iDEAL: Incentivized Dynamic Cellular Offloading via Auctions

Wei Dong†, Swati Rallapalli†, Rittwik Jana‡, Lili Qiu†, K. K. Ramakrishnan‡, Leo Razoumov‡, Yin Zhang†, Tae Won Cho‡
† The University of Texas at Austin ‡ AT&T Labs – Research

Abstract — The explosive growth of cellular traffic and its highly dynamic nature often make it increasingly expensive for a cellular service provider to provision enough cellular resources to support the peak traffic demands. In this paper, we propose iDEAL, a novel auction-based incentive framework that allows a cellular service provider to leverage resources from third-party resource owners on demand by buying capacity whenever needed through reverse auctions. iDEAL has several distinctive features: (i) iDEAL explicitly accounts for the diverse spatial coverage of different resources and can effectively foster competition among third-party resource owners in different regions, resulting in significant savings to the cellular service provider. (ii) iDEAL provides revenue incentives for third-party resource owners to participate in the reverse auction and be truthful in the bidding process. (iii) iDEAL is provably efficient. (iv) iDEAL effectively guards against collusion. (v) iDEAL effectively copes with the dynamic nature of traffic demands. In addition, iDEAL has useful extensions that address important practical issues. Extensive evaluation based on real traces from a large US cellular service provider clearly demonstrates the effectiveness of our approach. We further demonstrate the feasibility of iDEAL using a prototype implementation.

I. Introduction

The explosive growth of cellular traffic and its highly dynamic nature make it increasingly expensive for a cellular service provider to provision enough cellular resources to support all her consumers all the time. The current best practice is for service providers to augment the cellular network capacity by deploying alternative wireless technologies (*e.g.*, Wi-Fi and femtocells) on their own. While this approach is helpful in alleviating the stress at the busiest cellular regions in a short term, it alone is not sufficient in the long run due to the high deployment cost and excessive interference.

Our solution is to leverage resources *on demand* from third-party resource owners by buying capacity whenever needed. On-demand purchase of such resources can potentially lead to a win-win solution: the cellular service provider achieves significant savings by not having to provision for the peak traffic demands; the third-party resource owners gain additional revenue from the otherwise wasted spare capacity; the overall user experience is also improved. In order for this approach to be successful, however, it is essential to have an incentive framework that can effectively foster collaboration while guarding against non-truthful and collusive behavior.

Our approach: Incentivizing cellular offloading via auctions: We propose iDEAL, a novel auction-based incentive framework to enable dynamic offloading of cellular traffic. In iDEAL, a cellular service provider purchases bandwidth on demand from third-party resource owners, who may be a Wi-Fi hotspot owner, a femtocell owner, or another cellular service provider. This auction problem is naturally formulated as a *reverse auction*, where the goods of interest are bandwidth resources, third-party hotspot owners serve as sellers (*i.e.*, bidders or auctionees) and submit their bids while provider A or a trusted third party serves as an auctioneer, who evaluates the bids from all hotspot owners and makes decisions regarding whose services to purchase in order to satisfy A's traffic demands and minimize A's total cost. Each bidder submits a bid that specifies the total amount of bandwidth she offers in the next time interval and the unit price she asks for. After collecting all the bids, the cellular service provider determines (i) an *allocation*, *i.e.*, how to allocate her traffic between different third-party resource owners (depending on the region they cover) and her own cellular network, and (ii) a *price*, *i.e.*, how much she pays each third-party resource owner that offloads cellular traffic.

The use of reverse auction is motivated by the following observations. First, a key challenge in utilizing resources from third-party resource owners is that we do not know their cost function. Their cost function may be based on multiple considerations, some of which may not be revealed to the cellular service provider. Reverse auctions provide a formal framework for third-party resource owners to express the price they demand and for the cellular service provider to optimize the allocation based on the received bids. Second, by using reverse auctions, the cellular service provider avoids having to negotiate a long-term bi-lateral agreement with each individual third-party resource owner. Negotiating such long-term agreements is difficult and possibly inefficient due to dynamic traffic demands and resource availability. Instead, the cellular service provider can now establish short-term contracts with third party resource providers. It also potentially cuts costs by leveraging competition across third-party resource owners. Third, reverse auctions can be incrementally deployed today, yielding savings to the cellular service provider even when only a subset of third-party resource owners participate.

Unique challenges: While reverse auction has been applied to cellular offloading in the past (*e.g.*, [10]), our problem setting poses several unique challenges. Despite their importance, none of these challenges have been considered earlier.

- *Diverse spatial coverage.* Cellular resources can serve traffic anywhere in a cell sector (albeit at different rates depending on path loss etc.), whereas Wi-Fi hotspots and femtocells have a much more limited communication range, making it essential to consider the spatial coverage of different resources. However, one cannot simply partition resources into separate regions and launch independent reverse auctions within each region, because the longer-range cellular resource introduces coupling between the Wi-Fi hotspots or femtocells in different regions. For example, buying more resources from a cheaper Wi-Fi hotspot in one region frees up more cellular resources, which reduces the amount of cellular traffic to be offloaded in regions with more expensive Wi-Fi hotspots.
- *Traffic uncertainty.* Cellular traffic is highly dynamic and

unpredictable. Since the cellular service provider has to purchase third-party resources based on predicted traffic demands at a future time, it can easily result in under-provisioning or over-provisioning without an effective technique to cope with traffic uncertainties. In contrast, in conventional reverse auction settings, the total amount of goods that the buyer wants is typically known *a priori*.

- *Non-truthful bidding and collusion.* It is essential for us to explicitly guard against both non-truthful bidding and collusion. Due to the distributed nature of hotspot locations, collusion in our context is quite different from what was studied previously and calls for a new study to understand possible collusion strategies and mitigate them.

Contributions: Our paper makes three main contributions.

1. We design the iDEAL incentive framework to effectively address the above unique challenges. Compared with conventional mechanisms for reverse auctions, iDEAL has the following distinctive features: (i) iDEAL explicitly accounts for the spatial coverage of different resources and can effectively foster competition among third-party resource owners in different regions, resulting in significant savings to the cellular service provider. (ii) iDEAL incentivizes bidders (*i.e.*, third-party resource owners) to participate in the reverse auction and to be truthful in their bidding. (iii) iDEAL is provably efficient in that the winners are the bidders who have the lowest valuation of their resources. (iv) iDEAL can effectively mitigate collusion. (v) iDEAL can effectively cope with the highly dynamic nature of traffic demands.
2. We present useful extensions to iDEAL: (i) support femtocell offloading and dynamic roaming, and (ii) incorporate quality of service consideration (in addition to cost).
3. We extensively evaluate iDEAL using simulation based on real traces from one of the largest US cellular service providers. Our results clearly demonstrate the effectiveness of our approach. We further demonstrate the feasibility of our approach using a simple prototype implementation.

II. Problem Formulation

In this section, we formulate the problem of offloading cellular traffic as a *reverse auction*. The offloading is transparent to clients and does not affect cellular pricing (*i.e.*, users pay for the data usage regardless of whether it is carried by the cellular provider or third party resource owners).

Basic auction settings: Consider a cellular network A which is interested in purchasing and leveraging spare resources from third-party Wi-Fi hotspots to satisfy traffic demands from her customers. The third-party hotspot owners should be rewarded for opening up their services to A's customers. To facilitate such cooperation, provider A can set up an auction to let third-party hotspot owners submit bids to offer their network resources, *e.g.*, dollars per bit-rate for unit time (*e.g.*, 1 hour) that a third-party hotspot owner offers.

This problem is naturally formulated as a *reverse auction*. Since the demand changes over time, *e.g.*, due to diurnal variations [26], the auction takes place periodically or whenever demand changes. The auction frequency is chosen to balance the overhead and the accuracy of traffic demand estimation.

The cellular network is shared across a relatively large area typically called a cell site. A site is further sub-divided into three or more sectors. The sector can be considered to be divided into m small regions based on locations of Wi-Fi hotspots and Wi-Fi range. A Wi-Fi hotspot can satisfy traffic demands only in its region.

Naïve solution: A simple approach is to statically partition the cellular resource into different regions and determine the amount of Wi-Fi resource needed in each region (*e.g.*, based on the amount of user demand in the region). Then we conduct a local auction within a region to utilize the cellular resource and Wi-Fi resources dedicated to the region. We call it *static local auction*. While simple, this approach has several important limitations: 1) Due to limited Wi-Fi coverage, the number of hotspots in a region is limited, *i.e.*, the competition is limited. However, adequate competition is essential for an auction based approach to be effective. 2) This formulation treats different regions equally, however the service provider may view different regions differently because different regions may have different spectrum efficiencies due to different signal-to-interference-noise-ratio (SINR) from the base station. 3) The static allocation cannot effectively take into account the available Wi-Fi resources and their bids across different regions. For example, even when a region has higher traffic demand, we may or may not need to allocate more cellular resources to the region depending on (i) how many Wi-Fi hotspots are in the region, (ii) what are their prices, and (iii) how the Wi-Fi hotspots and their prices compare with those in other regions. If there are more Wi-Fi hotspots in a region offering cheaper bids than in the other regions, we can allocate less cellular resources.

Design goals: We seek an auction scheme to (i) *account for different spatial coverage of resources*, which has not been considered in existing work, (ii) *cope with dynamic traffic demands*, (iii) *achieve high efficiency*, where the winners in the auction are the hotspot owners who really can provide the service at a cheaper price, thereby improving the overall system efficiency and social welfare, (iv) *promote truthful bidding* to prevent bidders from gaming the system, effectively discover price to ensure that the overall system is efficient, and avoid unnecessary system fluctuation due to gaming, as unwanted switching between Wi-Fi and 3G can negatively impact user experience [14], (v) *low cost*, which is natural but is challenging to achieve simultaneously with truthfulness, and (vi) *guard against collusion*.

III. Our Solution: iDEAL

In this section we introduce our solution: iDEAL. We start by designing the auction setting that fosters more competition and captures the service provider's regional preferences. Then we describe the two stages of iDEAL: (i) *allocation*, *i.e.*, determine how to allocate traffic among third-party resource owners and the cellular network itself to minimize cost given the bids, (ii) *pricing*, *i.e.*, decide how much should be paid to individual third-party resource owners in order to provide enough incentives for them to be truthful. Table I summarizes the key notations.

m	number of regions in a cellular sector
n	number of sellers in a cellular sector
d_i	traffic demand in region i
c_i	cellular capacity in region i
e_i	spectrum efficiency of cellular network in region i
z	total cellular spectrum usage: $z = \sum_i^m c_i/e_i$
x_j	total capacity bought from seller j
p_j	the unit price seller j asks for
λ_j	the Wi-Fi capacity offered by seller j
$F(z)$	cellular cost function
$f(j)$	the region that seller j belongs to

TABLE I
NOTATIONS.

A. iDEAL Auction Setting

Third-party Wi-Fi resources and bids: Suppose n third-party hotspot owners offer their resources to the cellular service provider by submitting their bids. Let $A_j = \{\lambda_j, p_j\}$ denote hotspot owner j's bid, which indicates hotspot owner j wants to sell λ_j amount of bandwidth at a price p_j per bits-per-second. The bids are *non-atomic* (*i.e.*, a hotspot owner is willing to sell a part of the capacity she offers). Function $f(j)$ returns the region where hotspot owner j sells her capacity (*e.g.*, $f(j) = i$ means hotspot owner j sells her capacity in region i). For simplicity, we assume that each hotspot owner j sells capacity in a single region (relaxed in Section III-F). As Wi-Fi may not cover the whole sector, areas without Wi-Fi coverage can be treated as special regions with no Wi-Fi bids.

Cellular resources as a Virtual Bid: Let the traffic demand vector be $D = \{d_1, d_2, ..., d_m\}$, where d_i is the demand in region i. In order to effectively leverage both third party and cellular resources, we let the service provider also participate in the auction by submitting a *virtual bid*. The virtual bid is in the form of a cost function $F(z)$, where z is the total amount of spectrum used in the entire cellular sector. Let c_i be the cellular capacity in region i, let x_j be the total capacity bought from hotspot owner j. To satisfy the cellular traffic demand d_i in each region i, we must have: $c_i + \sum_{j:f(j)=i} x_j \geq d_i$. To allow us to capture the different spectrum efficiency, we denote the actual spectrum usage in region i as c_i/e_i, where e_i is the spectrum efficiency in region i. Thus, the total spectrum usage is $z = \sum_{i=1}^m c_i/e_i$.

We consider $F(z)$ to be a piecewise linear convex function, capturing the fact that below a certain value the cost (reflecting sunk cost [28]) is very low because the service provider has already invested in buying the spectrum and needs to keep the system running; as the cellular network becomes more loaded, the cost increases; and once it is overloaded, the cost increases sharply to capture the high cost of congestion. A similar convex cost function has been widely used in modeling congestion cost in the Internet (*e.g.*, [15], [25]).

Because the cellular resource in the virtual bid can be used in any region in the sector, it introduces coupling between the regions. The entire sector can now be viewed as one auction instead of several independent ones as in the naïve solution. Even if the number of hotspots in one region is small, its hotspots are not guaranteed to win since the auction may buy more Wi-Fi from other regions and save the cellular resource for this region, *i.e.*, hotspots compete not only within their regions, but also across regions. We now see a new type of competition, which we call *inter-region competition* in addition to *intra-region competition*.

Auction objective: The goal of the cellular service provider is to minimize the total Wi-Fi and cellular cost, while satisfying the customers' demands (*i.e.*, $c_i + \sum_{j:f(j)=i} x_j \geq d_i$) and offering appropriate incentives to the third-party Wi-Fi hotspot owners to share their resources.

B. Preparation: (Static) Global Allocation

We first ignore traffic variations and develop techniques to effectively utilize both cellular and Wi-Fi resources in serving user traffic demands.

$\triangleright$ $Input:\ d_i,\ e_i,\ \lambda_j,\ p_j,\ F(z)$
$\triangleright$ $Output:\ x_j,\ c_i,\ z$
minimize: $\sum_j p_j * x_j + F(z)$
subject to:
[C1] $\sum_{j:f(j)=i} x_j + c_i = d_i \quad \forall i = 1, 2, ..., m$
[C2] $\sum_{i=1}^m c_i/e_i = z$
[C3] $0 \leq x_j \leq \lambda_j \quad \forall j = 1, 2, ..., n$
[C4] $0 \leq c_i \quad \forall i = 1, 2, ..., m$

Fig. 1. Problem formulation to optimize allocation

We formulate a *global resource allocation problem* as a linear program in Figure 1. The formulation effectively captures global cellular resources and local Wi-Fi resources by treating the cellular resource as a single resource with a single bid. As shown, our goal is to minimize the sum of total Wi-Fi cost (based on their bids) plus cellular cost $F(z)$. The constraint [C1] ensures that we have enough Wi-Fi and cellular resources to satisfy traffic demands in each region i. The constraint [C2] relates the cellular capacity with the cellular spectrum. The constraints [C3] and [C4] put upper and lower bounds on x_j and c_i. Since there is no upper bound on z, there is always a feasible solution. When z increases beyond the available spectrum, $F(z)$ grows rapidly. This problem can be solved efficiently using linear program solvers, (*e.g.*, CPLEX).

C. iDEAL Dynamic Global Allocation

Traffic demand changes over time and is challenging to predict accurately. Based on the history of observed demand vectors, we can optimize for the representative demand vectors that are likely to occur in the next time interval. Our goal is to find the allocation to minimize the worst-case cost for these representative demand vectors.

Algorithm: Formally, suppose there are K historical demand vectors, denoted as $D_k = (d_{k1}, d_{k2}, \cdots, d_{km})$ ($k = 1, \cdots, K$), where d_{ki} denotes the k-th possible demand in region i ($i = 1, \cdots, m$). While it is difficult to predict accurately the demand vector for the next time interval, it is common in robust traffic engineering to assume that the demand vector for the next time interval is covered by the convex hull of all the historical demand vectors D_k [25]. Under this assumption, we can minimize the worst-case cost while satisfying all possible demands that may arise in the next time interval. We formulate this dynamic global allocation problem by modifying the LP formulation in Figure 1. In particular, we change [C1] and [C2] to the following:

[C1-dynamic] $\sum_{f(j)=i} x_j + c_{ki} \geq d_{ki} \quad \forall\ k$ and i

[C2-dynamic] $\sum_i (c_{ki}/e_i) = z \quad \forall\ k$

to ensure that we have enough cellular and Wi-Fi resources to satisfy all possible demand vectors. This is much more efficient than provisioning for the peak demand in each region.

From now on, we will refer to our dynamic global allocation algorithm as *iDEAL*, and the static global allocation algorithm as *iDEAL (static)*.

Property: A nice property of this dynamic global allocation is that it effectively leverages the global cellular resource on demand to satisfy different possible traffic demands. In particular, while the total cellular resource is fixed, the amount of cellular resource used in each region can change according to the real demand. When demand shifts from one region to another over time, the same global cellular resource can be used, rather than provisioning for the peak demand in each region. Therefore, global cellular resource has a distinctive advantage over local Wi-Fi resources in satisfying time-varying demand, which we explicitly leverage in our formulation.

D. iDEAL Pricing Solution

As discussed in section II, we want the pricing scheme to be truthful and efficient. Meanwhile, we want the pricing scheme to fully benefit from the inter-region competition. For example, when hotspots in one region lower their bids and offload more traffic, this would reduce the demand for third party resources in other regions and cause hotspots in other regions to sell less. To capture this unique interaction between intra-region and inter-region competition, we cannot treat auctions in different regions as separate auctions and compute pricing separately; instead we must consider them as a single auction and explicitly incorporate inter-region competition into the payment computation.

The Vickrey-Clarke-Groves [29] auction is well-known. It is both truthful and efficient. It pays a winner the opportunity cost that the presence of the winner introduces on the other players. VCG has a major weakness – its cost is generally high [5]. However, in our setting VCG is able to capture the inter-region competition, which lowers the cost. Thus to preserve the nice properties of VCG (*i.e.*, truthfulness and efficiency) while achieving low cost, we apply the VCG principle globally over the whole cellular sector and compute the *global opportunity cost* to capture both inter-region and intra-region competition.

Algorithm: We follow the general VCG principle and compute the global opportunity cost as follows. Let $V(D, N)$ denote the valuation consumed in the optimal allocation. D is a demand matrix containing K demand vectors $D_k = \{d_{k1}, d_{k2}, \cdots, d_{km}\}$ $(k = 1, \cdots, K)$, which specify the possible demands in each region. N is the set of bidders (including the cellular service provider). Given the result of the allocation scheme, if we buy t capacity from winner b in region r, the amount of money we pay to b will be $V(D, N \setminus \{b\}) - V(D^1, N \setminus \{b\})$ where D^1 is derived from D by setting $d_{kr} = max(0, d_{kr} - t)$ for each k and $N \setminus \{b\}$ is the set of remaining bidders after removing bidder b. Thus, $V(D^1, N \setminus \{b\})$ is the total value sold by other bidders under the current optimal allocation; $V(D, N \setminus \{b\})$ is the total value sold by the remaining bidders with b removed. The difference is the global opportunity cost b imposed on other bidders.

Next, we show that the inter-region competition can help reduce cost with an example. Consider 2 regions, each with 1 unit demand. Region 1 has 2 hotspots with valuations 1 and 3, respectively. Region 2 has 1 hotspot with valuation 2. Each hotspot has 1 unit resource. The cellular resource is 1 unit and is worth 1.5. The optimal allocation in this case is: 1 unit of Wi-Fi in region 1 with valuation 1 and 1 unit of cellular resource in region 2. To compute the global opportunity cost for the Wi-Fi winner, we remove this Wi-Fi winner and compute the optimal allocation without the winner. The new allocation should use all the cellular resource in region 1 and the Wi-Fi resource with valuation 2 in region 2. The total valuation sold by other bidders is thus $1.5+2 = 3.5$, while in the original allocation the number is 1.5. So the global opportunity cost we pay to the Wi-Fi winner is $3.5 - 1.5 = 2$. In comparison, with the same allocation, if we apply VCG in each region separately, the local opportunity cost is 3 since region 1 has only the Wi-Fi resource with valuation of 3 after we remove the Wi-Fi winner. This shows that global opportunity cost is lower since it effectively takes into account resources across all regions. Note that this notion of global opportunity cost and its computation work for both static and dynamic global allocations. The two versions only differ in the allocation (as described in Section III-B and III-C).

Properties: iDEAL inherits the following three important properties from VCG: (i) bidders have incentives to be *truthful*, (ii) the outcome of the auction is *efficient*, and (iii) the auction is *individually rational* meaning third-party resource owners have incentives to participate in the auction. Formally, we have the following three theorems.

Theorem 1: In iDEAL, truth-telling is an optimal strategy.

Theorem 2: iDEAL is efficient, which means when bidders are rational, the winners are the bidders whose valuation for their resources is the least.

Theorem 3: iDEAL is individually rational, *i.e.*, bidders of the auction will get non-negative utility, assuming a bidder does not bid lower than his valuation.

Theorem 1 indicates that it is beneficial for a bidder to bid truthfully regardless of other bidders' strategies. See [19] for formal proof. Theorem 2 follows from the truthfulness property and our allocation, which minimizes the total valuation assuming everyone bids truthfully. Theorem 3 guarantees that winners will be paid no less than their valuation.

While Theorem 3 is easy to see in normal settings, it is less straightforward with our dynamic allocation because in the dynamic allocation the total amount of resource we buy is not fixed. Specifically, when computing the opportunity cost, we remove a winner and compute a new allocation and use the bid(s) of the newly admitted winner(s) as the payment. While the unit prices of the newly admitted bids are not lower than the winner's, the total amount of capacity we buy in the new allocation might reduce. This is because the new allocation may buy more cellular resource, which can be used everywhere and may reduce the need for Wi-Fi in all regions. That makes it hard to tell if the opportunity cost is higher than the winner's valuation. We prove the theorem using contradiction: if we remove a winner w, and the amount of increased valuation we buy from others (*i.e.*, the opportunity cost) is less than what w sells, then w should not have won.

E. Understand and Guard Against Collusion

In this section, we first identify potential collusion strategies in iDEAL and show how they differ from those in normal VCG settings. We then discuss how to mitigate such strategies. We call a set of hotspots colluding together a *bidding ring*.

1) Collusion Strategies: Due to the distributed nature of hotspot locations, collusion in our context is quite different from collusion in normal settings, where the optimal collusion strategy is to let one proxy bidder buy (or sell in an reverse auction) for the whole bidding ring [6]. However, in our system each hotspot submits a separate bid. This forbids hotspots to collude optimally and thus may resort to other collusion strategies, identified below. In particular, we consider two types of collusion: (i) single seller collusion, whose objective is to maximize the total utility of all hotspots owned by this seller, and (ii) multi-seller collusion, where each seller colludes with other sellers, but tries to solely maximize her own utility.

In both types of collusion, a bidding ring can drive up the price and increase its utility by *Supply Reduction* (*i.e.*, drop losing bids or reduce the capacity offered in winning bids, which is equivalent to bidding an extremely high price for the capacity that is removed from bidding). Supply reduction can drive up price because it increases the opportunity cost, which is determined by the immediate losing bids.

2) Mitigating Collusion: We mitigate collusion as follows:

Dynamic demands: In order to benefit from supply reduction, a bidding ring needs to accurately predict which bids may lose and drop them. Without that, supply reduction can cause harm by letting the bidding ring miss opportunities to win. Making such predictions is challenging due to the dynamic nature of the traffic demand and Wi-Fi availability. Therefore, in practice supply reduction does not necessarily increase the utility of the hotspots, which can discourage them from colluding.

Bidding as a group: A single seller with multiple hotspots has an incentive to reduce supply because her hotspots submit separate bids. The opportunity cost of one hotspot can be affected by the price/availability of her other hotspots. So by strategically dropping some of her hotspots or raising their prices, she can increase her revenue. This strategy is especially harmful as it may also increase the opportunity cost of other sellers' hotspots. Ultimately, it incurs a higher cost to the service provider.

To address the issue, we let the hotspots owned by the same entity bid as a group, *i.e.*, the seller who owns multiple hotspots discloses all her hotspots and we consider them as a single bidder in the auction. The seller has an incentive to choose this option, since bidding truthfully is an optimal strategy (Theorem 1). It is also preferred from the service provider's perspective because it only removes competition within the group. The hotspots in this group still compete with hotspots of other sellers, which helps to bring down the cost.

Stability of multi-seller collusion: When multiple parties are involved in collusion, a natural question is whether the collusion is stable (*i.e.*, all members of the bidding ring have incentives to stay in the ring [6], [9]). [6] shows that in normal settings, collusion in VCG is stable under certain assumptions. However, their conclusion does not apply to our context because of the difference in collusion strategies. Specifically, we make the following two observations.

First, without utility sharing, members of a bidding ring have an incentive to leave the ring (*i.e.*, do not conduct supply reduction). Formally, we have the following lemma:

Lemma 1: Without utility sharing, for bidding ring members no supply reduction is a (weakly) dominant strategy (*i.e.*, no worse than supply reduction).

This follows from the truthfulness of VCG and the fact that different sellers submit separate, sealed bids and cannot pose as one entity in our system.

Second, the condition of "no utility sharing" is likely to hold in practice due to difficulties of estimating utility obtained from collusion in our system. One reason is that traffic demands and Wi-Fi availabilities are highly dynamic, which makes it hard to attribute utility changes to collusion. Moreover, using sealed bids makes it hard to validate the behavior of other members in the bidding ring. We can make it even harder through system design such as delayed payment (*e.g.*, paying the hotspots every week even though the auction is conducted hourly), which further obfuscates the utility.

F. Practical Considerations

Supporting offloading to femtocells and dynamic roaming: In addition to third-party Wi-Fi hotspots, femtocells and other cellular networks can also be used for offloading. Roaming to other cellular networks considered here is different from traditional roaming. Traditional roaming is enabled only outside the current cellular provider's coverage area whereas dynamic roaming in our context can take place within the coverage area to reduce congestion. In order to support offloading to different types of technologies, we need to effectively handle partially overlapping spatial coverage, as different resources have different coverage ranges. We extend our approach to support these scenarios by dividing overlapping regions into multiple non-overlapping regions and allowing one provider to belong to multiple regions. The constraint [C1] in Figure 1 is then replaced by the following two new constraints:

$$\sum\nolimits_{j:i\in f(j)} x_{ji} + c_i = d_i, \qquad \forall i = 1, 2, ..., m,$$
$$\sum\nolimits_{i} x_{ji} = x_j, \qquad \forall j = 1, 2, ..., n,$$

where x_{ji} is the amount of capacity bought from seller j and used in region i. This extension can not only support offloading to different types of networks, but also allow a hotspot provider to use her resources across different regions (*e.g.*, hotspots belonging to a single restaurant chain spread across different regions but sharing the same bottleneck capacity).

Incorporating quality score: The cellular service provider may prefer some hotspots over others due to different quality (*e.g.*, to avoid hotspots that do not guarantee the amount of capacity they offer). In this case, we can differentiate which hotspots to use based on the quality score q_i ($0 < q_i \le 1$) of hotspot i. The higher the score, the better the quality. To ensure the auction is still truthful and individually rational, we change the objective function in the allocation phase to $\sum_j (x_j \cdot p_j / q_j) + F(z)$ and change the payment for winner j to q_j times the opportunity cost, which is the quality weighted opportunity cost. It is not difficult to see that the auction is still truthful and individually rational.

IV. Evaluation Methodology

We evaluate our approach using trace-driven simulations. We first describe the traces and how they are used.

Traces: We use the following traces: (i) Locations of cell towers and femtocells from a large cellular provider in the US, and hotspot locations from [30]. (ii) Detailed network data with periodic (every 2 seconds) reports of which sectors mobile devices are using for their data communication. We use one-week data from 2011, and pick the busiest sectors out of thousands of sectors. We then use this data to estimate the number of users in a sector during one hour, and the amount of time they stay in that sector. (iii) 3G HTTP traces report detailed HTTP session information, such as HTTP duration, downloaded bytes and type of the download during all 24 hours on a single day in 2011. This is aggregated over several sectors and does not have information about which sector the user is currently in. (iv) The backhaul capacity of about 150 hotspots from a large service provider in the US. All of the network and trace data was anonymized, and no individual user information or identity was available or used. We only used the aggregate information for our evaluation.

Generating regions: We generate regions by clustering the Wi-Fi hotspots using k-means [22]. We use 6 regions as it minimizes partition index [8], which is a commonly used clustering metric. We run the clustering algorithm 100 times and pick the clustering that minimizes the average distance of Wi-Fi hotspots to the centers of their assigned regions.

Network configuration: Based on the typical cell tower spacing of 400-500m in busy urban areas, we use 250m as the communication range for a cell sector. The communication ranges for Wi-Fi and femtocell are set to 100 m and 40 m [4], respectively. To calculate spectrum efficiency, we use the distance between the centroid of the region and the cell tower, and the distance between the centroid of the region and the interfering cell towers, and compute path loss using Hata model [16]. We consider 6 nearest base-stations as interfering base-stations to calculate the $SINR$. We account for self-interference and compute the resulting $SINR'$ as: $SINR' = \frac{SINR}{1+\alpha * SINR}$ where $SINR'$ and $SINR$ denote the signal to interference and noise ratios with and without self interference, respectively, and $\alpha = 0.005$ [2]. We get the spectrum efficiency by applying the Shannon's Law. Since the Shannon capacity is an over-estimate of the real capacity, we scale down the result to match the maximum efficiency that is generally observed in a cellular network (2 bps/Hz).

Generating traffic demands: To generate the demand for an hour, we determine the number of users from the detailed network data during that hour, and pick all the HTTP requests of the corresponding number of users from the 3G HTTP trace. The detailed network data and the HTTP trace are both anonymized and we only use the aggregated demand information in our evaluation. We replace the data rate in the trace with the desired demand rate according to the application types: video 350 kbps, audio 128 kbps, application (*e.g.*, download binary files) 350 kbps, text 150 kbps, and image 165 kbps. We determine the rates of applications, text, and images according to the 90-th percentile rate that users receive from the 3G HTTP trace, and determine the video and audio rates using the data from a large service provider. The data rates in the traces are not used since they are limited by the current cellular capacity and may not indicate the real demand.

We place users randomly in the sector and assign them to regions according to their locations. When a single demand vector is used, we use the peak demand from each region as the final traffic demand. When dynamic allocation is used, we use all the demand vectors corresponding to the time when any region has peak demand. This way, both static and dynamic allocation schemes can sustain the peak loads in all regions.

Generating bids: We use the distribution of backhaul data-rates and pick the available data rate uniformly as being 25%-75% of the backhaul data-rate. The Wi-Fi bids are then generated based on the pricing plan of a major service provider. We uniformly choose 50%-150% of the price as a hotspot's valuation for a given backhaul capacity to capture varying costs from different service providers. We then determine the hourly Wi-Fi valuations according to its capacity and monthly bills assuming 30 days/month and 8 hours/day. The real bids depend on their bidding strategies and may differ from their valuations. The cellular bid $F(z)$ is set to 0 (reflecting the sunk costs) when z is below 80% of cellular capacity (which is set to 3 carriers *i.e.*, 3 times 3.84 MHz), and set to c times estimated maximum Wi-Fi valuation when z exceeds 80%. Note the Wi-Fi valuation is per bps whereas the spectrum cost is per Hz. Thus we translate the maximum Wi-Fi valuation to price per Hz using the lowest spectrum efficiency, such that Wi-Fi is always preferred when the cellular network is overloaded. We set c to 1.25 by default and vary it to evaluate its impact.

Performance metrics: We compare different schemes using efficiency and cost. Efficiency is measured as the total valuation of all resources consumed, whereas cost is the total price of cellular and Wi-Fi resources the service provider pays.

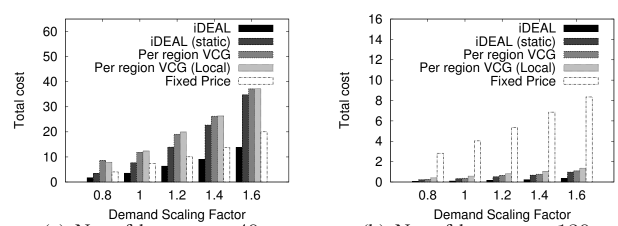

(a) No. of hotspots = 40 (b) No. of hotspots = 130

Fig. 2. Total cost comparison with truthful bids

V. Evaluation Results

A. Comparison of Truthful Auctions

We first compare the cost incurred under different auction schemes, including iDEAL, iDEAL (static), per region VCG with global allocation, and per region VCG with local allocation. All the auctions are truthful except per region VCG with global allocation, which is included to show how VCG will perform without inter-region competition. In addition, we also compare with fixed pricing, where the service provider pays the hotspots a fixed price and uses the global allocation to determine which hotspots to buy. A hotspot with higher valuation than the fixed price would not sell in this case so we use the maximum Wi-Fi price we may generate as the fixed price. The result of using average Wi-Fi price as the fixed price is similar and omitted for brevity.

Figure 2 shows the cost incurred under different schemes. We first observe that auction based approaches work much better than the fixed pricing when there is enough competition. In particular, iDEAL achieves lower cost than the fixed pricing even when the number of hotspots is 40. With 130 hotspots, iDEAL is almost an order of magnitude better than the fixed pricing. Second, iDEAL out-performs iDEAL (static), which out-performs both versions of per region VCG. Per region VCG fails to capture the inter-region competition and thus may suffer from limited competition and lead to high cost. In comparison, both versions of iDEAL fully benefit from inter-region competition. iDEAL further reduces cost by leveraging the flexibility of using cellular resource in different regions on demand, thus reducing the demands for third party resources. Therefore, iDEAL and iDEAL (static) out-perform per region VCG by 63%-80% and 10%-61%, respectively.

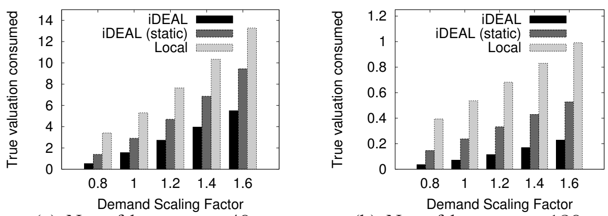

(a) No. of hotspots = 40 (b) No. of hotspots = 130

Fig. 3. Comparison of total true valuation consumed.

We further compare the efficiency of following allocations, all with truthful bids, namely (i) iDEAL, which can optimize allocation according to multiple possible demands, (ii) iDEAL (static), which optimizes allocation according to a single traffic demand, (iii) local allocation, which statically allocates cellular resources to different regions based on the traffic demands in these regions. Note here we omit the fixed pricing because it is not an auction and it makes allocation decisions solely based on the fixed price instead of the valuation. Figure 3 shows the total true valuation of different allocation schemes as we scale the traffic demands by a constant factor from 0.8 to 1.6 and vary the total number of hotspots participating in the auction. As before, iDEAL out-performs its static counterpart, iDEAL (static), which further out-performs the local allocation. iDEAL reduces the total valuation to only 8%-42% of local allocation since it can effectively adapt the cellular allocation to different regions based on real demand. Even iDEAL (static) performs very well: its total valuation consumed is only 34%-72% of local allocation.

Figure 4 further compares the cost of different auction schemes as we vary the cellular cost $F(z)$ by changing its parameter c from 1 to 2. The absolute cost increases with c, as we would expect. The relative performance across different schemes is similar for all values of c we use. The total cost reduces as competition increases (*i.e.*, when the number of hotspots goes up from 40 to 130).

B. Comparison with Non-truthful Auctions

In this section, we study the impact of individual hotspot gaming in non-truthful auctions. We compare iDEAL with the first price and regional uniform price, both of which are widely used [12], [18]. The first price pays winners the amount of their bids, and the regional uniform price pays all the winners in a region at the first losing bid in the region. We

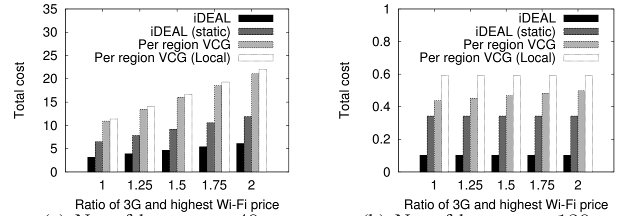

(a) No. of hotspots = 40 (b) No. of hotspots = 130

Fig. 4. Total cost comparison with varying cellular cost function

do not compare with GSP (Generalized second-price auction) because iDEAL does not differentiate between winning slots. If GSP were used, everyone would game to be the highest paid winner as in the first price. We use the static global allocation for all schemes, except that iDEAL uses dynamic global allocation. There are many possible gaming strategies. In our evaluation, we consider simple gaming strategies as examples and show that even these simple strategies can significantly degrade performance. In the first price, we assume a bidder can observe some fraction of bids from other bidders in his region. We call this fraction as Knowledge Factor (KF). He then uses that information to guide his bid in the next round by bidding the maximum between his valuation and the average of the lowest losing price he sees and the highest winning price he sees (including his own bid in the last round). In the first round, bidders start by bidding uniformly randomly between one time and two times their valuation. In the uniform price auction, bidders can game by supply reduction. So we let the winners who do not sell all their capacity reduce their capacity to slightly below the amount they sell in the hope of admitting new winners and potentially increasing the price. When they do sell all their capacity, they will try to increase their offered capacity. In reality, bidders can be more aggressive. For example, all bidders may attempt to reduce supply (*e.g.*, even when they sell all they offer, they can potentially still gain by supply reduction), which may harm the system even further. We conduct multiple runs, and show the results from one run since they are all similar.

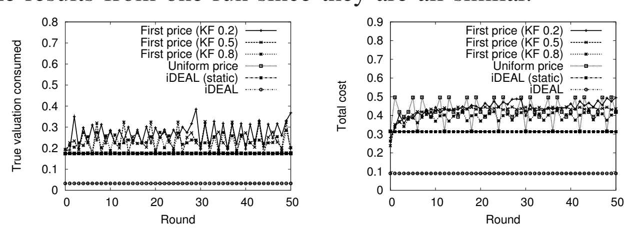

(a) True valuation consumed (b) Cost

Fig. 5. Cost of gaming

Figure 5 (a) shows how gaming affects efficiency. We make a few observations. First, both versions of iDEAL consume less total valuation. The total valuation of iDEAL is as low as 8% of the first price due to more effective use of cellular resources in presence of multiple demands. The total valuation of iDEAL (static) is only 45% of the first price. Second, both versions of iDEAL are stable as bidders are truthful. In comparison, the total value consumption fluctuates considerably in the first price auction because the bidders adapt their bids according to the others' bids. The uniform price performs close to iDEAL (static), because the bidders in our simulation only reduce supply slightly and they do not game by asking higher. In reality, the damage can only be worse.

Figure 5 (b) further compares the total cost to the provider.

Similar to the case of total valuation, both iDEAL versions yield significantly lower cost. Specifically, iDEAL reduces the cost to 18% of the first price and regional uniform price. Moreover, even iDEAL (static) reduces the total cost to 63% of first price and regional uniform price. This result shows that with the help of inter-region competition, using VCG does not incur higher cost than first price or regional uniform price.

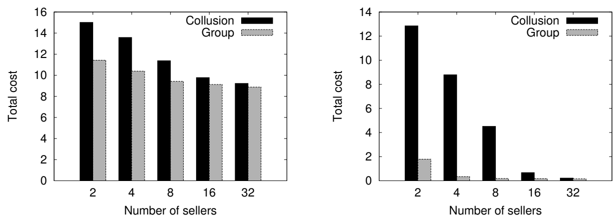

(a) No. of hotspots = 40 (b) No. of hotspots = 130

Fig. 6. Auction cost under collusion and with group option

C. Collusion

Collusion under dynamic demands: We first study how often a bidding ring can improve its utility by supply reduction. We use two different sizes of the bidding rings: 20 and 50 out of 144 hotspots. For each size, we run the experiment 10 times with different random sets of hotspots. Each run consists of 50 rounds. In each round, the bidding ring drops all losing hotspots from the previous round. If there is no losing hotspot, it brings back the cheapest previous dropped hotspot. We vary the demand during each round, but we keep Wi-Fi bids constant. We confirm the degree of traffic variation in the hourly traffic traces in multiple cellular sectors from a major cellular provider is comparable to the traces used for our evaluation. We find that for the bidding ring of size 50, collusion reduces the hotspots' utility for 13% of time and improves the utility for 28% of the time. For the bidding ring of size 20, the numbers are 20% and only 5%. When collusion reduces utility, it reduces by 79% on average, while the number for improvement is only 30%. These results suggest dynamic demand significantly reduces the incentive to collude. In reality, when Wi-Fi bids are also dynamic, it is even harder to predict which set of hotspots will lose.

Bidding as a group: Next we compare bidding as a group with collusion using the same strategy mentioned above. Figure 6 plots the average cost as we vary the total number of sellers and the total number of hotspots they own and perform 100 random runs for each configuration, where each configuration generates 10 sets of sellers and 10 sets of hotspots. The results are consistent with our expectation: a single seller collusion does not always improve utility, but it always incurs a higher cost to the service provider, especially when each seller has a large number of hotspots. In comparison, the group option, which is preferred by sellers, reduces the total cost by as much as 36% and 96% when the number of hotspots is 40 and 130, respectively. The damage of collusion reduces as the number of sellers increases since there are more sellers in competition and each seller controls fewer hotspots.

D. Extensions

Supporting femtocell offload: In Figure 7 (a), we let both Wi-Fi hotspots and femtocells participate in the auction. We vary the number of Wi-Fi bidders while keeping 16 femtocells. As expected, the benefit of femtocells is larger when we have fewer Wi-Fi hotspots. For example, the femtocells reduce the cost by 32% when there are only 40 hotspots. As the number of hotspots increases, the additional benefit from femtocells becomes marginal since Wi-Fi has a higher communication range and is more effective in offloading.

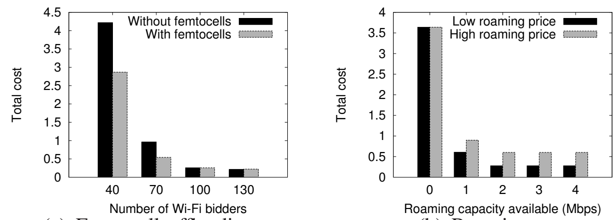

(a) Femtocell offloading (b) Roaming

Fig. 7. Benefit of femtocell offloading and roaming

Supporting Dynamic Roaming: Figure 7 (b) shows the total cost as the roaming capacity available varies from 0 to 4 Mbps, where 0 corresponds to no roaming. The evaluation has 40 hotspots. In this case, since the Wi-Fi resource is insufficient, even having 1 Mbps of available roaming capacity (around 10% total cellular traffic in the sector) can significantly cut down cost. Dynamic roaming reduces the cost to 17% of that when only Wi-Fi is used with the low roaming price (which is set to the maximum winning Wi-Fi bid we observe in the default settings), and 25% under the high roaming price (which is the maximum Wi-Fi bid we may generate based on the distribution we use). Further increasing roaming capacity leads to an even lower cost but the improvement tapers off as the capacity increases beyond 2 Mbps.

VI. Implementation

Offloading involves the following three challenges: (i) identifying a network to offload, (ii) automatic authentication, and (iii) seamless offload so that existing sessions are maintained during the offload. iDEAL already solves the first issue. To address (ii), Hotspot 2.0 can be used to discover hotspot information and support authentication with externally owned hotspots. To support dynamic offloading in this paper, the 'roaming' partners are updated dynamically according to the offloading decisions of iDEAL. To address (iii), Dual Stack Mobile IP (DSMIP), DSMIPv6 [3], [1] have been proposed and various implementations (*e.g.*,[20], [27]) show that there is a low switching overhead.

We develop a prototype implementation on Linux machines using a NetGear WAG511 NIC to demonstrate the feasibility of our solution. Figure 8 shows our system architecture. Through a simple web interface, bidders *i.e.*, hotspot owners, can submit their bids to the service provider machine, who runs the auction. Hotspots are configured using hostap [17]. The service provider sends a message to the winning hotspot with the ssid and password it should use in the current round and also sends this information to the mobile client so that it can connect to the winning hotspot. This is achieved using TCP sockets. Authentication between mobile client and hotspot is done using WPA PSK through WPA Supplicant [31]. We collect performance statistics from the mobile client for billing and keeping track of hotspot quality score. We measure the upload and download statistics on the wireless interface using the Collectl tool [11] periodically (*e.g.*, every 10 secs) and send

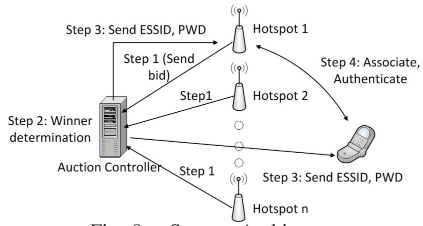

Fig. 8. System Architecture.

back the data to the service provider PC for bookkeeping. The running time of allocation and pricing is only around 100ms, which is small enough for practical use.

We measured the association time in our implementation. After getting the scan results, it takes 18 ms to associate, 103 ms to perform a 4-way handshake (*i.e.*, defining individual keys for unicast transmission), and 3 ms to perform the group handshake (*i.e.*, defining keys for broadcasts). The authentication and scan times can be further reduced (*e.g.*, using techniques in [23], [24]). To further enhance efficiency, after deciding how much traffic to offload to each third party resource owner (which is the focus of this paper), we can strategically select users to offload to minimize the switching time and avoid offload users who will soon leave the hotspots.

VII. Related Work

The need to complement cellular networks with other forms of connectivity has been considered in the past. The authors in [7] conduct measurements in a vehicular testbed, and report that Wi-Fi is available only 11% of the time and 3G is available 87% of the time. Moreover, they find that 3G and Wi-Fi availability are negatively correlated, *e.g.*, Wi-Fi is available 50% of the times that 3G is not available. Lee *et al.*, in [21] use daily mobility patterns of 100 iPhone users to measure the amount of data Wi-Fi can offload. They find that Wi-Fi can offload 65% of data traffic without any delay; if 1 hour or longer delay can be tolerated, the offload traffic increases further by 29%. Zhuo *et al.*, in [32] leverage VCG based auction mechanism to incentivize mobile users to wait until they come in contact with a Wi-Fi AP. Authors in [13] quantify city-wide Wi-Fi offloading gain. They show that even a sparse Wi-Fi network improves performance. Different from the above existing works, our paper focuses on how to incentivize third party resource owners to offload cellular traffic. The work in [10] is closest to ours. It proposes a VCG reverse auction framework to buy femtocell resources. As mentioned in Section I, it does not address the three unique challenges we focus on, namely, diverse spatial coverage, traffic uncertainty, and collusion. The scheme is similar to the local allocation in spirit in that it statically determines the amount of third-party resource to buy in each region.

VIII. Conclusion

We propose iDEAL to enable a cellular service provider to purchase and leverage third-party resources on demand through reverse auctions. iDEAL has several important features: (i) explicitly accounts for diverse spatial coverage of different resources, (ii) copes with the dynamic nature of traffic demands, (iii) effectively incentivizes third-party resource owners to be truthful to their true valuation in the bidding process, (iv) is provably efficient by choosing the bidders with the lowest valuations as the winners, and (v) mitigates potential collusion. Our trace-driven simulation shows that iDEAL effectively reduces cost and is robust against collusion. Our prototype implementation demonstrates its feasibility.

Acknowledgements: This work is supported in part by NSF Grants CNS-0916106 and CNS-1017549.

References

[1] 3G Wi-Fi Seamless Offload. http://www.qualcomm.com/documents/files/3g-wifi-seamless-offload.pdf.
[2] 3GPP2 C.R1002-A CDMA2000 evaluation methodology revision A, version 1.0. May 2009.
[3] *Access Network Discovery and Selection Function (ANDSF) Management Object (MO)*. 3GPP TS 24.312, 2011.
[4] Femto-cell range. http://www.att.com/shop/wireless/devices/3gmicrocell.jsp?fbid=jZ_3W-pHS%6d.
[5] L. M. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. *Combinatorial Auctions*, pages 1–37, 1961.
[6] Y. Bachrach. Honor among thieves: collusion in multi-unit auctions. In *Proc. of AAMAS*, Richland, SC, 2010.
[7] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3G using WiFi. In *Proc. of MobiSys*, 2010.
[8] J. C. Bezdek. Numerical taxonomy with fuzzy sets. *Journal of Mathematical Biology*, 1974.
[9] O. Biran and F. Forges. Core-stable rings in auctions with independent private values. CESifo Working Paper Series 3067, CESifo Group Munich, 2010.
[10] Y. Chen, J. Zhang, Q. Zhang, and J. Jia. A reverse auction framework for access permission transaction to promote hybrid access in femtocell network. In *Proc. of IEEE INFOCOM*, 2012.
[11] Collectl tool. http://collectl.sourceforge.net/.
[12] P. Cramton. Competitive bidding behavior in uniform-price auction markets. *Hawaii International Conference on System Sciences*, 2, 2004.
[13] S. Dimatteo, P. Hui, B. Han, and V. O. Li. Cellular traffic offloading through WiFi networks. In *Proc. of IEEE MASS*, 2011.
[14] W. Dong, Z. Ge, and S. Lee. 3G meets the Internet: understanding the performance of hierarchical routing in 3G networks. In *Proc. of ITC*, 2011.
[15] B. Fortz, J. Rexford, and M. Thorup. Traffic engineering with traditional IP routing protocols. *IEEE Comm. Magazine*, 40(10), Oct. 2002.
[16] M. Hata. Empirical formula for propagation loss in land mobile radio services. *IEEE Trans. Vehicular Tech.*, VT-29(3):317–325, Aug. 1980.
[17] Hostap. http://hostap.epitest.fi/hostapd/.
[18] V. Krishna. *Auction theory*. Academic Press, 2002.
[19] V. Krishna and M. Perry. Efficient mechanism design. Technical report, 1998.
[20] R. Kuntz and J. Lorchat. Versatile IPv6 mobility deployment with dual stack mobile IPv6. In *Proc. of MobiArch*, 2008.
[21] K. Lee, I. Rhee, J. Lee, S. Chong, and Y. Yi. Mobile data offloading: how much can WiFi deliver? In *Proc. of CoNEXT*, 2010.
[22] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129 – 137, 1982.
[23] A. M. Min-Ho, M. ho Shin, and W. A. Arbaugh. Proactive key distribution using neighbor graphs. *IEEE Wireless Comm.*, 2004.
[24] S.-H. Park, H.-S. Kim, C.-S. Park, J.-W. Kim, and S.-J. Ko. Selective channel scanning for fast handoff in wireless LAN using neighbor graph. *Proc. of PWC*, 2004.
[25] M. Roughan, M. Thorup, and Y. Zhang. Traffic engineering with estimated traffic matrices. In *Proc. of IMC*, Oct. 2003.
[26] M. Shafiq, L. Ji, and A. Liu. Characterizing and modeling Internet traffic dynamics of cellular devices. In *Proc. of ACM SIGMETRICS*, 2011.
[27] S. Sharma, I. Baek, and T.-C. Chiueh. Omnicon: a mobile IP-based vertical handoff system for wireless LAN and GPRS links. *Softw. Pract. Exper.*, 37, 2007.
[28] Sunk costs. http://en.wikipedia.org/wiki/Sunk_costs.
[29] W. Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *Journal of Finance*, pages 8–37, 1961.
[30] Wireless Geographic Logging Engine. http://wigle.net.
[31] WPA Supplicant. http://hostap.epitest.fi/wpa_supplicant/.
[32] X. Zhuo, W. Gao, G. Cao, and Y. Dai. Win-coupon: An incentive framework for 3G traffic offloading. In *Proc. of ICNP*, 2011.