

A Token-based Admission Control and Request Scheduling in Lane Reservation Systems

Kai Liu*

Sang H. Son[†]

Victor C.S. Lee*

Krasimira Kapitanova[†]

* Department of Computer Science, City University of Hong Kong, Kowloon, Hong Kong
Email: kevin.liu@student.cityu.edu.hk csvlee@cityu.edu.hk

[†] Department of Computer Science, University of Virginia, Charlottesville VA, USA
Email:son@cs.virginia.edu krasi@virginia.edu

Abstract—In many parts of the world, the ever-expanding traffic congestion problem has become a major source of wasted fuel, economic burden, and environmental pollution. Alleviating traffic congestion is not only a matter of expanding the transportation capacity, such as adding more lanes or building new roads, but also the problem of providing good traffic management and polices. Recently, the concept of road reservation systems has been widely discussed. With this system, in order to use the lanes and roads controlled by the transportation operators, drivers make reservations in advance. The goal of the system is to provide certain quality of services to the drivers, e.g., guaranteed end to end travel delay. In this work, we focus on the admission control and request scheduling for a reservation system. Specifically, we consider a high-priority lane reservation scenario, in which the system has to determine which vehicles could be allowed to enter the high-priority lane, so that the traffic workload does not exceed the lane capacity. We present a token-based admission control policy that implements the reservation scheme. We have also designed an on-line scheduling algorithm that selects which reservation requests can be allowed based on the admission control policy. Using a simulation model to evaluate the system performance in a variety of scenarios, we show that the proposed algorithm can achieve efficient utilization of the high priority lane.

I. INTRODUCTION

Traffic congestion becomes a global problem since it seriously undermines the efficiency of current road transportation systems. Conventional wisdom holds that expanding transportation capacity, such as adding additional lanes or constructing new roads, is the most straightforward solution to traffic congestion. Unfortunately, the growing demand for road network capacities is far beyond what we can afford in terms of both physical resources and economic cost. There are nearly a billion vehicles on the road today and this number is expected to double over the next decade. [1]. The problem is even more serious in large cities. For example, as reported by the Beijing municipal taxation office, the number of new cars registered in Beijing in the first four months of 2010 increased by 23.8% and reached 248,000. Therefore, simply expanding road network capacities to cater to this dramatic increasing demand is not a sustainable solution.

Transportation planners and government have both recognized the importance of traffic congestion management and policy. Therefore, a great effort has been made to improve the efficiency of the existing road infrastructure. The Department of Transportation in Hong Kong has set up websites with maps showing congestion in some hotspots

like Cross Harbor Tunnel. To deal with the rapid growth of the number of vehicles, Beijing announced a policy to limit the number of new plates issued to passenger cars in the end of 2010. In Singapore, transportation controllers receive real-time data through sensors to model and predict future traffic flows. Besides the government efforts, a variety of traffic demand management (TDM) strategies have been employed to alleviate traffic congestion, including high occupancy vehicle lanes (HOV) [2], congestion pricing [3], and ramp metering [4]. While these strategies are helpful in relieving the congestion problem to some extent, they are far from sufficient. Additional efforts are needed in order to further improve the efficiency of the current transportation systems.

The cutting-edge technologies in vehicular networks enable us to try more innovative and sophisticated solutions to tackle traffic problems and improve transportation efficiency. As envisioned in [5], today's technologies can be used to enable the communication between highways and vehicles, thus making the highway system aware of the drivers travel plans. The roadway reservation system is one of the many applications that would benefit from these technologies. Several researchers have been working on this problem over the past few years. A lane reservation system has been presented in [6]. It describes a system which allows drivers to reserve time slots on a high-priority lane by paying a premium price, while the high-priority lane provides drivers with congestion free travel. Similar studies are also presented in [7] and [8], where the concept of Highway Space Inventory Control System (HSICS) is proposed. The basic idea of HSICS is that all road users have to make reservations in advance to enter the highway. In order to achieve certain system-wide objectives, highway operators have to make real-time decisions whether to accept or reject travelers requests.

The majority of previous work had mainly investigated the roadway reservation from the perspective of system modeling. Our work, however, focuses on the issue of admission control and request scheduling in a reservation system. This paper has three main contributions: i) we have developed a token-based system that implements the reservation scheme, ii) we have developed an online scheduling algorithm which makes admission decisions by selecting the most rewarding requests, iii) we evaluate our system through simulations. The simulation results show that our algorithms achieves efficient use of the high priority lane.

The rest of this paper is organized as follows. Section II outlines the motivation of this study. In Section III, we present our token-based admission control policy. Section IV describes the on-line scheduling algorithm which makes decisions based on the admission control policy. Section V shows the results of our performance evaluation and Section VI summarizes the work and presents future directions.

II. MOTIVATION

Traffic congestion becomes a major source of both economic and environmental problems. In the U.S., 439 urban areas are faced with the congestion problem. In 2009, urban Americans wasted 4.8 billion hours in total during their travel due to congestion. It caused the purchase of extra 3.9 billion gallons of fuel at cost of \$115 billion [9]. In Japan, traffic congestion causes the loss of about 3.8 billion man-hours (worth about 12 trillion yen) each year [10]. About 11% of the fuel consumed by automobiles is wasted during congestion [11]. In addition to extra fuel consumption, traffic congestion also causes the emission of higher levels of air pollutants and greenhouse gases, including carbon monoxide (CO), carbon dioxide (CO₂), hydrocarbons (HC) and oxides of nitrogen (NO_x) [12]. As reported by the U.S. Environmental Protection Agency in [13], transportation accounted for approximately 33% of total CO₂ emissions from fossil fuel combustion, which is 7% higher than the share of industry. A more efficient roadway system is urgently needed in order to alleviate the costs of traffic congestion.

The idea of adopting reservation schemes into the roadway system is inspired from the success of many existing applications, such as reserving a seat in airlines, booking a room in hotels, and making doctor's appointments. However, in road traffic management, the reservation scheme is not yet a well explored topic. There are many issues that could have previously hindered the development of roadway reservation systems, such as the high cost on infrastructure deployment, the lack of reliable real-time communication capability, as well as possible contradictions in public opinions.

Nowadays, with the advent of emerging vehicular networking technologies, we believe that most of the critical issues in developing the roadway reservation system could be gracefully addressed. First, the demands for real-time wireless communication in vehicular networks can be well supported by the DSRC (Dedicated Short Range Communication) [14] in parallel with alternate communication technologies such as GSM (Global System for Mobile Communications) [15], CDMA2000 [16] and WiMAX (Worldwide Interoperability for Microwave Access) [17]. In addition to adequate technical support, the reservation system also needs active user participation. Therefore, it cannot be deployed without acceptance from the public. Kim et al. studied the public opinion on deploying highway reservation systems [18]. Their work showed that, although some factors, such as perceived travel time, age, and travel distance, may influence people's opinions, the majority is willing to accept the reservation policy, albeit inconvenient, to avoid possible severe traffic congestion.

Given the necessity and feasibility to deploy roadway reservation systems, this study focuses on the design of an admission control scheme and a scheduling algorithm. While the proposed scheme could theoretically be applied to any type of roadway reservation systems, we currently focus on the reservation of a specific lane along the freeway. It is likely that this would be the first actual deployment of the roadway reservation system since it would be easier to implement the access control, traffic control, and enforcement on freeways for only one lane. In addition, controlling and restriction on just one particular lane instead of the whole freeway would be a more graceful migration to the implementation of roadway reservation systems.

III. TOKEN-BASED ADMISSION CONTROL

A high-priority lane consists of N segments, denoted by $S = \{s_1, s_2, \dots, s_N\}$. Each segment is a part of the lane between two adjacent entrance/exit points along the freeway. The *capacity of a segment* is adopted to restrict the traffic workload, which is defined as the maximum number of vehicles allowed to be in this segment concurrently. Drivers who wish to enter the high-priority lane have to submit a request to make a reservation in advance. The drivers who successfully make the reservation are expected to enter the reserved segments at the specified time. Otherwise, either early or late arrival may result in the revocation of the reservation. Also, if they decide not to use their reservation, drivers can cancel it in advance.

To enable the high-priority lane described above, we present a *token-based admission control* policy to determine whether a reservation can be accepted or not. The basic idea is that the system maintains a number of tokens to monitor and trace the workload of each segment so that it will not be overloaded. Given a request, the system determines whether it should be accepted or rejected by checking the status of relevant tokens. We transform the admission control problem into a token-based policing mechanism in the following way. The set of tokens maintained in the system is denoted by T where $T = \{T_1, T_2, \dots, T_N\}$. T_i ($1 \leq i \leq N$) is the set of tokens for segment s_i , which is further described as $T_i = \{T_{i_1}, T_{i_2}, \dots, T_{i_{|T_i|}}\}$, where $|T_i|$ is the capacity of segment s_i , i.e. the maximum number of tokens for that segment. Each token T_{i_j} ($1 \leq j \leq |T_i|$) is characterized by a 2-tuple: $\langle status, reservation_table \rangle$. The *status* of a token could be either *IDLE* or *BUSY*. *IDLE* indicates that the token is not currently occupied by any vehicle, while *BUSY* means that a certain vehicle is holding this token and running in the corresponding segment. The *reservation_table* of a token records the time slots for which the token has been reserved and the corresponding vehicles that have reserved them. The table is dynamic and each entry consists of three elements: $[v_id, in_time, out_time]$, which corresponds to the vehicle's ID, arrival time, and departure time respectively. For a vehicle with a requested time slot, it can be granted with a token only if the requested time slot is not overlapped with any reserved time slots recorded in the *reservation_table* of this token.

v_id	in_time	out_time
001	1:30pm	2:00pm
002	2:10pm	2:40pm

(a) at 1:00pm

v_id	in_time	out_time
001	1:30pm	2:00pm
002	2:10pm	2:40pm
004	3:00pm	3:30pm

(b) at 1:10pm

v_id	in_time	out_time
001	1:30pm	2:00pm
004	3:00pm	3:30pm

(c) at 1:20pm

v_id	in_time	out_time
004	3:00pm	3:30pm

(d) at 2:00pm

Fig. 1. An example of token-based admission control

Figure 1 shows an example of how the admission policy is used. Suppose at 1:00pm, a token in IDLE status has been granted to two vehicles (ID: 001 and 002), which reserved the time slots of 1:30pm~2:00pm and 2:10pm~2:40pm, respectively. The *reservation_table* of this token is shown in Figure 1(a). At 1:10pm, two other vehicles (ID: 003 and 004) request the time slots of 2:30pm~3:00pm and 3:00pm~3:30pm, respectively. In this case, the token can be only granted to vehicle 004, because the time slot requested by vehicle 003 overlapped with the time slot reserved by vehicle 002 (the overlapping is at 2:30pm~2:40pm). Note that each token can be only held by one vehicle at a time. In other words, all the reserved time slots in a token must be non-overlapped. When the token is granted to vehicle 004, a new entry is inserted into the *reservation_table* as shown in Figure 1(b). At 1:20pm, vehicle 002 cancels its reservation and the *reservation_table* is updated by removing the corresponding entry as shown in Figure 1(c). Note that if the request of vehicle 003 is still pending in the service queue, this token could be granted to it now. Suppose vehicle 001 enters the segment on time. It obtains the token since 1:30 pm and the status of the token is set to *BUSY*. When the vehicle leaves the segment at 2:00pm, the token's status is set back to *IDLE* and the corresponding entry is removed, as shown in Figure 1(d).

When a request is submitted, it gets added to the service queue. The pending requests at time t are denoted by $Q(t)$ and $Q(t) = \{Q_1, Q_2, \dots, Q_{|Q(t)|}\}$, where $|Q(t)|$ is the current number of pending requests. Each request Q_m ($1 \leq m \leq |Q(t)|$) may require one or multiple segments: $Q(m) = \{s'_1, s'_2, \dots, s'_{|Q_m|}\}$, where $|Q_m|$ ($1 \leq |Q_m| \leq N$) is the total number of required segments. Each request specifies the vehicle's ID and the expected reservation time slot for each segment. To facilitate the discussion, we assume that both the arrival time (t_{in}) and the departure time (t_{out}) are specified by the request, although practically, t_{out} can be derived from t_{in} and the estimated travel time. In addition, each request is associated with a deadline, which is the time-bound for receiving the result. In practice, this time-bound could be set by either the drivers (e.g. maximum time a driver is willing to wait) or the system (e.g. the time to announce the result cannot be later than the expected arrival time of the

vehicle). Any request which cannot get the reservation before its deadline would be regarded as *expired* and removed from the pending queue.

IV. A NEW SCHEDULING ALGORITHM

In order to achieve the best utilization of the high-priority lane, we have designed an on-line scheduling algorithm based on the admission control policy. This algorithm determines which requests should be granted a reservation, as well as which tokens should be granted to the selected requests. Before going into the details of the new algorithms, we give several definitions.

Although the estimated arrival time (t_{in}) and departure time (t_{out}) are specified for each requested segment, in reality, it is unlikely that every vehicle will be able to arrive at the reserved segment on time. Therefore, in order to tolerate early and late arrival cases to some extent, we introduce the notion of a *time window*.

Definition 4.1: Time window: Given a specified t_{in} , a vehicle is allowed to enter the reserved segment between $t_{in} - \omega$ and $t_{in} + \omega$, where ω is the time window.

Note that given a time window ω , the time slot to be reserved is extended to $[t_{in} - \omega, t_{out} + \omega]$. The larger the value of ω is, the more tolerant the system is to the early/late arrival rate. However, disadvantage is that a longer time slot has to be reserved for each segment, which may decrease the utilization efficiency of the high-priority lane.

Since the time slots of a token granted to different requests cannot be overlapped with each other, we define the *available token* to a request as follows.

Definition 4.2: A token T_{ij} is considered available to a request Q_m asking for segment s_i in time slot $[t_{in}(s_i), t_{out}(s_i)]$ if for any time slot pair $[in_time, out_time]$ in T_{ij} 's reservation table, there is no overlap with the newly requested time slot $[t_{in}(s_i) - \omega, t_{out}(s_i) + \omega]$. For the requested segment s_i , the set of tokens available to Q_m is denoted by $T_i(Q_m)$, where $T_i(Q_m) \subseteq T_i$.

With the above definition, we can determine whether a segment is available to a request.

Definition 4.3: Available segments: Given a request Q_m which requires segment s_i , s_i is considered to be available to Q_m if it has at least one available token, namely, $T_i(Q_m) \neq \emptyset$. The set of available segments to Q_m is denoted by $S(Q_m)$, where $S(Q_m) \subseteq Q_m$.

If a request asks for multiple segments, it is possible that not all of the requested segments are available. We define the *service ratio* to a request as follows.

Definition 4.4: For a request Q_m , where $|Q_m|$ is the total number of requested segments, and $S(Q_m)$ is the number of available segments, the request service ratio is defined as the number of available segments over the total number of required segments, which is calculated by $SR_{Q_m} = \frac{|S(Q_m)|}{|Q_m|}$.

Using SR_{Q_m} as the sole metric to evaluate lane utilization may result in poor performance. For instance, requests that only ask for a small number of segments are likely to have higher service ratio. Nevertheless, systems that choose to prioritize such small requests can hardly achieve the best




Fig. 2. Infeasible time slots of token T_i caused by a new reservation

utilization of the entire high-priority lane. We introduce the following metrics in order to enhance lane utilization.

An unreserved time slot can be granted to a new request without overlapping with other slots only when the length of an unreserved time slot is at least $travel_time + 2 \cdot \omega$, where $travel_time = t_{out} - t_{in}$. We use this observation to define an infeasible time slot of a token as follows.

Definition 4.5: Infeasible time slot of a token: A time slot is considered infeasible if its length is less than $travel_time_i + 2 \cdot \omega$, where $travel_time_i = t_{out}(s_i) - t_{in}(s_i)$.

The infeasible time slots of a token are caused by new reservations, which divide the unreserved time slots of a token into smaller pieces. The pieces shorter than $travel_time_i + 2 \cdot \omega$ become infeasible. Figure 2 shows such an example. Suppose the time slots $[a, b]$ and $[c, d]$ have been reserved for token T_i , while time slot $[b, c]$ is unreserved. The length of $[b, c]$ satisfies: $(travel_time_i + 2 \cdot \omega) < c - b < 2 \cdot (travel_time_i + 2 \cdot \omega)$. When a new reservation is made on $[b, c]$, infeasible time slots will be created as shown in Figure 2.

Based on this example, we can conclude that in order to achieve better lane utilization, the scheduling algorithm should select tokens, such that the length of the newly generated infeasible time slots remains minimal. Accordingly, we define the *most fitting token* as follows.

Definition 4.6: Most fitting token: Token T_{ij} is the most fitting token for a request Q_m asking for segment s_i if: 1) T_{ij} is among the set of available tokens $T_i(Q_m)$, and 2) the length of the newly caused infeasible slot(s) on T_{ij} is the minimum among those for all other tokens in $T_i(Q_m)$.

In addition to minimizing the number of the infeasible time slots of tokens, a scheduling algorithm is also expected to maximize the total length of the reserved time slots.

Definition 4.7: Request profit ratio: Given a set of available segments $S(Q_m)$ and the corresponding most fitting tokens for request Q_m , the profit ratio of Q_m is represented as follows: $PR_{Q_m} = \frac{\sum_{s_i \in S(Q_m)} R(T_{ij})}{\sum_{s_i \in S(Q_m)} R(T_{ij}) + \sum_{s_i \in S(Q_m)} I(T_{ij})}$, where T_{ij} is the most fitting token for s_i ; $R(T_{ij})$ represents the length of reserved time slots on T_{ij} by Q_m , while $I(T_{ij})$ represents the length of infeasible time slots on T_{ij} caused by Q_m .

Each submitted request is associated with a deadline, beyond which the request expires. The *slack time* is defined to represent the request urgency.

Definition 4.8: Request slack time: Given a request Q_m with absolute deadline, $submission_time_{Q_m} + \epsilon$, where ϵ is the time-bound for system response, the slack time of Q_m is calculated by $ST_{Q_m} = submission_time_{Q_m} + \epsilon - t$, where t is

TABLE I
DEFAULT SETTINGS

Parameter	Default	Description
N	10	Number of segments
Capacity	200	Capacity of each segment
Arrival Rate	20 vehicles/min	Mean vehicle arrival rate
		Exponential distribution
Travel Time	30 min	Travel time in each segment
Request Size	5	Average number of requested segments per vehicle

the current time.

Based on the above discussion, intuitively, a request is more rewarding if it has a higher service ratio, a higher profit ratio and less slack time. Accordingly, we define the priority of a request Q_m as follows.

Definition 4.9: Request priority: Given a request Q_m , its priority is calculated by: $Pri_{Q_m} = \frac{SR_{Q_m} \cdot PR_{Q_m}}{ST_{Q_m}}$

Given the above definitions, at each scheduling point, the procedure of the algorithm is presented as follows.

- 1) Determine if there is an available segment for any of the pending requests based on Def.4.1 ~ Def.4.3.
- 2) Calculate the priority of each request based on Def.4.4 ~ Def.4.9.
- 3) Select the request with the highest priority.
- 4) Select the most fitting token (Def.4.6) for each available segment and grant these tokens to the selected request.
- 5) Update the *reservation_table* of each granted token.
- 6) Remove the selected request from the pending queue.
- 7) Repeat step 1) ~ 6) until there are no pending requests or no available segments for any of the pending requests.

V. PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed admission control policy and the scheduling algorithm. The model is implemented by CSIM19 [19]. Table I shows the default parameter settings used to produce challenging conditions for the performance evaluation. Unless stated otherwise, the simulations are conducted under these default settings. Specifically, the vehicle arrival follows a Poisson process with mean arrival rate of 20 vehicles per minute. The average number of requested segments for each vehicle is 5. The deadline of a request is uniformly distributed between 1 and 10 minutes. For simplicity, we assume that each segment has the same capacity and the travel time of vehicles on each segment is the same. In practice, the estimation of travel time can be determined by the techniques discussed in [20] and [21]. Note that the high-priority lane is used only if there is a token which is in its BUSY status. The overall utilization of the high-priority lane is estimated using the average percentage of time tokens spend in their BUSY status. In this regard, we adopt the following metric for performance evaluation.

Lane Utilization Ratio: We denote the duration of time a token spends in its *IDLE* or *BUSY* status as t_{idle} and t_{busy} ,

respectively. The Lane Utilization Ratio (LUR) is defined as the percentage of time all tokens in the system spend in their BUSY status. It is evaluated with the formula: $LUR = \left(\sum_{i=1}^N \sum_{j=1}^{|T_i|} \frac{t_{Busy}}{t_{Idle} + t_{Busy}} \right) / \sum_{i=1}^N |T_i|$, where N is the total number of segments and $|T_i|$ is the number of tokens for segment s_i .

Different scheduling algorithms can be used as part of the token-based admission control policy. First-come first-served (FCFS) is included for comparison in the first set of simulation experiments shown in Figures 3~5. These experiments examine lane utilization under different system workloads. Two observations can be made based on the first set of simulation results. First, the proposed scheduling algorithm outperforms FCFS in all the examined scenarios in terms of maximizing the lane utilization ratio. Second, increase in the system workload, such as higher vehicle arrival rate, larger request size, or longer travel time, leads to higher lane utilization ratio. However, rather than growing linearly to the system workload, the system utilization grows slower. This implies that although high workload leads to better lane utilization, it might also cause higher number or rejected requests.

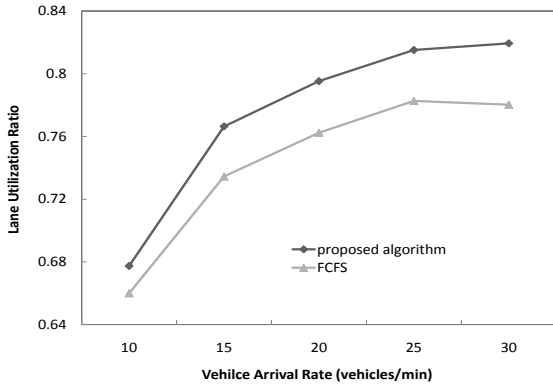


Fig. 3. Lane utilization ratio under different vehicle arrival rates

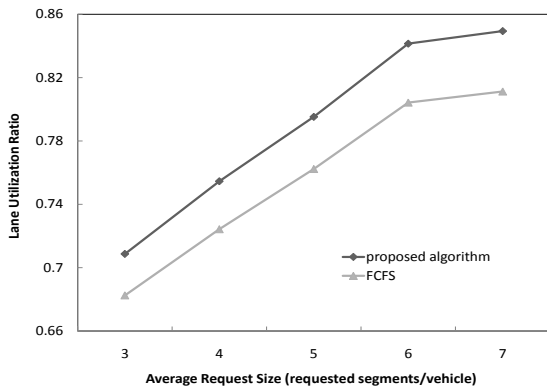


Fig. 4. Lane utilization ratio under different request size

Figures 6~8 show the results of the second set of experiments which examine the influence of exceptional cases, such as withdrawal of the reservation and early/late arrival,

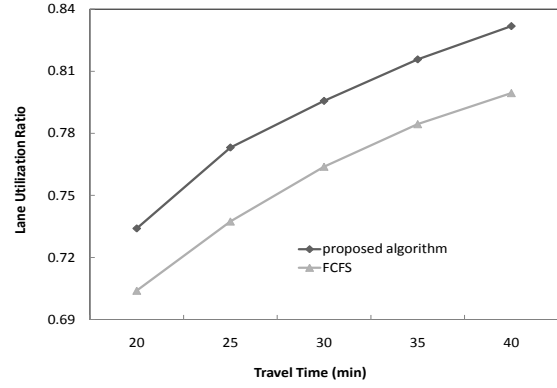


Fig. 5. Lane utilization ratio under different travel time

on the system performance. Although the system will revoke the granted time slots of tokens when a vehicle withdraws its reservation, the chance to re-grant these released time slots is slim because it is not likely to receive new requests asking for these released time slots. Therefore, high withdrawal probability results in poor performance as shown in Figure 6.

To alleviate the influence of reservation withdrawal, we investigate a scenario where the system attempts to negotiate the arrival time of vehicles within a given maximum adjustment range. Specifically, when there is not an available time slot for a specified arrival and departure time, the system will continue searching for other time slots. If another available time slot within the adjustable time range is found, the request will still get a reservation. The capability to adjust the arrival time within a certain range will improve the chances that a time slot will be utilized. Figure 7 shows the lane utilization ratio under different adjustment ranges while the withdrawal probability is kept at 15%. We can see that higher lane utilization can be achieved when the adjustment range is larger.

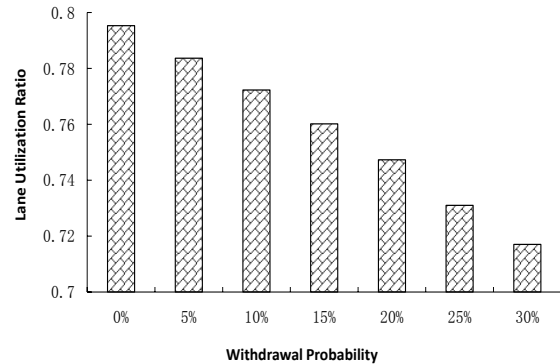


Fig. 6. Lane utilization ratio under different withdrawal probability

Last, we evaluate the effect of the time window, ω , on the system performance. In this simulation, the deviation of the actual arrival time follows an Exponential distribution with mean value of 5 (mins). Note that the length of the reserved time slot equals $travel_time + 2 \cdot \omega$, while the

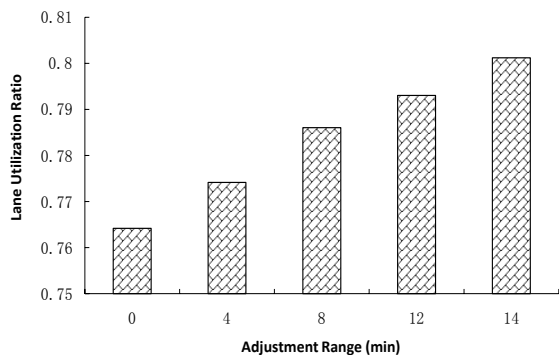


Fig. 7. Lane utilization ratio under different adjustment range




Fig. 8. Lane utilization ratio under different time window

token is actually utilized only for the duration of $travel_time$. Therefore, a large value of ω will cause inefficient lane utilization. On the other hand, a large value of ω provides high tolerance for the early/late arrival cases and increases the chance of a vehicle to enter the reserved segment. Consequently, there is a trade-off in setting the value of ω . The results in Figure 8 give an insight into this argument and demonstrate that a suitable value of ω is needed in order to achieve the best performance.

VI. CONCLUSIONS AND FUTURE WORK

The ever-growing traffic congestion problem results in a series of serious consequences, including transportation inefficiency, wasted fuel, and commuter frustration. The cutting-edge technologies in vehicular networks enable us to develop more sophisticated systems to improve the transportation efficiency. First, we propose a token-based admission control policy for a lane reservation system. Second, we have designed an on-line scheduling algorithm which enhances the system performance by identifying the most rewarding requests and granting the most fitting tokens to these requests. Last, the performance evaluation of our algorithm not only shows the effectiveness of the proposed solution, but also gives insights into future research on the lane reservation system.

There are several other interesting yet important problems related to the practicality and soundness of the lane reservation system to be explored in our future work. For example,

the social benefit should be further evaluated by examining the performance improvement of the proposed model compared to a regular system without roadway reservation options. Besides, more efforts are needed in the evaluation of user perspectives. The opinion of the users is especially important for this system since its development is not just a matter of designing rules, but also getting users to participate.

ACKNOWLEDGMENT

The work described in this paper was partially supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU 115407].

REFERENCES

- [1] D. Sperling and D. Gordon, *Two billion cars: driving toward sustainability*. Oxford University Press, USA, 2009.
- [2] J. Dahlgren, "High-occupancy/toll lanes: where should they be implemented?" *Transportation Research Part A: Policy and Practice*, vol. 36, no. 3, pp. 239–255, 2002.
- [3] X. Zhang and H. Yang, "The optimal cordon-based network congestion pricing problem," *Transportation Research Part B: Methodological*, vol. 38, no. 6, pp. 517–537, 2004.
- [4] M. Papageorgiou and A. Kotsialos, "Freeway ramp metering: An overview," in *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC'00)*, 2000, pp. 228–239.
- [5] L. Iftode, S. Smaldone, M. Gerla, and J. Misener, "Active highways (position paper)," in *IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'08)*, 2008.
- [6] N. Ravi, S. Smaldone, L. Iftode, and M. Gerla, "Lane reservation for highways (position paper)," in *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC'07)*, 2007, pp. 795–800.
- [7] D. Teodorovic and P. Edara, "Highway space inventory control system," *Transportation and Traffic Theory*, vol. 16, pp. 43–62, 2005.
- [8] P. Edara and D. Teodorovic, "Model of an advance-booking system for highway trips," *Transportation Research Part C: Emerging Technologies*, vol. 16, no. 1, pp. 36–53, 2008.
- [9] "Texas Transportation Institute's 2010 Urban Mobility Report," http://tti.tamu.edu/documents/mobility_report_2010.pdf, dec., 2010.
- [10] "Japan ITS Hand Book 2006-2007," <http://www.mlit.go.jp/road/ITS/index/indexHBook2006.html>.
- [11] S. Tsugawa and S. Kato, "Energy ITS: another application of vehicular communications," *IEEE Communications Magazine*, vol. 48, no. 11, pp. 120–126, 2010.
- [12] O. Servin, K. Boriboonsomsin, and M. Barth, "Preliminary design of speed control strategies in dynamic intelligent speed adaptation system for freeways," in *Transportation Research Board 87th Annual Meeting (TRB'08)*, no. 08-0378, 2008.
- [13] "Draft inventory of U.S. greenhouse gas emissions and sinks," <http://www.epa.gov/climatechange/emissions/usinventoryreport.html>, 2011 Draft U.S. Greenhouse Gas Inventory Report, Feb. 2011.
- [14] "FCC 03-024," Federal Communications Commission, FCC Report and Order, Feb. 2004.
- [15] "GSM world," <http://www.gsmworld.com>.
- [16] "CDMA development group," <http://www.cdg.org>.
- [17] "4G products, companies and resources," <http://www.wimax.com>.
- [18] Y. Kim and S. Kang, "An innovative traffic demand management strategy: expressway reservation system," in *Transportation Research Board 90th Annual Meeting (TRB'11)*, 2011.
- [19] H. Schwetman, "CSIM19: a powerful tool for building system models," in *Proceedings of the 33rd IEEE Winter Simulation Conference (WSC'01)*, 2001, pp. 250–255.
- [20] P. Mohan, V. Padmanabhan, and R. Ramjee, "Nericell: rich monitoring of road and traffic conditions using mobile smartphones," in *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems (SenSys'08)*, 2008, pp. 323–336.
- [21] A. Thiagarajan, L. Ravindranath, K. LaCurts, S. Madden, H. Balakrishnan, S. Toledo, and J. Eriksson, "VTrack: accurate, energy-aware road traffic delay estimation using mobile phones," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09)*, 2009, pp. 85–98.