

# Maximizing the Contact Opportunity for Vehicular Internet Access

Zizhan Zheng<sup>†</sup> Zhixue Lu<sup>†</sup> Prasun Sinha<sup>†</sup> Santosh Kumar<sup>§</sup>  
<sup>†</sup> The Ohio State University      <sup>§</sup> University of Memphis  
{zhengz,luz,prasun}@cse.ohio-state.edu    santosh.kumar@memphis.edu

**Abstract**—With increasing popularity of media enabled hand-helds, the need for high data-rate services for mobile users is evident. Large-scale Wireless LANs (WLANs) can provide such a service, but they are expensive to deploy and maintain. Open WLAN access-points (APs), on the other hand, need no new deployments, but can offer only opportunistic services with no guarantees on short term throughput. In contrast, a carefully planned *sparse* deployment of roadside WiFi provides an economically scalable infrastructure with quality of service assurance to mobile users. In this paper, we propose to study deployment techniques for providing roadside WiFi services. In particular, we present a new metric, called *Contact Opportunity*, as a characterization of a roadside WiFi network. Informally, the contact opportunity for a given deployment measures the fraction of distance or time that a mobile user is in contact with some AP when moving through a certain path. Such a metric is closely related to the quality of data service that a mobile user might experience while driving through the system. We then present an efficient deployment method that maximizes the worst case contact opportunity under a budget constraint. We further show how to extend this concept and the deployment techniques to a more intuitive metric – the average throughput – by taking various dynamic elements into account. Simulations over a real road network and experimental results show that our approach achieves more than 200% higher minimum contact opportunity, 30%-100% higher average contact opportunity and a significantly improved distribution of average throughput compared with two commonly used algorithms.

## I. INTRODUCTION

With increasing popularity of media enabled hand-helds, the need for high data-rate services for mobile users is evident. WiFi hotspots are rapidly mushrooming in every city to meet this demand. But, their primary target is static users. These networks fail to provide any assured level of service to a mobile user. Although large deployments of WLANs can be used to provide high data-rate services over large areas, the cost becomes prohibitive due to the sheer number of access-points (APs) required. For instance, to cover a 2km x 2km area in Mountain View, Google needed to deploy 400 access points [12] to barely provide coverage at the base data rate. In addition to the deployment cost, the maintenance and management complexity has led to abandonment or scaling back of several WLAN projects from San Francisco to Philadelphia [3].

New Wireless Wide-Area Networking (WWAN) technologies such as 3GPP LTE (Long Term Evolution) and mobile WiMAX are expected to provide *either* long range coverage *or* high data rates, but practical numbers are far from the

promised levels. For example WiMAX is intended to support data rates as high as 75 Mbps per 20 MHz channel, or a range of 30 miles [23]. However, one of the first deployments of WiMAX in US is reported to provide a downlink bandwidth of 3 Mbps [3], which is only within a factor of 2 better than the current 3G networks. Note that these resources will potentially be shared by a large number of active users within the respective sector of the antenna. Given the resistance from majority of users to pay high monthly fees for mobile data access, which is essential for supporting expensive new deployments, ubiquitous service from such new deployments could take several years, and possibly decades.

On the other hand, evaluation of wireless data access by mobile users using “in situ” (or “open”) WiFi networks [4], [7], [8], [18], and in various controlled environments [4], [10], [18], [20] have confirmed the feasibility of WiFi based vehicular Internet access for non-interactive applications. The possibility and challenges to support certain interactive applications, such as Web browsing, have also been studied [4], [5]. Most existing work, however, consider an unplanned deployment of APs based on open-APs [4], [7], [8], [10], [18], [20]. Consequently, these solutions *fail to provide any throughput assurance* to a mobile user; they can only provide opportunistic services to mobile users.

The two objectives – an economically scalable infrastructure and quality of service assurance – can be achieved by a carefully planned sparse deployment of WiFi APs at roadside. In this paper, we propose to study deployment techniques for providing roadside WiFi services. In particular, we present a new metric, called **Contact Opportunity**, as a characterization of a roadside WiFi network. Informally, the contact opportunity for a given deployment measures the fraction of distance or time that a mobile user is in contact with some AP when moving through a certain path. Such a metric is closely related to the quality of data service that a mobile user might experience while driving through the system. We then present an efficient deployment method that maximizes the worst case contact opportunity under a budget constraint. We further show how to extend this concept and the deployment techniques to a more intuitive metric – the average throughput – by taking various dynamic elements into account.

This is the first work that addresses the challenges in achieving sparse deployment of APs that provides quality of service assurance to the mobile users. The deployment issues with respect to roadside WiFi networks have not received

much attention in the past. The only study in this area is our previous work on Alpha Coverage [25]. In Alpha Coverage, each AP is viewed as a point, and an optimal deployment is found such that one contact is guaranteed for every  $\alpha$ -distance that a mobile user moves in a road network. In other words, only the *number* of contacts is considered but not the *quality* of each contact. In contrast, contact opportunity is more closely related to the real performance that a mobile user experiences by taking various static and dynamic parameters into account, such as the coverage region of each potential location for deploying an AP, driving speed, the data rate of APs (when average throughput is considered). Consequently, finding an optimal deployment in terms of contact opportunity is significantly more challenging.

We make the following contributions in this paper.

- We present a new metric, called Contact Opportunity, as a characterization of roadside WiFi deployment, which is closely related to the quality of data service that a mobile user might experience when driving through the network.
- We present an efficient deployment method that maximizes the worst case contact opportunity given a budget constraint by utilizing submodular optimization techniques.
- We extend the concept of contact opportunity and the corresponding deployment techniques to average throughput by modeling various dynamic parameters.
- Simulations over a real road network [1] and experimental results show that our approach achieves more than 200% higher minimum contact opportunity, 30%-100% higher average contact opportunity and a significantly improved distribution of average throughput compared with two commonly used algorithms.

A number of applications can benefit from a roadside WiFi network deployed using our techniques. Remote monitoring and tracking of shipments is one such application. For example, Walmart currently depends on a satellite based system [2] for tracking its trailers, which is an expensive solution. Similarly, businesses with mobile workforce can benefit from media-rich communication over such a system. Recently, the feasibility and usefulness of a system that provides road condition updates has been studied [13], which is another use case of the system. Our techniques can be applied to either plan a new deployment or to improve an existing one. For instance, a new set of WiFi APs can be added to a set of open APs to improve the quality of service.

We present the formal definition of contact opportunity, the deployment problem statement and our solution in Section II. Several extensions of the concept are presented in Section III. Evaluation based on simulations and experiments are presented in Sections IV and V, respectively, followed by a summary of related work in Section VI, and conclusions in Section VII.

## II. CONTACT OPPORTUNITY OPTIMIZATION

Ideally, we would like to have an economic deployment of APs that is able to serve mobile users moving through a road network with guaranteed performance in terms of some

intuitive metric such as average throughput. Such an objective is complicated by various uncertainties in the system, such as unpredictable traffic conditions, unknown moving patterns of mobile users, and the dynamics involved in the performance of APs. To this end, we use an incremental approach; we introduce a performance metric for roadside AP deployment that is closely related to average throughput while avoiding the uncertainties such that an efficient solution can be obtained. In Section III, several extensions that consider more intuitive performance metrics and more practical system models are introduced.

### A. System Model

We model a road network as a connected *geometric graph*, where vertices represent points where road centerline segments and road intersections meet, and edges represent road centerline segments connecting road intersections. For a curved road segment, we introduce artificial road intersections, so that each edge represents a straight line segment. Without loss of generality, the road network graph is assumed to be undirected.

We assume that  $A$  is a set of known candidate locations in the 2D region covering the road network where APs can be deployed. Associated with each candidate location  $a \in A$ , there is a fixed cost  $w_a \in \mathbb{R}^+$  for installing an AP at  $a$ , and a coverage region  $C_a$ , which is a connected region in the 2D space consisting of the set of points where the received SNR from an AP deployed at  $a$  is higher than a fixed threshold. The coverage regions  $C_a, \forall a \in A$  partition the road network graph into smaller segments called *subsegments*. Figure 1 shows a road network with four roads (lines) and three candidate locations with coverage regions shown as disks that partition the roads into subsegments such as  $va, vb, vc, vd, ce$ , etc. Although the coverage regions are plotted as disks in Figure 1, our problem definitions and solutions are independent of the shape of the coverage region. Also note that such a partition is with respect to a particular deployment and each subsegment is either fully covered or not covered at all.

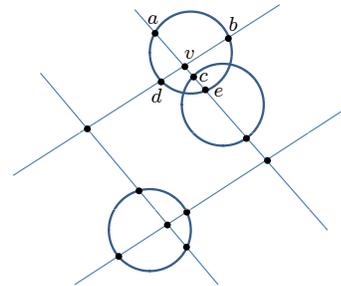


Fig. 1: A road network with four roads (lines) and three candidate locations with coverage regions shown as disks that partition the roads into subsegments such as  $va, vb, vc, vd, ce$ , etc.

Let  $\mathcal{L}$  denote the set of all the subsegments in a road network graph. For each  $l \in \mathcal{L}$ , let  $d_l \in \mathbb{R}^+$  denote the length of the corresponding road centerline segment. For any deployment  $S \subseteq A$ , let  $\mathcal{L}_S \subseteq \mathcal{L}$  denote the set of subsegments covered by  $S$ , that is,  $\mathcal{L}_S = \{l \in \mathcal{L} | l \subseteq \cup_{a \in S} C_a\}$ .

TABLE I: Symbols used in the system model

$A$	The set of all candidate locations for deploying APs
$w_a$	The cost of installing an AP at location $a \in A$
$C_a$	The coverage region of an AP deployed at $a \in A$
$\mathcal{L}$	The set of subsegments
$d_l$	The length of a subsegment $l \in \mathcal{L}$
$\mathcal{L}_S$	The set of subsegments covered by $S \subseteq A$
$A_l$	The set of candidate locations that cover $l \in \mathcal{L}$
$\mathbb{P}$	The set of movements
$\mathcal{L}_p$	The set of subsegments constituting $p \in \mathbb{P}$
$A_p$	The set of candidate locations that cover $p \in \mathbb{P}$

A movement on a road network is modeled as a simple path on the corresponding graph. We assume that there is a set of movements, denoted as  $\mathbb{P}$ , given as part of the input to the deployment decision maker. The concrete definition of  $\mathbb{P}$  is independent of our problem definitions and solutions, while the size of the set  $\mathbb{P}$  impacts the computational complexity and performance guarantee of our solutions as discussed below. For instance,  $\mathbb{P}$  could be a set of shortest (or fastest) paths or a set of most frequently traveled paths. Such information can be learned from a road network database [1] and historical traffic data [11]. For each  $p \in \mathbb{P}$ , let  $\mathcal{L}_p \subseteq \mathcal{L}$  denote the set of subsegments that constitute  $p$ .

### B. Problem Statement

We now define a performance metric for roadside deployment that does not require any information about the dynamics of the system. Given a deployment  $S \subseteq A$ , the **Contact Opportunity in Distance** of a path  $p \in \mathbb{P}$ , denoted as  $\eta_p^d$ , is defined as the fraction of distance on  $p$  that is covered by some AP in  $S$ . Formally,

$$\eta_p^d(S) = \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} d_l}{\sum_{l \in \mathcal{L}_p} d_l}. \quad (1)$$

When a mobile user travels at a constant speed where each AP has the same data rate, and there is only one user in the system, contact opportunity in distance can be directly translated into average throughput that the user will experience. We show in Section III how to extend this concept by taking care of various dynamic elements. Given a budget  $B$ , we are looking for a deployment where the total cost of the APs deployed does not exceed the budget, and the minimum contact opportunity over all movements in  $\mathbb{P}$  is maximized. Such a deployment provides a worst case guarantee and hence does not require statistics about the movement patterns of mobile users. Furthermore, simulation results (in Section IV) show that even though our solution is designed for the worst case, it works well in the average case. Formally, if we let  $w(S)$  denote the cost of a deployment  $S \subseteq A$ , that is  $w(S) = \sum_{a \in S} w_a$ , the optimization problem becomes

$$\max_{S \subseteq A} \min_{p \in \mathbb{P}} \eta_p^d(S), \text{ subject to } w(S) \leq B. \quad (2)$$

### C. Approximation via Submodular Set Cover

In this section, we first show that Problem (2) is an instance of a budgeted version of the submodular set covering problem first studied in [24] and recently extensively explored in [15] and hence allows an efficient bicriterion approximation [15]. We then describe the detailed steps for solving the problem followed by various general and problem specific techniques for accelerating the computation. The following theorem summarizes our results in this section.

**Theorem II.1.** *Let  $OPT(B)$  denote the minimum contact opportunity in distance of an optimal solution to (2). There is a polynomial time algorithm that finds a solution with the minimum contact opportunity to be at least  $OPT(B/\epsilon)$ , where  $\epsilon$  is a logarithmic function of problem parameters defined below.*

To show that (2) can be reduced to the budgeted submodular set covering problem, we note that the set function  $\eta_p^d : 2^A \rightarrow [0, 1]$  satisfies the following properties: (1) nondecreasing, i.e.,  $\eta_p^d(S) \leq \eta_p^d(T)$  whenever  $S \subseteq T \subseteq A$ , (2) normalized, i.e.,  $\eta_p^d(\emptyset) = 0$ , and (3) *submodular*, i.e., for all  $S \subseteq T \subseteq A$  and  $a \in A \setminus T$ ,  $\eta_p^d(S \cup \{a\}) - \eta_p^d(S) \geq \eta_p^d(T \cup \{a\}) - \eta_p^d(T)$ . The last property is formally proved below, which essentially says that *adding a new AP to a small set helps more than adding it to a large set*. It captures our intuition that the total coverage that two APs can provide to a mobile user is reduced if their communication regions overlap with each other.

**Lemma II.1.**  $\eta_p^d$  is submodular.

*Proof:* For any  $S \subseteq T \subseteq A$ ,  $a \in A \setminus T$ ,  $\eta_p^d(S \cup \{a\}) - \eta_p^d(S) = \frac{\sum_{l \in \mathcal{L}_p \cap (\mathcal{L}_{\{a\}} \setminus \mathcal{L}_S)} d_l}{\sum_{l \in \mathcal{L}_p} d_l}$ , and  $\eta_p^d(T \cup \{a\}) - \eta_p^d(T) = \frac{\sum_{l \in \mathcal{L}_p \cap (\mathcal{L}_{\{a\}} \setminus \mathcal{L}_T)} d_l}{\sum_{l \in \mathcal{L}_p} d_l}$ . Since  $S \subseteq T$ ,  $\mathcal{L}_{\{a\}} \setminus \mathcal{L}_S \supseteq \mathcal{L}_{\{a\}} \setminus \mathcal{L}_T$ . Therefore,  $\eta_p^d(S \cup \{a\}) - \eta_p^d(S) \geq \eta_p^d(T \cup \{a\}) - \eta_p^d(T)$ . ■

It follows that (2) is an instance of the budgeted submodular set covering problem, which does not have a polynomial time approximation algorithm unless  $P=NP$  as shown in [15] using a reduction from the hitting set problem. The same argument can be applied to (2) with slight modification. Fortunately, an efficient bicriterion approximation can be achieved by relaxing both the requirement on the objective function and that on the budget. The solution framework proposed in [15] requires first solving the following variant of the problem, which is interesting by itself: given a required minimum contact opportunity  $\lambda \in [0, 1]$  over all the movements, find a deployment of minimum cost. Formally,

$$\min_{S \subseteq A} w(S), \text{ subject to } \min_{p \in \mathbb{P}} \eta_p^d(S) \geq \lambda. \quad (3)$$

A binary search of  $\lambda \in [0, 1]$  is then applied. For each  $\lambda$ , an instance of (3) is solved until a close to optimal solution to (2) is found. Although the budgeted version (2) is hard to approximate, the subproblem (3) allows an efficient approximation since it can be reduced to the submodular set covering problem as follows. Given  $\lambda$ , define:

$$\eta^d(S) = \sum_{p \in \mathbb{P}} \min\{\eta_p^d(S), \lambda\} \quad (4)$$

We note that  $\eta^d$  is also a submodular function since (a)  $\min\{\eta_p^d(S), \lambda\}$  as a set function on  $A$  is submodular when  $\eta_p^d$  is submodular [19] and (b) the sum of submodular functions is submodular. Note that a subset  $S \subseteq A$  is a feasible solution to (3) iff  $\eta^d(S) = \eta^d(A) = |\mathbb{P}|\lambda$ . Therefore, (3) can be reformulated as a submodular set covering problem [24]:

$$\min_{S \subseteq A} w(S), \text{ subject to } \eta^d(S) = \eta^d(A) \quad (5)$$

Due to the submodularity of  $\eta^d$ , (5) allows an efficient greedy approximation. See Algorithm II.1. The algorithm starts with an empty set and in each iteration picks and adds a new candidate location that is most cost effective until the required contact opportunity is achieved. This simple greedy procedure outputs a subset  $S \subseteq A$ , the cost of which never exceeds the cost of the optimal solution by more than a logarithmic factor [24]. In particular, by multiplying all  $\eta_p^d$  by  $10^n$  for some  $n \geq 2$ ,  $\eta^d$  can be made an integer valued function without loss of much accuracy, then an approximation factor  $\epsilon = O(1) + \ln(\max_{a \in A} \eta^d(a))$  can be achieved [24].

---

**Algorithm II.1** Minimum Cost Contact Opportunity
 

---

Input:  $A, \mathbb{P}, \lambda$

Output: A subset  $S \subseteq A$

- 1:  $S \leftarrow \emptyset$ ;
  - 2: **while**  $\eta^d(S) < \eta^d(A)$  **do**
  - 3:   Find  $a \in A \setminus S$  that maximizes  $\frac{\eta^d(S \cup \{a\}) - \eta^d(S)}{w_a}$ ;
  - 4:    $S \leftarrow S \cup \{a\}$ ;
- 

A binary search of  $\lambda \in [0, 1]$  is then applied to solve (2). Let  $B(\lambda)$  denote the total cost of a deployment that achieves  $\lambda$  computed by Algorithm II.1. Starting at  $\lambda = \min_{p \in \mathbb{P}} \eta_p^d(A)$ , that is the minimum contact opportunity when all the candidate locations are used, if  $B(\lambda) > B$ , a lower  $\lambda$  is selected; otherwise, a higher  $\lambda$  is selected. The procedure continues until  $B(\lambda) \leq B$  and  $B(\lambda') > B$  for any  $\lambda' : \lambda' - \lambda \geq \delta$ , where  $\delta$  can be adjusted to control the accuracy. Given a budget  $B$ , such a binary search finds a subset  $S \subseteq A$  that does not exceed the budget and has the minimum contact opportunity to be at least what an optimal solution with budget  $B/\epsilon$  can achieve, that is,  $\min_{p \in \mathbb{P}} \eta_p^d(S) \geq \max_{T \subseteq A, w(T) \leq B/\epsilon} \min_{p \in \mathbb{P}} \eta_p^d(T)$ . Furthermore, this result has shown to be hard to improve [15].

The above procedure can be naturally extended to the case of improving an existing deployment by adding new APs, by substituting all the evaluations of  $\eta_p^d(S)$  with  $\eta_p^d(S \cup A_0)$ , where  $A_0$  denotes the set of APs previously deployed.

#### D. Techniques to Accelerate the Computation

The number of different  $\lambda$  values evaluated by the binary search procedure is  $O(\log(\min_{p \in \mathbb{P}} \eta_p^d(A)/\delta))$ , and for each  $\lambda$ , Algorithm II.1 requires  $O(|A|)$  iterations (line 2 to line 4) where each iteration involves  $|A|$  evaluations of  $\eta^d$ . Hence the

above solution requires  $O(|A|^2 \log(\min_{p \in \mathbb{P}} \eta_p^d(A)/\delta))$  evaluations of  $\eta^d$  and each evaluation involves computing  $\eta_p^d$  for each  $p \in \mathbb{P}$ , which is very time consuming for a large road network, large  $|A|$  and  $|\mathbb{P}|$ . In this section, we describe several techniques that can be applied to accelerate the procedure.

**Techniques to reduce the number of evaluations of  $\eta^d$ :** First, the accelerated greedy algorithm proposed in [17] is applied, which requires significantly fewer evaluations to find a candidate location that maximizes the marginal improvement in each iteration (line 3 in Algorithm II.1) by making use of the submodularity of  $\eta^d$ . Second, whenever the subset found (line 4 in Algorithm II.1) already violates the budget constraint, the procedure moves on to a new  $\lambda$ .

**Techniques to accelerate each evaluation of  $\eta^d$ :** First, we note that for any path  $p \in \mathbb{P}$ , if  $p$  can be divided at certain road intersection into two sub-movements  $p_1, p_2 \in \mathbb{P}$  such that  $\eta_{p_1}^d(S) \geq \lambda, \eta_{p_2}^d(S) \geq \lambda$ , then  $\eta_p^d(S) \geq \lambda$  as well. It follows that only a subset of  $\mathbb{P}$ , denoted as  $\hat{\mathbb{P}}$ , need to be considered in the evaluation of  $\eta^d$ . For instance, suppose  $\mathbb{P}$  is the set of shortest paths in the road network graph of length at least  $\alpha$ , and  $\kappa \ll \alpha$  where  $\kappa$  denotes the maximum edge length in the graph, then  $\hat{\mathbb{P}}$  contains mainly the shortest paths of length from  $\alpha$  to  $2\alpha$  since every shortest path of length greater than  $2\alpha$  can be divided into shortest paths of length between  $\alpha$  and  $2\alpha$  with high chance. Second, we note that each candidate location only contributes to a small subset of  $\hat{\mathbb{P}}$ , and therefore an incremental calculation is more efficient, where  $\eta^d(S \cup \{a\})$  is obtained from  $\eta^d(S)$  by updating only  $\eta_p^d$  for those  $p$  covered by  $C_a$ .

We observe that these techniques improve the performance of our algorithm significantly in practice. For the  $6 \times 6$  road network used in our simulations (in Section IV), the running time to find a solution to (2) is reduced from days to about half an hour under the same machine configuration.

### III. FROM CONTACT OPPORTUNITY TO AVERAGE THROUGHPUT

The concept of contact opportunity in distance discussed above ignores several complexities involved in a real system and does not correspond directly to the real performance a mobile user might experience when driving through the system. Therefore, we seek to design performance metrics that are more intuitive to mobile application designers and end users. To this end, the various uncertainties involved in the system need to be taken into account. The strategy we use is to consider the worst case, and design for it whenever possible.

#### A. Contact Opportunity in Time

The first natural extension to contact opportunity in distance is to replace distance with time, which informally is defined to be the fraction of time when a mobile user moving through a path is within the range of some AP. The main difficulty with such a definition is that both the contact time and the travel time are not fixed due to the uncertainties of traffic conditions such as traffic jams, accidents and stop signs.

To model these uncertainties, we follow the interval based modeling approach [14] and assume that for each  $l \in \mathcal{L}$ , the possible travel time  $t_l$  over  $l$  varies within an interval  $[b_l, u_l]$ ,  $0 < b_l \leq u_l$ , which can be learned from historical traffic data. We define a *scenario*  $k$  to be an assignment of travel time  $t_l(k) \in [b_l, u_l]$  to each  $l \in \mathcal{L}$ . Let  $K$  denote the set of all possible scenarios. Given a deployment  $S \subseteq A$  and a scenario  $k \in K$ , the **Contact Opportunity in Time** of a path  $p \in \mathbb{P}$  is defined to be:

$$\eta_p^t(S, k) = \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} t_l(k)}{\sum_{l \in \mathcal{L}_p} t_l(k)}. \quad (6)$$

Although there are infinitely many scenarios, we seek to find a deployment that performs well even in the worst case. Fortunately, a worst case scenario for a given deployment  $S \subseteq A$  can be easily determined: let  $k_S$  denote the scenario where  $t_l(k_S) = b_l$  if  $l$  is covered by some AP in  $S$ , that is if  $l \in \mathcal{L}_S$ , and  $t_l(k_S) = u_l$  otherwise. We claim that  $k_S$  is a worst case scenario for  $S$ , that is:

**Proposition III.1.** *Every  $\eta_p^t(S, k)$ ,  $p \in \mathbb{P}$  is minimized across all the scenarios if  $k = k_S$ .*

*Proof:* For every  $p \in \mathbb{P}$ ,  $\eta_p^t(S, k) = \frac{t_1}{t_2 + t_1}$  where  $t_1$  denotes the travel time over the set of subsegments in  $p$  that are covered by  $S$  and  $t_2$  denotes the travel time over the other subsegments in  $p$ . Note that  $\eta_p^t(S, k)$  is minimized when  $t_1$  is minimized and  $t_2$  is maximized, which happens at  $k_S$ . ■

Given a budget  $B$ , we would like to find a deployment that does not exceed the budget while maximizing the minimum contact opportunity in time over all the movements in the worst case scenario. Formally, we want to solve

$$\max_{S \subseteq A} \min_{p \in \mathbb{P}} \eta_p^t(S, k_S), \text{ subject to } w(S) \leq B. \quad (7)$$

Note that  $\eta_p^t(S, k_S)$  can be viewed as a set function since whenever a deployment  $S \subseteq A$  is given,  $\eta_p^t(S, k_S)$  can be evaluated. However,

**Proposition III.2.**  *$\eta_p^t$  as a set function is nondecreasing and normalized, but not submodular.*

A formal proof of the proposition is given in the Appendix. It follows that the solution framework for (2) when applied here (by replacing  $\eta_p^d(S)$  by  $\eta_p^t(S, k_S)$ ) does not guarantee the same approximation factor.

We note that for any fixed scenario  $k$ ,  $\eta_p^t(S, k)$  is submodular and hence the above framework can still be applied. The general problem can be viewed as a robust optimization problem and is more challenging. In fact, it has been observed that the robust counterparts of many polynomial time solvable optimization problems are NP-hard [14]. Although an efficient solution with guaranteed performance to the general problem remains open, we propose the following two approaches as first steps that work well in many practical cases.

First, we study the potential loss if we only design for a fixed scenario. Let  $k_0$  denote the ‘‘mean’’ scenario where  $t_l(k_0) = (b_l + u_l)/2, \forall l \in \mathcal{L}$ . It turns out that, if  $u_l/b_l$  is

small for all  $l \in \mathcal{L}$ ,  $k_0$  can be used as a good approximation. More concretely, let  $S$  denote an optimal deployment with respect to  $k_0$  and let  $T$  denote an optimal solution to (7), we have the following theorem, where to simplify the notion, we write  $\tilde{\eta}(S, k)$  to denote the minimum contact opportunity in time across all the movements under the deployment  $S$  and scenario  $k$ . The theorem says that if  $u_l/b_l$  is bounded by  $\beta \geq 1$ , then the minimum contact opportunity in time (with respect to the worst case scenario) achieved by an optimal deployment computed using the ‘‘mean’’ scenario is at most a factor  $\beta$  worse than the optimal solution. A formal proof is given in the Appendix.

**Theorem III.1.** *If  $u_l/b_l \leq \beta$  for all  $l \in \mathcal{L}$ , then  $\tilde{\eta}(S, k_S) \leq \tilde{\eta}(T, k_T) \leq \beta \tilde{\eta}(S, k_S)$ .*

We then propose a solution to the general problem without making assumptions on  $u_l/b_l$ , which works well if the set of candidate locations that cover  $p$ , denoted as  $A_p$ , has small cardinality, for every  $p \in \mathbb{P}$ . For each pair of path  $p$  and scenario  $k$ , write  $\eta_{p,k}^t(S) = \eta_p^t(S, k)$ , where we move  $k$  to the subscript to emphasize that we are now considering  $\eta_{p,k}^t : 2^A \rightarrow [0, 1]$  as a set function for each  $(p, k)$  pair. Notice that the objective in (7) is equivalent to

$$\max_{S \subseteq A} \min_{p \in \mathbb{P}, k \in K} \eta_{p,k}^t(S), \text{ subject to } w(S) \leq B. \quad (8)$$

This formulation essentially says that we are looking for a deployment that maximizes the worst case contact opportunity across all the movements as well as all the scenarios. Since each  $\eta_{p,k}^t$  is nondecreasing and submodular, we can apply the framework discussed before. The only issue is that the possible scenarios could be a huge set. The good news is that for each  $p$ , we only need to consider the set of scenarios  $\{k_S : S \subseteq A_p\}$ . The remaining scenarios are either not distinguishable to  $p$  or will never be a worst case scenario. Hence the total number of relevant  $(p, k)$  pairs is  $\sum_{p \in \mathbb{P}} 2^{|A_p|}$ . When  $|A_p|$  is small (e.g.  $< 10$ ) for every  $p$ , this formulation gives an efficient solution with guaranteed performance.

### B. Average Throughput for a Single Mobile User

We next extend the concept of contact opportunity in time to average throughput for a single mobile user. To this end, we need to consider a more complex system model. First, we assume each access point deployed at a candidate location  $a \in A$  has a fixed worst case data rate  $r_a \in \mathbb{R}^+$ . Second, let  $A_l \subseteq A$  denote the set of candidate locations that cover  $l \in \mathcal{L}$ . We assume that a user moving through  $l$  always selects the AP with the highest rate among  $A_l$  to associate with. Third, we assume that connection setup time when a user starts to associate with an AP and the handoff time when a user switches between APs are small enough to be ignored. Then for a given deployment  $S \subseteq A$  and a traffic condition scenario  $k$ , the average throughput for a user moving through a path  $p$  can be stated as follows:

$$\eta_p^s(S, k) = \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} t_l(k) \max_{a \in A_l \cap S} r_a}{\sum_{l \in \mathcal{L}_p} t_l(k)}. \quad (9)$$

We again want to find a deployment that is within the budget while maximizing the minimum average throughput in the worst case scenario across all the movements. We note that  $\eta_p^s$  is not submodular in general, but is submodular for any fixed scenario, and the solutions proposed for contact opportunity in time can be applied as well.

### C. Average Throughput for Multiple Mobile Users

Finally, we discuss an approach to take multiple mobile users into account by making two additional assumptions. First, for any subsegment  $l \in \mathcal{L}$ , the maximum number of mobile users moving through  $l$  at any time instant is known, and is denoted as  $\nu_l$ , which can be learned from historical traffic data. For any candidate location  $a \in A$ , let  $\nu_a$  denote the maximum number of users within  $C_a$ , that is  $\nu_a = \sum_{l \in \mathcal{L}_{\{a\}}} \nu_l$ . Second, define the normalized data rate of an AP deployed at  $a \in A$  to be  $\hat{r}_a = r_a/\nu_a$ . We assume that a user moving through  $l$  always selects the AP with the highest normalized rate among  $A_l$  to associate with. Note that this is consistent with the scheme defined in the previous section where  $\nu_a = 1$ . With these assumptions, for a given deployment  $S \subseteq A$  and a traffic condition scenario  $k$ , the estimated average throughput in the worst case traffic condition scenario for *each* user moving through a path  $p$  can be stated as follows:

$$\eta_p^m(S, k) = \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} t_l(k) \max_{a \in A_l \cap S} \hat{r}_a}{\sum_{l \in \mathcal{L}_p} t_l(k)}. \quad (10)$$

Similar optimization problem as before can then be studied with respect to  $\eta_p^m$ , which is submodular for any fixed scenario. This problem formulation, however, has some limitations. First, the simplified association protocol does not fully model the reality and hence the estimated average throughput could be far from the real throughput for any particular user. Second, fairness among mobile users has been ignored. We propose to study more sophisticated joint deployment and association models that take fairness into account as part of the future work.

## IV. SIMULATIONS

In this section, we evaluate our solution to the roadside AP deployment problem via simulations over a real road network retrieved from the 2008 Tiger/Line shapefiles [1] to understand its performance, including the minimum and average contact opportunity in time or distance, and the distribution of contact opportunity and average throughput across a set of movements under various budget constraints.

### A. Baseline Algorithms

Let  $\hat{A} \subseteq A$  denote the set of candidate locations covering  $\mathbb{P}$ . The following two baseline algorithms are used:

**Uniform Random Sampling** (Rand for short), which at each step randomly picks a new element from  $\hat{A}$  as long as the total cost of the selected elements does not exceed the budget.

**Max-Min Distance Sampling** [22] (Dist for short), which starts at a randomly selected location in  $\hat{A}$ , and at each step finds a new element from  $\hat{A}$  that maximizes the minimum

graph distance (in terms of shortest paths) from the elements already selected, as long as the total cost of the selected elements does not exceed the budget.

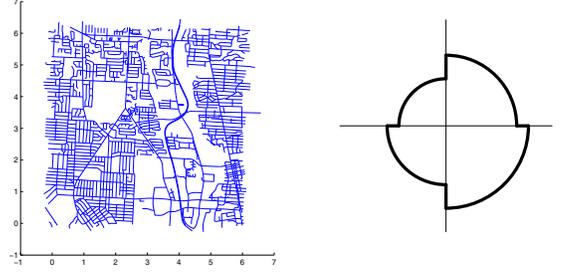


Fig. 2: Left: A road network spanning an  $6 \times 6 \text{ km}^2$  region used in simulations. Right: An instance of AP's coverage region with its boundary highlighted.

### B. Simulation Settings

Figure 2(left) shows the road network used in simulations, where we only consider the largest connected component of the corresponding graph, which has 1802 road intersections and each of them is assumed to be a candidate location for deploying APs. Each edge is associated with an interval of travel speed:  $[(\tau - 5) \text{ m/s}, \tau \text{ m/s}]$ , where  $\tau$  is randomly selected from  $[10, 20]$ . The coverage region at each candidate location is modeled using a simplified version of a sector based approach recently proposed in [21], where each region is composed of 4 sectors of  $90^\circ$  with radius randomly selected from  $[100\text{m}, 200\text{m}]$ , as shown in Figure 2(right). Each candidate location is assigned a unit cost. The set of movements considered are the set of all the shortest paths of length at least 2km between any two road intersections (there are about  $1.4 \times 10^6$  such paths in the road network).

The following variants of our algorithms are used in the evaluation. For contact opportunity in distance, the algorithm presented in Section II-C is used. The parameter  $\delta$  used in the binary search is set to 0.5%. For contact opportunity in time, the same algorithm is extended to consider time instead of distance and uses the “mean” traffic scenario based approach. We also extend the same algorithm to the average throughput case by again using the “mean” scenario, where the data rate of each AP is set to 1Mbps and a large movement file is generated (see the next paragraph) to estimate user density for multiple user case. For each of the baseline algorithms, 100 different deployments are evaluated and the average or cumulative results across all the deployments are plotted.

We further carry out  $ns$ -2 based packet level simulations to evaluate the performance in terms of average throughput for both a single mobile user and multiple user scenarios. Each user moves in the network following a restricted random waypoint model composed of randomly selected shortest paths of length at least 2km, and the travel speed over each edge is randomly selected from the corresponding interval. We did not start simulations at steady-state since it is too time consuming for large graphs. The physical layer data rate of each AP is

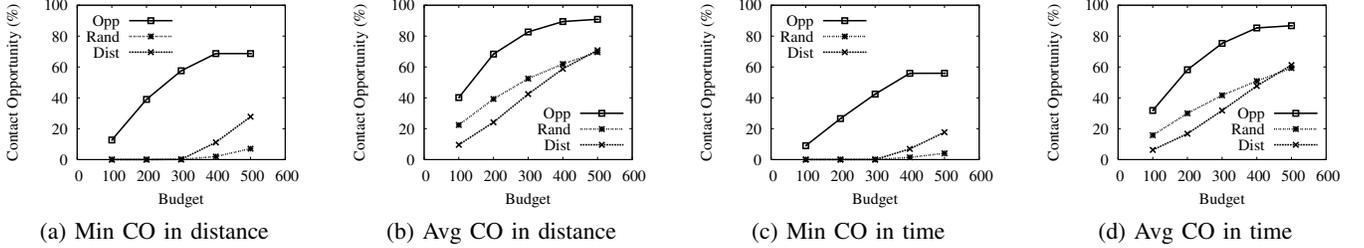


Fig. 3: The minimum and average contact opportunity (CO) across all the movements. For (c) and (d), the contact opportunity in time corresponding to the worst case traffic condition is plotted. For the two baseline algorithms, the results are averaged over all the deployments.

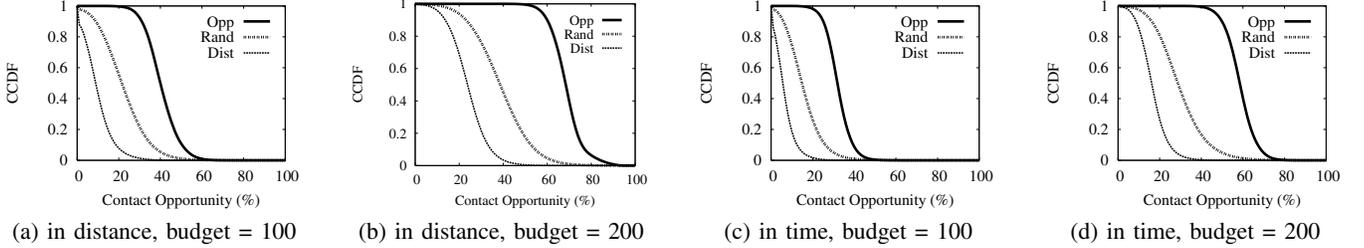


Fig. 4: The complementary cumulative distribution of contact opportunity (for all the shortest paths  $\geq 2$ km) in distance and time. For (c) and (d), the contact opportunity in time corresponds to the worst case traffic condition is plotted.

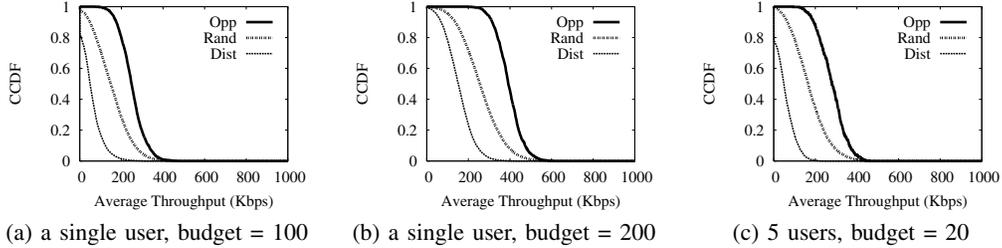


Fig. 5: The complementary cumulative distribution of average throughput (over every 400s interval). (a) and (b): a single mobile user in the large network; (c): 5 mobile users in the small network. The results are cumulative across all the users.

set to 1Mbps. A mobile user selects one of the APs in range (if there is one) with minimum load to associate and then downloads CBR data (1Mbps) from that AP, where the load of an AP is the number of users currently associated with it. To avoid frequent switching between APs, a node switches from AP  $a$  to  $b$  only if the current load of  $a$  is higher than the current load of  $b$  plus 1. For the single user case, a movement file describes a continuous move for  $10^4$  minutes (more than 1500 paths are traveled) is generated, and the deployment computed by our algorithm and the first 10 of the 100 deployments computed by each baseline algorithm are evaluated. For the multiple user case, a large movement file is first generated as the training data to estimate the worst case user density in the coverage region of each candidate location. Due to the time required to simulate multiple user scenarios, we use a smaller road network (a  $2\text{km} \times 2\text{km}$  subregion in the same area) and simulate 5 mobile users where each user moves continuously for 1000 minutes.

### C. Simulation Results

Figure 3 shows the minimum and average (across all the movements) contact opportunity in distance and time achieved by our solution (Opp for short) and the two baseline algorithms

under various budgets, where a budget is simply the number of APs allowed to use since each candidate location is assigned a unit cost. For contact opportunity in time, the value for the worst case traffic condition is plotted. Our algorithm performs significantly better than the baselines in all the cases. The minimum contact opportunity is more than 200% higher. In fact, when the budget is low, the minimum contact opportunity of the two baseline algorithms is very close to 0. Even if our solution is designed for the worst case, it still achieves 30%-100% higher average contact opportunity. Furthermore, we observe that (not shown in the figure) the contact opportunity when deploying one AP at each of the 1802 candidate locations is very close to the one achieved by deploying only 400 APs computed by our solution. We also observe that the max-min sampling performs better than random sampling in terms of minimum contact opportunity but worse in terms of average performance unless the budget is very high. One explanation is that the max-min sampling distributes APs in a more uniform way and when the budget is low, it does not provide enough coverage to shorter movements.

Figure 4 shows the complementary cumulative distribution of contact opportunity in distance and time for the three algorithms, and our algorithm performs clearly better. We

then perform  $ns-2$  simulations over the set of deployments used in Figures 4(c) and 4(d), and compute the average throughput that a single mobile user experiences for every 400 seconds interval, as shown in Figures 5(a) and 5(b). Even if the user only visits a small subset of all the shortest paths, our solution still performs much better than the baselines. Figure 5(c) shows the result for the 5 user case. We observe that the performance of different users are similar and hence the cumulative results across all the users are plotted.

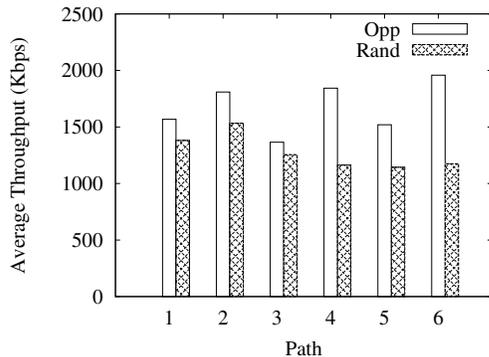


Fig. 6: The average throughput of the 6 paths under evaluation, where Rand represents the average of 5 random deployments.

## V. EXPERIMENTAL EVALUATION

We set up a small scale controlled experiment to better understand the performance of our approach. The experiment was carried out in a  $180\text{m} \times 120\text{m}$  parking lot located at the west campus of OSU and is free of potential interference from other WiFi networks. The experiment was usually carried out at night when the parking lot was empty. We artificially divided the parking lot area into a 6 by 4 grid and use it as a small road network. All the 24 intersections are treated as candidate locations for deploying APs.

A single mobile node carried by a car and 4 APs are used in the experiment. Each AP is a laptop equipped with an Orinoco 802.11b/g PC card and an external antenna mounted on a 1.7m high tripod so that the signal will not be blocked by the car in the test. The single mobile node is a laptop equipped with a Ubiquiti Networks SRC 802.11a/b/g PC card and two external antennas fixed at the two sides of the car. The transmission power of each AP is set to 6 dBm, which is tested to give an effective transmission distance of no more than 50 meters. Each node runs Ubuntu Linux with Linux 2.6.24 kernel and madwifi device driver for the 802.11 interface. The physical layer data rate of each node is fixed at 54Mbps.

A total of 5 random deployments are evaluated and compared with a deployment computed by our algorithm that maximizes the minimum contact opportunity in distance across the set of shortest paths between intersections of length at least 200m (there are 30 such paths in total). The algorithm assumes that each AP has a unit cost and the coverage region of each AP is a disk of a radius 50m.

Because of the large volume of driving work and limited availability of that place, we picked 6 representative shortest paths that go through different parts and directions of the

parking lot, and drove through each of them 3 times for each deployment. The moving speed is kept at about 10mph. When moving through a path, the mobile node attempts to associate with an AP with the strongest signal. Once associated, it downloads UDP packets from the AP until it is disconnected from that AP. The mobile node then finds another AP with the strongest signal to associate. Figure 6 shows the average throughput of each of the 6 paths. For random sampling, the average results across the five deployments are plotted. We observe that our solution achieves up to 66.7% higher throughput, and across all the 6 paths the average improvement in throughput is 26.4%.

## VI. RELATED WORK

The idea of Drive-thru Internet by connecting to existing roadside Access Points is introduced in [20], which shows that a single moving vehicle connected via 802.11b with an AP located at roadside of an empty street can access several megabytes of TCP or UDP traffic, even when the velocity is as high as 180 km/h. Subsequently, evaluations in various controlled environments [4], [10], [18], [20] and in situ WiFi networks [4], [7], [8], [18] have been conducted, further confirming the feasibility of WiFi-based Vehicular Internet Access for non-interactive applications. However, the deployment issues with respect to WiFi-based Vehicular Internet Access have not been carefully studied so far. Instead, an unplanned deployment of APs is commonly assumed in most previous works. A simple non-uniform strategy that places more stationary nodes in the network core was considered in a recent work [6]. However, it was completely based on intuition without providing any performance guarantees. The only exception is our previous work on Alpha Coverage [25], which however, only considers the number of contacts but ignores the quality of each contact.

Submodular functions play a critical role in combinatorial optimization. The theory was first developed half a century ago. Since then, various submodular optimization problems have been intensively studied [9], [19]. Although the submodular minimization problem is polynomial time solvable, the submodular maximization problem is NP-hard. However, a simple greedy algorithm gives an  $(1 - 1/e)$  approximation ratio when the function is also nondecreasing and normalized. The submodular set covering problem was first studied in [24]. Recently, submodular optimization has been applied to several network deployment problems, including placing sensors to efficiently detect outbreak [16], or to provide robust observations [15]. The former is formulated as a submodular maximization problem while the latter is formulated as a budgeted submodular set covering problem.

## VII. CONCLUSION

Existing solutions for wide-area data services fail to provide an economical solution for high data-rate services. This paper takes one more step towards this target by presenting Contact Opportunity in terms of distance or time as a new metric for the characterization of roadside WiFi networks and an efficient

deployment approach that maximizes the worst case contact opportunity for a given budget. We further extend this metric to average throughput, which is more intuitive for mobile users and application designers. Using simulations and experiments, we show that our approach achieves much better performance in terms of both contact opportunity and average throughput compared with two commonly used deployment techniques.

### VIII. ACKNOWLEDGMENT

This material is based upon work supported by NSF grants CNS-0546630 (CAREER Award), CNS-0403342, CNS-0721983, CNS-0910878, CCF-0728928, NIH Grant U01DA023812 from National Institute for Drug Abuse (NIDA), and Fedex Institute of Technology (FIT) at the University of Memphis.

### APPENDIX

**Proof of Proposition III.2:**  $\eta_p^t$  is normalized by definition. For every  $S \subseteq T \subseteq A$ , we have  $\mathcal{L}_p \cap \mathcal{L}_S \subseteq \mathcal{L}_p \cap \mathcal{L}_T$  and  $\mathcal{L}_p \setminus \mathcal{L}_S \supseteq \mathcal{L}_p \setminus \mathcal{L}_T$ . Therefore,  $\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l \leq \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_T} b_l$  and,  $\sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_S} u_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l = \sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_T} u_l + \sum_{l \in \mathcal{L}_p \cap (\mathcal{L}_T \setminus \mathcal{L}_S)} u_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l \geq \sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_T} u_l + \sum_{l \in \mathcal{L}_p \cap (\mathcal{L}_T \setminus \mathcal{L}_S)} b_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l = \sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_T} u_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l$ . Hence  $\eta_p^t(S) = \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l}{\sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_S} u_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_S} b_l} \leq \frac{\sum_{l \in \mathcal{L}_p \cap \mathcal{L}_T} b_l}{\sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_T} u_l + \sum_{l \in \mathcal{L}_p \cap \mathcal{L}_T} b_l} = \eta_p^t(T)$ . So  $\eta_p^t$  is nondecreasing.

To see that  $\eta_p^t$  is not submodular, it suffices to give a counter-example. Let  $p$  be a path from  $(0, 0)$  to  $(3, 0)$  in  $\mathbb{R}^2$  and there is no other road intersection in  $p$  other than the two endpoints. Assume that there are three candidate locations at points  $(i + 0.5, 0), i = 0, 1, 2$ , where each of them covers an interval  $[(i, 0), (i + 1, 0)], i = 0, 1, 2$  in  $p$ , and  $p$  is partitioned into 3 subsegments by them. Suppose the travel time of each subsegment varies over  $[1, 2]$ . Let  $S = \{(0.5, 0)\}$ ,  $T = \{(0.5, 0), (1.5, 0)\}$ , and  $a = (2.5, 0)$ . Then  $S \subseteq T, a \notin T$ .  $\eta_p^t(S) = 1/(2 \times 2 + 1) = 0.2$ ,  $\eta_p^t(S \cup \{a\}) = \eta_p^t(T) = 2/(1 \times 2 + 2) = 0.5$ , and  $\eta_p^t(T \cup \{a\}) = 1$ . Therefore,  $\eta_p^t(S \cup \{a\}) - \eta_p^t(S) = 0.3$ , while  $\eta_p^t(T \cup \{a\}) - \eta_p^t(T) = 0.5$ .

**Proof of Theorem III.1:** First, we have  $\tilde{\eta}(S, k_S) \leq \tilde{\eta}(T, k_T) \leq \tilde{\eta}(T, k_0) \leq \tilde{\eta}(S, k_0)$ , where the first inequality follows from the optimality of  $T$  with respect to worst case scenarios, and the second inequality follows from the fact that  $k_T$  is a worst case scenario with respect to  $T$ , and the third inequality follows from the optimality of  $T$  with respect to  $k_0$ . To show that  $\tilde{\eta}(T, k_T) \leq \beta \tilde{\eta}(S, k_S)$ , it suffices to show that  $\tilde{\eta}(S, k_0) \leq \beta \tilde{\eta}(S, k_S)$ . For every path  $p \in \mathbb{P}$ , let  $\mathcal{L}_0 \subseteq \mathcal{L}_p$  denote the set of subsegments covered by  $S$ . Let  $t_1$  and  $t_2$  denote the travel time over covered subsegments in scenario  $k_0$  and  $k_S$ , respectively, that is,  $t_1 = \sum_{l \in \mathcal{L}_0} (b_l + u_l)/2$  and  $t_2 = \sum_{l \in \mathcal{L}_0} b_l$ . Let  $T_1$  and  $T_2$  denote the total travel time over  $p$  in scenario  $k_0$  and  $k_S$ , respectively, that is,  $T_1 = \sum_{l \in \mathcal{L}_p} (b_l + u_l)/2$  and  $T_2 = \sum_{l \in \mathcal{L}_p \setminus \mathcal{L}_0} u_l + \sum_{l \in \mathcal{L}_0} b_l \leq \sum_{l \in \mathcal{L}_p} u_l$ . Then

$\eta_p^t(S, k_0) = \frac{t_1}{T_1}$ ,  $\eta_p^t(S, k_S) = \frac{t_2}{T_2}$ , and  $\frac{\eta_p^t(S, k_0)}{\eta_p^t(S, k_S)} = \frac{t_1}{t_2} \frac{T_2}{T_1}$ . When  $u_l/b_l \leq \beta$  for all  $l \in \mathcal{L}$ , we have  $\frac{t_1}{t_2} \leq \frac{1+\beta}{2}$  and  $\frac{T_2}{T_1} \leq \frac{2\beta}{1+\beta}$ . Hence  $\frac{\eta_p^t(S, k_0)}{\eta_p^t(S, k_S)} \leq \beta$ . Since this inequality holds for every  $p \in \mathbb{P}$ ,  $\tilde{\eta}(S, k_0) \leq \beta \tilde{\eta}(S, k_S)$ .

### REFERENCES

- [1] U.S. Census Bureau - TIGER/Line <http://www.census.gov/geo/www/tiger/>.
- [2] VeriWise Asset Intelligence. <http://www.ge.com/equipmentservices/assetintelligence/>.
- [3] Taking Wireless to the Max. *Business Week (businessweek.com/go/techmaven)*, pages 101–102, Nov. 2008.
- [4] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan. Interactive WiFi Connectivity For Moving Vehicles. In *Proc. of ACM SIGCOMM*, Sept. 2008.
- [5] A. Balasubramanian, Y. Zhou, W. B. Croft, B. N. Levine, and A. Venkataramani. Web Search From a Bus. In *Proc. of CHANTS*, Sept. 2007.
- [6] N. Banerjee, M. Corner, D. Towsley, and B. Levine. Relays, Base Stations and Meshes: Enhancing Mobile Networks with Infrastructure. In *Proc. of ACM MOBICOM*, Sept. 2008.
- [7] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden. A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. In *Proc. of ACM MOBICOM*, Sept. 2006.
- [8] J. Eriksson, H. Balakrishnan, and S. Madden. Cabernet: A WiFi-Based Vehicular Content Delivery Network. In *Proc. of ACM MOBICOM*, Sept. 2008.
- [9] S. Fujishige. *Submodular Functions and Optimization, Volume 58, Second Edition (Annals of Discrete Mathematics)*. Elsevier Science, 2005.
- [10] R. Gass, J. Scott, and C. Diot. Measurements of In-Motion 802.11 Networking. In *Proc. of WMCSA*, Apr. 2006.
- [11] H. Gonzalez, J. Han, X. Li, M. Myslinska, and J. P. Sondag. Adaptive Fastest Path Computation on a Road Network: a Traffic Mining Approach. In *Proc. of VLDB*, Sept. 2007.
- [12] Google WiFi. Google's Mountain View WiFi Network. <http://wifi.google.com/>.
- [13] Joint Nokia Research Project Captures Traffic Data using GPS-enabled Cell Phones. [http://www.berkeley.edu/news/media/releases/2008/02/08\\_gps.shtml](http://www.berkeley.edu/news/media/releases/2008/02/08_gps.shtml), 2008.
- [14] P. Kouvelis and G. Yu. *Robust Discrete Optimization and Its Applications*. Springer, 1996.
- [15] A. Krause, B. McMahan, C. Guestrin, and A. Gupta. Robust Submodular Observation Selection. *Journal of Machine Learning Research (JMLR)*, 9:2761–2801, 2008.
- [16] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and N. Glance. Cost-effective Outbreak Detection in Networks. In *Proc. of ACM SIGKDD*, 2007.
- [17] Michel Minoux. Accelerated Greedy Algorithms for Maximizing Submodular Set Functions. *Optimization Techniques, LNCS*, 7:234–243, 1978.
- [18] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-giel, and S. R. Das. MobiSteer: Using Steerable Beam Directional Antenna for Vehicular Network Access. In *Proc. of MOBISYS*, June 2007.
- [19] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1988.
- [20] J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for "Automobile" Users. In *Proc. of INFOCOM*, Mar. 2004.
- [21] J. Robinson, R. Swaminathan, and E. W. Knightly. Assessment of Urban-Scale Wireless Networks with a Small Number of Measurements. In *Proc. of IEEE MOBICOM*, 2008.
- [22] S.-H. Teng. Mutually repellent sampling. *Minimax and its Applications, Ding-Zu Du and Panos M. Pardalos ed., Kluwer Academic Publishers*, pages 129–140, 1995.
- [23] WiMax.com FAQ. <http://www.wimax.com/education/faq/>, 2008.
- [24] L. A. Wolsey. An Analysis of the Greedy Algorithm for the Submodular Set Covering Problem. *Combinatorica*, 2(4):385–393, 1982.
- [25] Z. Zheng, P. Sinha, and S. Kumar. Alpha Coverage: Bounding the Interconnection Gap for Vehicular Internet Access. In *Proc. of IEEE INFOCOM Mini-Conference*, Apr. 2009.