

Traffic-Driven Dynamic Spectrum Auctions

Xia Zhou, Shravan Mettu[†], Heather Zheng, Elizabeth M. Belding
Department of Computer Science

[†]Department of Electrical and Computer Engineering
University of California, Santa Barbara

xiazhou, htzheng, ebelding@cs.ucsb.edu [†]smetu@umail.ucsb.edu

Abstract—Wireless growth has been limited by the shortage of radio spectrum. While the spectrum assigned to legacy technologies remain unused, new prominent technologies such as Mesh/WiFi networks are forced to crowd into a small unlicensed band, suffering from significant interference and degraded performance. Using economic incentives, dynamic spectrum auctions redistribute spectrum to make it available to new technologies while providing financial benefits to legacy owners. In this paper, we investigate the performance of dynamic spectrum auctions under traffic dynamics. Using measured traffic traces from deployed WiFi access points, we evaluate the advantages and artifacts of dynamic auctions over plain channel sharing, and investigate the impact of bidding formats and auction intervals. Our results show that short-term dynamic auctions with traffic-aware bidding can significantly improve system throughput and provide bidders with cost-effective spectrum usage.

I. INTRODUCTION

As WiFi networks become a pervasive last mile connectivity tool, wireless users are suffering from poor performance at crowded hotspots. The fundamental observation is that increasing user participation leads to greater variability in traffic density and demands, and consequently more unpredictable user experiences. As bandwidth-hungry devices such as AppleTV and iPhone join the fray, the level of variability further increases.

To improve user experience, WiFi access points (AP) must adjust their allocated bandwidth based on varying traffic demands. Varying APs' spectrum allocation would be a natural and highly effective approach – APs can adjust the frequency location and the amount of spectrum to accommodate traffic dynamics. However, even with this flexibility, APs' spectrum usages are critically limited by the amount of frequency available in the unlicensed bands as well as the elevated competition from other consumer devices.

To obtain additional spectrum in a cost-effective manner, we propose a dynamic spectrum auction framework [6] where APs use the unlicensed bands to main-

tain basic connectivity but bid for additional external spectrum to improve user experience. For simplicity, we assume a single auctioneer (or spectrum owner) who auctions a single channel among APs. The auction is dynamic and repeated in short cycles, adapting to AP traffic dynamics. In each cycle, the auctioneer determines the set of winning APs based on their bids submitted for this cycle, exploiting time and spatial multiplexing to improve spectrum utilization.

In this paper, we focus on investigating the impact of dynamic spectrum auctions on network performance. While previous work in spectrum auctions [6] relies on static and synthetic traffics, we use the publicly available measurement traces from Dartmouth campus WiFi networks [7] on AP traffic demands. Using these traffic traces, we investigate the advantages and artifacts of obtaining additional channels through auctions, as well as the impact of bidding behaviors and the size of auction cycles.

Our results reveal the following findings:

- The measurement traces indicate that AP traffic demands vary significantly over time and differ significantly across APs. Such independent dynamics offer the potential for time and spatial multiplexing.
- As spectrum becomes scarce, dynamic auctions favor bidders (APs) with better financial return from their desired spectrum usage, granting spectrum to those who value them the most.
- The size of auction cycles can significantly impact the system performance. In the measurement traces, APs experience fast traffic dynamics in the rate of 50 minutes. Performing auctions every 50 minutes can double the system throughput compared to those using 300 minutes.
- To stay competitive, bidders with heavy traffic dynamics should plan their bids intelligently, adapting to traffic dynamics. By spreading bids over

time proportional to traffic volume, a traffic-aware bidding format enables bidders to obtain spectrum when it is needed the most.

II. THE NEED FOR DYNAMIC SPECTRUM AUCTIONS

In this section, we present the characteristics of the measured AP traffic traces from Dartmouth [7], which motivate the use of dynamic spectrum auctions. Using the Dartmouth campus WiFi AP traces, we extract the volume of per-AP traffic demand by 5-minute intervals. Figure 1 shows the traffic demands for two randomly selected APs in one day. The volume of the traffic demand fluctuates significantly over time and differs significantly across APs. Most APs carry heavy and bursty traffic during the day, and become idle at night. A small set of APs have very bursty and spiky traffic demands that are 100 times higher than the average of all the other APs.




Fig. 1. Traffic demands of two randomly selected APs from Dartmouth Campus WiFi traces.

The above results clearly indicate a trend of bursty and unpredictable traffic dynamics. As the network population increases, the original unlicensed spectrum is insufficient to accommodate the requests. Driven by economic incentives, spectrum auctions serves as an efficient and fair approach for bidders to obtain additional spectrum based on their needs.

Intuitively, the design of auctions can significantly impact the usage of spectrum. The two simplest formats would be to auction spectrum in *static long-term leases* driven by each bidder's peak traffic demand or by mean traffic demand. However, because of the inherent traffic dynamics, the peak-demand based auctions will lead to severe under-utilization of spectrum while the mean-

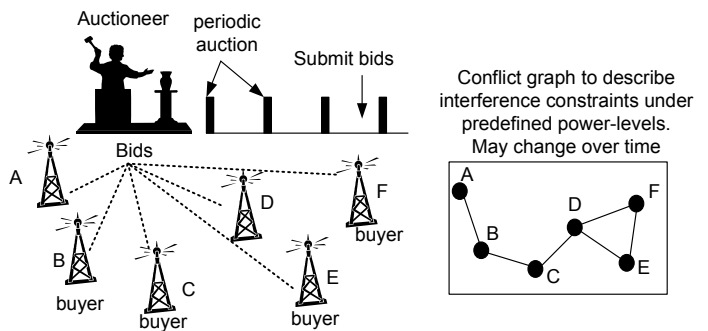


Fig. 2. An illustrative example of dynamic spectrum auctions. The interference constraints among bidders are represented by a conflict graph.

demand based auctions will lead to severe contentions among users and lack of allocation guarantee.

On the other hand, *dynamic* spectrum auctions, if designed properly, offer the potential to adapt to traffic dynamics and utilize spectrum efficiently. By auctioning spectrum in short intervals matching traffic dynamics, this new approach can exploit both time and spatial multiplexing to make the best use of spectrum. Finally, bidders can adjust their bidding behaviors based on their instantaneous traffic demands, and acquire spectrum in a cost-effective manner.

III. DYNAMIC SPECTRUM AUCTION DESIGN

In this section, we present the problem model and the auction clearing algorithm that determines winners and their spectrum usage. We consider a seal-bid auction system with private bids. In each auction interval, an auctioneer broadcasts the set of channels to be auctioned in this interval to APs. APs submit their bids to the auctioneer privately. The auctioneer determines the set of winning APs and their channel allocations. The winning APs set up their communications using both the auctioned spectrum and the original unlicensed spectrum while the losing APs share the unlicensed spectrum to set up their communications.

The fundamental difference between spectrum auctions and conventional auctions is the spatial reusability of spectrum. While conventional auction items (*e.g.* paintings, stocks) cannot be reused among multiple bidders, the same spectrum band can be reused by multiple bidders if they do not conflict with each other. Therefore, an efficient spectrum auction must consider interference constraints and exploit spatial reuse to maximize its usage. In this paper, we assume that the interference constraints can be represented by a conflict-graph, or binary pair-wise constraints among APs (see Figure 2).

That is, two APs either conflict or not. We also assume that the conflict constraints are known to the auctioneer. Extensions to sophisticated interference models incur much higher complexity and will be addressed in a separate study.

A. Problem Model

Consider N APs indexed from 0 to $N-1$ compete for the M extra channels. Assume F_i and F_j denote the set of channels assigned to i and j , $f_i = |F_i|/M$. Let $\mathbb{C} = \{c_{i,j}\}_{N \times N}$, a N by N array, represent the interference constraints among N APs, where $c_{i,j} = 1$ if i and j conflict. Given \mathbb{C} , a channel assignment $\mathbb{A} = \{F_i\}_{i=0}^{N-1}$ is a *conflict-free channel assignment* if

$$F_i \cap F_j = \emptyset, \text{ if } c_{i,j} = 1, \forall i, j, 0 \leq i, j < N \quad (1)$$

Each participating AP expresses its preference of the channels by a piece-wise linear price-demand bid [6], which defines the per-channel price p_i of bidder i over its demand f_i . In its general form, the demand f_i is divided into m sections with the maximal f_i^m and minimum demand f_i^0 . The price p_i in each section is defined as:

$$p_i(f_i) = \begin{cases} -a_1 f_i + b_1, & f_i \in [f_i^0, f_i^1), a_1 \geq 0, b_1 > 0 \\ \dots \\ -a_m f_i + b_m, & f_i \in [f_i^{m-1}, f_i^m), a_m \geq 0, b_m > 0 \\ 0, & \text{otherwise} \end{cases}$$

A simple example is the linear demand price bid with a single price section: $p_i(f_i) = -a_i f_i + b_i$, $a_i \geq 0, b_i > 0$.

Definition 1: A revenue-maximizing spectrum auction with the outcome \mathbb{A}^* is the one that provides a conflict-free channel assignment for the bidders and maximizes the auction revenue. That is,

$$\begin{aligned} \mathbb{A}^* &= \operatorname{argmax}_{\mathbb{A}} \sum_{i=0}^{N-1} f_i \cdot p_i(f_i) \\ \text{subject to} & \quad f_i \leq 1 \\ & \quad \mathbb{A} \text{ is a conflict-free assignment.} \end{aligned} \quad (2)$$

B. Auction Clearing Algorithms

Given a set of bids and the conflict graph, the auctioneer applies a clearing algorithm to produce a conflict-free channel allocation that maximizes the revenue. When $M > 1$, and $f_i > 1/M$, *i.e.* each bidder demands more than one channels, the problem is NP-hard because of the combinatorial interference constraints [11]. Our previous work has provided efficient heuristics for the auction clearing [6].

When $f_i = 1/M$, *i.e.* each bidder demands one channel, this problem becomes a variant of the weighted independent set problem [17], which is also NP-hard. We utilize the greedy algorithms proposed in [14] to solve the auction clearing. This algorithm allocates spectrum in multiple iterations. In each iteration, the algorithm computes a label for each bidder based on its bid and conflict degree, assigns the lowest indexed channel to the bidder with the highest label, and updates the conflict graph by eliminating the winning bidders. This algorithm runs in polynomial time and provides a provable worst-case bound [14].

IV. EXPERIMENTS

Using Qualnet and Matlab simulations, we evaluate the performance of dynamic spectrum auctions using real traffic traces. For simplicity, we assume $M = 1$, *i.e.* every bidder requests one channel. Under this assumption, we investigate the advantages and artifacts of dynamic auctions and the impact of bidder behaviors and the size of auction intervals.

A. Evaluation Methodology

We use Dartmouth traces collected from Jan. 2004 to Mar. 2004 on AP traffic demands [7]. Using 28 APs' SNMP files on Feb. 12, 2004, we extract the set of packet sequence numbers for each 5-minute interval and calculate the number of packets sent during this interval. Because auctions are performed in intervals ≥ 5 minutes, we aggregate the set of packets in each auction interval.

We build a simulation platform by integrating Matlab and Qualnet. We use Qualnet to deploy APs and their subscribers, and simulate traffic transmissions over the original unlicensed channel and the auctioned channel. Losing bidders will share the original unlicensed channel using the 802.11 MAC protocol, while winning bidders use the auctioned channel to support high-throughput transmissions in the current interval. In each auction interval, we use Matlab to run the auction clearing algorithm and determine the revenue-maximizing channel assignment. The channel assignment is then used as an input for Qualnet simulations.

To investigate the impact of economic incentives, we consider two budget scenarios: (1) uniform budget where each bidder has the same budget, (2) traffic-dependent budget where each bidder's total budget scales linearly with its total traffic. We assume that bidders do not accumulate their unused budgets to increase their future bids.

We consider three bidding strategies:

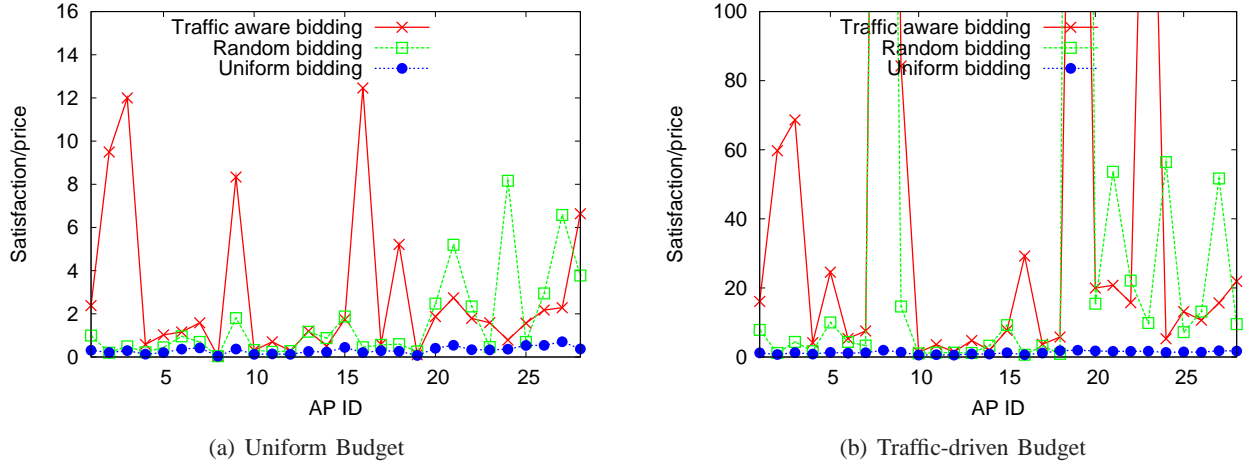


Fig. 3. Distribution of satisfaction/price under different bidding behaviors.

- *uniform bidding* – A bidder i uniformly distributes its budget X_i across auction intervals when it has adequate traffic demands. That is, bidder i 's bid at auction interval j is $p_i(j) = X_i/R_i$ where R_i is the number of auction intervals where i has adequate traffic demands.
- *random bidding* – A bidder randomly distributes its budget X_i among auction intervals when it has adequate traffic demand.
- *traffic-aware bidding* – A bidder distributes its budget X_i among auction intervals proportionally to its traffic demand. At interval j , the bid of AP i is

$$p_i(j) = X_i * \frac{t_{ij}}{\sum_{j=1}^{R_i} t_{ij}} \quad (3)$$

where t_{ij} is the amount of traffic demands at bidder i in auction interval j .

We use the following performance metrics. Let *AP satisfaction* represent the ratio between the amount of traffic served by each AP to the total amount of its demands over time. Let *AP satisfaction/price* represent the cost-effectiveness of spectrum bidding. Let *throughput* or *spectrum utilization* represent the auction efficiency.

B. Impact of Bidding Behaviors


We first study the impact of bidding behaviors by examining the user satisfaction/price, which refers to the time-average satisfaction per unit price. We compare the performance of uniform, random and traffic-aware bidding in Figure 3 under different budget settings. Under the uniform budget, traffic-aware bidding achieves significantly higher satisfaction/price tradeoffs compared to uniform bidding. This is because APs can concentrate

their budgets into intervals with heavy traffic and hence can ensure the success of auctions in these intervals. We also observe that the effectiveness of traffic-aware bidding differs greatly among APs. After examining the traces, we observe that APs with higher degree of traffic burstyness benefit more from traffic-aware bidding. Next, under the traffic-driven budget, the variance of the tradeoff increases. This is because APs with heavy traffic now have higher chances in winning the auction and use the auctioned channel to improve its satisfaction.

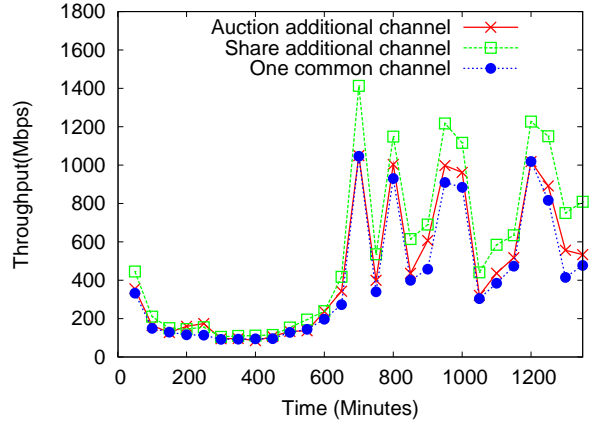
C. Auction Efficiency

We then examine the spectrum efficiency of the auction from the total system throughput in each auction interval. We compare (1) the basic system without auctions where APs share a single (unlicensed) channel, (2) the auction system where APs bid for one extra channel and (3) the basic system without auctions but APs share 2 channels. Results in Figure 4 show that under the traffic-driven budget, the auction system achieves similar performance as the basic sharing system with 2 channels. This result demonstrates the effectiveness of time-multiplexing achieved by auctions.

On the other hand, under the uniform budget, the auction system achieves less throughput compared to the 2-channel sharing system. This is because APs with lighter traffic now have equal opportunity in winning the auction, leading to spectrum under-utilization. One possible solution to eliminate such inefficiency is to partition the channel auctioned into many sub-channels.




(a) Traffic-driven Budget

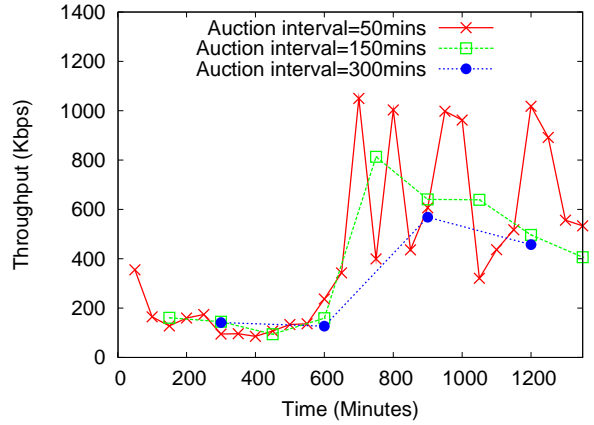


(b) Uniform Budget

Fig. 4. The total AP throughput as a function of time. We compare three systems (1) APs share one channel, (2) APs share one channel and bid for one additional channel and (3) APs share two channels.



(a) Traffic-driven Budget



(b) Uniform Budget

Fig. 5. The impact of auction interval on system throughput, assuming traffic-award bidding, under both uniform and traffic-driven budgets.

D. The Choice of Auction Cycle

Next we investigate the impact of the size of auction interval. Intuitively, large auction intervals produce coarse spectrum allocation and fail to respond to fluctuating traffic demands while small auction intervals lead to significant overhead in computations and communications. After parsing the traffic traces, we show that the rate of traffic dynamics is roughly 50 minutes (Figure 1), which is the baseline value for the experiments above. In the following, by varying the auction interval, we study its impact on the overall system throughput.

Figure 5 compares the impact of auction intervals in terms of the throughput in each auction interval. For a fair comparison, we assume that bidders have *perfect knowledge* of the amount of traffic in each

interval¹. Results show that auctions at a 50-minute interval can quickly respond to traffic dynamics, leading to the best throughput results. Auctions at 150 and 300 minutes intervals cannot fully exploit the benefits of time multiplexing, and hence suffer from large throughput degradations.

V. RELATED WORK

There are multiple complementary ways to design spectrum auctions [12], [13], each applicable to different scenarios. First, the system can auction transmit power among bidders to minimize interference [8], [9]. All winning bidders operate on the same spectrum channel.

¹In practice, the accuracy of the knowledge decreases with the size of auction interval.

By enforcing a price proportional to the amount of power emitted and the relative degree of competition (interference) in a given area, this approach controls power emissions to provide each winner their desired bandwidth. Second, the system can allocate conflicting users with orthogonal channels to avoid interference, and compute appropriate prices and allocations to maximize system utility. Prior work in this category mainly follows a cellular network model. The work in [10] proposed a demand responsive pricing framework, and applied iterative bidding to maximize the social welfare for small scale networks. In [4], the authors proposed a general concept of dynamic auctions among cellular networks and developed centralized heuristics for small scale networks. The work in [16] introduces a hybrid pricing model to reduce the frequency of auctions. The system will use auctions during peak periods and apply an uniform pricing mechanism during off-peak periods. Finally, targeting on fast auction clearing over large-scale networks, the work in [6] proposed light-weight auction clearing algorithms by simplifying the combinatorial interference constraints into linear constraints.

We note that prior work on dynamic spectrum auctions mainly focused on designing allocation algorithm for each auction interval. In this paper, we use real measurement traces to examine the impact of auctions on network throughput and the impact of bidding behaviors and the size of auction interval.

There has also been a parallel stream of work on allocating spectrum to users [15] without economic incentives. Many have proposed centralized and distributed allocation mechanisms to distribute spectrum efficiently [2], [3], [5]. A good survey on existing works in dynamic spectrum access can be found in [1].

VI. CONCLUSION

In this paper, we focus on investigating the impact of dynamic spectrum auctions on WiFi user performance. Using the measured AP traffic traces from Dartmouth campus WiFi networks, we investigate advantages and artifacts of accessing spectrum through dynamic auctions. We show that by choosing appropriate auction intervals, dynamic auctions can provide similar spectrum utilization as the system where nodes share spectrum, and yet apply economic incentives to redistribute spectrum to nodes who value it the most. We also show that traffic-aware channel bidding can significantly improve user satisfaction/price ratio, leading to the best tradeoffs between economic incentives and spectrum efficiency.

REFERENCES

- [1] AKYILDIZ, I. F., LEE, W.-Y., VURAN, M. C., AND MOHANTY, S. Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey. *Comput. Networks* 50, 13 (2006), 2127–2159.
- [2] BRIK, V., ROZNER, E., BANERJEE, S., AND BAHL, P. A protocol for coordinated spectrum access. In *Proc. of IEEE DySPAN* (2005), pp. 611–614.
- [3] BUDDHIKOT, M. M., KOLODY, P., MILLER, S., RYAN, K., AND EVANS, J. DIMUSUMNet: new directions in wireless networking using coordinated dynamic spectrum access. In *Proc. of IEEE WoWMoM* (2005), pp. 78–85.
- [4] BUDDHIKOT, M. M., AND RYAN, K. Spectrum management in coordinated dynamic spectrum access based cellular networks. In *Proc. of IEEE DySPAN* (November 2005).
- [5] ETKIN, R., PAREKH, A., AND TSE, D. Spectrum sharing for unlicensed bands. In *IEEE Journal on Selected Areas in Communications* (2005), pp. 517–528.
- [6] GANDHI, S., BURAGOHAIN, C., CAO, L., ZHENG, H., AND SURI, S. A general framework for wireless spectrum auctions. In *Proc. of IEEE DySPAN* (2007).
- [7] HENDERSON, T., KOTZ, D., AND ABYZOV, I. The changing usage of a mature campus-wide wireless network. In *Proc. of MobiCom* (Sept. 2004).
- [8] HUANG, J., BERRY, R., AND HONIG, M. Auction mechanisms for distributed spectrum sharing. In *Proc. of 42nd Allerton Conference* (September 2004).
- [9] HUANG, J., BERRY, R., AND HONIG, M. Spectrum sharing with distributed interference compensation. In *Proc. of IEEE DySPAN* (November 2005).
- [10] ILERI, O., SAMARDZIJA, D., SIZER, T., AND MANDAYAM, N. B. Demand responsive pricing and competitive spectrum allocation via a spectrum server. In *Proc. of DySpan' 05* (November 2005).
- [11] JAIN, K., PADHYE, J., PADMANABHAN, V. N., AND QIU, L. Impact of interference on multi-hop wireless network performance. In *Proc. of MobiCom* (2003), pp. 66–80.
- [12] NOAM, E. M. Taking the next step beyond spectrum auctions: Open spectrum access. *IEEE Communications* 33, 12 (1995), 66–73.
- [13] PEHA, J. M. Approaches to spectrum sharing. *IEEE Communications Magazine* 43 (February 2005), 10–12.
- [14] PENG, C., ZHENG, H., AND ZHAO, B. Y. Utilization and fairness in spectrum assignment for opportunistic spectrum access. *Mobile Networks and Applications (MONET)* 11 (May 2006), 555–576.
- [15] RODRIGUEZ, V., MOESSNER, K., AND TAFAZOLLI, R. Market driven dynamic spectrum allocation over space and time among radio-access networks: Dvb-t and b3g cdma with heterogeneous terminals. *ACM Mobile Networks and Applications (MONET)* 11, 6 (2006), 847–860.
- [16] RYAN, K., ARAVANTINOS, E., AND BUDDHIKOT, M. M. A new pricing model for next generation spectrum access. In *Proc. of TAPAS* (August 2006).
- [17] SAKAI, S., TOGASAKI, M., AND YAMAZAKI, K. A note on greedy algorithms for the maximum weighted independent set problem. *Discrete Appl. Math.* 126, 2-3 (2003), 313–322.