity of on-ramp i, whereas when $g_{i+1} > T_{max}$, it also depends on the downstream traffic condition. Naturally, higher values of T_{max} make the policy more "decentralized", but it may take longer to reach the free flow state.

Algorithm 2 Update rule for the minimum time gap between release of vehicles under the DisDRR policy

Require: design constants: $T_{\text{per}} > 0$, $T_{\max} > 0$, $\alpha_{>}0$, $\gamma_2 > 0$, $\theta_i^\circ > 0$, $\beta > 1$
initial condition: $(g_i(0), g_i(T_{\text{per}})) = (0,0)$, $(\theta_i(0), \theta_i(T_{\text{per}})) = (\theta_i^\circ, \theta_i^\circ)$, $X_f^i(0) = X_{f_1}^i(0)$, $i \in [m]$
for $t = 2T_{\text{per}}, 3T_{\text{per}}, \cdots$ **do** the following for each on-ramp $i \in [m]$
  if $X_f^i(t) \le \max\{X_f^i(t - T_{\text{per}}) - \alpha_{,}0\}$ **then**
    if $g_{i+1}(t - T_{\text{per}}) \le T_{\max}$ **or** $\sum_{j>i} X_f^j(t) \le \max\{\sum_{j>i} X_f^j(t - T_{\text{per}}) - \alpha_{,}0\}$ **then**
      $\theta_i(t) \leftarrow \theta_i(t - T_{\text{per}})$
      $g_i(t) \leftarrow \max\{g_i(t - T_{\text{per}}) - \gamma_2, 0\}$
    else
      $\theta_i(t) \leftarrow \beta\theta_i(t - T_{\text{per}})$
      $g_i(t) \leftarrow g_i(t - T_{\text{per}}) + \theta_i(t)$
    end if
  else
    if $X_f^i(t - T_{\text{per}}) \le \max\{X_f^i(t - 2T_{\text{per}}) - \alpha_{,}0\}$ **then**
      $\theta_i(t) \leftarrow \beta\theta_i(t - T_{\text{per}})$
    else
      $\theta_i(t) \leftarrow \theta_i(t - T_{\text{per}})$
    end if
    $g_i(t) \leftarrow g_i(t - T_{\text{per}}) + \theta_i(t)$
  end if
end for

Proposition 1. *For any initial condition, $T \in \mathbb{N}$, and design constants in Algorithm 2, the DisDRR policy keeps the network under-saturated if $(\frac{\tau_i}{\tau} - 1)\rho_i < 1$ for all $i \in [m]$.*

V2I communication requirements: This policy uses information about $|Q|$ (for $T > 1$) and X. Its worst-case communication cost is calculated similar to the DRR policy, except that at the end of each update period T_{per}, X is not communicated to all on-ramps. Instead, for all $i \in [m]$, only the part X associated with the vehicles between the i-th and $i+1$-th on-ramps is communicated to on-ramp i. The contribution of the queue size to the communication cost is the same as the DRR policy because of the same cycle mechanism. Thus, the worst-case communication cost is reduced to $(n_c + n_a)/T_{\text{per}} + n_m + \overline{Q}_2$.

5.2.4 Dynamic Space Gap Policy

In this policy, on-ramps require an additional space gap, in addition to the safe gaps in (M3)-(M4), before releasing a vehicle. This additional space gap is updated periodically based on the state of all vehicles. Recall that the DRR policy enforces an additional time gap between release of successive vehicles, which is updated based on the current state of the vehicles as well as their state in the past. The dynamic space gap policy only requires the current state of the vehicles. However, it requires additional assumptions on the vehicle-level rules: consider n vehicles over a time interval during which at least one vehicle is in the speed tracking mode, and no vehicle leaves the freeway. Then, during this time interval:

(VC4) each vehicle changes mode at most once.

(VC5) if no vehicle changes mode, then $X_g := X_{g_1} + X_{g_2}$ converges to zero globally exponentially, where X_{g_1} and X_{g_2} are defined as follows: let $X_{g_1}(t) := \sum_{e\in[n]}(|v_e(t)-V_f|+|a_e(t)|)J_e(t)$, where J_e is a binary variable which is equal to zero if the ego vehicle has been in the speed tracking mode at all times since being released from an on-ramp, and one otherwise. Moreover,

$$X_{g_2} := \sum_{e\in[n]} \left(|y_e(t) - S_e(t)| \mathbb{1}_{\{y_e(t)<S_e(t)\}} + \max_{s\in[t',t_f]} |\hat{y}_e(s|t) - \hat{S}_e(s|t)| \mathbb{1}_{\{\hat{y}_e(s|t)<\hat{S}_e(s|t)\}} \right),$$

where the second term in the sum is set to zero if the ego vehicle is not in a merging area. If the ego vehicle is in a merging area but has not yet crossed the merging point, then $t' = t_m$. Otherwise, $t' = t$. Furthermore, recall from Section 5.1.1 that the acceleration lane of an on-ramp ends either at the merging point (if the merging speed is V_f), or on the mainline (if the merging speed is less than V_f). The term $t_f|t$ is the time the ego vehicle predicts to cross the endpoint of the acceleration lane based on the information available at time $t \le t_f$.

Definition 4. ***(Dynamic Space Gap (DSG) ramp metering policy)*** *The policy works in cycles of fixed length $T\tau$, where $T \in \mathbb{N}$. At the beginning of the k-th cycle at $t_k = (k-1)T\tau$, each on-ramp allocates itself a "quota" equal to the queue size at that on-ramp at t_k. At time $t \in [t_k, t_{k+1}]$ during the k-th cycle, on-ramp i releases the ego vehicle if (M1)-(M2), and the following condition is satisfied:*

(M6) at time t, $y_e(t) \ge K_1(X(t))$, $\hat{y}_e(t_m|t) \ge K_2(X(t))$, and $\hat{y}_{\hat{f}}(t_m|t) \ge K_3(X(t))$, where $\hat{y}_{\hat{f}}$ is the predicted gap between the ego vehicle and its virtual following vehicle, $K_1(\cdot)$ is the safety distance $S_e(\cdot)$ plus an additional gap $f_1(\cdot)$, $K_2(\cdot)$ and $K_3(\cdot)$ are the gaps in (M4) required to ensure safety between merging and exiting the acceleration lane, plus additional gaps $f_2(\cdot)$ and $f_3(\cdot)$, respectively.

Once an on-ramp reaches its quota, it does not release a vehicle during the rest of the cycle. The additional gaps $f_i(X(t))$, $i = 1, 2, 3$, are piecewise constant and updated periodically at each time step according to some rule which will be determined in the proof of Theorem 3.

Theorem 3. *There exists $f_i(\cdot)$, $i = 1, 2, 3$, such that for any initial condition, and $T \in \mathbb{N}$, the DSG policy keeps the network under-saturated if $(\frac{\tau_i}{\tau} - 1)\rho_i < 1$ for all $i \in [m]$.*

V2I communication requirements: This policy uses information about $|Q|$ (for $T > 1$) and X. Except for an initial finite time, X is communicated to every on-ramp at each time step. Moreover, the contribution of the queue size to the communication cost is the same as the DRR policy because of the same cycle mechanism. Thus, the communication cost is upper bounded by $(n_c + n_a)m + \overline{Q}_2$.

Remark 10. *The choice of the additional gaps f_i, $i = 1, 2, 3$, in the DSG policy is not limited to the expressions found in the proof of Theorem 3. An alternative expression is simulated in Section 5.3.4*

5.2.5 Local and Greedy Policies

The performance analysis of the three policies in Sections 5.2.2-5.2.4 can be divided into two phases. The first phase concerns the transient from the initial condition to the free flow state where each vehicle occupies some slot, and the second phase is from this free flow state onward. Since the performance index of throughput is an asymptotic notion, it is natural to examine the policies specifically in the second phase. Indeed, in the second phase, the actions of all the three policies can be shown to be equivalent to the following policy:

Definition 5. ***(Fixed-Cycle Quota (FCQ) ramp metering policy)*** *The policy works in cycles of fixed length $T\tau$, where $T \in \mathbb{N}$. At the beginning of the k-th cycle at $t_k = (k-1)T\tau$, each on-ramp allocates itself a "quota" equal to the queue size at that on-ramp. During a cycle, the i-th on-ramp releases the ego vehicle if (M1)-(M4) are satisfied. Once an on-ramp reaches its quota, it does not release a vehicle during the rest of the cycle.*

V2I communication requirements: This policy uses information about $|Q|$ (for $T > 1$) and part of X. At each time step during a cycle, the vehicles in the merging area of every on-ramp communicate their state to that on-ramp. Moreover, the contribution of the queue size to the communication cost is the same as the DRR policy because of the same cycle mechanism. Hence, C is upper bounded by $n_m + \overline{Q}_2$.

Note that the FCQ policy is local, except for synchronizing the beginning of the cycles which can be done once offline. In the special case of $T = 1$, the FCQ becomes the simple *greedy* (and local) policy under which the on-ramps do not need to synchronize cycles with other on-ramps or keep track of their quota. Therefore, the communication cost of the greedy policy is upper bounded by n_m. One can see from the proof of the DRR policy that, starting from the aforementioned second phase, the freeway is under-saturated using the FCQ policy if $(\tau_i/\tau - 1)\rho_i < 1$ for all $i \in [m]$. It is natural to wonder if such a result holds if we start from *any* initial condition, or under *any* other greedy policy (not just the slot-based). In Section 5.3.3, we provide intuition as to when the former statement might be true.

5.2.6 An Outer Estimate

We now provide an outer estimate to the under-saturation region of any policy, against which we benchmark the inner estimates derived in the previous sections for the proposed policies. This outer estimate can be thought of as the network analogue of the flow capacity of a single on-ramp. To formalize this, let $\bar{D}_{\pi,p}(k\tau)$ be the cumulative number of vehicles that has crossed point p on the mainline up to time $k\tau$, $k \in \mathbb{N}_0$, under the RM policy π. Then, the crossing rate at point p is defined as $\bar{D}_{\pi,p}(k\tau)/k$ and the "long-run" crossing rate is $\limsup_{k\to\infty} \bar{D}_{\pi,p}(k\tau)/k$. Note that $\bar{D}_{\pi,p}((k+1)\tau) - \bar{D}_{\pi,p}(k\tau)$ represents the *traffic flow* at point p in terms of the number of vehicles per τ seconds. Macroscopic traffic models suggest that the flow is no more than the mainline capacity, which is 1 vehicle per τ seconds. Hence, $\bar{D}_{\pi,p}((k+1)\tau) - \bar{D}_{\pi,p}(k\tau) \leq 1$ for all $k \in \mathbb{N}_0$ and any RM policy π. This implies that

$$\limsup_{k\to\infty} \bar{D}_{\pi,p}(k\tau)/k \leq 1, \tag{13}$$

for any RM policy π. We take (13) as an assumption to the next Theorem.

Theorem 4. *If a policy π keeps the freeway under-saturated and satisfies* (13) *for at least one point on each link, then the demand must satisfy $\rho \leq 1$.*

Remark 11. *In all of the policies studied in previous sections, vehicles move at the constant speed V_f at some point p_i on the i-th link, $i \in [m]$, after a finite time. This and the constant time headway rule imply that the number of vehicles that can cross p_i at each time step is no more than one. Therefore, $\limsup_{k\to\infty} \bar{D}_{\pi,p_i}(k\tau)/k \leq 1$ for all $i \in [m]$, and the long-run crossing rate condition in* (13) *holds.*

Remark 12. *If the merging speed at all the on-ramps is V_f, then the minimum time headways τ_i, $i \in [m]$, are all equal to 2τ. In this case, one can check that the inner estimate of the under-saturation region given in Theorem 1, Theorem 2, and Proposition 1 become $\rho < 1$. Comparing with*

Theorem 4, this implies that the Renewal, DRR, and DisDRR policies give the maximum possible throughput for all practical purposes.

5.2.7 Discussion of the Straight Road Geometry

In this section, we discuss how the previous results can be extended to the straight road geometry with $m-1$ on- and off-ramps. The on- and off-ramps are placed alternatively, and they are numbered in an increasing order along the direction of travel. The on-ramps are numbered from 2 to m, and the off-ramps are numbered from 1 to $m-1$. The upstream entry point and the downstream exit point are numbered 1 and m, respectively; see Figure 7a.

The vehicle-level rules discussed in Section 5.1.1 remain unchanged, except that, at the free flow state, we require the vehicles from the upstream entry point to enter at the constant speed V_f. One can also relax (VC3) as follows:

(VC3) there exists T_{empty} such that for any initial condition, if no vehicle is released from on-ramps $2, \dots, j$ for some $j \in \{2, \dots, m\}$, and all the vehicles downstream of link j are at the free flow state, then all the vehicles reach the free flow state after at most T_{empty} time.

Note that condition (VC3) imposes macroscopic-level constraints on the inflow from the upstream entry point before reaching the free flow state. This is justified if, e.g., the on-ramps upstream of the entry point are also controlled. We also need to specify vehicle arrivals and their destination at the microscopic level after reaching the free flow state. To this end, we may adopt an i.i.d Bernoulli process with parameter λ_1 and an i.i.d Bernoulli routing process that are independent of the on-ramp processes discussed in Section 5.1.2. The routing matrix, the cumulative routing matrix, and the induced loads are defined accordingly. However, this choice of demand model may be generalized to include practical cases as noted in Remark 2.

Since the vehicle arrivals from the upstream entry point is not controlled, e.g., by a ramp meter, the mainline may never become empty even if all the on-ramps pause release. Hence, we need to modify the description of the Renewal policy. In particular, the first cycle of the Renewal policy begins when the vehicles reach the free flow state. Thereafter, the k-th cycle, $k > 1$, begins when all quotas of the $(k-1)$-th cycle reach their destination. The description and performance of the DRR policy remains the same. In the description of the DisDRR policy, X_f^2 is modified to include all vehicles between the upstream entry point and on-ramp 2. The performance of this policy also remains the same.

5.3 Simulations

The following setup is common to all the simulations in this section. We consider a ring road with $P = 1860$ [m] and $m = 3$ on- and off-ramps. Let $h = 1.5$ [s], $S_0 = 4$ [m], $L = 4.5$ [m], and $V_f = 15$ [m/s]. For these parameters, we obtain $n_c = 60$. The on-ramps are located symmetrically at 0, $P/3$, and $2P/3$; the off-ramps are also located symmetrically 155 [m] upstream of each on-ramp. The initial queue size of all the on-ramps is set to zero. Vehicles arrive at the on-ramps according to i.i.d Bernoulli processes with the same rate λ; their destinations are determined by

$$R = \begin{pmatrix} 0.2 & 0.7 & 0.1 \\ 0 & 0.8 & 0.2 \\ 0.5 & 0 & 0.5 \end{pmatrix}.$$

Note that, on average, most of the traffic wants to exit from off-ramp 2. Thus, one should expect that on-ramp 3 finds more safe gaps on the mainline than the other two on-ramps. As a result, on-ramp 3's queue size is expected to be less than the other two on-ramps.

5.3.1 Greedy Policy for Low Merging Speed

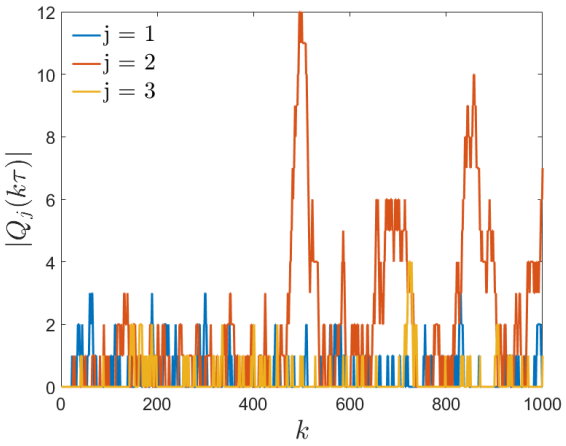

Figure 13: On-ramp queue size profiles when all the on-ramps are long, under the FCQ policy.

The mainline and acceleration lanes are assumed to be initially empty in this section and Section 5.3.2. The on-ramps use the greedy policy, i.e., the FCQ policy with $T = 1$, from Section 5.2.5. When the merging speed of all the on-ramps is V_f, then for the given R, the throughput is given by $\lambda = 5/9$ [veh/time step]. Note that since the arrival rates are the same, the throughput is a single point. The queue size profiles for $\lambda = 0.5$ [veh/time step], which corresponds to $\rho = 0.9$ (heavy demand), is shown in Figure 13. As expected, $|Q_3(\cdot)|$ is generally less than the other two on-ramps.

We next consider the case when the merging speed of on-ramps 1 and 3 are V_f, i.e., $\tau_1 = \tau_3 = 2\tau$, and is approximately $V_f/3 = 5$ [m/s] for on-ramp 2 which corresponds to $\tau_2 = 3\tau$. In the heavy demand regime, i.e., $\rho = 0.9$, $|Q_2(\cdot)|$ increases steadily which suggests that the freeway becomes saturated even though $\rho < 1$. The throughput in this case is estimated from simulations to be $\lambda = 0.44$ [veh/time step]. The estimates of the throughput of the Renewal and FCQ policies, given by Theorems 1 and 2 are $\lambda = 0.5$ [veh/time step] and $\lambda = 0.33$ [veh/time step], respectively. Combining this with the simulation results, we see that the Renewal policy performs better than the greedy policy in terms of its throughput.

5.3.2 Effect of Cycle Length on Queue Size

The long-run expected queue size, i.e., $\overline{Q} := \limsup_{k\to\infty} \mathbb{E}\sum_i |Q_i(k\tau)|$, are compared under the FCQ policy for different T. The expected queue size is computed using the *batch means* approach, with warm-up period of 10^5, i.e., the first 10^5 observations are not used, and batch size of 10^5. In each case, the simulations are run until the margin of error of the 95% confidence intervals are 1%.

Figure 14a shows $\overline{Q}$ vs. ρ for different T when the merging speed of all the on-ramps is V_f. The plot suggests that for all ρ, $\overline{Q}$ increases monotonically with T. However, this does not hold true when the merging speeds are low. For example, Figure 14b shows $\overline{Q}$ in log scale vs. T for $\lambda = 0.455$, which corresponds to $\rho \approx 0.82$, when the merging speed of on-ramps 1 and 3 is V_f, and

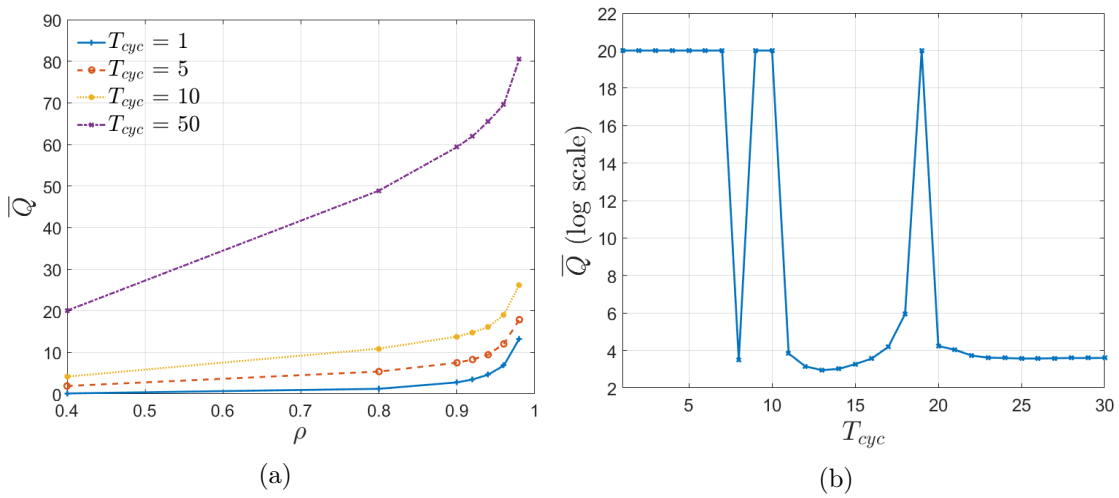

(a) (b)

Figure 14: Effect of cycle length T on the long-run expected queue size (a) for different ρ when both on-ramps are long, (b) for a fixed ρ when on-ramp 2 is short. Both plots are under the FCQ policy. In plot (b), the logarithm of the expected queue size is set to 20 whenever the freeway becomes saturated.

is $V_f/3$ for on-ramp 2. The plot shows that the freeway may become saturated depending on the choice of T, and that the dependence of $\overline{Q}$ on T is not monotonic.

5.3.3 Relaxing the V2X Requirements

In this scenario, we evaluate the performance of the greedy policy for $\rho = 0.9$ (heavy demand) when the merging speed of all the on-ramps is V_f, and the mainline is initially not empty. The safety mode only consists of the following submode: in a non-merging scenario, the ego vehicle follows the leading vehicle by keeping a safe time headway as in, e.g., [29]. In a merging scenario, it applies the aforementioned control law with respect to both its leading and virtual leading vehicles. The most restrictive acceleration is used by the ego vehicle. For this control law, platoons of vehicles are *stable* and *string stable* ([29]).

The initial number of vehicles on the mainline, their location, and their speed are chosen at random such that the safety distance is not violated. The initial acceleration of all the vehicles is set to zero. We conducted 10 rounds of simulation with different random seed for each round. It is observed in all scenarios that vehicles reach the free flow state after a finite time using the greedy policy. We conjecture that this observation holds for *any* initial condition if platoons of vehicles are stable and string stable. The intuition behind this statement is as follows: using the greedy policy, vehicles are released only if there is enough space gap on the mainline, and successive releases are at least τ seconds apart. In between releases, vehicles on the mainline try to reach the free flow state because of platoon stability. Once a vehicle is released, it may create disturbance, e.g., in the acceleration of the upstream vehicles. However, string stability implies that such disturbance is attenuated. Therefore, it is natural to expect that vehicles will reach the free flow state after a certain finite time. Thereafter, by Theorem 2, the greedy policy keeps the freeway under-saturated if $(\tau_i/\tau - 1)\rho_i < 1$ for all $i \in [m]$.

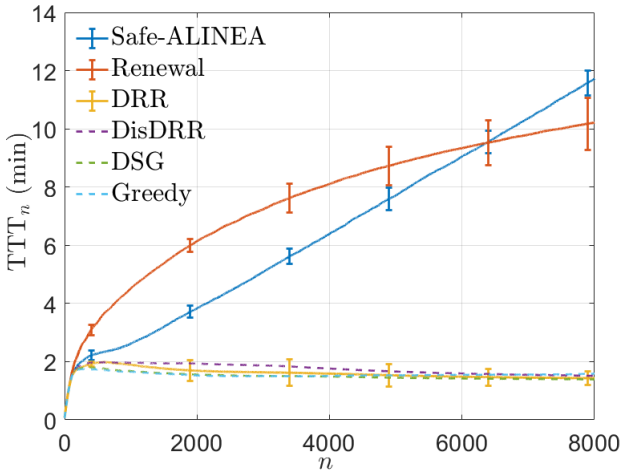

Figure 15: The total travel time (TTT_n), where n is the number of completed trips, under the fixed demand $\rho = 0.82$. The vertical lines in the Safe-ALINEA, Renewal, and DRR policies represent the standard deviation. The standard deviation of the DisDRR, DSG, and greedy policies are similar to the DRR policy and are omitted for clarity.

5.3.4 Comparing the Total Travel Time

We evaluate the total travel time under the Renewal, DRR, DisDRR, DSG, and greedy policies and compare them with the ALINEA ramp metering policy under vehicle safety constraints. The ALINEA policy was introduced in [19] and, for a given on-ramp, can be expressed as follows:

$$r(k) = r(k-1) + K_r(\hat{o} - o(k)).$$

Here, $r(\cdot)$ is the on-ramp outflow, K_r is a positive design constant, $o(\cdot)$ is the occupancy of the mainline downstream of the on-ramp, and $\hat{o}$ is the desired occupancy. For the ALINEA policy, we use time steps of size 60 [s], $K_r = 70$ [veh/h] ([19]), and $\hat{o} = 13$ % corresponding to the capacity flow. We also add the following safety filter on top of ALINEA: the ego vehicle is released only if it predicts to be at a safe gap with respect to its virtual leading and following vehicles at the moment of merging. We use the name Safe-ALINEA to refer to his policy. For the DRR and DisDRR policies, we use the following parameters: $T = 12$, $T_{\mathrm{per}} = 2\tau$, $\alpha_{=}50$, $\gamma_2 = 10$, $\theta_i^\circ = 0.1$, $\beta = 1.01$, and $T_{\max} = 100$ [s]. Informally, γ_1 is set to a high value so that the release times $g_i(\cdot)$ increase if the traffic condition, i.e., the value of X_f, does not improve significantly in between the update periods; γ_2 and θ_i° are set to low values so that the jump sizes in $g_i(\cdot)$ are small; $T_{\max}$ is set to a high value so that on-ramps update their release time in a completely decentralized way. For the DSG policy, we set $T = 12$ and use the additional gap $\kappa(X_{f_1} + \sum_{e\in[n]} |y_e - (hv_e + S_0)|I_e)$ on top of (M3)-(M4), where $\kappa = 0.01$ is chosen so that the additional gap is relatively small (≈ 8 [m] for the initial condition described next).

We let the merging speed of on-ramps 1 and 3 be V_f, and $V_f/3$ for on-ramp 2. We let $\lambda = 0.455$, and consider a congested initial condition where the initial number of vehicles on the mainline is $100 > n_c$; each vehicle is at the constant speed of 6.7 [m/s], and is at the distance $h \times 6.7 + S_0 \approx 14.1$ [m] from its leading vehicle.

We evaluate the Total Travel Time (TTT_n), which is the average travel time of the first n completed trips. In order to show consistency, for each policy we conducted 10 rounds of simulations

with different random seed for each round. In each round, we used the same seed across different policies. We use the average and standard deviation of the 10 simulation rounds for illustration.

For the aforementioned setup, the freeway becomes saturated under the Safe-ALINEA and greedy policies, but remains under-saturated under the other policies. Figure 15 and Table 2 show the simulation results. From the results, we can see that the DRR policy provides significant improvement in the TTT as compared to the Safe-ALINEA (approximately 83%) and Renewal (approximately 80%) policies. The DisDRR, DSG, and greedy policies show similar improvements. However, as shown in Section 5.3.2, the freeway becomes saturated under the greedy policy for long enough simulation time (not shown in Figure 15). This implies that the current improvement in the TTT will disappear as the number of completed trips grows. On the other hand, the choice of $T = 12$ in the DRR, DisDRR, and DSG policies will show steady improvement.

Table 2 also provides the time average queue size, i.e., $\overline{Q}(K\tau) = \frac{1}{mK}\sum_{i,j}|Q_i(j\tau)|$, where $|Q_i|$ is the average queue size over the 10 simulation rounds, and $K\tau$ is the simulation time. It can be seen that the Safe-ALINEA policy has the largest queue size as compared to the other policies. This is mainly because of the safety filter added on top of the ALINEA policy. Without the safety filter, ALINEA releases vehicles more frequently, which results in shorter on-ramp queues on average. However, the safety distance is violated more often and more congestion is created on the mainline, which results in higher travel time compared to the DRR policy. In conclusion, with proper choice of design parameters, our proposed policies create and maintain the free flow state on the mainline, which improves the travel time.

Table 2: Performance of the policies.

Ramp Metering Policy	Renewal	DRR	DisDRR	DSG	Greedy	Safe-ALINEA
$\max_n \mathrm{TTT}_n$ [min]	10.2	2	2	1.8	1.7	11.7
Average queue size	101	6	10	7	12	127

6 Conclusion

In this project, we exploited the departure time-based equilibrium formulation to design pricing profiles under budget and price cap constraints that achieve the system optimum for a single bottleneck. Our approach for the single bottleneck case under inelastic demand is based on the bilevel non-convex optimization problem that has a special structure, which can be analytically solved under a reasonably weak assumption on the schedule delay cost function. In the case of the single bottleneck under the elastic demand, we converted the bilevel non-convex optimization problem into a single level mathematical program with equilibrium constraints, which was approximately solved. The designed pricing profile allows for decentralized control of the commuters' departure times from their homes and further allows for effective management of the transportation infrastructure with a single bottleneck under an elastic or inelastic demand setting.

We provided a RM framework at the microscopic level subject to vehicle following safety constraints. This allows to explicitly take into account the V2X communication scenarios into the RM design, and study the impact on the freeway performance. We specifically provided RM policies for a few such scenarios, and analyzed performance in terms of the throughput of the freeway. There are several avenues for generalizing the setup and methodologies initiated in this project. Of

immediate interest would be to consider a general network structure, and to derive sharper inner estimates on the under-saturation region for the low merging speed case. We are also interested in expanding performance analysis to include travel time, possibly by leveraging waiting time analysis from queuing theory.

7 Implementation

Not applicable.

8 References

References

[1] Chen-Hsiu Laih. Queueing at a bottleneck with single- and multi-step tolls. *Transportation Research Part A: Policy and Practice*, 28(3):197–208, May 1994.

[2] Zhi-Chun Li, Hai-Jun Huang, and Hai Yang. Fifty years of the bottleneck model: A bibliometric review and future research directions. *Transportation Research Part B: Methodological*, 139(C):311–342, 2020.

[3] Richard Arnott, André de Palma, and Robin Lindsey. A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *American Economic Review*, 83:161–79, 02 1993.

[4] Moshe Ben-Akiva, Andre de Palma, and Pavlos Kanaroglou. Dynamic Model of Peak Period Traffic Congestion with Elastic Arrival Rates. *Transportation Science*, 20(3):164–181, August 1986.

[5] Vincent A.C. van den Berg. Step-tolling with price-sensitive demand: Why more steps in the toll make the consumer better off. *Transportation Research Part A: Policy and Practice*, 46(10):1608–1622, 2012.

[6] Mogens Fosgerau and Kurt Van Dender. Road pricing with complications. *Transportation*, 40(3):479–503, May 2013.

[7] Charles Raux and Stéphanie Souche. The acceptability of urban road pricing: A theoretical analysis applied to experience in lyon. *Journal of Transport Economics and Policy (JTEP)*, 38(2):191–215, 2004.

[8] Jan Rouwendal, Erik T Verhoef, and Jasper Knockaert. Give or take? rewards versus charges for a congested bottleneck. *Regional Science and Urban Economics*, 42(1-2):166–176, 2012.

[9] Richard Arnott, André de Palma, and Robin Lindsey. Economics of a bottleneck. *Journal of Urban Economics*, 27(1):111–130, 1990.

[10] G. Rostomyan, K. Savla, and P. A. Ioannou. Bottleneck management using pricing under constraints. In *American Control Conference*, 2023. Under review.

[11] Hai Yang and Huang Hai-Jun. Analysis of the time-varying pricing of a bottleneck with elastic demand using optimal control theory. *Transportation Research Part B: Methodological*, 31(6):425–440, November 1997.

[12] Kenneth A Small. The scheduling of consumer activities: work trips. *The American Economic Review*, 72(3):467–479, 1982.

[13] Markos Papageorgiou and Apostolos Kotsialos. Freeway ramp metering: An overview. *IEEE transactions on intelligent transportation systems*, 3(4):271–281, 2002.

[14] Jackeline Rios-Torres and Andreas A Malikopoulos. Automated and cooperative vehicle merging at highway on-ramps. *IEEE Transactions on Intelligent Transportation Systems*, 18(4):780–789, 2016.

[15] Chris Lee, Bruce Hellinga, and Kaan Ozbay. Quantifying effects of ramp metering on freeway safety. *Accident Analysis & Prevention*, 38(2):279–288, 2006.

[16] Joseph A Wattleworth. Peak-period analysis and control of a freeway system. Technical report, Texas Transportation Institute, 1965.

[17] Ioannis Papamichail, Apostolos Kotsialos, Ioannis Margonis, and Markos Papageorgiou. Coordinated ramp metering for freeway networks–a model-predictive hierarchical control approach. *Transportation Research Part C: Emerging Technologies*, 18(3):311–331, 2010.

[18] Markos Papageorgiou, Christina Diakaki, Vaya Dinopoulou, Apostolos Kotsialos, and Yibing Wang. Review of road traffic control strategies. *Proceedings of the IEEE*, 91(12):2043–2067, 2003.

[19] Markos Papageorgiou, Habib Hadj-Salem, Jean-Marc Blosseville, et al. Alinea: A local feedback control law for on-ramp metering. *Transportation research record*, 1320(1):58–67, 1991.

[20] Markos Papageorgiou, Jean-Marc Blosseville, and Habib Haj-Salem. Modelling and real-time control of traffic flow on the southern part of boulevard périphérique in paris: Part ii: Coordinated on-ramp metering. *Transportation Research Part A: General*, 24(5):361–370, 1990.

[21] Gabriel Gomes and Roberto Horowitz. Optimal freeway ramp metering using the asymmetric cell transmission model. *Transportation Research Part C: Emerging Technologies*, 14(4):244–262, 2006.

[22] Li Li, Ding Wen, and Danya Yao. A survey of traffic control with vehicular communications. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):425–432, 2013.

[23] Jennie Lioris, Ramtin Pedarsani, Fatma Yildiz Tascikaraoglu, and Pravin Varaiya. Platoons of connected vehicles can double throughput in urban roads. *Transportation Research Part C: Emerging Technologies*, 77:292–305, 2017.

[24] M Karimi, C Roncoli, C Alecsandru, and M Papageorgiou. Cooperative merging control via trajectory optimization in mixed vehicular traffic. *Transportation Research Part C: Emerging Technologies*, 116:102663, 2020.

[25] Yuki Sugiyama, Minoru Fukui, Makiko Kikuchi, Katsuya Hasebe, Akihiro Nakayama, Katsuyoshi Nishinari, Shin ichi Tadaki, and Satoshi Yukawa. Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam. 2008.

[26] Raphael E Stern, Shumo Cui, Maria Laura Delle Monache, Rahul Bhadani, Matt Bunting, Miles Churchill, Nathaniel Hamilton, Hannah Pohlmann, Fangyu Wu, Benedetto Piccoli, et al. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89:205–221, 2018.

[27] Yang Zheng, Jiawei Wang, and Keqiang Li. Smoothing traffic flow via control of autonomous vehicles. *IEEE Internet of Things Journal*, 7(5):3882–3896, 2020.

[28] Milad Pooladsanj, Ketan Savla, and Petros Ioannou. Vehicle following over a closed ring road under safety constraint. In *2020 IEEE Intelligent Vehicles Symposium (IV)*, pages 413–418. IEEE, 2020.

[29] Milad Pooladsanj, Ketan Savla, and Petros Ioannou. Vehicle following on a ring road under safety constraints: Role of connectivity and coordination. *IEEE Transactions on Intelligent Vehicles*, pages 1–1, 2022.

[30] P.A. Ioannou and C.C Chien. Autonomous intelligent cruise control. *IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY*, 42(4):657–672, 1993.

[31] M. Pooladsanj, K. Savla, and P. A. Ioannou. Ramp metering to maximize throughput under vehicle safety constraints. *Transportation Part C: Emerging Technologies*, 2022. Under review.

[32] Xuexiang Jin, Yi Zhang, Fa Wang, Li Li, Danya Yao, Yuelong Su, and Zheng Wei. Departure headways at signalized intersections: A log-normal distribution model approach. *Transportation research part C: emerging technologies*, 17(3):318–327, 2009.

[33] Leonidas Georgiadis, Wojciech Szpankowski, and Leandros Tassiulas. A scheduling policy with maximal stability region for ring networks with spatial reuse. *Queueing Systems*, 19(1):131–148, 1995.

[34] Mor Armony and Nicholas Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing systems*, 44(3):209–252, 2003.

9 Data Management Plan

Products of Research

No new data was collected for this project. Only simulation data was generated.

Data Format and Consent

All the simulations were done in Matlab.

Data Access and Sharing

The input/output data for the simulations in this report is available at `https://viterbi-web.usc.edu/~ksavla/code.html`.

Reuse and Redistribution

The data can be reused freely for non-commercial purposes. Its usage, in original or after modification, in publications is to be done with due acknowledgement to the authors of this report and by citation of relevant publications by the authors.