Placement Optimization for Advertisement Dissemination in Smart City

Huanyang Zheng and Jie Wu, *Fellow, IEEE*

Abstract—This paper studies a promising application in Vehicular Cyber-Physical Systems (VCPS) called *roadside advertisement dissemination*. Its application involves three elements: the drivers in the vehicles, Roadside Access Points (RAPs), and shopkeepers. The shopkeeper wants to attract as many customers as possible by using RAPs to disseminate advertisements to the passing vehicles. Upon receiving an advertisement, the driver might detour towards the shop, depending on the detour distance. Given a fixed number of RAPs and the traffic distribution, our goal is to optimize the RAP placement for the shopkeeper to maximally attract potential customers. This application is a non-trivial extension of traditional coverage problems, the difference being that RAPs are used to cover the traffic flows. RAP placement algorithms pose complex trade-offs. If we place RAPs at locations that can provide small detour distances to attract more customers, these locations may not necessarily be located in heavy traffic regions. While heavy traffic regions cover more flows, they might cause large detour distances, making shopping less attractive to customers. To balance the above trade-off, bounded RAP placement algorithms are proposed with respect to submodular and non-submodular scenarios. Since real-world traffic distributions exhibit unique patterns, here we further consider the Manhattan grid scenario and propose improved solutions. Extensive real trace-driven experiments validate the competitive performances of the proposed algorithms.

Index Terms—Vehicular cyber-physical systems, placement problem, coverage problem, submodularity and optimization.

1 Introduction

Vehicular Cyber-Physical Systems (VCPS) refer to a new generation of vehicular systems; VCPS integrate computational and physical capabilities that can interact with humans through many new modalities [1]. While traditional vehicular systems were generally considered to be a common component of the physical world, VCPS actively interact with humans through communications which yield a very tight coordination between cyber and physical resources [2]. Therefore, it is critical for VCPS to consider the perceptions and reactions of humans (i.e., the drivers in the vehicles). The effectiveness and efficiency of VCPS depend on how humans could benefit from such a system [3, 4].

This paper addresses a novel and promising application in VCPS called *roadside advertisement dissemination* [5–7]. Its application involves three basic elements: the drivers in the vehicles (the human factor), Roadside Access Points (RAPs), and shopkeepers. The shopkeeper wants to attract as many customers as possible by using RAPs to disseminate electronic advertisements to passing vehicles. The drivers may decide to go shopping or not upon receiving advertisements, depending on the detour distance. An example is shown in Fig. 1(a), where commuters drive home after work. During their trip home, they receive an advertisement from an RAP, and then, decide to detour to the shop. We observe that the driver may not shop if the detour distance to the shop is too large. This is because the desire to shop does not outweigh the cost of the journey. If the shop is on the driver's way home, he or she may stop by due to convenience. We focus on the scenario with only one shop; however, multiple shops can be easily extended.

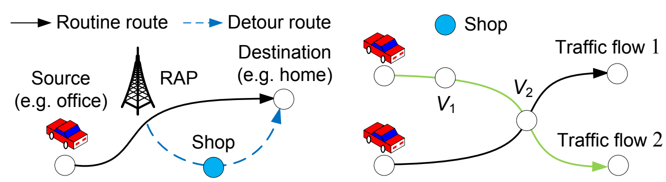

(a) The detour model. (b) The RAP placement.

Fig. 1. The scenario of the roadside advertisement dissemination.

Our roadside advertisement dissemination application is motivated by roadside e-displays. Although today's e-displays are popular, its effectiveness is not clear because of its inability to cater to individual's shopping preferences. By comparison, online targeted advertising [8] can deliver ads to attract customers based on their search/email contents, web browsing, etc, but it mainly targets people at home rather than on the road. Additionally, very little is known about how such targeted ads may affect people's online behaviors, let alone how targeted ads delivered to commuters may affect their traveling routes and activities. In comparison, our application avoids these drawbacks at a similar cost of e-displays.

Fig. 1(b) shows some traffic flows on the streets. Each traffic flow represents a group of commuters traveling home from work. Given a fixed number of RAPs and the traffic distribution that can be obtained from the previous records, we focus on optimizing the RAP placement for the shopkeeper, as to maximally attract potential customers. Our problem is a non-trivial extension of traditional coverage problems, the difference being that we use RAPs to cover the

- *H. Zheng and J. Wu are with the Department of Computer and Information Sciences, Temple University, Philadelphia, PA 19121, USA. E-mail: {huanyang.zheng, jiewu}@temple.edu*

Manuscript received April 19, 2017; revised August 26, 2017.

traffic flows. A trade-off exists between the traffic density and the detour probability for the RAP placement problem. Let us consider the placement of only one RAP in Fig. 1(b). If we place the RAP at v_1 near the shop, then this RAP can only cover traffic flow 2, leaving traffic flow 1 not covered. On the other hand, if we place the RAP at v_2, although both traffic flows 1 and 2 are covered, the nearby drivers are not likely to go shopping due to a large detour distance. Our problem becomes more challenging when placing multiple RAPs. Let us consider the placement of two RAPs in Fig. 1(b). Suppose these two RAPs are placed at v_1 and v_2, respectively. Then, the RAP at v_2 is meaningless for the drivers in traffic flow 2, since the RAP at v_1 provides a smaller detour distance for those drivers. Redundant advertisements do not provide additional shopping incentives. If a driver decides not to go shopping despite a smaller detour distance, they would not go shopping with a larger detour distance. Accordingly, the geographical distribution of RAPs should also be controlled.

Furthermore, the real-world traffic distributions exhibit certain patterns, which can be utilized for the RAP placement. For example, the streets in Manhattan are mapped as a grid, meaning that all the vehicles have only four possible moving directions. This means that the RAP placement is more controllable. Another interesting observation is that multiple shortest paths exist, connecting a pair of locations in the Manhattan grid streets. These properties pose several unique challenges regarding RAP placement optimization.

Our main contributions are summarized as follows:

- We address a novel and promising application called roadside advertisement dissemination, which could follow the design principle of VCPS. It is a non-trivial extension of traditional coverage problems.
- Several utility functions are used to model the driver's detour probability. Three greedy solutions with different approximation ratios are proposed for those utility functions, respectively.
- Since the real-world traffic distributions exhibit certain patterns, we further study the RAP placement problem in the Manhattan grid street scenario, where we propose solutions with tightened bounds.
- Extensive experiments are conducted to evaluate the proposed solutions. The results are provided from different perspectives for insightful conclusions.

The remainder of this paper is organized as follows. Section 2 mainly surveys the related work. Section 3 describes the model and formulates the problem. Section 4 analyzes the problem. Section 5 gives out several greedy solutions. Section 6 discusses the Manhattan grid scenario. Section 7 includes the experiments. Finally, Section 8 concludes the paper and suggests future works.

2 Related Work

Cyber-physical systems are engineered systems whose operations are monitored, coordinated, controlled, and integrated by a computing and communication core [9, 10]. VCPS are special types of cyper-physical systems designed for vehicles. While traditional protocols are imperceptible by humans, VCPS take humans' perceptions into account [11]. For example, the data dissemination mechanism in [12] considers the data to be location-dependent, but humans are not aware of this mechanism. In contrast, Li et al. [13] considered a human-oriented service scheduling in VCPS, where a driver cannot receive multiple services in a short time. Their goal is to deliver a limited number of services, each having a time-dependent (and non-increasing) utility, to a subset of intended drivers so as to minimize the system-wide total utility loss due to unsuccessful service deliveries. Wagh et al. [14] proposed that the data composition in VCPS should be flexible for drivers. Given a limited number of data elements that can be transmitted, their objective is to maximize the system-wide utility at the receiver by choosing to transmit data elements with different utilities.

The practical implementation of VCPS usually involves RAPs in addition to the Dedicated Short Range Communications (DSRC) / IEEE 802.11p units within the vehicles [15]. These RAPs can supplement a sparse network in a low-density scenario, and help coordinate and move data in dense scenarios. The drawbacks are that RAPs are extremely costly in the past decades. A 2012 industry survey by Michigans DoT and the Center for Automotive Research reiterated that "one of the biggest challenges respondents see to the broad adoption of connected vehicle technology is funding for roadside infrastructure." [15]. However, recent technology developments bring down the cost of RSUs, and thus, a VCPS implementation becomes acceptable. The latest DSRC and Wireless Access Vehicular Environments (WAVE) are based on IEEE 802.11p PHY/MAC, DSRC wireless communication, and messaging protocols. For example, Redpine Signals provides fully certified 802.11p DSRC/WAVE wireless module along with complete software solution including 1609.x/DSRC stack. The RSU module from Redpine Signals only costs about $30 [16].

Currently, the advertisement dissemination is considered as a novel and promising application in VCPS [17], since advertisements belong to the practically useful data. While traditional studies focus on online advertisements [18–20], Li et al. [8] first considered the advertisement dissemination in VCPS as a bandwidth allocation problem with pre-fixed locations of RAPs. We optimize the RAP placement for the advertisement dissemination. This application is also studied from different perspectives. Shen et al. [21] studied the message authentication problem for safety advertisements. A cooperative message authentication protocol is developed to alleviate vehicles' computation burden. All the vehicles would share their verification results with each other in a cooperative way, so that the number of safety messages that each vehicle needs to verify reduces significantly. Liu et al. [22] presented a study on real-time data services via roadside-to-vehicle communication by considering both the time constraint of data dissemination and the freshness of data items. A temporal data dissemination problem is proposed by introducing the snapshot consistency requirement on serving real-time requests for temporal data items.

Our RAP placement problem also relates to the existing maximum coverage problem [23–25], since RAPs are used to cover the passing vehicles. The maximum coverage problem is based on some sets defined over a domain of elements associated with weights. Its objective is to select k sets, such that the total weight of elements within the selected sets is maximized. However, our problem cannot be solved by

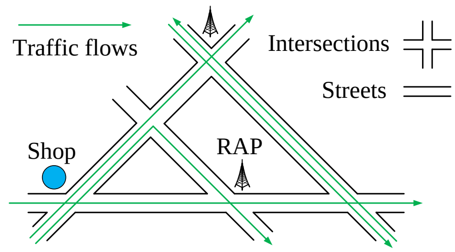

Fig. 2. The scenario for the RAP placement problem.

existing techniques, since the detour distance is considered. The location chosen for a driver to receive the advertisement is critical for the driver's detour decision. Our problem is a variation of the maximum coverage problem, in terms of the detour distance. The maximum coverage problem has been well studied through a greedy approximation algorithm that guarantees a ratio of $1-\frac{1}{e}$. However, it requires the coverage function to be submodular [26] and fails to deal with non-submodular coverage functions [27, 28].

Our problem brings more unique challenges, since we focus on both submodular coverage and non-submodular coverage. Note that the problem of non-submodular function maximization [29] has not been perfectly solved in the literature [30]. This is because certain properties of the objective function are required to design approximation algorithms. Although the problem of supermodular function maximization can be optimally solved by the minimum-norm-point algorithm [31], non-submodular functions are not the same. The latest approach is based on the curvature [32], which typically assumes that the marginal gain of the non-submodular function varies within a given curvature. This approach is based on modifications of the continuous greedy algorithm and non-oblivious local search, and allows us to approximately maximize the sum of a nonnegative, nondecreasing submodular function and a (possibly negative) linear function. Meanwhile, it has been proved that these approximation results are the best possible in the value oracle model, even in the case of a cardinality constraint [33].

3 Model and Formulation

This section starts with the graph model and the driver attraction model. The problem is also formulated.

3.1 Graph and Notations

As illustrated in Fig. 2, our advertisement dissemination scenario is based on a directed graph $G=(V,E)$, where $V=\{v_i\}$ is a set of nodes (i.e., street intersections), and $E\subseteq V^2$ is a set of directed edges (i.e., streets). Some traffic flows exist on the streets. We assume that all cars start from and stop at intersections. Let $T_{i,j}$ denote the traffic flow from intersections v_i to v_j (e.g., vehicles that return home from the office). The traveling path for $T_{i,j}$ is unique and is known a priori (a shortest path in general). If $T_{i,j}$ goes through an intersection v, we say $v\in T_{i,j}$. Let $n_{i,j}$ denote the number of drivers in $T_{i,j}$. $T=\{T_{i,j}\}$ is the set of traffic flows that are targeted for the advertisement dissemination, and $|T|$ is the number of traffic flows in T. Traffic flows with insufficient potential customers are not counted in T. Let S denote the set of street intersections with placed RAPs. S is the only variable in this paper.

3.2 Driver Attraction Model

This subsection describes the driver attraction model, which assumes that repeated advertisements do not provide additional shopping incentives. Drivers will make a decision of whether or not to shop upon first receiving the advertisement, depending on the detour distance. We start with the scenario with only one shop. For each traffic flow, $T_{i,j}\in T$, its detour distance is denoted by $d_{i,j}(S)$. As shown in Fig. 3, suppose that $T_{i,j}$ goes through an RAP, say $v\in S$. When the driver receives the advertisement, the shortest path distance from the current location to the shop is d'_v, from the shop to the destination v_j is d''_v, and from the current location to the destination v_j is d'''_v. The detour distance is calculated as:

$$d_{i,j}(S)=\begin{cases}\min\limits_{v\in S, v\in T_{i,j}} d'_v+d''_v-d'''_v & \text{if } \exists\, v, v\in S, v\in T_{i,j}\\ \infty & \text{otherwise}\end{cases} \quad (1)$$

Once the traffic flow $T_{i,j}$ goes through multiple RAPs, then the detour distance is the minimum detour distance among all these RAPs. Interestingly, the RAP, which provides the minimum detour distance, is the first RAP met by the driver. It is consistent that drivers will make a decision of whether or not to shop upon first receiving the advertisement:

Theorem 1. For a specified traffic flow $T_{i,j}$ that goes from v_i to v_j, the first RAP on its path always provides a smaller detour distance than all the other RAPs on the path.

Proof: As shown in Fig. 4, let us select two arbitrary RAPs (say v and u) on the traffic flow $T_{i,j}$. The difference between the detour distances of v and u is:

$$[d'_v+d''_v+d'''_v]-[d'_u+d''_u+d'''_u]=d'_v-\{d'_u+[d'''_v-d'''_u]\}<0 \quad (2)$$

In Eq. 2, $\{d'_u+[d'''_v-d'''_u]\}$ is the distance from v to the shop via u. Since d'_v is the shortest path from v to the shop, it is smaller than $\{d'_u+[d'''_u-d'''_u]\}$. Hence, the detour distance at v is smaller than the detour distance at u. Since v and u are arbitrarily selected, the proof completes. ■

Let $f_{i,j}(S)$ denote the number of drivers attracted from $T_{i,j}$, when RAPs are placed at S. This paper models $f_{i,j}(S)$ as a utility function with respect to $d_{i,j}(S)$ and some constants. Note that only $d_{i,j}(S)$ in $f_{i,j}(S)$ depends on S. An example of $f_{i,j}(S)$ is the threshold utility function:

$$f_{i,j}(S)=\begin{cases}\alpha_{i,j}\times n_{i,j} & \text{if } d_{i,j}(S)\le D\\ 0 & \text{otherwise}\end{cases} \quad (3)$$

$\alpha_{i,j}$ is a constant advertisement attractiveness for the drivers in the traffic flow $T_{i,j}$ ($0\le\alpha_{i,j}\le 1$). It is known a priori through previous statistics. $n_{i,j}$ is the number of drivers in $T_{i,j}$. D is a constant threshold. Another example of $f_{i,j}(S)$ is the decreasing utility function:

$$f_{i,j}(S)=\begin{cases}\alpha_{i,j}\times n_{i,j}\times[1-d_{i,j}(S)/D] & \text{if } d_{i,j}(S)\le D\\ 0 & \text{otherwise}\end{cases} \quad (4)$$

Note that our driver attraction model can be extended to scenarios with multiple shops. For those cases, the result depends on the shop that provides the smallest detour distance among all the shops. We do not consider the commercial competition among different shops, in terms of attracting the drivers. All shops belong to the same shopkeeper.

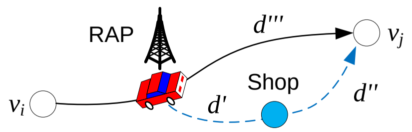

Fig. 3. The calculation of the detour distance.

TABLE 1
All notations.

$G = (V, E)$	Graph of intersections and streets
$V = v_i$	Set of intersections (v_i denotes i-th intersection)
E	Set of directed streets connecting intersections
$T_{i,j}$	Traffic flow that goes from intersections v_i to v_j ($v \in T_{i,j}$ if $T_{i,j}$ goes through intersection v)
$n_{i,j}$	Number of drivers in $T_{i,j}$ as potential customers
$T = \{T_{i,j}\}$	Set of traffic flows with $\lvert T\rvert$ to be its cardinality
S	Set of intersections with placed RAPs with $\lvert S\rvert$ to be its cardinality and S' to be its subset
$d_{i,j}(S)$	Detour distance for $T_{i,j}$ under S (d'_v, d''_v, and d'''_v are parameters to compute it)
$f_{i,j}(S)$	Number of drivers attracted from $T_{i,j}$ under S (e.g., threshold or decreasing utility function)
$\alpha_{i,j}$	Constant advertisement attractiveness for $T_{i,j}$
k	A given bound for number of RAPs
$f(S)$	Number of attracted drivers as our objective
M_v	Modularity set for v, i.e., all intersections that might increase the marginal gain of v for $f(S)$
M'_v	Subset of M_v used in algorithms and proofs
Δ	Modularity computed by the maximum cardinality among all modularity sets
Δ'	A parameter that satisfies $1 \le \Delta' \le \Delta$

3.3 Problem Formulation

To attract customers, the shopkeeper places, at most, k RAPs at street intersections for the advertisement dissemination. Since the set of street intersections with placed RAPs is S, we have $|S| \le k$. A placed RAP will send electronic advertisements to all passing vehicles. Since $f_{i,j}(S)$ is the number of attracted drivers from $T_{i,j}$, let $f(S) = \sum_{T_{i,j}\in T} f_{i,j}(S)$ be the total number of attracted drivers from all traffic flows. This paper studies the RAP placement problem. Given a fixed number of RAPs and the traffic distribution, our goal is to optimize the RAP placement for the shopkeeper to maximally attract potential customers:

$$\text{maximize} \quad f(S) = \sum_{T_{i,j}\in T} f_{i,j}(S)$$
$$\text{s.t.} \quad |S| \le k \ \text{ and } \ S \subseteq V$$

The RAP placement problem may pose complex trade-offs. If we place RAPs at locations that can provide small detour distances to attract more customers, these locations may not necessarily be located in heavy traffic regions. While heavy traffic regions cover more flows, they can cause large detour distances, making shopping less attractive to customers. As a result, the relationship between $f_{i,j}(S)$ and $d_{i,j}(S)$ is critical. We will analyze this relationship in the next section. Finally, all notations are listed in Table 3.3.

4 Problem Analysis

We start with the problem hardness:

Theorem 2. The RAP placement problem is NP-hard.

Proof: We prove by reduction from the weighted maximum coverage problem [24], which is NP-complete. The maximum coverage problem is based on some sets defined over a domain of elements associated with weights. Its objective is to select k sets, such that the total weight of elements within the selected sets is maximized. Let us consider the RAP placement problem, using the threshold utility function in Eq. 3. In this case, the elements reduce to the traffic flows, and the sets reduce to the intersections with placed RAPs. The weight of a traffic flow is its number of expected drivers that detour to the shop, if the detour distance is smaller than D. The selection of a set corresponds to the placement of an RAP. As a result, the maximum coverage problem reduces to the RAP placement problem, and the proof completes. ∎

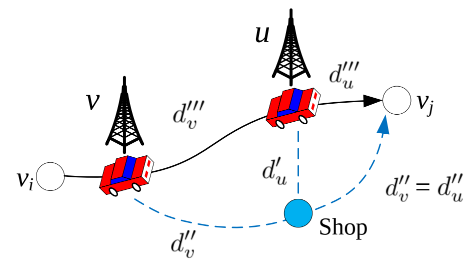

Fig. 4. An illustration for the proof of Theorem 1.

Theorem 2 shows that the RAP placement problem is essentially a variation of the maximum coverage problem, under a certain relationship between $f_{i,j}(S)$ and $d_{i,j}(S)$. To further analyze the RAP problem, we show that $d_{i,j}(S)$ is non-increasing with respect to S:

Theorem 3. For any given traffic flow $T_{i,j}$, its detour distance $d_{i,j}(S)$ is non-increasing with respect to S, meaning that $d_{i,j}(S) \le d_{i,j}(S')$ for $\forall S' \subseteq S \subseteq V$.

Proof: This Theorem is straightforward based on Theorem 1: for a specified traffic flow $T_{i,j}$, the first RAP on its path always provides a smaller detour distance than all the other RAPs on the path. Let v be the first RAP in S for $T_{i,j}$. Since $S' \subseteq S$, we have $v \in S'$. However, v may not be the first RAP in S' for $T_{i,j}$. This is because $S'\backslash S$ may include another RAP that is before v. As a result, $d_{i,j}(S) \le d_{i,j}(S')$ for $\forall S' \subseteq S$, and the proof completes. ∎

The insight of Theorem 3 is that more RAPs will never bring a larger detour distance. This property leads to another property with respect to our objective function:

Theorem 4. If $f_{i,j}(S)$ is non-increasing with respect to the detour distance $d_{i,j}(S)$, then $f_{i,j}(S)$ is submodular with respect to S: $f_{i,j}(S\cup\{v\}) - f_{i,j}(S) \le f_{i,j}(S'\cup\{v\}) - f_{i,j}(S')$ for $\forall v \in V$ and $\forall S' \subseteq S \subseteq V$.

Proof: Since $S' \subseteq S$, we have $d_{i,j}(S) \le d_{i,j}(S')$ according to Theorem 3. Since $f_{i,j}(S)$ is non-increasing with respect to the detour distance $d_{i,j}(S)$, we have $f_{i,j}(S) \ge f_{i,j}(S')$. We prove Theorem 4 by exhausting all cases:

- In the first case, v is the first RAP in both $S \cup \{v\}$ and $S' \cup \{v\}$ for $T_{i,j}$. According to Theorem 1, we have $d_{i,j}(S\cup\{v\}) = d_{i,j}(S'\cup\{v\})$ and consequently $f_{i,j}(S\cup\{v\}) = f_{i,j}(S'\cup\{v\})$. Since $f_{i,j}(S) \ge f_{i,j}(S')$, Theorem 4 holds.
- In the second case, v is the first RAP in $S' \cup \{v\}$ for $T_{i,j}$, but not the first RAP in $S \cup \{v\}$ for $T_{i,j}$. We

Algorithm 1 Naive Greedy

Input: The directed graph G; the set of traffic flows T; The number of RAPs to place (i.e., k).
Output: The RAP placement.

1: Initialize $S = \emptyset$.
2: **for** $i = 1$ to k **do**
3: Find $v = \arg\max_{v \in V} f(S \cup \{v\}) - f(S)$.
4: Update $S = S \cup \{v\}$.
5: **return** S as the RAP placement.

have $d_{i,j}(S' \cup \{v\}) \leq d_{i,j}(S')$ and $d_{i,j}(S \cup \{v\}) = d_{i,j}(S)$. Since $f_{i,j}(S)$ is non-increasing with respect to $d_{i,j}(S)$, we have $f_{i,j}(S' \cup \{v\}) \geq f_{i,j}(S')$ and $f_{i,j}(S \cup \{v\}) = f_{i,j}(S)$. Theorem 4 holds.

- In the third case, v is not the first RAP in both $S \cup \{v\}$ and $S' \cup \{v\}$ for $T_{i,j}$. We have $d_{i,j}(S' \cup \{v\}) = d_{i,j}(S')$ and $d_{i,j}(S \cup \{v\}) = d_{i,j}(S)$. As a result, we have $f_{i,j}(S' \cup \{v\}) \leq f_{i,j}(S')$ and $f_{i,j}(S \cup \{v\}) = f_{i,j}(S)$. Theorem 4 also holds.

Since Theorem 4 holds in all cases, the proof completes. ■

The insight of Theorem 4 is the diminishing return of the RAP placement. If a smaller detour distance can always attract more drivers, then the marginal gain of one more RAP decreases with respect to the number of existing RAPs. Examples of non-increasing functions include Eqs. 3 and 4. Since the summation of submodular functions is also submodular, we have:

Corollary 1. If $f_{i,j}(S)$ is non-increasing with respect to the detour distance $d_{i,j}(S)$ for any given $T_{i,j}$, then $f(S) = \sum_{T_{i,j} \in T} f_{i,j}(S)$ is submodular.

A notable point is that a smaller detour distance is likely to attract more drivers, but the assumption of the non-increasing function does not always hold. For example, if the detour distance is 0, the driver may not go to the shop, since he or she can go to the shop later as a weekly routine, ignoring the advertisement.

5 RAP Placement Solutions

This section shows solutions to the RAP placement problem, in terms of both submodular and non-submodular objective functions. Three approximation algorithms are proposed to deal with the RAP placement problem.

5.1 Naive Greedy Maximization

This subsection describes a naive greedy approach to solve the RAP placement problem. We start with two definitions:

Definition 1. An intersection *includes* a traffic flow, if an RAP placed at this intersection can attract more than zero drivers from this traffic flow to the shop.

Definition 2. An RAP *covers* a traffic flow, if this RAP is placed at an intersection that includes this traffic flow.

The naive greedy algorithm is shown in Algorithm 1. In line 1, it initializes $S = \emptyset$ as the set of intersections for the RAP placement. Lines 1 to 4 include k iterations to select the intersections one by one. In each iteration, the intersection v, which greedily maximizes $f(S \cup \{v\}) - f(S)$, is selected for

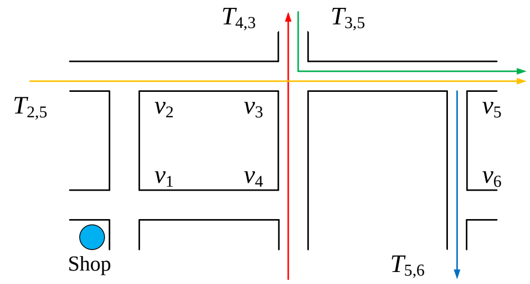

Fig. 5. An example of the RAP placement ($k = 2$ and $D = 6$).

RAP placement (lines 2 and 3). After k iterations, selected intersections are turned in line 4. The algorithm complexity is $O(|V|^3 + k|V||T|)$, where k is the number of RAPs, $|V|$ is the total number of intersections, and $|T|$ is the total number of traffic flows. $O(|V|^3)$ results from the calculation of detour distances, since we need to calculate the shortest paths between all pairs of intersections. $O(k|V||T|)$ comes from the greedy approach. Each greedy step takes $O(|V||T|)$ to examine all the intersections. Examining each intersection takes $O(|T|)$ to check all the traffic flows.

For a better explanation, an example is shown in Fig. 5, where we have two RAPs ($k = 2$) to place. This example is based on the threshold utility function in Eq. 3 with $D = 6$. The shop is located at v_1. $\alpha_{i,j}$ in Eq. 3 is set to be 1 for all the traffic flows. Meanwhile, we have $T_{2,5} = T_{4,3} = 6$, $T_{3,5} = 3$, and $T_{5,6} = 1$. This example involves four traffic flows,. In the first iteration, v_3 is picked to place an RAP, since it can attract the maximum amount of drivers from uncovered traffic flows ($n_{2,5}+n_{3,5}+n_{4,3}$=15). The remaining uncovered traffic is $T_{5,6}$. Consequently, the second RAP is placed at v_5 to cover $T_{5,6}$. Algorithm 1 terminates for this example, since $k = 2$. Note that v_6 does not include $T_{5,6}$, since its detour distance is 8 (the path changes from v_5v_6 to $v_5v_6v_5v_3v_2v_1v_2v_3v_5v_6$). Since the detour distance is larger than the threshold D, the driver would not detour to the shop, upon receiving the advertisement at v_6.

We have the following property for Algorithm 1:

Theorem 5. If $f_{i,j}(S)$ is non-increasing with respect to $d_{i,j}(S)$ for any given $T_{i,j}$, then Algorithm 1 can guarantee an approximation ratio of $1 - \frac{1}{e}$ to the optimal algorithm.

Theorem 5 is from the literature, i.e., the submodular function maximization [26]. If $f_{i,j}(S)$ is non-increasing with respect to $d_{i,j}(S)$, $f(S)$ becomes submodular, meaning that a greedy marginal maximization algorithm has an approximation ratio of $1 - \frac{1}{e}$. Algorithm 1 is essentially a greedy marginal maximization algorithm, which could be applied to other driver attraction models that are submodular. For example, if we allow drivers to be influenced by subsequent advertisements other than the first one, Algorithm 1 still works as long as the objective function $f(S)$ is submodular. However, the performance of Algorithm 1 is not guaranteed once $f(S)$ is not submodular. The next three subsections will discuss more general solutions for a non-submodular $f(S)$.

5.2 Generalized Approach and Modularity

Theorem 5 points out that the naive greedy algorithm can guarantee an approximation ratio of $1 - \frac{1}{e}$, when $f_{i,j}(S)$

Algorithm 2 Improved Greedy (IG)

Input: The directed graph G; the set of traffic flows T; The number of RAPs to place (i.e., k).

Output: The RAP placement.

1: Initialize $i = 0$ and $S_0 = \emptyset$.
2: **while** $|S_i| < k$ **do**
3: Find out $\arg\max_{v \in V, M'_v \subseteq M_v} f(S_i \cup \{v\} \cup M'_v) - f(S_i)$ constrained by $|S_i \cup \{v\} \cup M'_v| \leq k$.
4: Update $S_{i+1} = S_i \cup \{v\} \cup M'_v$.
5: Update i to be $i+1$.
6: **return** $S = S_i$ as the RAP placement.

Algorithm 3 Capped Greedy (CG)

Input: The directed graph G; the set of traffic flows T; The number of RAPs to place (i.e., k).

Output: The RAP placement.

Initialize $S = \emptyset$.
for each intersection $u \in V$, each Δ' from 1 to Δ, and each set $S_0 \subseteq M_u$ with $|S_0| = k \mod (\Delta' + 1)$ **do**
  Initialize $i = 0$.
  while $|S_i| < k$ **do**
    Find $\arg\max_{v \in V, M'_v \subseteq M_v} f(S_i \cup \{v\} \cup M'_v) - f(S_i)$ constrained by $|S_i \cup \{v\} \cup M'_v| \leq k$ and $M'_v \leq \Delta'$.
    Update $S_{i+1} = S_i \cup \{v\} \cup M'_v$.
    Update i to be $i+1$.
  if $f(S_i) > f(S)$ **then**
    Update S to be S_i.
return S as the RAP placement.

is non-increasing with respect to $d_{i,j}(S)$. The assumption of the non-increasing function means that a smaller detour distance should attract more drivers, leading to the submodularity of $f(S)$. However, this assumption may not always hold. For example, if the detour distance is 0, the driver may not go to the shop, since he or she can go to the shop later as a weekly routine, ignoring the advertisement. Therefore, we need a more generalized approach for the RAP placement problem, when $f(S)$ is not submodular.

In the following part of this section, we assume that $f(S)$ is monotone, but may not be submodular. The monotonicity means that $f(S) \geq f(S')$ for $\forall S' \subseteq S \subseteq V$, i.e., more RAPs attract more drivers. To analyze the non-submodularity, we introduce Feldman's approach and algorithms [29]:

Definition 3. Given a monotone objective function $f(\cdot)$, the modularity set of an intersection v is defined as $M_v = \{u \,|\, f(S \cup \{u, v\}) - f(S \cup \{u\}) > f(S \cup \{v\}) - f(S)$ for $\exists\, v \in V, u \in V, S \subseteq V\}$, which includes all intersections that might increase the marginal gain of v.

Definition 4. The modularity, Δ, is the maximum cardinality among all modularity sets, i.e., $\Delta = \max_v |M_v|$.

For an intersection v, only intersections in M_v might increase the marginal gain of v for the objective function $f(\cdot)$. In contrast, intersections that are not in M_v can never increase the marginal gain of v. If v is locally submodular for $f(\cdot)$, then $M_v = \emptyset$. The modularity Δ measures the degree to which $f(\cdot)$ violates the submodularity. Note that $f(\cdot)$ gets closer to the submodularity for a smaller Δ, and is submodular when $\Delta = 0$. In the following part of this section, we assume Δ to be a small constant. This is because a small detour distance is likely to attract more drivers through common sense and there will be few violations.

5.3 Improved Greedy

This subsection extends the naive greedy algorithm to tolerate the non-submodularity of $f(S)$. The key idea is that, when the intersection v is selected into S for the RAP placements, related intersections in M_v should be further considered, since they can improve v's marginal gain. Instead of selecting one intersection at each greedy iteration, we can select a small set of intersections at each greedy iteration to optimize the RAP placement. This observation can improve Algorithm 1 for a non-submodular $f(S)$.

Algorithm 2 is proposed as another greedy algorithm. It extends Algorithm 1 by considering place a set of RAPs instead of only one RAP in each greedy iteration. In line 1, it initializes $i = 0$ and $S_0 = \emptyset$. Lines 2 to 5 describe greedy iterations. While Algorithm 1 iteratively selects one intersection, Algorithm 2 iteratively selects a set of intersections, in order to mitigate the negative impact resulting from the non-submodularity. In line 3, once v is selected into S for the RAP placement, partial intersections in M_v (denoted as M'_v) are jointly selected into S. The greedy criterion is that v and M'_v together can maximize the marginal gain of the current intersections for the RAP placement, i.e., maximize $f(S_i \cup \{v\} \cup M'_v) - f(S_i)$. The constraint is that $|S_i \cup \{v\} \cup M'_v| \leq k$, i.e., at most k intersections are selected for the RAP placement. Lines 4 and 5 update the set S_i and the index i. The greedy iteration terminates once k intersections are selected. Finally, the best S_i is returned.

Computing $f(S)$ for a given S takes $O(k|T|)$. This is because we need to exhaust each traffic flow for each intersection in S. Algorithm 2 has at most k greedy iterations, and each iteration it exhausts v and M_v in line 3. Note that a time complexity of $O(2^{\Delta})$ is needed to exhaust v and M_v. Consequently, the time complexity of Algorithm 2 becomes $O(|V|^3 + 2^{\Delta} k^2 |V||T|)$. According to [29], Algorithm 2 is bounded with respect to the modularity Δ:

Theorem 6. Algorithm 2 has an approximation ratio of $\frac{1}{\Delta+2}$ to the optimal algorithm.

The key insight of Theorem 6 is that Algorithm 2 ignores at most $1 + \Delta$ optimal intersections at each greedy iteration, leading to a bound. However, Algorithm 2 still has critical drawbacks with respect to Δ. Once Δ becomes large, the performance of Algorithm 2 becomes poor, in terms of the time complexity and the approximation ratio. For example, even if $\Delta = 10$, the approximation ratio becomes $\frac{1}{12}$, which is fairly small and useless in practice. Meanwhile, the time complexity of Algorithm 2 grows exponentially with respect to Δ. Moreover, when $\Delta = 0$ and $f(S)$ becomes submodular, Algorithm 2 has a worse approximation ratio than Algorithm 1, i.e., $\frac{1}{2} < 1 - \frac{1}{e}$. These drawbacks come from the fact that Algorithm 2 may select an overly large set of intersections in each greedy iteration. The next subsection will improve Algorithm 2 by considering this important observation through a designed cap.

5.4 Capped Greedy

Since Δ has a critical impact on Algorithm 2, we need to further identify its role in the algorithm design. The key idea is that, although $|M_v|$ could be large, not all intersections in M_v have huge impacts on the marginal gain of v for $f(\cdot)$. Intuitively, we only need to consider the important intersections in M_v for v's marginal gain. Meanwhile, the optimal set of intersections for the RAP placement are not able to include all intersections in M_v if $|M_v| > k$. Therefore, capping the number of selected intersections in M_v might lead to a better performance, since low impact intersections in M_v can be replaced by high impact intersections outside M_v. To find out the best cap, we could simply exhaust all possible scenarios (using parameter Δ').

As a result, Algorithm 3 is proposed as an extension of Algorithm 2, in terms of capping the number of selected intersections in each greedy iteration. In line 1, it initializes $S = \emptyset$. Line 2 includes a loop statement to exhaust all possible scenarios, in terms of all possible combinations of each intersection $u \in V$, each Δ' from 1 to Δ, and each set $S_0 \subseteq M_u$ constrained by $|S_0| = k \bmod (\Delta' + 1)$. Instead of Δ, Δ' is used as the cap. Lines 3 to 7 are basically the same as Algorithm 2, except for the cap. This part embeds Algorithm 2 to search the set of intersections for the RAP placement in each possible scenario. The cap is added at the end of line 5 ($M'_v \leq \Delta'$), while Algorithm 2 uses $M'_v \leq \Delta$ by default ($\Delta = \max_v |M_v|$ by definition). Lines 8 and 9 record the best set of intersections sought among all possible scenarios (specified by line 2). Finally, S is returned.

Note that the total number of all possible scenarios for Algorithm 3 is $O(2^{\Delta}\Delta|V|)$, since it exhausts each intersection $u \in V$, each Δ' from 1 to Δ, and each set $S_0 \subseteq M_u$ constrained by $|S_0| = k \bmod (\Delta' + 1)$. For each possible scenario, Algorithm 3 does the same complexity behavior as Algorithm 2. Meanwhile, the time complexity of Algorithm 2 is $O(|V|^3 + 2^{\Delta}k^2|V||T|)$. As a result, the time complexity of Algorithm 3 is $O(|V|^3 + 4^{\Delta}\Delta k^2|V|^2|T|)$. Although Algorithm 3 has a worse time complexity than Algorithm 2, its approximation ratio becomes better by capping the modularity set [29]:

Theorem 7. Algorithm 3 has an approximation ratio of $1 - e^{-\frac{1}{\Delta+1}}$ to the optimal algorithm.

Proof: Let S^* denote the optimal set of intersections for the RAP placement, in terms of maximizing $f(\cdot)$. Since Algorithm 3 exhausts all possible u, Δ', and S_0 in line 2, there must exist a scenario in which $u = \arg\max_u |M_u \cap S^*|$, $\Delta' = |M_u \cap S^*|$, $S_0 \subseteq M_u \cap S^*$, and $|S_0| = k \bmod (\Delta' + 1)$. All of the following proof is based on the above scenario, although Algorithm 3 eventually has the best effort among all scenarios (lines 8 and 9).

In the above scenario, we claim that $f(S_i)$ in each greedy iteration of Algorithm 2 has a lower bound to $f(S^*)$:

$$f(S_i) \geq (1 - \frac{1}{k'})^i \times f(S_0) + [1 - (1 - \frac{1}{k'})^i] \times f(S^*) \quad (5)$$

Here, k' is defined as $k - [k \bmod (\Delta' + 1)]$. In other words, k' is the largest multiple of $\Delta' + 1$ constrained by $k' \leq k$. Eq. 5 is proved by induction. It is trivial that Eq. 5 holds when $i = 0$, since $(1 - \frac{1}{k'})^0 = 1$. Assume that Eq. 5 holds for i, and we prove that Eq. 5 holds for $i + 1$. Since $S_0 \subseteq (M_u \cap S^*) \subseteq S^*$ and $|S_0| = k \bmod (\Delta' + 1)$, we have:

$$|S^* \setminus S_0| = |S^*| - |S^0| = k - [k \bmod (\Delta' + 1)] = k' \quad (6)$$

Similarly, let us order intersections in $|S^* \setminus S_0|$ in an arbitrary order (say $v_1, v_2, ..., v_{k'}$), and let $S^*_j = \{v_1, v_2, ..., v_j\}$ for $1 \leq j \leq j'$ ($S^*_0 = \emptyset$). For each j, we have:

$$\begin{aligned}
&f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*)) - f(S_i) \\
\geq &f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*_{j-1})) - f(S_i) \\
\geq &f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*_{j-1})) - f(S_i \cup (M_{v_j} \cap S^*_{j-1}))) \\
\geq &f(S_i \cup \{v_j\} \cup S^*_{j-1}) - (S_i \cup S^*_{j-1})
\end{aligned} \quad (7)$$

The first and second inequalities are from the monotonicity in Theorem 3, since $S^*_{j-1} \subseteq S^*$ and $S_i \subseteq S_i \cup (M_{v_j} \cap S^*_{j-1})$). So, $f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*)) \geq f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*_{j-1}))$ and $f(S_i) \leq f(S_i \cup (M_{v_j} \cap S^*_{j-1})))$. The third inequality is from the definition of the modularity set, since only intersections in M_{v_j} can increase the marginal gain of v_j (other intersections might decrease the marginal gain of v_j). By accumulating Eq. 7 among j, we have the following inequality to bound $f(S^*) - f(S_i)$:

$$\begin{aligned}
&\sum\nolimits_{j=1}^{k'} [f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*)) - f(S_i)] \\
\geq &\sum\nolimits_{j=1}^{k'} [f(S_i \cup \{v_j\} \cup S^*_{j-1}) - (S_i \cup S^*_{j-1})] \\
= &f(S_i \cup S^*) - f(S_i) \geq f(S^*) - f(S_i)
\end{aligned} \quad (8)$$

The first inequality is from Eq. 7. The equality results from the definition of S^*_j, since $\{v_j\} \cup S^*_{j-1} = S^*_j$. Note that, since $S^*_{k'} \subseteq S^* \setminus S_0$ and $S_0 \subseteq S_i$, we have $S_i \cup S^*_{k'} = S_i \cup S^*$. We have $S_i \cup S^*_0 = S_i$ since $S^*_0 = \emptyset$. The last inequality is from the monotonicity, since $S^* \subseteq S_i \cup S^*$. We have:

$$\begin{aligned}
&f(S_{i+1}) - f(S_i) = f(S_i \cup \{v\} \cup M'_v) - f(S_i) \\
= &\frac{1}{k'} \sum\nolimits_{j=1}^{k'} [f(S_i \cup \{v\} \cup M'_v) - f(S_i)] \\
\geq &\frac{1}{k'} \sum\nolimits_{j=1}^{k'} [f(S_i \cup \{v_j\} \cup (M_{v_j} \cap S^*)) - f(S_i)] \\
\geq &\frac{1}{k'} [f(S^*) - f(S_i)]
\end{aligned} \quad (9)$$

The first inequality is because line 5 in Algorithm 3 always selects the maximum marginal gain in each greedy iteration. $M_{v_j} \cap S^*$ is also constrained by $|(M_{v_j} \cap S^*)| \leq \Delta'$, since the scenario sets $u = \arg\max_u |M_u \cap S^*|$ and $\Delta' = |M_u \cap S^*|$. The second inequality is from Eq. 8. We rewrite Eq. 9 as:

$$\begin{aligned}
&f(S_{i+1}) \geq \frac{1}{k'}[f(S^*) - f(S_i)] + f(S_i) \\
= &\frac{1}{k'} \times f(S^*) + (1 - \frac{1}{k'}) \times f(S_i) \\
\geq &(1 - \frac{1}{k'})^{i+1} \times f(S_0) + [1 - (1 - \frac{1}{k'})^{i+1}] \times f(S^*)
\end{aligned} \quad (10)$$

The first equality is from Eq. 9 and the last equality is from the induction hypothesis (substituting $f(S_i)$ in Eq. 5). As a result, Eq. 5 is proved by induction.

Since each greedy iteration of Algorithm 3 selects at most $\Delta' + 1$ intersections into S, Algorithm 3 has at least $\lfloor k/(\Delta' + 1) \rfloor = k'/(\Delta' + 1)$ greedy iterations. If we use $i = k'/(\Delta' + 1)$ for Eq. 5, the proof completes:

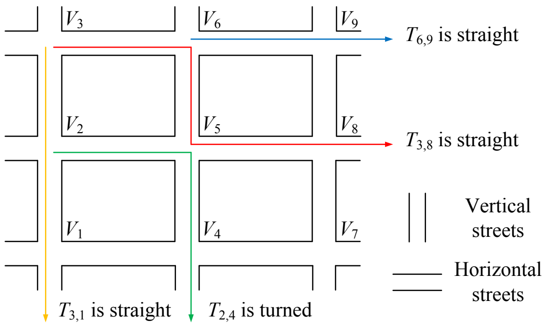

Fig. 6. An illustration of the Manhattan grid streets.

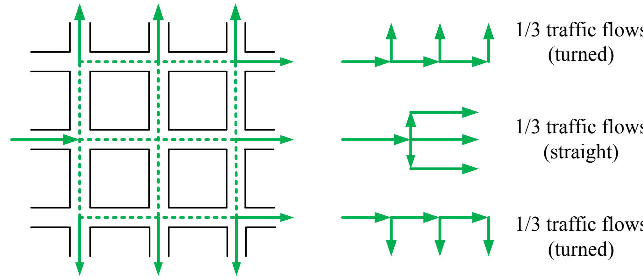

Fig. 7. Traffic flow distribution.

$$\begin{aligned}&f(S) \geq f(S_{k'/(\Delta'+1)})\\ \geq&(1-\frac{1}{k'})^{\frac{k'}{\Delta'+1}} \times f(S_0) + [1-(1-\frac{1}{k'})^{\frac{k'}{\Delta'+1}}] \times f(S^*)\\ \geq&[1-(1-\frac{1}{k'})^{\frac{k'}{\Delta'+1}}] \times f(S^*)\\ \geq&(1-e^{-\frac{1}{\Delta'+1}}) \times f(S^*) \geq (1-e^{-\frac{1}{\Delta+1}}) \times f(S^*) \qquad (11)\end{aligned}$$

The approximation ratio is $1-e^{-\frac{1}{\Delta'+1}} \geq 1-e^{-\frac{1}{\Delta+1}}$. ■

5.5 Time Complexity Reduction

Although Algorithm 3 has a better bound than Algorithm 2, its time complexity is much larger. However, we claim that the time complexity of Algorithm 3 can be reduced. This is because we do not need to exhaust all possible scenarios for practical usage. Rather than exhausting Δ' from 1 to Δ, we can simply stop at a small constant. For example, we only exhaust Δ' from 1 to 3. This is because we only need to consider the most important intersections for v's marginal gain. Similarly, we do not need to exhaust each intersection v for the initialization of S_0. Instead, we can focus on the intersections with most numbers of passing drivers. This is because the optimal set of intersections for the RAP placement is not likely to exclude those intersections. Using this approach, the time complexity of Algorithm 3 can be reduced to $O(|V|^3 + k|V||T|)$, which is asymptotically the same as Algorithm 1. Our experiments show that this approach only slightly hurts the performance of Algorithm 3.

6 RAP Placement for Manhattan Grid

Considering that the real-world traffic distributions have some patterns, this section studies a special case of the RAP placement problem, using the Manhattan grid scenario that has a grid street layout. The RAP placement problem is re-explored under the deadline and decreasing utility functions, as shown in Eqs. 3 and 4. The RAP placement problem in this special case remains NP-hard by reduction from the geometric maximum coverage problem [34].

6.1 Properties of Manhattan Grid

The Manhattan grid streets plan is a type of city plan in which streets run at right angles to each other. In this city plan, vehicles can only move in four given directions, as shown in Fig. 6. We classify the streets into vertical streets and horizontal streets, based on their orientations. Multiple shortest paths between pairs of intersections may exist in the Manhattan grid streets. For example, in Fig. 6, the shortest path from v_1 to v_6 could be $v_1v_2v_3v_6$, $v_1v_2v_5v_6$, or $v_1v_4v_5v_6$. We relax the constraint used in the previous section, where the traveling path for a traffic flow is unique and is known a priori. In this section, the traveling path for a traffic flow is not pre-fixed. Let us consider a driver in $T_{1,6}$ that travels from v_1 to v_6. His/Her traveling path would be randomly chosen from among the three shortest paths. At this time, what could happen if an RAP is placed at v_3? This driver would definitely choose $v_1v_2v_3v_6$ as his/her traveling path, since it is one of the shortest paths with a free additional advertisement. We consider a traffic flow, $T_{i,j}$, to travel along one of the shortest paths from i to j; if an RAP is placed in one of the shortest paths, then the traffic flows would choose that path to obtain a free additional advertisement. Locations of placed RAPs are assumed to be known by all the drivers, i.e., they are published on the internet for drivers to access.

Considering the above property, we reformulate the RAP placement problem under the Manhattan grid scenario, as follows. The shop is located within a square region. D is large enough such that vehicles would detour to the shop once receiving an advertisement. All the traffic flows travel through their shortest paths. For a specified traffic flow, if an RAP is placed in one of its shortest paths, this traffic flow would travel through that path to obtain a free advertisement. For simplicity, the Manhattan grid scenario is discussed under two special forms for $f_{i,j}(S)$, i.e., the deadline and decreasing utility functions in Eqs. 3 and 4. We do not consider a general form for $f_{i,j}(S)$ in this section.

6.2 Manhattan Grid with Threshold Utility Function

Let us start with the Manhattan RAP placement problem under the threshold utility function in Eqs. 3. An intersection can still include multiple traffic flows, while a traffic flow can be included by multiple intersections. We start with the following definition:

Definition 5. A traffic flow is *turned* if it has exactly one turn within the grid. Otherwise, it is *straight*.

For example, in Fig. 6, the traffic flows of $T_{3,1}$ and $T_{6,9}$ are straight, while $T_{2,4}$ is turned. Note that $T_{3,8}$ is straight, since it has two turns at v_5 and v_6, respectively. All the traffic flows have at most two turns within the scenario, otherwise, the corresponding traveling path is not a shortest path. Then, we observe that a turned traffic flow has multiple shortest traveling paths. Therefore, the locations of the RAPs may have some impact on its actual traveling path.

Algorithm 4 A two-stage solution

Input: The grid graph G; the set of traffic flows T;
The number of RAPs to place (i.e., k).
Output: The RAP placement.
1: **if** $k \leq 5$ **then**
2: **return** the optimal solution by exhaustive search.
3: Initialize $S = \emptyset$.
4: **for** each corner intersection, v, of the square region **do**
5: Update $S = S \cup \{v\}$.
6: **for** $i = 1$ to $k - 4$ **do**
7: Find $v = \arg\max_{v \in V} f(S \cup \{v\}) - f(S)$.
8: Update $S = S \cup \{v\}$.
9: **return** the RAP placement.

Following the above intuition, a two-stage RAP placement algorithm is proposed as Algorithm 4. The RAPs are placed for turned and straight traffic flows, respectively. Algorithm 4 assumes that the traffic flows are uniformly distributed. We start with the case in which all the traffic flows go through the scenario. For this case, turned and straight traffic flows have fractions of $\frac{2}{3}$ and $\frac{1}{3}$ with respect to the total traffic flows, respectively. The reason is illustrated in Fig. 7. We have two observations. (i) Four RAPs at the corners of the grid can cover all the turned traffic flows. This is because corresponding drivers could go to the corner for a free advertisement without extra traveling distances. (ii) The RAP placement for the straight traffic flows can be obtained through the same idea as Algorithm 1. This placement keeps a ratio of $1-\frac{1}{e}$ to the optimal solution. The cases are similar, when traffic flows may start from or stop at an intersection within the scenario. Consequently, we have:

Theorem 8. If traffic flows are uniformly distributed, then Algorithm 4 is expected to have a ratio of $1-\frac{1}{3e}-\frac{4}{3k}$ to the optimal solution, under the Manhattan grid scenario with the threshold utility function.

Proof: Our proof includes two parts. The first part is for the turned traffic flows and the second part is for the straight traffic flows. In the first part, we claim that four RAPs at the corners of the grid are enough to cover all the turned traffic flows that go through the scenario (lines 4 and 5 in Algorithm 4). Let us consider a traffic flow that enters the grid via the west boundary of the grid, and exits the grid via the south boundary of the grid. Such a traffic flow could be $T_{2,4}$ in Fig. 6. The shortest paths for this traffic flow only include two kinds of orientations: going Eastward or going Southward at an intersection. If this traffic flow goes Southward to the end and then goes Eastward, it would result in a shortest path that goes through the southwest corner of the grid. For example, such a shortest path for $T_{2,4}$ in Fig. 6 is $v_2v_1v_4$ that goes through the corner v_1. Since $v_2v_1v_4$ is a shortest path for $T_{2,4}$ with a free advertisement, drivers in $T_{2,4}$ would choose $v_2v_1v_4$ as their traveling paths. By enumerating all the possibilities, our claim is true. On expectation, four RAPs at the corners of the grid can cover $\frac{2}{3}$ of the total traffic flows. The above analysis is also valid, when traffic flows may start from or stop at an intersection within the scenario (i.e., not go through the scenario).

In the second part, we focus on the straight traffic flows. The greedy placement in lines 6 to 8 of Algorithm 4 has a ratio of $1-\frac{1}{e}$ to the optimal solution using $k-4$ RAPs. This is similar to Algorithm 1. Furthermore, it has a ratio of $\frac{k-4}{k}(1-\frac{1}{e})$ to the optimal solution using k RAPs. This is because four RAPs are placed at the corner of the scenario for the turned traffic flows. Note that straight traffic flows have a fraction of $\frac{1}{3}$ with respect to all the traffic flows on expectation. Considering both turned traffic flows and straight traffic flows, the total fraction of the traffic flows that are covered by Algorithm 4 is:

$$\frac{2}{3}+\frac{1}{3}\cdot\frac{k-4}{k}\cdot(1-\frac{1}{e}) \geq 1-\frac{1}{3e}-\frac{4}{3k} \tag{12}$$

Eq. 12 completes the proof of Theorem 8. ∎

When k becomes larger, $1-\frac{1}{3e}-\frac{4}{3k}$ becomes larger, meaning that Algorithm 4 has a better performance. $1-\frac{1}{3e}-\frac{4}{3k}$ is larger than $1-\frac{1}{e}$ when $k>5$.

Algorithm 5 A modified two-stage solution

Input: The square region; the set of traffic flows T;
The number of RAPs to place (i.e., k);
Output: The RAP placement;
1: Same as Algorithm 4, except the change in lines 4 and 5: Instead of each corner intersection of the square region, S selects the intersection at the middle of that corner and the center of the square;

6.3 Manhattan Grid with Deceasing Utility Function

This subsection discusses the Manhattan RAP placement problem under the decreasing utility function. Similarly, we place RAPs for turned and straight traffic flows, respectively. However, the overlaps among RAPs bring some performance degradations. Algorithm 5 is proposed as an extension of Algorithm 4. It has a prerequisite in which the decreasing utility function must be the one in Eq. 4. The performance of Algorithm 5 is also guaranteed:

Theorem 9. If traffic flows are uniformly distributed, then Algorithm 5 is expected to have a ratio of $\frac{1}{2}-\frac{1}{6e}-\frac{2}{3k}$ to the optimal solution, under the Manhattan grid scenario with the decreasing utility function.

Proof: This proof is similar to that of Theorem 8. First, we prove that the four RAPs in the middle of the corner and the shop can attract half of the maximum drivers from the turned traffic flows. This is because the turned traffic flows have an average detour distance of $\frac{D}{2}$, while the four RAPs cover half of the turned traffic flows with the detour distance $\frac{D}{2}$. Then, we show that the remaining $k-4$ RAPs can attract half of the maximum drivers from straight traffic flows. This is because Algorithm 5 can attract no fewer drivers than the maximum drivers from $k-4$ traffic straight flows (either vertical or horizontal). Through a similar argument, it can be seen that Algorithm 5 achieves a ratio of $\frac{1}{2}(1-\frac{4}{k})=\frac{1}{2}-\frac{2}{k}$ to the optimal solution. ∎

7 EVALUATIONS

This section conducts extensive experiments to evaluate the performances of the proposed algorithms. After presenting the settings, the evaluation results are shown from different perspectives to provide insightful conclusions.

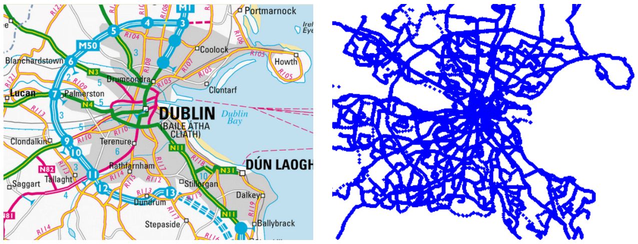

(a) The Dublin map. (b) The bus traces.

Fig. 8. The map and bus traces for Dublin's central area.

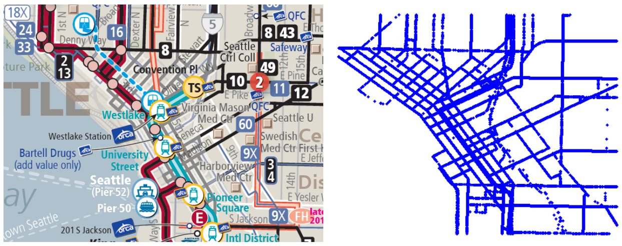

(a) The Seattle map. (b) The bus traces.

Fig. 9. The map and bus traces for Seattle's central area.

7.1 Real Trace-driven Datasets and Basic Settings

This section conducts extensive experiments based on two real traces, i.e., the Dublin bus trace [35] and the Seattle bus trace [36]. The city plan of Dublin is not grid-based, and thus the Dublin bus trace is used to test our Algorithms 1, 2, and 3 for the general scenario in Section 5. The city plan of Seattle is partially grid-based, and thus the Seattle bus trace is used to test all of our algorithms for both the general scenario in Section 5 and the Manhattan grid scenario in Section 6.

For the Dublin bus trace, we focus on the part within Dublin's central area, which is a $80,000 \times 80,000$ square foot area, as shown in Fig. 8. The Dublin bus trace includes the bus ID, longitude, latitude, and vehicle journey ID. The vehicle journey is a given run on a journey pattern, which corresponds to our concept of the traffic flow. Buses with the same vehicle journey ID have similar routing paths in terms of longitude and latitude. To obtain the number of attracted customers, we assume that each bus in Dublin carries 100 people (who are potential customers) per day on average. The Dublin bus trace has 411 traffic flows with 5,119 buses.

For the Seattle bus trace, we focus on the part within Seattle's central area, which is a $10^4 \times 10^4$ square foot area, as shown in Fig. 9. The Seattle bus trace includes the bus ID, x-coordinate, y-coordinate, and route ID. Each route is regarded as a traffic flow. Buses with the same route ID have similar routing paths in terms of x and y coordinates. To obtain the number of attracted customers, we assume that each bus in Seattle carries 200 people per day on average. The Seattle bus trace has 236 traffic flows with 1,163 buses. Although the Seattle bus trace has fewer traffic flows and buses than the Dublin bus trace, its area is also smaller.

According to the amount of passing traffic flows, all the street intersections in both traces are classified into the city's center, city, or suburb. This is used to observe the impact of the shop location. Our experiments are based on three utility functions. The first one is the threshold utility function in Eq. 3. The second one is the decreasing utility function in Eq. 4, which decays linearly. The third one is the sine utility function, as defined in the following:

TABLE 2
Parameters for experiments.

	Dublin	Seattle
Trace size	80,000×80,000 ft^2	10,000×10,000 ft^2
Bus/Vehicle load	100 per day	200 per day
# of traffic flows	411	236
# of buses/vehicles	5,119	1,163
Utility function	Threshold/Decreasing/Sine function	
$\alpha_{i,j}$	0.001 for all traffic flows	

$$f_{i,j}(S) = \begin{cases} \alpha_{i,j} \times n_{i,j} \times \sin \frac{\pi \times d_{i,j}(S)}{D} & \text{if } d_{i,j}(S) \leq D \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

A notable point is that the sine utility function is no longer non-increasing. Therefore, $f(S)$ is not submodular when the sine utility function is used. Meanwhile, under the same detour distance, d, and the same threshold, D, the detour probability of the threshold utility function is larger than that of the decreasing utility function. In these three utility functions, $\alpha_{i,j}$ is set to be 0.001 for all the traffic flows [8]. All parameters for experiments are shown in Table 7.1.

7.2 Comparison Algorithms and Metrics

In our experiments, four baseline algorithms (MaxCardinality, MaxVehicles, MaxCustomers, and Random) are used for comparisons, as in the following:

- MaxCardinality ranks the intersections by the number of passing traffic flows, and then places the RAPs at the top-k intersections.
- MaxVehicles ranks the intersections by the number of passing buses (a traffic flow has multiple buses), and then places the RAPs at the top-k intersections.
- MaxCustomers ranks the intersections by the number of attracted customers if only one RAP is placed there without other RAPs. Then, MaxCustomers also places RAPs at the top-k intersections.
- Random places RAPs uniform-randomly at the intersections within the $D \times D$ square region centered at the shop. It is the baseline.

Our experiments focus on the relationship between the number of placed RAPs and the number of attracted customers, under different settings (utility functions, threshold D, and shop locations). Street intersections are classified into the city's center, city, or suburb, depending on the amount of passing traffic flows. In the following experiments, if we say that the shop is located in the city, it means that the intersections with city tags are randomly selected as the shop locations. All the results are averaged over 1,000 times.

7.3 Evaluation Results in the Dublin Bus Trace

This subsection focuses on Algorithms 1, 2, and 3 in the Dublin bus trace, in terms of different numbers of RAPs, utility functions, shop locations, and utility function thresholds. We also analyze the running time and complexity reduction for Algorithms 2 and 3 in this trace.

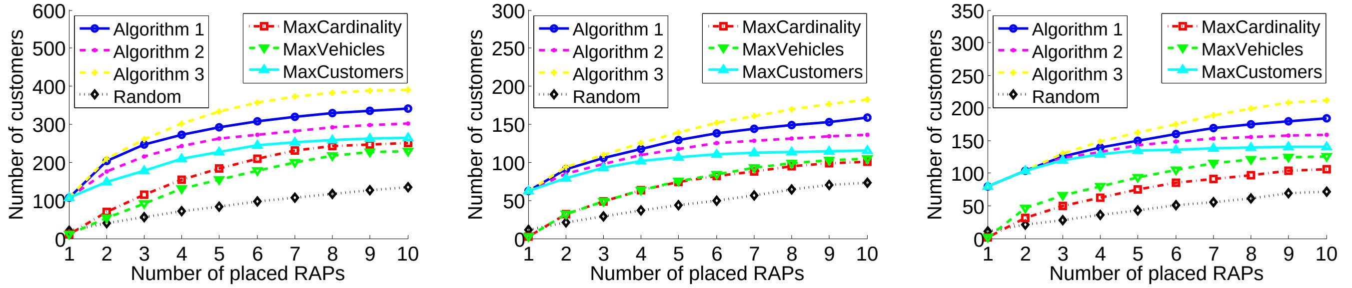

(a) Threshold utility function. (b) Decreasing utility function. (c) Sine utility function.

Fig. 10. The experimental results for the Dublin bus traces with different utility functions. The shop is located in the city where $D = 20,000$ feet.

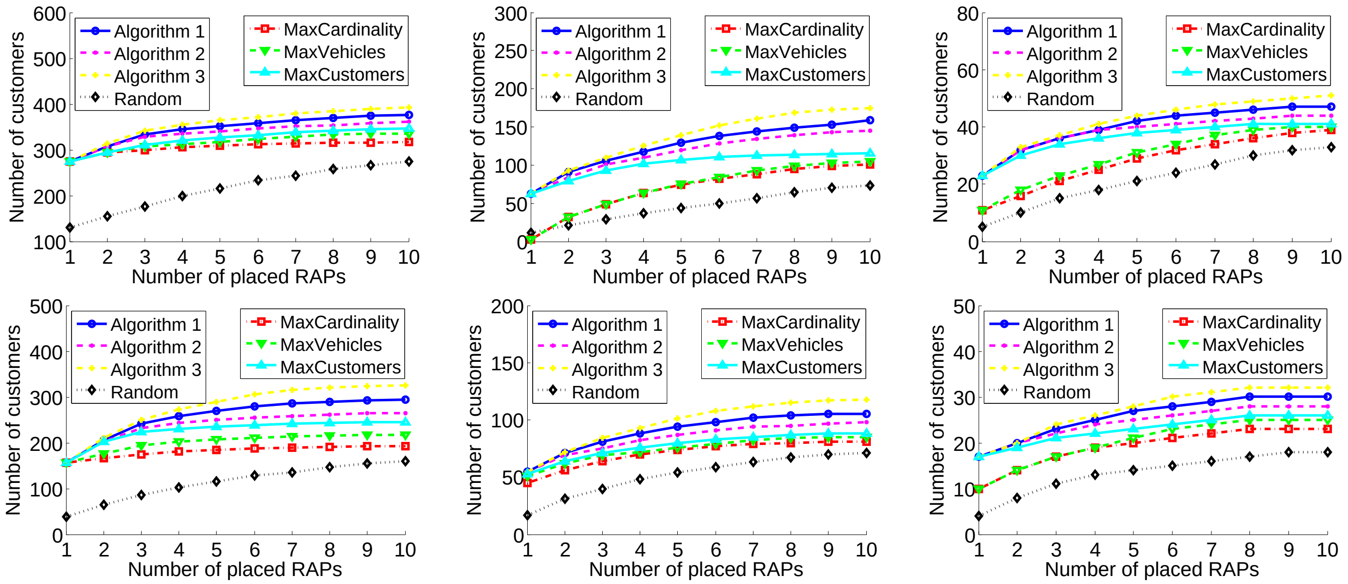

(a) The shop is located in the city's center. (b) The shop is located in the city. (c) The shop is located in the suburb.

Fig. 11. The experimental results for the Dublin bus traces with different shop locations. The decreasing utility function is used with a different D.

7.3.1 *Number of RAPs and Utility Function*

Fig. 10 shows the impact of the number of RAPs and the utility function in the Dublin bus trace under the general scenario. The shop is located in the city with the threshold $D = 20,000$ feet. Figs. 10(a), 10(b), and 10(c) show the results for the threshold utility function, the decreasing utility function, and the sine utility function, respectively. It can be seen that the performance gap between the our algorithms and the other algorithms is significant. For example, when we use the threshold utility function with $k = 10$, Algorithm 3 attracted more than 50% drivers to the shop than MaxCustomers.

Algorithm 3 could always outperform Algorithms 1 and 2. This is because Algorithms 1 and 2 are special cases of Algorithm 3. Algorithm 3 reduces to Algorithm 1 by only using the cap $\Delta' = 1$. Similarly, Algorithm 3 reduces to Algorithm 2 by only using the cap $\Delta' = \Delta$. As a trade-off, Algorithm 3 has a higher time complexity than Algorithms 1 and 2. We also find that Algorithm 2 is outperformed by Algorithm 1. Although Algorithm 2 guarantees an approximation ratio when the objective function is non-submodular, its intersection selection is overly aggressive. In each greedy iteration, Algorithm 2 may select too many intersections for the RAP placement, ignoring the benefit-to-cost ratio. Meanwhile, the comparison algorithms (MaxCardinality, MaxVehicles, and MaxCustomers) perform poorly, since they only focus on top-k intersections with different ranking criteria. They do not consider the relationships between different intersections to maximize the number of attracted drivers.

For all algorithms, the impact of the number of RAPs is basically similar. Due to the monotonicity, more RAPs always bring more attracted drivers to the shop. In contrast, the impact of the utility function is more significant. All algorithms attract more drivers under the threshold utility function than the decreasing utility function and the sine utility function. This is because the detour probability of the threshold utility function is the largest, under the same d and D. Moreover, the utility function also determines the sharpness of the performance curve with respect to the number of RAPs. For Algorithm 3, the threshold utility function brings more diminishing return effects.

7.3.2 *Shop Location and Utility Function Threshold*

Fig. 11 shows the impact of the shop location and the threshold D in the Dublin bus trace under the general scenario. The decreasing utility function is used. Figs. 11(a), 11(b), and 11(c) show the results for different shop locations (city's center, city, and suburb). For each subfigure, the top

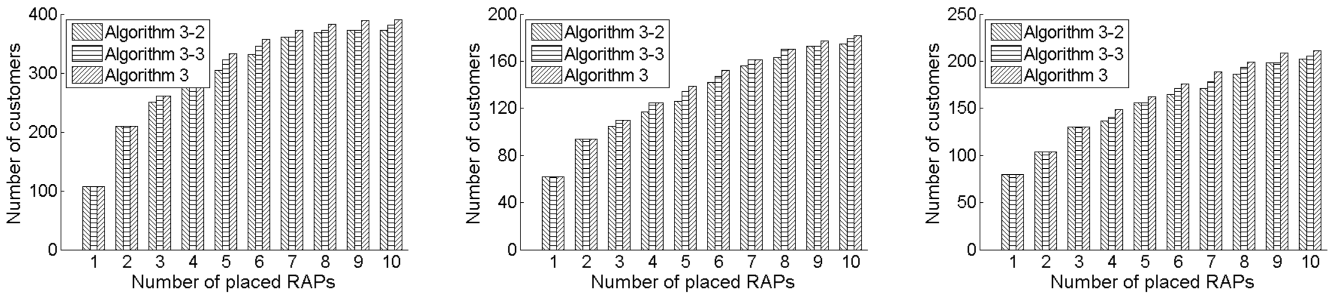

(a) Threshold utility function. (b) Decreasing utility function. (c) Sine utility function.

Fig. 12. Running time reduction for Algorithm 3.

and bottom parts show the results with $D = 20,000$ feet and $D = 10,000$ feet, respectively.

The impact of the shop's location is very significant. Under the same settings, many more drivers can be attracted to the shop, when the shop is located in the city's center than the suburb. This is because there are more traffic flows in the city's center than the suburb. If the shop is located in the city's center, the RAPs can cover more traffic flows at small detour distances. Meanwhile, the performance gain of Algorithm 3 is relatively small, when the shop is located in the city's center or suburb. If the shop is located in the city's center, randomly placing RAPs around the shop can already cover most traffic flows with small detour distances. On the other hand, if the shop is located in the suburb, none of the placement strategies can cover too many traffic flows.

The utility function threshold D is also critical. A larger D means that the drivers are more likely to detour to the shop, and thus, the shop can attract more drivers. When the shop is located in the city's center, a large D does not bring too many additional drivers, since most traffic flows are already near the shop. When the shop is located in the suburbs, a large D still does not bring too many additional drivers, since the detour distances are large. However, when the shop is located in the city, as shown in Fig. 11(b), a large D brings many more drivers, since more traffic flows are covered with small detour distances.

7.3.3 Running Time and Complexity Reduction

This subsection evaluates the running time of the proposed algorithms. Comparison algorithms are the same. All codes are implemented in Matlab and are executed on Dell Inspiron i15RN-3647BK laptop, which includes a 2.5GHz Intel Core i5 2450M processor. In addition, we further introduce two variations of Algorithm 3 by using different maximum cap sizes. The first variation is Algorithm 3-2, using 2 as its maximum cap size (Δ' ranges from 1 to 2). In each greedy iteration, Algorithm 3-2 selects at most 2 intersections. The second variation is Algorithm 3-3, using 3 as its maximum cap size. We evaluate the impact of the maximum cap size, in terms of both the performance and running time.

MaxCardinality	MaxVehicles	MaxCustomers
3 seconds	4 seconds	87 seconds
Random	Algorithm 1	Algorithm 2
1 second	13 minutes	191 minutes
Algorithm 3	Algorithm 3-2	Algorithm 3-3
16 hours	73 minutes	231 minutes

We start with the running time, as shown in the above table. The shop is located in the city, and we set $k = 10$ and $D = 20,000$ feet. The deadline utility function is used here, although other utility functions do not significantly change the running time. Random is fastest, since it does not care about the performance at all. As a trade-off, it has the worst performance. MaxCardinality, MaxVehicles, and MaxCustomers take seconds, since they only rank intersections by different criteria. The running time of MaxCustomers is larger, since it needs to calculate the number of attracted drivers for each intersection. The proposed algorithms take minutes to hours. Algorithm 1 takes more time than MaxCustomers due to its greedy loop for calculating marginal gain. Algorithm 2 takes significantly more time than Algorithm 1, since it needs to consider the combinations of intersections for the RAP placement. Algorithm 3 takes even more time (i.e., 16 hours) to exhaust all possible caps. However, if the maximum cap size is used, the running time of Algorithm 3 can be significantly reduced to minutes.

Fig. 12 shows the impact of the maximum cap size for Algorithm 3 under these three utility functions. Algorithm 3-2, Algorithm 3-3, and Algorithm 3 have very close performances (Algorithm 3-2 $\leq$ Algorithm 3-3 $\leq$ Algorithm 3). This is because Algorithm 3 does not need to select a large set of intersections in each iteration (there are few important combinations of these intersections). Algorithm 3-3 has almost the same performance as CG, especially when there are few RAPs. After considering the running time, capping Δ' to 3 in Algorithm 3 is a practical strategy for large-scale RAP placements.

7.4 Evaluation Results in the Seattle Bus Trace

This subsection evaluates all algorithms in the Seattle bus trace, under the general scenario and the Manhattan grid scenario. For the general scenario, we still evaluate Algorithms 1, 2, and 3 through a similar setting as the previous subsection. Their performances in the Dublin and Seattle traces are compared. For the Manhattan grid scenario, we only evaluate Algorithms 4 and 5, since they are specially designed for this case. Note that Algorithms 4 and 5 require a grid street plan and special utility functions. Consequently, they are not general solutions for the RAP placement.

7.4.1 General Scenario

Fig. 13 shows the evaluation results in the Seattle bus trace under the general scenario. The shop is located in

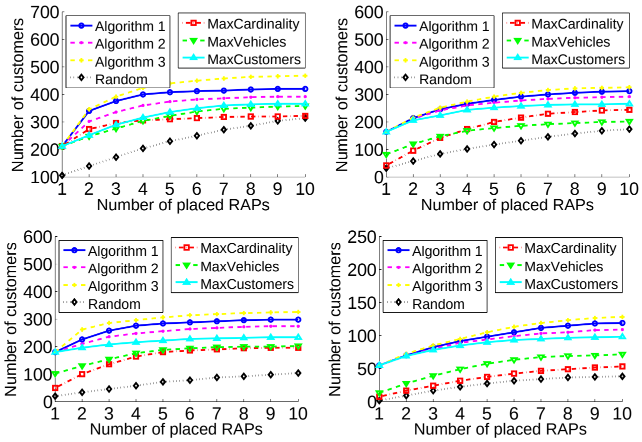

(a) Threshold utility function. (b) Decreasing utility function.

Fig. 13. The experimental results for the Seattle bus traces under the general scenario in Section 5. The shop is located in the city. Different utility functions with different threshold D are used.

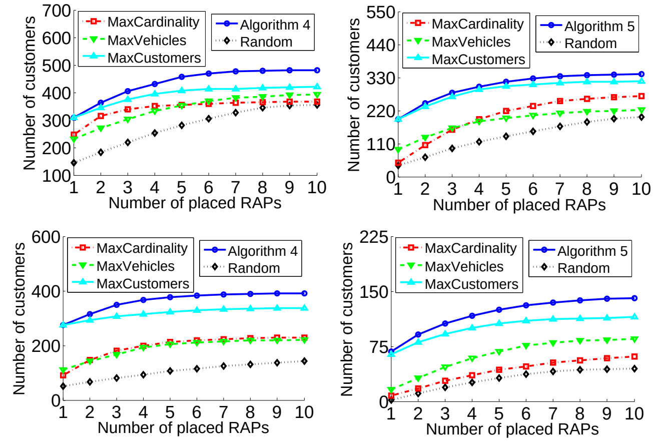

(a) Threshold utility function. (b) Decreasing utility function.

Fig. 14. The experimental results for the Seattle bus traces under the Manhattan grid scenario in Section 6. The shop is located in the city. Different utility functions with different threshold D are used.

the city. We focus on the impacts of the utility function and the threshold D. The threshold utility function and the decreasing utility function are used in Figs. 13(a) and 13(b), respectively. For each subfigure, the top and bottom parts show the results under $D = 2,500$ feet and $D = 1,000$ feet, respectively. It can be seen that all algorithms attract more customers under the threshold utility function than the decreasing utility functions, since the former one brings higher detour probabilities. D remains critical, especially when the shop is located in the city. The number of attracted customers with $D = 2,500$ feet is 30% more than that with $D = 1,000$ feet. The evaluation results are similar for the Dublin bus trace and the Seattle bus trace under the general scenario. The slight algorithmic performance differences are caused by area density, in terms of the area size, the number of traffic flows, and the number of buses.

7.4.2 *Manhattan Scenario*

Fig. 14 shows the evaluation results in the Seattle bus trace under the Manhattan grid scenario. The settings (i.e., shop location, utility function, and threshold D) are the same as the general scenario. Compared with the results under the general scenario in Fig. 13, more customers are attracted under the Manhattan grid scenario. This is mainly because the traveling paths of all the traffic flows are pre-fixed under the general scenario, the assumption of which is relaxed in the Manhattan grid scenario. Algorithm 4 is designed under the threshold utility function in Fig. 14(a). It performs much better than the comparison algorithms (MaxCardinality, MaxVehicles, and MaxCustomers). For example, when $D = 1,000$, Algorithm 4 attracts almost a doubled number of drivers to the shop than MaxCardinality and MaxVehicles. We also find that a larger threshold D brings more customers to the shop. Algorithm 5 is designed under the decreasing utility function in Fig. 14(b). The performance gap between Algorithm 5 and the comparison algorithms is smaller than the performance gap between Algorithm 4 and the comparison algorithms. This is because Algorithm 5 has a worse approximation ratio than Algorithm 4. Algorithm 5 cannot attract drivers from all turned traffic flows.

8 Conclusions

This paper addresses a novel roadside advertisement dissemination problem that involves three elements: the drivers, RAPs, and shopkeepers. The shopkeeper uses RAPs to disseminate advertisements to the drivers, in order to attract customers. Upon receiving an advertisement, the driver may detour to the shop, depending on the detour distance. Our goal is to optimize the RAP placement for the shopkeeper to maximally attract potential customers. Three bounded RAP placement algorithms are proposed for the general scenario. As a special case, the Manhattan scenario is also discussed. Real trace-driven experiments validate the competitive performance of our algorithms.

9 Acknowledgments

This research was supported in part by NSF grants CNS 1629746, CNS 1564128, CNS 1449860, CNS 1461932, CNS 1460971, and CNS 1439672.

References

[1] N. Nikookaran, G. Karakostas, and T. D. Todd, "Combining capital and operating expenditure costs in vehicular roadside unit placement," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 8, pp. 7317–7331, 2017.

[2] J. Liu, J. Wan, D. Jia, B. Zeng, D. Li, C.-H. Hsu, and H. Chen, "High-efficiency urban traffic management in context-aware computing and 5g communication," *IEEE Communications Magazine*, vol. 55, no. 1, pp. 34–40, 2017.

[3] Y. Zhou, S. Chen, Z. Mo, and Q. Xiao, "Point-to-point traffic volume measurement through variable-length bit array masking in vehicular cyber-physical systems," in *Proceedings of IEEE ICDCS 2015*, pp. 51–60.

[4] Y. Hou, Y. Zhao, A. Wagh, L. Zhang, C. Qiao, K. F. Hulme, C. Wu, A. W. Sadek, and X. Liu, "Simulation-based testing and evaluation tools for transportation cyber–physical systems," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 3, pp. 1098–1108, 2016.

[5] C. Huang, R. Lu, and K.-K. R. Choo, "Vehicular fog computing: architecture, use case, and security and forensic challenges," *IEEE Communications Magazine*, vol. 55, no. 11, pp. 105–111, 2017.

[6] C. T. Ngo and O. Hoon, "A roadside unit based hybrid routing protocol for vehicular ad hoc networks," *IEICE Transactions on Communications*, vol. 98, no. 12, pp. 2400–2418, 2015.
[7] J. Qin, H. Zhu, Y. Zhu, L. Lu, G. Xue, and M. Li, "Post: Exploiting dynamic sociality for mobile advertising in vehicular networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 6, pp. 1770–1782, 2016.
[8] X. Li, C. Qiao, Y. Hou, Y. Zhao, A. Wagh, A. Sadek, L. Huang, and H. Xu, "On-road ads delivery scheduling and bandwidth allocation in vehicular CPS," in *IEEE INFOCOM 2013*, pp. 2571–2579.
[9] Y. Li, L. Zhang, H. Zheng, X. He, S. Peeta, T. Zheng, and Y. Li, "Nonlane-discipline-based car-following model for electric vehicles in transportation-cyber-physical systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 1, pp. 38–47, 2018.
[10] S. M. Khan, K. C. Dey, and M. Chowdhury, "Real-time traffic state estimation with connected vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 7, pp. 1687–1699, 2017.
[11] K. Lakshmanan, D. De Niz, R. Rajkumar, and G. Moreno, "Resource allocation in distributed mixed-criticality cyber-physical systems," in *IEEE ICDCS 2010*, pp. 169–178.
[12] S. Ishihara, N. Nakamura, and Y. Niimi, "Demand-based location dependent data dissemination in vanets," in *ACM MobiCom 2013*, pp. 219–222.
[13] X. Li, X. Yu, A. Wagh, and C. Qiao, "Human factors-aware service scheduling in vehicular cyber-physical systems," in *IEEE INFOCOM 2011*, pp. 2174–2182.
[14] A. Wagh, X. Li, R. Sudhaakar, S. Addepalli, and C. Qiao, "Data fusion with flexible message composition in Driver-in-the-Loop VCPS," *Ad Hoc Network*, vol. 11, no. 7, pp. 2083–2095, 2013.
[15] A. B. Reis, S. Sargento, and O. K. Tonguz, "Parked cars are excellent roadside units," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 9, pp. 2490–2502, 2017.
[16] http://www.redpinesignals.com/.
[17] H. Zheng and J. Wu, "Optimizing roadside advertisement dissemination in vehicular cyber-physical systems," in *Proceedings of IEEE ICDCS 2015*, pp. 41–50.
[18] J. Chen and J. Stallaert, "An economic analysis of online advertising using behavioral targeting," *Mis Quarterly*, vol. 38, no. 2, pp. 429–449, 2014.
[19] Y. Li, D. Zhang, and K.-L. Tan, "Real-time targeted influence maximization for online advertisements," *Proceedings of the VLDB Endowment*, vol. 8, no. 10, pp. 1070–1081, 2015.
[20] Y. Zhang, L. Song, C. Jiang, N. H. Tran, Z. Dawy, and Z. Han, "A social-aware framework for efficient information dissemination in wireless ad hoc networks," *IEEE Communications Magazine*, vol. 55, no. 1, pp. 174–179, 2017.
[21] W. Shen, L. Liu, X. Cao, Y. Hao, and Y. Cheng, "Cooperative message authentication in vehicular cyber-physical systems," *IEEE Transactions on Emerging Topics in Computing*, vol. 1, no. 1, pp. 84–97, 2013.
[22] K. Liu, V. C. S. Lee, J. K.-Y. Ng, J. Chen, and S. H. Son, "Temporal data dissemination in vehicular cyber–physical systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 6, pp. 2419–2431, 2014.
[23] B. Kar, E. H.-K. Wu, and Y.-D. Lin, "The budgeted maximum coverage problem in partially deployed software defined networks," *IEEE Transactions on Network and Service Management*, vol. 13, no. 3, pp. 394–406, 2016.
[24] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a monotone submodular function subject to a matroid constraint," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1740–1766, 2011.
[25] Z. Zhang, J. Willson, Z. Lu, W. Wu, X. Zhu, and D.-Z. Du, "Approximating maximum lifetime k-coverage through minimizing weighted k-cover in homogeneous wireless sensor networks," *IEEE/ACM Transactions on Networking*, vol. 24, no. 6, pp. 3620–3633, 2016.
[26] N. Buchbinder and M. Feldman, "Deterministic algorithms for submodular maximization problems," in *Proceedings of ACM-SIAM SODA 2016*, pp. 392–403.
[27] S. Im, V. Nagarajan, and R. V. D. Zwaan, "Minimum latency submodular cover," *ACM Transactions on Algorithms*, vol. 13, no. 1, p. 13, 2016.
[28] A. Badanidiyuru, C. Papadimitriou, A. Rubinstein, L. Seeman, and Y. Singer, "Locally adaptive optimization: Adaptive seeding for monotone submodular functions," in *Proceedings of ACM-SIAM SODA 2016*, pp. 414–429.
[29] M. Feldman and R. Izsak, "Constrained monotone function maximization and the supermodular degree," in *Proceedings of ACM-SIAM SODA 2014*, pp. 1–23.
[30] S. Dughmi, "Algorithmic information structure design: a survey," *ACM SIGecom Exchanges*, vol. 15, no. 2, pp. 2–24, 2017.
[31] S. Fujishige and S. Isotani, "A submodular function minimization algorithm based on the minimum-norm base," *Pacific Journal of Optimization*, vol. 7, no. 1, pp. 3–17, 2011.
[32] M. Sviridenko, J. Vondrák, and J. Ward, "Optimal approximation for submodular and supermodular optimization with bounded curvature," in *Proceedings of ACM-SIAM SODA 2015*, pp. 1134–1148.
[33] ——, "Optimal approximation for submodular and supermodular optimization with bounded curvature," *Mathematics of Operations Research*, vol. 42, no. 4, pp. 1197–1218, 2017.
[34] A. Pananjady, V. K. Bagaria, and R. Vaze, "Optimally approximating the coverage lifetime of wireless sensor networks," *IEEE/ACM Transactions on Networking*, vol. 25, no. 1, pp. 98–111, 2017.
[35] https://data.dublinked.ie/dataset/dublin-bus-gtfs-data.
[36] https://crawdad.cs.dartmouth.edu/rice/ad_hoc_city/20030911/bus_mobility/.

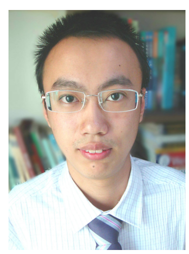

Huanyang Zheng received his B.Eng. degree in Telecommunication Engineering from Beijing University of Posts and Telecommunications, China, in 2012. He is currently a Ph.D. candidate in the Department of Computer and Information Sciences, Temple University, USA. His research focuses on wireless and mobile networks, social networks and structures, and cloud systems.

Jie Wu is the Associate Vice Provost for International Affairs at Temple University. He also serves as the Chair and Laura H. Carnell professor in the Department of Computer and Information Sciences. Prior to joining Tempe University, he was a program director at the National Science Foundation and was a distinguished professor at Florida Atlantic University. His current research interests include mobile computing and wireless networks, routing protocols, cloud and green computing, network trust and security, and social network applications. Dr. Wu regularly publishes in scholarly journals, conference proceedings, and books. He serves on several editorial boards, including IEEE Transactions on Service Computing and the Journal of Parallel and Distributed Computing. Dr. Wu was general co-chair/chair for IEEE MASS 2006, IEEE IPDPS 2008, IEEE ICDCS 2013, and ACM MobiHoc 2014, as well as program co-chair for IEEE INFOCOM 2011 and CCF CNCC 2013. He was an IEEE Computer Society Distinguished Visitor, ACM Distinguished Speaker, and chair for the IEEE Technical Committee on Distributed Processing (TCDP). Dr. Wu is a CCF Distinguished Speaker and a Fellow of the IEEE. He is the recipient of the 2011 China Computer Federation (CCF) Overseas Outstanding Achievement Award.