

Double Auctions for Dynamic Spectrum Allocation

Wei Dong, Swati Rallapalli, Lili Qiu, *Senior Member, IEEE*, K. K. Ramakrishnan, *Fellow, IEEE*, and Yin Zhang, *Fellow, IEEE*

Abstract—Wireless spectrum is a precious resource and must be allocated and used efficiently. Conventional spectrum allocations let a government agency (e.g., FCC) sell a portion of spectrum to one provider. This is not only restrictive, but also limits spectrum reuse and may lead to significant under-utilization of spectrum. In this paper, we develop a novel truthful double-auction scheme to let any resource owner (e.g., a cellular provider), who has spare spectrum at a given time period, sell to one or more providers that need additional spectrum at that time. Spectrum auctions are fundamentally different from conventional auction problems since spectrum can be reused and competition among buyers is complex due to wireless interference. Our proposal is the first double-auction design for spectrum allocation that explicitly decouples the buyer-side and seller-side auction design while achieving: 1) truthfulness; 2) individual rationality; and 3) budget-balance. To accurately capture wireless interference and support spectrum reuse, we partition the conflict graph so that buyers with strong direct and indirect interference are put into the same subgraph, and buyers with no interference or weak interference are put into separate subgraphs. Then, we compute pricing independently within each subgraph. We then develop a scheme to combine spectrum allocation results from different subgraphs and resolve potential conflicts. We further extend our approach to support local sellers whose spectrum can only be sold to buyers within certain regions, instead of all buyers. Using conflict graphs generated from real cell tower locations, we extensively evaluate our approach and demonstrate that it achieves high efficiency, revenue, and utilization.

Index Terms—Radio spectrum management, wireless networks,

I. INTRODUCTION

A. Motivation

THE EXPLOSIVE growth and dynamic nature of wireless traffic makes it costly for a single wireless service provider to buy sufficient spectrum, based on a long-term contract, to sustain its peak load, which may only last for brief periods of time. This long-term spectrum allocation, which is prevalent today, is highly inefficient. This motivates the need for dynamic spectrum access—wherein a wireless service provider only obtains sufficient spectrum to support the “typical” traffic

Manuscript received August 08, 2014; revised January 25, 2015 and May 28, 2015; accepted August 03, 2015; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor H. Zheng. Date of publication October 05, 2015; date of current version August 16, 2016. This work was supported in part by the NSF under Grants CNS-1017549 and CNS-1343383.

W. Dong was with the The University of Texas at Austin, Austin, TX 78712 USA. He is now with Google, Inc., San Francisco, CA 94105 USA (e-mail: weid@google.com).

S. Rallapalli is with IBM Research, Yorktown Heights, NY 10598 USA.

L. Qiu and Y. Zhang are with the The University of Texas at Austin, Austin, TX 78712 USA (e-mail: weid@google.com).

K. K. Ramakrishnan is with the University of California, Riverside, Riverside, CA 92521 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2015.2476497

demands and can: 1) purchase additional spectrum on-demand from other providers to satisfy higher traffic demands; or 2) offer spare spectrum to other providers for profit when there is lower demand.

While dynamic spectrum access is attractive and it is now technically feasible to dynamically change the spectrum to use on-the-fly [6], [20], [28], an important open issue remains—how to allocate spectrum across multiple parties. It is essential to have an incentive framework that can effectively foster collaboration while guarding against dishonest behaviors.

We aim to achieve the following properties for the spectrum double auction, where the first three properties are necessary economic properties for a good double auction and the remaining three quantify the effectiveness of the auction:

- 1) *Truthfulness*: Bidders cannot benefit from bidding differently from their true valuation.
- 2) *Individual rationality*: Bidders get nonnegative utilities, i.e., sellers are paid no less than their asks and buyers do not pay more than their bids.
- 3) *Budget-balance*: The total amount paid to the sellers is no more than the total amount received from the buyers. This prevents the auctioneer, who runs the auction, from losing money
- 4) *Efficiency*: It is the difference between the sum of the winning buyers' valuations and the sum of the winning sellers' valuations. To achieve good efficiency, the goods should be sold to the buyers that value them the most and be sold by the sellers that value them the least. Reference [15] shows that it is impossible to simultaneously achieve truthfulness, budget-balance, and maximum efficiency. Thus, in this paper we aim to achieve truthfulness, budget-balance, and high (but not maximum) efficiency.
- 5) *Revenue*: It is the total amount of payment from all winning buyers. A winning buyer may pay a different amount from its bid, depending on the auction design. A higher revenue gives sellers a stronger incentive to participate.
- 6) *Utilization*: Unique to spectrum auctions, we seek to maximize the spectrum utilization by allowing as many buyers as possible to reuse the spectrum. We quantify the utilization based on the total number of buyers that are assigned spectrum. In addition, we also optimize fairness: Payments should capture the real competition that each individual auction participant faces, where the competition is determined by the wireless interference experienced by each buyer.

B. Challenges

Designing a good double auction for spectrum reuse poses the following significant challenges: 1) How to accurately capture wireless interference among the buyers? This is necessary to support spectrum reuse, where multiple buyers may share

the same spectrum. Achieving spectrum reuse is the key property that distinguishes spectrum auctions from traditional auctions, which assume all entities compete against each other for an item. 2) How to design a truthful double auction? Simply applying truthful auctions for the sellers and for the buyers does not lead to a truthful double auction. It is critical to ensure participants cannot gain by manipulating the interaction between the two sides. This is especially challenging in a spectrum auction because the buyer-side auction differs from the seller side in that we need to take into account wireless interference between buyers in order to support spectrum reuse. 3) How to maximize spectrum utilization while preserving budget-balance? To maximize spectrum utilization, we need to sell as many channels as possible and let them be concurrently used by as many buyers as possible. However, this goal conflicts with budget-balance, as selling more channels means a lower average price and possibly lower revenue (as the price is determined by the loser's price). Lower revenue means that fewer channels may be sold due to the need to preserve budget-balance. We will further elaborate these challenges in Section II-B.

C. Limitations of Existing Spectrum Auctions

The many advantages offered by double-sided spectrum auctions have resulted in strong research attention to the problem. Despite considerable previous work, significant challenges still remain. In particular, some works fail to support spectrum reuse while preserving truthfulness [23], [26]. The first work that satisfies all economic properties and supports spectrum reuse is TRUST [31]. It follows the classic McAfee's double auction [13], which jointly computes the auction result for buyers and sellers. However, to apply the classic design, it makes several simplifications that could sacrifice performance. Specifically, it randomly groups noninterfering buyers and requires buyers in the same group to win or lose together. Hence, the fate of a group is determined by the lowest bidding buyer. A group can lose even if it has many high-bidding buyers, which results in unfairness and low efficiency. It also enforces uniform pricing for all buyers that win the same channel, in which case the price can be no more than the lowest bid of these buyers. This limits total revenue and further hurts the auction performance because only a few channels can be traded in order to preserve budget-balance. Several follow-up works share similar problems [27], [2].

D. Our Approach

We develop a novel double auction for dynamic allocation of spectrum (DA^2) in Section III. DA^2 uses separate designs of the buyer and seller-side auctions while achieving truthfulness, budget-balance, and individual rationality. Compared to the classic joint design strategy (*i.e.*, McAfee [13]) used by most existing spectrum double-auction designs [31], [27], a unique advantage of our separate auction designs for the buyers and sellers is that this enables flexible combinations of different buyer- and seller-side designs. This significantly increases the design space and can immediately benefit from future enhancements to auction design on either the buyer or seller side. Moreover, it also allows different properties of both sides to be captured accurately, which is especially important for spectrum double auctions since the buyer side is much more complicated due to wireless interference. We show how to

combine the design of the two sides to ensure budget-balance and identify the necessary properties of a single-side auction in order for the double auction to be truthful.

DA^2 consists of three components: 1) seller-side auction design, assuming N channels are sold; 2) buyer-side auction design, assuming N channels are sold; and 3) determining the number of channels, N , to sell to satisfy budget-balance.

The buyer-side design is more challenging due to the complications of wireless interference. We propose to partition the graph into subgraphs based on the wireless interference among the buyers. Graph partitioning enables us to compute allocations independently in each subgraph and then combine the results. Therefore, it reduces the randomness of independent set construction and allows us to better capture competition in each subgraph and achieve higher revenue by computing prices in each subgraph independently.

We identify important requirements of a good partition algorithm: cutting as few edges as possible while yielding balanced partitions. These two requirements correspond to the well-known metric, called the Normalized Cut. We apply spectral clustering [16], which is an effective approximation algorithm to minimize the Normalized Cut. We then use the group bid definition in [27] to sort the groups and determine potential winners in a subgraph. We develop a novel pairwise merge procedure to combine the allocation results from subgraphs and preserve truthfulness. It merges the results from two subgraphs at a time and reduces the efficiency and revenue loss due to conflicts between subgraphs.

In Section IV, we further extend DA^2 to support localized sellers, *i.e.*, sellers that sell spectrum that can only be used by buyers in certain regions. To optimize performance and maximize the number of trades to carry out, we develop *cross-region budget-balance* to allow regions to help each other and let a seller win in all regions as long as it is budget-balanced in any one region. In addition, we discuss practical issues of applying DA^2 in Section V, such as incorporating quality of service, leveraging prior knowledge about bid distributions, and how to avoid starvation.

We extensively evaluate our approaches using conflict graphs generated from real cell tower locations derived from a major US cellular provider (Section VI). Our results show that DA^2 consistently yields high efficiency, revenue, and utilization and significantly outperforms the existing approaches: It outperforms by orders of magnitude when compared to TRUST and improves all three metrics to up to $4\times$ when compared to TDSA [27]. When localized sellers are considered, DA^2 improves efficiency to up to $13\times$ of that of District [24], revenue to up to $5.5\times$, and utilization to up to $21.5\times$.

II. BACKGROUND

A. Spectrum Double Auction

A double auction implements a two-sided market. The market consists of three groups of entities: *buyers*, *sellers*, and an *auctioneer*. Buyers submit *bids* that specify the item they are interested in and their maximum willingness to pay. Sellers submit *asks* that include the item they offer for sale and their asking price. The auctioneer evaluates the bids to determine the winners and the items that are traded. The auctioneer also determines the amount to pay the sellers and the amount to charge the buyers. All payments from buyers are paid to the auctioneer,

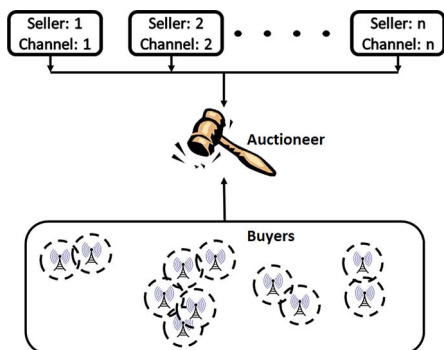


Fig. 1. Double auction: Spectrum market.

and the auctioneer pays the sellers. A buyer's utility is then the difference between his valuation of the item that he wins and the amount that he pays to the auctioneer. A similar utility applies for the sellers.

Fig. 1 shows a spectrum double auction. Here, the sellers are spectrum resource owners (*e.g.*, cellular service providers or organizations that own spectrum resource). The buyers can be entities that need more spectrum. To capture the interference relationships among the buyers, measurement can be conducted by the buyers, sellers, auctioneer, or a third party to derive the conflict graph. We can leverage existing approaches to either actively or passively measure the conflict graph (*e.g.*, [1], [18], [19], and [30]).

After winning an auction, buyers then start using the new spectrum by switching some of their clients onto the new spectrum. Techniques in intracell handover (*e.g.*, [4] and [5]) and spectrum virtualization (*e.g.*, [6] and [20]) can be applied to efficiently and seamlessly switch users to the new spectrum.

Basic Auction Setting: We consider a two-sided spectrum market where the goods of interest are wireless channels (frequency bands) and the players are spectrum resource users, such as cellular service providers. We assume a seller offers one channel for sale and a buyer seeks to buy one channel as considered in previous work [31], [27]. A seller can sell to any buyer in the market. We relax this requirement in Section IV. The auction runs periodically to enable spectrum reallocation in a dynamic fashion. The multiunit double-spectrum auction is more challenging and is left for future work.

We define *Critical value* as the lowest value that a buyer can bid and still win. A *monotonic allocation* means that a buyer who wins by bidding x will still win if it bids y , where $y > x$. The definitions are similar for the seller side.

B. Challenges in Spectrum Double Auction

Challenges in Supporting Spectrum Reuse: Spectrum auction is fundamentally different from a conventional auction in that an item in a conventional auction can be won by one buyer, whereas spectrum can be reused by multiple buyers as long as they do not interfere with each other.

Wireless interference dictates how spectrum can be reused, so it has significant impact on the design of spectrum auction. A *conflict graph* [9] is commonly used to capture the interference relationship between the buyers, where each node in the conflict graph denotes a buyer and there is an edge between two nodes if the corresponding buyers interfere.

Wireless interference has significant impact on spectrum auction. First, conventional auctions do not allow an item to be bought by multiple buyers. However it is important to support spectrum reuse in spectrum auctions to maximize the utilization of the spectrum resource. Second, the conflict graph significantly complicates the competition among buyers. A buyer only directly competes with nearby interfering buyers, but at the same time a buyer's competition is coupled with the competition seen by other buyers, depending on their interference. As an example, consider the conflict graph: $a(7) - c(10) - b(5)$, where “-” represents the interference between two adjacent buyers and the numbers are their bids. In this example, buyer a or buyer b individually does not outbid buyer c . However, since the buyers a and b do not conflict with each other and can share the channel, letting them win together is a better choice than letting buyer c win alone. In other words, a buyer's auction outcome not only depends on its competitors' bids, but also depends on its competitors' competitors' bids and so on. This significantly complicates the auction design.

Challenges in Designing a Truthful Double Auction: As mentioned in Section I, it is challenging to achieve truthfulness in a double auction. Combining a truthful seller-side auction and a truthful buyer-side auction does not necessarily lead to a truthful double auction. This is because the single-sided auctions do not need to consider budget-balance, whereas in a double auction, a participant may try to improve his utility by manipulating the point where budget-balance is achieved and changing the number of traded items. We illustrate this using the following example.

Consider the buyer-side design that tries to maximize efficiency on the buyer side by maximizing the sum of winning buyers' bids and uses critical value pricing. When only one side is considered, the auction is truthful, but this design is not truthful in a double auction. Suppose there are four buyers a , b , c , and d , where a , b , and c do not conflict with each other and they all conflict with d . The valuation of a , b , and c is 3, and the valuation of d is 7. Assuming we have 1 channel to allocate, the winners are a , b , and c because the sum of their valuation is 9, which is greater than 7. The critical value for each winner is 1 because one of the three nodes, say a , can win even when its bid is as low as 1 (as the sum of a , b , c is still as high as d 's bid). Hence, the total revenue is 3 (considering that buyers pay their critical price). However, it is possible that on the seller side, when one channel is sold, the selling price is higher than 3, in which case budget-balance is not satisfied and this trade cannot happen. Hence, no buyer gets a channel, and everyone has 0 utility.

Now a lies and lowers its bid to 1. a , b , and c still win, but the critical values of b and c both become 3 and a 's critical value is 1 in order for the sum of their critical values to still be as high as d 's bid. The total revenue thus increases to 7. If this is higher than the price on the seller side, the channel will be traded and a receives a positive utility of $3 - 1 = 2$. Thus, the double auction is not truthful since a 's utility increases when a lies about its valuation.

Challenges in Maximizing Utilization Under the Constraints of Auction Design: In an idealized world where everyone is altruistic, we would design the system solely to maximize resource utilization, *i.e.*, reuse whenever possible and always allocate all channels. However, in reality there is a

cost associated with the resource, and service providers need incentives for sharing it. Though an auction is considered an efficient way of providing such incentives, it comes with constraints. First, auction design becomes complicated in order to support spectrum reuse. One extreme is no reuse at all, in which case it becomes an auction for normal goods and traditional approaches can be applied. Our goal is to allow reuse while still having a sound auction design. Second, always allocating all available channels means there will be no losing seller, *i.e.*, no competition. However, competition is the major reason that makes auction effective. In a double auction where all sellers always win, the sellers can ask for whatever price they want, and the market would never work. Thus, the auction design needs to be highly effective in boosting competition so that the level of competition is sufficient to support the market, while many channels get allocated. This is especially important in a spectrum auction, where the spectrum resource is precious.

C. Related Work

There is previous work on this topic. Researchers have considered both single-sided auction [29] and double auction for allocating spectrum resource. We focus on double auctions in our work. McAfee's double auction [13] is one of the earliest double-auction designs for single-unit, homogeneous goods that satisfies all the desired economic properties. Reference [3] then proposes a more generalized approach to design double auctions that achieve these economic properties. It introduces the concept of substitutability and removes the requirement for the goods to be homogeneous. It then provides a way to evaluate every single trade, determine whether it should take place, and finally determine the price.

However, it is challenging to apply these existing approaches to spectrum auctions due to the unique properties of the spectrum resource. Unlike traditional goods, where all buyers compete with each other, a buyer in a spectrum auction competes with only those that interfere with it. Thus, defining the appropriate substitutability is critical to the economic properties of spectrum auction. The pioneering and most representative work in spectrum double auction is TRUST [31]. It uses a simple way to define substitutability by generating noninterfering groups at the cost of auction performance. In comparison, our work only establishes substitutability within each subgraph, and deviates from [3] as we cannot evaluate an individual trade alone to determine if it can be sacrificed but need to defer the decision until we have the complete picture. Our approach is an instance of [3] only when there is no conflict between subgraphs and every trade can potentially be evaluated independently. For the general case, we need a new solution like ours to make the allocation decision.

Next, we discuss TRUST in more detail. TRUST divides buyers into multiple independent sets or groups, where the buyers in each independent set do not interfere with each other. For each group, TRUST computes the group bid as the lowest bid in the group multiplied by the group size. Then, McAfee is applied by treating each group as one virtual buyer. As in McAfee, all winning groups pay the group bid of the group in the sacrificed match. Within the group, all members share the price equally. TRUST enables spectrum reuse by grouping

buyers and achieves truthfulness by applying McAfee's design. We believe it has several weaknesses:

- Low efficiency: The independent sets are constructed randomly. It is possible that a high bidding buyer loses just because it is put into a bad independent set. As a result, the final winning groups can be suboptimal.
- Low revenue: The group bid only depends on the lowest bid in the group, and all higher bids are ignored. Thus, the group payment is significantly limited by the lowest bid, which can be much lower than other bids in the group.
- Unfairness: TRUST uses uniform price for the buyers that are assigned the same channel. However, buyers assigned the same channel may have different interference (competition) patterns. The uniform pricing forces some winners to pay for the competition of other winners.

We illustrate the above weaknesses using the following example. The conflict graph and bids are: $a(60) - b(2) - c(30) - d(100)$. Assume two sellers each offering one channel: channels 1 and 2, respectively. Let (a, c) and (b, d) be the independent sets generated by TRUST. The group bid of (a, c) is $30 \times 2 = 60$, while the group bid of (b, d) is $2 \times 2 = 4$. Thus (a, c) is matched with seller 1 and (b, d) is matched with seller 2. (b, d) is the last profitable match, so only (a, c) wins. a and c together pay the group bid of (b, d) , which is 4. Hence, a and c each pay 2.

In this example, because d is in the same group as b , his bid is ignored even though it is the highest. a and d winning together gives the highest efficiency. However, TRUST fails to consider that possibility because of the group bid computation. TRUST also yields low revenue in this example since it only uses the lowest bid, while the winning bidders can potentially pay much more. TRUST is also unfair to d in this example because he cannot win, no matter how high he bids.

There have been several works on improving TRUST [7], [25], [27] and alleviating some of the problems in TRUST. However, they are still similar in spirit in that they still use random independent sets and enforce uniform price for buyers assigned the same channel. Thus, they still suffer from low efficiency and poor fairness. For example, [25] improves on TRUST to sacrifice fewer buyers. However, it uses a single-sided auction, a much simpler problem. Reference [7] builds on [25], but focuses on the privacy aspect instead of the auction design. Reference [27] proposes a new definition of group bid that achieves better auction performance and preserves all economic properties. Our evaluation also compares with [27].

There are also some spectrum double-auction designs that are not based on TRUST. However, they are either not truthful or do not support general conflict graphs. For example, [23] proposes an online spectrum double auction where buyers can request for resource at different times. It assumes a complete conflict graph (*i.e.*, everyone interferes with everyone else), which simplifies the competition pattern and disables spectrum reuse. Reference [26] proposes a single-unit double auction with discriminatory pricing, but their auction is not proven to be truthful. Several papers consider practical issues in spectrum auction. Reference [2] proposes a double auction similar to TRUST, but takes the buyers' frequency preference into consideration while forming groups. This is complementary to our work. Reference [24] focuses on incorporating market locality. Our design

also supports this feature and achieves much better performance as shown in Section VI. In addition, [8] proposes a double-auction-based approach for mobile data offloading, but it does not consider wireless interference as in spectrum auctions.

III. OUR SOLUTION

A. Overview

Design Strategy: We develop a novel double auction for dynamic spectrum allocation. It consists of the following three parts: 1) seller-side auction design assuming N channels are sold; 2) buyer-side auction design assuming N channels are sold; and 3) a procedure to determine the number of channels, N , to sell to satisfy budget-balance.

To the best of our knowledge, our solution represents the first double-auction design for spectrum allocation that explicitly decouples the buyer-side and seller-side auction design while achieving: 1) truthfulness; 2) individual rationality; and 3) budget-balance. Previously, seller- and buyer-side auctions were designed jointly in order to satisfy all three properties (e.g., McAfee's principle used in TRUST and all its variants).

Decoupling seller-side and buyer-side design is crucial for dynamic spectrum allocation for two reasons. First, by using separate auction designs for the buyers and sellers, we no longer require groups of buyers that share the same channel to be formed in advance to match with sellers one by one. Instead, we can make more informed decisions based on the bids to select a stronger set of winners and improve efficiency, revenue, and utilization. Second, it enables flexible combination of different buyer/seller-side designs. This is especially beneficial in our context since the two sides have rather different properties. For example, it is commonly assumed that a seller can sell to any buyer in the auction. Competition between sellers is similar to traditional auctions, whereas the buyers' competition is much more complicated due to complex wireless interference.

Proof Strategy: The key to designing the two sides separately is a new proof strategy to establish truthfulness in a (decoupled) double auction. Formally, we have the following theorem.

Theorem 1: A double auction for dynamic spectrum allocation is truthful if the following two conditions hold: i) both seller-side and buyer-side auctions are truthful when the number of channels that are sold, denoted as N , is fixed; and ii) no seller or buyer can improve its own utility by unilaterally modifying its own bid and causing N to change.

The correctness of this theorem is easy to see: When a bidder lies but does not change N , it cannot gain because both sides are truthful when N is fixed. When a bidder lies and changes N , it cannot gain either because no seller or buyer can improve its own utility by unilaterally modifying its own bid and causing N to change. Thus, a bidder never gains by lying, and the auction is truthful.

We use the following theorem from [10] to assist our design.

Theorem 2: If every losing bidder (*i.e.*, unallocated) in an auction pays nothing, it is truthful if and only if its allocation algorithm is monotonic and it uses critical value as payment.

Design Overview: Below is an overview of our design. The buyer-side auction is complicated since buyers' competition is determined by the complex conflict graph and buyers without interference can share a channel. Therefore, this paper focuses on the buyer-side auction design.

The seller-side auction is a standard auction since a seller can sell to any buyer in the auction. Therefore, we can apply uniform

pricing or VCG for single-unit auction. We prove it satisfies the properties specified in Theorem 1.

To determine the number of channels N , we start by setting N as the total number of channels that sellers collectively have. We run seller- and buyer-side auctions separately, and then check if budget-balance is satisfied (*i.e.*, the payment collected from the buyers is no less than the payment to be paid to the sellers). Note the payment is determined by our seller-side and buyer-side auction design described below and not the sum of winning bids/asks. If budget-balance is already satisfied, we terminate. Otherwise, we decrease N by 1, and run the auction again. By reducing N , it requires buyers to bid even higher to win and sellers to ask even lower to sell, thus reducing the gap between revenue and payment. We stop when the budget-balance is satisfied and N channels are then traded. This procedure always terminates because budget-balance is satisfied when N drops to 0. Thus, we guarantee budget-balance, which is stated in the following theorem.

Theorem 3: Our design satisfies budget-balance.

Note that McAfee's design is a special case of our framework, where both the buyer-side design and the seller-side design use uniform pricing. This shows that our framework is general and under our framework it is possible to design a double auction that satisfies all three economic properties. We describe our design in detail.

B. Seller-Side Design

The seller-side design is a standard auction since there is no need to consider interference among sellers. Thus, we simply use the traditional uniform price design. Assuming N channels are sold, the N sellers with the lowest asking prices win, and they each get paid at the $(N + 1)$ th seller's asking price. When there is no $(N + 1)$ th seller, a predetermined "reserve" price can be used. This uniform price design is known to be truthful in a single unit auction when N is fixed. Moreover, since sellers are paid higher than their asking prices, individual rationality is satisfied. In the Appendix, we further prove Theorem 4 to ensure this design is truthful when applied to a double auction.

Theorem 4: This seller-side design, when applied to double auctions, does not allow a seller to unilaterally manipulate N (*i.e.*, the number of channels that can be sold) and gain.

C. Buyer-Side Design

Overview: To ensure fairness and achieve high revenue, it is important for the buyer-side auction to explicitly take into account the fact that different buyers face different levels of competition depending on their locations and interference patterns. Existing approaches do not account for such individuality and apply uniform pricing, where all buyers that share the same channel pay the same amount. This is not only unfair, but also reduces revenue since a buyer's payment is limited by the lowest bids among the buyers that share the channel.

Our key idea is to divide-and-conquer. Specifically, we first partition the conflict graph into subgraphs. Independent sets are then constructed within each subgraph, and pricing is computed independently in each subgraph. We then design a merge strategy to combine the winners from different subgraphs while preserving truthfulness. Note that when computing the allocation within each individual subgraph, we ignore the inter-subgraph conflict edges. During the merge procedure, we need




Fig. 2. Benefit of graph partition. (a) Without partition. (b) With partition.

to add back the intersubgraph conflict edges. As a result, some winners from individual subgraphs may have to be removed due to the intersubgraph conflict edges. The challenge is how to do so without compromising truthfulness.

Benefit of Graph Partition: Before going to the details of our design, we first show the two major benefits of graph partition: 1) better independent set construction, and 2) higher revenue due to the removal of the uniform pricing constraint.

We illustrate these benefits using an example in Fig. 2. It shows a conflict graph involving 6 nodes. We ignore the seller side in this example and assume we want to allocate 1 channel. To see the benefit 1), Fig. 2(a) shows the groups constructed by TRUST: $G1: (a, e, f)$ and $G2: (b, c, d)$. The group bids of $G1$ and $G2$ are $20 \times 3 = 60$ and $10 \times 3 = 30$, respectively, which are the lowest bid multiplied by the group size. Hence, in TRUST, $G1$ wins and pays the first losing group's bid, which is 30. The efficiency on the buyer side is the sum of all winning buyers' bids, which is $20 + 20 + 30 = 70$, and the revenue is 30. Now if we treat each subgraph independently as shown in Fig. 2(b): $a, b,$ and c win in their subgraphs, and each pays the first losing buyer's bid, i.e., 10, 20, and 30, respectively. Then, we get the efficiency of $20 + 30 + 40 = 90$ and revenue of $10 + 20 + 30 = 60$, both of which are higher than TRUST.

To demonstrate the benefit 2), we consider the case when TRUST happens to construct the same groups as ours, namely $G1': (a, b, c)$ and $G2': (d, e, f)$. The group bids of $G1'$ and $G2'$ are $20 \times 3 = 60$ and $10 \times 3 = 30$, respectively. So $G1'$ wins, and pays the first losing group's group bid, which is 30. This is still lower than the revenue in our scheme because the revenue of TRUST is limited by the uniform pricing (i.e., the lowest bid multiplied by the number of buyers). The difference in revenue may have an even bigger impact when we take the seller side into consideration. If the revenue is lower than the amount sellers should get, some channels cannot be traded because it violates budget-balance, which further reduces efficiency and revenue.

Essentially, partitioning allows us to decouple nodes with no or weak interference into different subgraphs to improve the independent set construction and avoid unnecessary coupling in different buyers' pricing when they interfere with different sets of nodes. Note the example shows a disconnected conflict graph, but the idea also works on realistic, densely connected conflict graphs as we show in Section VI.

Design Questions: Several important questions should be addressed in order to realize the benefits of graph partition:

- How to partition the graph to retain important interference relationships within a subgraph and decouple nodes with weak and no interference into different subgraphs?

- How to compute auction results within each partition?
- How to merge the auction results from different partitions to achieve truthfulness and budget-balance?

We answer each question in turn.

1) *Graph Partitioning:* Given a graph A , a graph partitioning algorithm strives to find a partitioning A_1, A_2, \dots, A_k that minimizes a certain objective. For the purpose of dynamic spectrum allocation through double auction, a good graph partitioning algorithm should balance two key requirements: 1) the number of intersubgraph edges should be small, and 2) different subgraphs should be similar in size. If a subgraph is too small (e.g., having only one buyer), a small number of channels can satisfy all its buyers, leaving no losing buyers and no revenue from the subgraph. On the other hand, too big a subgraph may lead to poor independent set construction and poor performance due to the uniform pricing in each group.

Two common objectives have been proposed for graph partition: RatioCut and Normalized cut (NCut). The former normalizes the weights of the edges on the cut by the number of vertices in each partition (i.e., minimizes $\sum_{i=1}^k \frac{W(A_i, \bar{A}_i)}{|A_i|}$, where $W(A_i, \bar{A}_i)$ is the total weight of all edges between A_i and the remaining nodes \bar{A}_i , and $|A_i|$ denotes the number of vertices in A_i). The latter normalizes the weights of the edges on the cut by the sum of node degrees in each partition (i.e., minimizes $\sum_{i=1}^k \frac{W(A_i, \bar{A}_i)}{vol(A_i)}$, where $vol(A_i)$ denotes the sum of degrees of all nodes in the partition A_i). It is easy to see $\frac{1}{|A_i|}$ and $\frac{1}{vol(A_i)}$ are minimized when either the number of vertices or the sum of node degrees within each partition is the same. This captures our goal of finding balanced cuts while minimizing the weights of edges on the cut.

Minimizing either RatioCut or NCut is an NP-hard problem. Spectral clustering is a well-known effective scheme to find approximate solutions to this NP-hard problem (see [12] for a nice tutorial). There are many variants of spectral clustering [12], [14], [16]. In this paper, we use the Meila–Shi algorithm [14] to minimize NCut, which is the recommended algorithm in [12] due to its excellent performance and solid mathematical foundation. Let W be the adjacency matrix with weight w_{ij} on its i th row and j th column. Let D be the degree matrix, which is a diagonal matrix with the node degree $d_i = \sum_j w_{ij}$ on the diagonal. The Meila–Shi algorithm takes the eigenvectors corresponding to the k smallest eigenvalues of the normalized graph Laplacian matrix $L_{rw} = I - D^{-1}W$ (where I is the identity matrix), and then invokes another algorithm (e.g., k -means clustering [11]) to cluster points by their respective k components in these eigenvectors.

To automatically determine the number of clusters to create (i.e., k), we follow the suggestion of [12] and apply the eigengap heuristic. Specifically, let $\lambda_1 \leq \lambda_2 \leq \dots$ be the eigenvalues of the normalized graph Laplacian matrix L_{rw} sorted in an ascending order. The eigengap heuristic computes all the eigengaps (i.e., difference between two successive eigenvalues) and chooses the number of clusters k such that $(\lambda_{k+1} - \lambda_k)$ is the largest eigengap.

Note that when the conflict graph is disconnected, we first divide it into multiple connected components since nodes in different connected components have no competition at all. Then, we apply spectral clustering to each connected component to further partition the connected component.

2) *Allocation and Pricing Within a Subgraph*: When a good partitioning is found, we first compute the allocation in each subgraph independently. For that we can apply existing algorithms, such as TRUST. To further improve the performance, in our implementation we apply the allocation algorithm proposed in TDSA [27]. It is similar to TRUST but it defines a new group bid and allows a subset of a group to win while the rest lose. Consider a group k and assume its members are sorted in a decreasing order of their bids. Denote members as 1 to m , and the group bid as $\max\{b_i \times i \mid i = 1 \dots m\}$. This group bid quantifies the maximum potential payment of a group if we allow a subset of this group to win. For each subset, the maximum potential payment is the lowest bid in the subset times the size of the subset. Thus, this group bid finds the highest potential payment by enumerating all possible sizes of the subset. For example, if a group contains bids $(1, 3, 5)$, its potential payment could be either letting 5 win alone, which corresponds to payment of $5 \times 1 = 5$; or letting 3 and 5 win together, which yields a payment of $3 \times 2 = 6$; or letting all of them win, which yields a payment of $1 \times 3 = 3$. The maximum payment achieved is 6 in this example.

To allocate N channels, the groups with the top N group bids win. A winning group is then charged the first losing group's group bid and all members in the group share the price equally. If a member i cannot afford its fair share (*i.e.*, $b_i * i$ is smaller than the first losing group bid, which is possible due to TDSA-based group bid computation), i does not win, and the price is shared among the remaining group members. Since the winning group has a higher group bid than the losing group, there always exists a subset of the winning group such that they all bid no lower than their fair share. This automatically guarantees individual rationality on the buyer side. We prove that this procedure finds the critical value of a buyer under our partition and merge framework in Theorem 6 described in the Appendix.

We also make a temporary assignment by sorting the channels in increasing order of their asking prices and assign the first channel to the first group and second channel to the second group, and so on. Note if there is no conflict between the subgraphs, then this is a good final allocation and we can compute the total revenue by summing up the revenue from each subgraph and determine if N needs to be reduced or not. In this case, the scheme looks like a very simple generalization of McAfee/TRUST/TDSA in that it finds and sacrifices the last efficient trade such that the payment is sufficient to support the remaining sales. However, it often requires removing edges for better and more balanced partition, in which case the assignment is subject to change in the merge procedure below.

3) *Merge Strategy*: Next, we describe how to merge allocation results from different subgraphs. We merge two subgraphs at a time. The input to the merge are the winners selected from the two subgraphs, including their channel assignment. The purpose of the merge is to find one way to reorder the channel assignment in one subgraph such that winners on the cut do not conflict. If such a reordering does not exist, certain nodes on the cut may be dropped. Note all these happen without affecting the pricing within each subgraph. We propose the following bid-independent merge strategy to select a node to drop. We prove this merge strategy preserves truthfulness.

```

1 for both  $A$  and  $B$ 
2   Find all winning buyers on the cut
3   Merge buyers sharing the same channel into a virtual buyer
4   Put all winning buyers into  $Candidate$ 
5 for each  $i$  in  $Candidate$ 
6   Let  $degree_i$  be number of  $i$ 's edges on the cut
7    $flag =$  a feasible reordering function  $reorder(x)$  can be found,
   where  $x$  can be any channel currently used in  $B$ .
   Note  $reorder(x)$  is only defined for channels that need to be replaced.
8 while  $flag$  is  $false$ 
9   Sort  $Candidate$  in a decreasing order of  $degree_i$ 
10  Drop  $Candidate[0]$ 
11  Update effective degrees
12   $flag =$  a feasible reordering function  $reorder(x)$  can be found
13 for each winning buyer  $i$  in  $B$ 
14   if  $reorder(assign_B(i))$  is defined
15      $assign_B(i) = reorder(assign_B(i))$ 
16  $assign_{AB} = assign_A \cup assign_B$ 
17 return  $assign_{AB}$ 

```

Fig. 3. Pseudo-code for bid-independent merge.

Fig. 3 gives the pseudo-code of the merge procedure for two subgraphs A and B . In lines 1–4, we preprocess the buyers on the cut by combining the buyers that share the same channel on each side into one virtual node because these buyers are always assigned the same channel no matter how the assignment is reordered. Lines 5–6 compute the *effective degree* for each buyer i , which denotes the number of edges that i has on the cut. This information is used in lines 7–12 to determine which buyer to drop when a feasible reordering does not exist. We search for a feasible reordering by enumerating all possibilities. The cost of this procedure is acceptable because its complexity depends on the number of channels (which is small) instead of number of buyers. We drop buyers in a decreasing order of their effective degrees in order to minimize the number of dropped buyers and reduce loss in efficiency and revenue. Every time a buyer is dropped, the effective degree of the remaining nodes is updated. We iterate until a feasible reordering is found, and then use it to derive a joint assignment for the union of A and B . Intuitively, our partition and merge framework preserves truthfulness because both operations by themselves are bid-independent and there is no incentive to lie. It is still possible to lie to change the allocation in subgraphs, which changes the input of the merge operation, but a buyer cannot gain in this way due to the use of critical value pricing. We formally prove truthfulness in Theorem 5.

Fig. 4 shows an example of merging two subgraphs. Let A and B be the left and right subgraphs, respectively, and $c1$ and $c2$ denote the channels to be allocated. The labels next to nodes represent the channel assignments to the nodes. Fig. 4(a) shows the assignment to the winners in each subgraph. If in one subgraph, some winning buyers on the cut share the same channel, they should be merged to form a virtual buyer before we proceed. There are no such buyers in this case, so we go to the next step by adding back the removed edges. In Fig. 4(b), the three previously removed edges are added back. There is conflict between nodes 3 and 4 and conflict between nodes 2 and 7. In Fig. 4(c), the algorithm tries to reorder the assignment in B by swapping $c1$ and $c2$, but there is still conflict between nodes 3 and 7. Thus, in Fig. 4(d), node 3, which is the node on the cut with highest degree, is dropped to resolve the conflict.

Theorem 5: DA^2 is truthful.




Fig. 4. Simple example of the merge procedure. (a) Subgraphs with temporary assignment. (b) Removed edges added back. (c) Assignment reorder fails to resolve conflict. (d) Node 3 dropped.

Proof: To prove this theorem, we first prove Theorems 6 and 7 in the Appendix, and then apply Theorems 1 and 4. ■

IV. INCORPORATING SELLER LOCALITY

Motivation: So far, our solution assumes a seller can sell to all buyers. In reality, a seller may own channels in a limited area and cannot sell to all buyers. This problem is considered in [24], but it requires each buyer–seller pair to be budget-balanced, which is a very strong requirement and can significantly limit spectrum reuse because the new buyer may not be able to afford the entire cost of a channel. An important advantage of a spectrum double auction is that we can leverage payment from multiple buyers that reuse the same channel to pay to a seller.

Challenges: A naive approach is to divide the global area into regions where each region is covered by a few sellers and in each region sellers can sell to all buyers in that region. Then, our allocation can be applied to each region independently, assuming each buyer only belongs to one region. However, treating each region as an individual auction can prevent feasible trades and limit spectrum reuse because the revenue in a single region may not be enough to satisfy the sellers' price (*i.e.*, budget-balance is violated within one region), while the revenue from multiple regions together may be sufficient to pay for the sellers (*i.e.*, budget-balance is satisfied globally). Moreover, uniform pricing for sellers may not be truthful since only the sellers whose coverage overlap compete against each other. Different sellers face different competition.

To address these challenges, we change the seller-side design of the auction to use *Cross-region budget-balance* and *Nonuniform seller-side pricing*. In *Cross-region budget-balance*, we allow a region that has sufficient revenue to help other regions. Note in this design we treat the auction across all the regions as a single auction, and a service provider still only places one ask for selling its resource in the regions that it is available. We illustrate the idea using a simple example. Consider a region with three sellers 1, 2, and 3, and sellers 2 and 3 are already budget-balanced in other regions (*i.e.*, sellers 2 and 3 get enough payment from the buyers in other regions). Then, we can allocate seller 2 and 3's resource in this region without considering their asking price when checking budget-balance, and this region is budget-balanced as long as the revenue is high enough to pay for seller 1. Note this not only improves spectrum reuse,

```

1 for each seller  $i$ 
2    $SellerBalanced_i = 0$ 
3    $Lose_i = 0$ 
4 for each region  $m$ 
5    $RegionBalanced_m = 0$ 
6 while not all regions are budget balanced
7   find seller  $j$  with the highest ask price  $a_j$ 
8   and  $SellerBalanced_j = 0$ 
9   Drop  $j$  by setting  $Lose_j = 1$ 
10  for each region  $m$  where  $RegionBalanced_m = 0$ 
11    Compute allocation assuming all remaining sellers in  $m$  win
12    Compute buyer pricing and the total revenue  $R$ 
13    Compute the total selling price:
14     $S = \sum_{i \in m, SellerBalanced_i = 0} a_j$ 
15    if  $R \geq S$ 
16       $RegionBalanced_m = 1$ 
17      for each seller  $i$  in  $m$  where  $Lose_i = 0$ 
18         $SellerBalanced_i = 1$ 
19 Allocation completes and all sellers with  $SellerBalanced_i = 1$  win

```

Fig. 5. Pseudo-code for allocation with cross-region budget-balance.

but also strictly improves auction efficiency, as we allow more buyers to win, increasing buyer utility without affecting seller utility. In *Nonuniform seller-side pricing*, we compute the critical value of winning sellers under our new allocation algorithm to ensure truthfulness. We describe these in detail as follows.

Cross-Region Budget-Balance: To generalize the intuition from the previous simple example, we design the scheme outlined in Fig. 5. At a high level, it finds regions that are budget-balanced and uses them to lower the budget requirement for the other regions that overlap with the budget-balanced regions. As shown in Fig. 5, lines 1–5 initialize all regions and all sellers to not be budget-balanced, and all sellers are initially included in the auction. In lines 6–17, we drop the seller with the highest asking price that is still not budget-balanced until all regions are budget-balanced. Because the selling price is determined by the losing sellers' prices, dropping a seller lowers the selling price and makes it easier to satisfy budget-balance. Every time a seller is dropped, as shown in lines 9–13, we compute allocation in each region that is still not budget-balanced and check if budget-balance is satisfied with the new allocation, assuming any seller that is not already budget-balanced will be paid at the last dropped seller's asking price. For sellers that are already budget-balanced, we consider the payment as 0. In lines 14–17, if a region becomes budget-balanced in this round, we label all its sellers left as budget-balanced, and we also label this region as budget-balanced. Finally, all the sellers that are budget-balanced will win.

With cross-region budget-balance, the point where budget-balance is satisfied is no longer characterized by a single N . Instead, it is a vector of numbers N_i , where N_i is the first losing seller's position in region i .

Nonuniform Seller-Side Pricing: It is easy to see that the above allocation is monotonic for sellers. Then, in order to achieve truthfulness, we pay the winning sellers their critical values. If a seller becomes budget-balanced after seller j is dropped, its critical value is the j 's asking price. Different sellers' critical values may differ, which is desirable, since it captures their different competition patterns. We prove that it is indeed the critical value in Theorem 8 in the Appendix.

According to Theorem 2, we conclude that the seller side is truthful if the budget-balance point does not change. To show that the double auction is truthful, we require that no seller can change the budget-balance point and gain by lying. The proof is similar to the proof of Theorem 4 and is omitted.

V. PRACTICAL ISSUES

Next, we discuss practical issues involved in applying DA^2 .

Auction Cost: Auction can be conducted periodically or on demand whenever the inputs change significantly. The frequency of auction needs to be chosen to balance the cost of running the auction and changing the spectrum allocation versus the optimality of the spectrum allocation. In case of a sudden change in demand or conflict graph, a new auction can be conducted to reallocate the spectrum resources. The auction cost includes gathering inputs for the auction (*e.g.*, bids, asks, conflict graph), computing the auction solution, disseminating auction outputs, and collecting payments. On average, given 3–7 sellers and with the NYC downtown conflict graph, which has 230 buyers (see Section VI for the details of the data), the auction takes 48 s to run on a desktop with Intel E8200 CPU and 2 GB memory. As part of this, spectral clustering, computing allocation within subgraphs, and merging take 10, 23, and 15 s, respectively. This is fast enough for practical use. When the conflict graph is big, allocation in different subgraphs and the merge between different pairs of subgraphs can be easily parallelized to save time. Spectral clustering only needs to be recomputed when the conflict graph changes. The cost of changing spectrum allocation involves notifying user devices, changing parameters on towers and devices, and handoffs. Handoffs can be achieved efficiently owing to recent advances in handoff and spectrum virtualization techniques.

Monitoring Quality of Service: The auction-based approach is only useful if its results are honored. That is, sellers stop using channels that are sold in the auction area, and buyers only use the channels they win in the appropriate area. To ensure this, the auctioneer monitors the quality of service by letting the buyers report performance anomalies and the locations where the anomalies occur. Upon receiving an anomaly report, the auctioneer identifies the potential causes of the anomaly and verifies the real coverages of the involved participants. Using unauthorized spectrum resource is illegal. If that happens, the auctioneer can take legal action. If the anomaly is caused by careless parameter settings or resource management, the involved participants can be punished by making it harder for them to win in future auctions by using the following *Reputation Score*.

Reputation Score: The auctioneer maintains a reputation score r_i ($0 < r_i \leq 1$) for both buyers and sellers to reflect their quality. For a seller, the score can be based on his channel quality. For a buyer, it can be based on whether this buyer uses wireless resources carefully without causing extra interference to other buyers. The higher the score, the better the reputation. When computing the allocation, we divide every seller's asking price by its reputation score so that it is harder for a seller with a bad reputation to win. We then multiply a seller's reputation score with its critical value to compute its final payment. Since the critical value is computed based on the reputation-weighted asking prices and is greater than a winner's asking price divided by its reputation, individual rationality is achieved. Similarly, for buyers we can multiply their bids by their reputations when computing the allocation, and then divide its critical value by its reputation score to get the real price they need to pay to ensure individual rationality. Similar ideas have been used in sponsored search auctions and online advertisement auctions [21], [22]. Truthfulness is preserved as the reputation scores do not depend on bids. Moreover,

budget-balance is satisfied regardless of whether we consider the reputation score adjustment when determining N because the real selling price is always less than or equal to the price before adjustment and real payment is always greater than or equal to its corresponding unadjusted value. Using the adjusted values may allow a higher N .

Leveraging Prior Knowledge: The independent set construction is critical to the performance. Our solution mitigates the randomness but the construction can be further improved if the auctioneer has prior knowledge, *e.g.*, distributions of buyers' valuations. Specifically, we can maximize the expected valuations in winning groups, which directly relates to auction efficiency by formulating the construction as a *maximum weight independent set* problem (MWIS). The expected valuation of an independent set is simply the sum of expected valuations of all the members. The MWIS problem on conflict graphs can be approximated in polynomial time, *e.g.*, [17]. When the error of prior knowledge with respect to the actual bids is up to 10%, the efficiency improves by 16% over our current scheme on average; when the error increases to up to 50%, the average improvement becomes 6%. The auction is still truthful as long as the prior knowledge is bid-independent and cannot be manipulated.

Avoiding Starvation: Our conflict graph partition is designed to minimize the number of intersubgraph edges while balancing the subgraphs. However, it is still possible that well-connected nodes are on the cut of two subgraphs. These nodes may be always dropped even if they win in their subgraph due to their large degrees and get starved. To address the starvation issue, we can introduce randomness in the merge procedure by dropping a buyer with a probability proportional to its degree. Thus, well-connected nodes are more likely to be dropped, but they are not always dropped and will not starve. This new procedure is still bid-independent, so truthfulness is preserved.

VI. EVALUATION

Simulation Setup: In order to experiment with realistic conflict graphs, we use the location data of cell towers from a large US service provider. We consider there are buyers at each tower/location looking for spectrum resources at that location and that they do not collude. We construct the conflict graph for three cities: New York City (NYC), San Francisco (SF), and Chicago. In each city, we pick a grid (approximately 5×5 km²) encompassing the downtown area and use all the cell towers in the grid to generate the conflict graphs. We consider that two nodes conflict if the internode distance is smaller than 500 m, which is considered a typical cell range. We also vary the range to see how the network density impacts performance. We only present results with conflict graphs generated based on real locations, but we also experimented with random conflict graphs and observed similar benefits.

We use 5 sellers by default, and also vary the number from 3 to 7 to see the impact. The small number of sellers is consistent with the reality where there are only a few large spectrum owners. The number of winning sellers is constrained by the conflict graph and increasing the number of sellers further has little impact on the result. For the sellers' asks and buyers' bids, the absolute values do not matter and only their ratio matters. We generate the buyers' bids drawn from a uniform distribution between 0–100. We also use a uniform distribution to generate

the asking prices. Since each grid can cover at most 25 buyers and we assume buyers and sellers value the spectrum resource similarly such that the price from the two sides compare fairly, we generate the asking prices to be between 0–2500 such that the mean is 25 times the mean of a buyer's bid. We also scale the sellers' asking prices to see its impact in our evaluation.

We compare our scheme to TRUST and TDSA [27] in terms of the following three metrics:

- 1) *Efficiency*: This is widely used to quantify auction performance. It is defined as the difference between the sum of all winning buyers' bids and the sum of all winning sellers' asks.
- 2) *Revenue*: It is defined as the total payment from all winning buyers. Revenue is different from efficiency on the buyer side since revenue depends on critical values, whereas efficiency depends on bids. A higher revenue gives a stronger incentive for sellers to participate.
- 3) *Utilization*: This is defined as the number of winning buyers. This is a unique metric in spectrum auctions because the spectrum resource is precious but reusable, and a higher utilization means more winners can utilize the spectrum at the same time, which is preferred. Utilization cannot be derived from efficiency since the winning buyers may not be the ones with the highest bids, and it is impossible to tell the number of winning buyers just from the efficiency value.

For every setting, we repeat the experiment 20 times with different random asking prices and bids and report the average.

Performance at Different Locations: We first compare the performance in all three cities. As shown in Fig. 6, the performance varies significantly across cities. Specifically, SF gives the highest value for all three metrics, while Chicago gives the lowest and NYC lies in between. The performance difference is primarily due to the difference in the number of buyers. SF has 16% more buyers than NYC, while the number of buyers in Chicago is only around half of that in NYC. With the same interference range, more buyers means more spectrum reuse, and this tends to generate higher revenue, thus more channels can be traded. We also evaluate the impact of network density with fixed location by varying interference range in subsequent experiments. Among all the schemes, DA^2 achieves the best performance for all three metrics and for all three cities. Compared to TRUST, DA^2 's efficiency is $22\times$ – $62\times$, its revenue is $27\times$ – $126\times$, and its utilization is $42\times$ – $65\times$. TRUST does not perform well because it is limited by the use of uniform pricing and the lowest bid in each group. Only few channels can be sold, and their revenue is low. TDSA performs better than TRUST because it searches for the best subset of a group to win and is thus more robust to the lowest bid in a group. DA^2 still outperforms TDSA by 51% to 101% in efficiency, 57% to 115% in revenue, and 47% to 93% in utilization. This is because DA^2 : 1) decouples pricing in different subgraphs to better capture the competition; and 2) combines the top groups from different subgraphs to reduce randomness and find a set of winners with higher valuations.

For the rest of our study, we use the NYC conflict graph as default and study the impact of other parameters.

Impact of the Number of Sellers: Next, we vary the number of sellers. Fig. 7 shows that as the number of sellers increases,

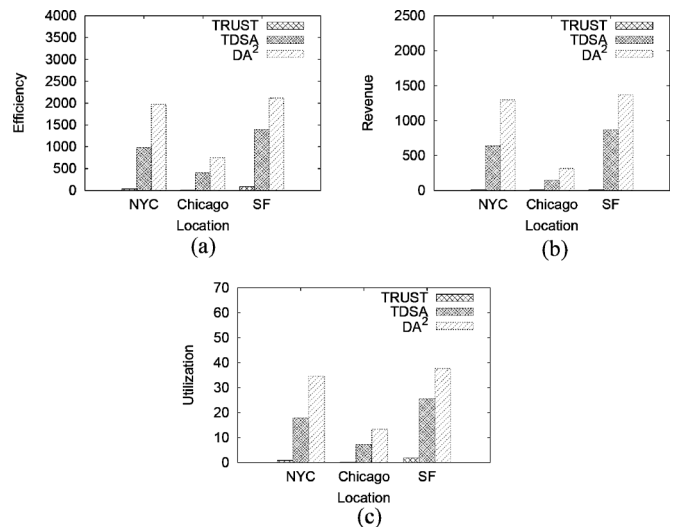


Fig. 6. Performance at different locations. (a) Efficiency. (b) Revenue. (c) Utilization.

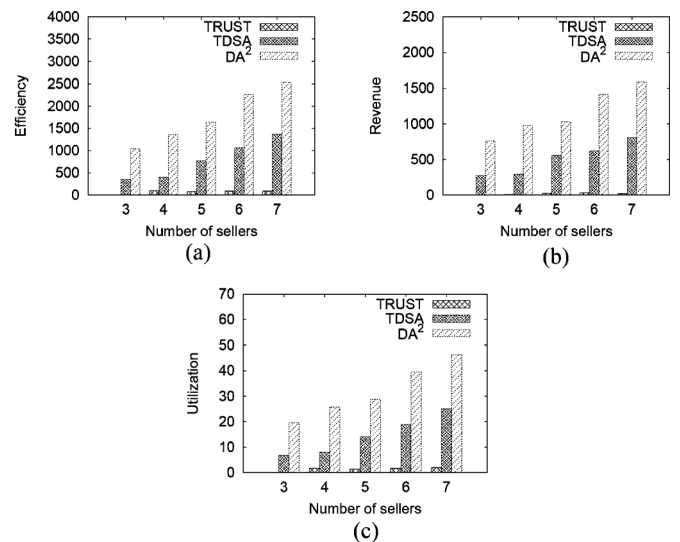


Fig. 7. Impact of the number of sellers. (a) Efficiency. (b) Revenue. (c) Utilization.

more channels become available and the price reduces due to increased seller-side competition. As a result, more channels can be traded, and all performance metrics improve. DA^2 consistently outperforms the existing approaches. Its improvement is highest when there are only 3 sellers, in which case its performance is $3\times$ of TDSA in all three metrics. TRUST does not sell a single channel in the 20 runs with 3 sellers because its revenue is low and it is more challenging to sell a channel while ensuring budget-balance when the number of sellers is small. When the number of sellers increases to 7, compared to TRUST, DA^2 's efficiency is $28\times$, its revenue is $72\times$, and its utilization is $24\times$; DA^2 outperforms TDSA by 85% in efficiency, 97% in revenue, and 86% in utilization.

Impact of Network Density: Next, we vary the network density by changing the buyer communication range from 250 to 750 m. A longer range indicates more competition among buyers and fewer buyers can reuse a channel. As shown in Fig. 8, the benefit of our scheme increases with the range since it is harder to sell a channel, and a good auction design

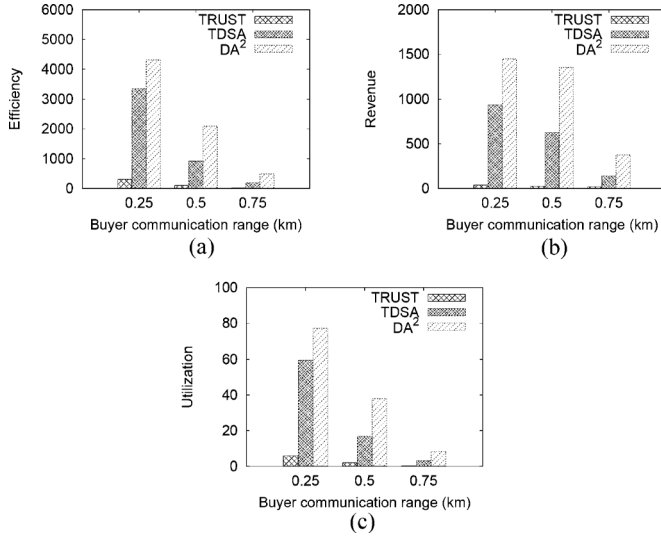


Fig. 8. Impact of network density. (a) Efficiency. (b) Revenue. (c) Utilization.

becomes even more important. For example, when the range is 250 m, our scheme outperforms TDSA by 29% in efficiency, 55% in revenue, and 30% in utilization. The corresponding numbers for a range of 750 m are 152%, 172%, and 173%, respectively. The performance trend in TRUST is less clear because TRUST heavily depends on the random independent set construction. However, under all three ranges, DA^2 is much better than TRUST: Its efficiency is $14\times$, its revenue is $22\times$, and its utilization is $13\times$.

Impact of Bid Distribution: To understand the impact of bid distribution, we scale all sellers' asking prices by a factor between 0.5 to 1.5 after the initial ask is drawn from a uniform distribution between 0 and 2500. This changes the ratios between sellers' valuation and buyers' valuation. A higher asking price requires a higher revenue from the buyer side in order to sell the channels, and is more challenging for an auction. This trend is shown in Fig. 9. In all cases, DA^2 outperforms the existing schemes. The benefit of DA^2 increases with the asking price. For example, when we scale the asking price to 1.5 times, our approach improves all three metrics to around $4x$ of those in TDSA. The improvement over TRUST is even larger, as TRUST does not sell anything and all its metrics are 0.

Seller Locality: To study localized sellers, we compare DA^2 to the District algorithm described in [24], which is specifically designed to incorporate seller locality. For fair comparison, we compare with the solution in [24] that does not require prior knowledge. In order to take seller locality into consideration, it takes a conservative approach and requires every single seller-buyer pair to be budget-balanced, and then finds the assignment for buyers one by one.

We randomly generate three regions and also randomly assign buyers to one of these regions with the same probability. District requires a preselected parameter that specifies the maximum number of winning buyers, and picks that many highest bid buyers and uses the bid of the first buyer that is not selected as the amount to charge the winning buyers. We vary that parameter from 10 to 160.

Fig. 10 summarizes the result. We make three observations: First, our scheme is significantly better in all the experiments.

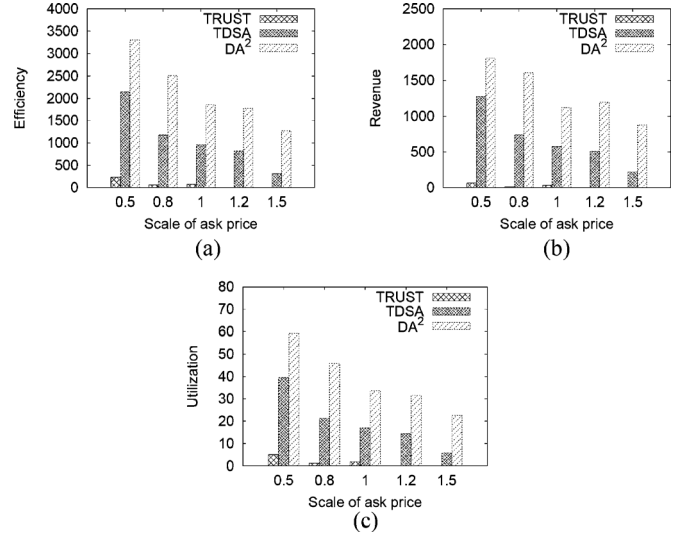


Fig. 9. Impact of bid distribution. (a) Efficiency. (b) Revenue. (c) Utilization.

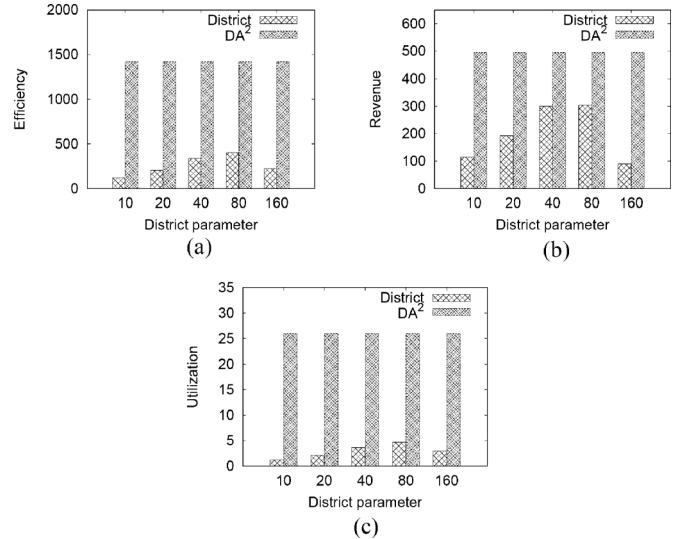


Fig. 10. Local sellers. (a) Efficiency. (b) Revenue. (c) Utilization.

For example, DA^2 improves efficiency to $3.5\times$ to $13\times$ of that of District, revenue to 162% to $5.5\times$, and utilization to $5.5\times$ to $21.5\times$. The main reason is that District requires each seller-buyer pair to be budget-balanced. This artificial restriction makes it hard to complete a trade. Our scheme not only benefits from considering all the buyers from all regions, but also uses cross-region budget-balance to allow regions with higher revenue to help weaker regions, thus achieving much better performance. Second, the preselected parameter (*i.e.*, the maximum number of winning buyers in District) has significant impact on the performance of District, and the best parameter is difficult to determine. If the parameter is too small, too few buyers get selected; if the parameter is too big, the payment from winning buyers becomes too low, and only a few sellers have their budget-balance requirement satisfied as a result. In comparison, our scheme does not require a preselected parameter. Third, compared to the default setting as in the previous evaluation where a seller can sell to all buyers, the absolute efficiency, revenue, and utilization are all lower. This is as what

we would expect since the sellers here only sell in a specific region, which may not cover all buyers.

VII. CONCLUSION

In this paper, we develop a truthful spectrum double auction to allow different providers to buy and sell spectrum to each other dynamically. We explicitly decouple buyer-side and seller-side designs to capture different properties of the two sides. To accommodate complicated interference patterns on the buyer side, we partition the conflict graph to decouple nodes with no or weak interference while putting nodes with strong interference in the same subgraph. In this way, group bids and critical values can accurately capture the real competition among different buyers. We design a novel merge strategy that combines the auction results from different subgraphs to achieve truthfulness. We further extend our approach to support localized sellers' spectrum resources. Using real cell tower topology traces from a large US wireless provider, we show our approach achieves high efficiency, revenue, and fairness and significantly outperforms existing work.

APPENDIX

Proof of Theorem 4

Proof: To change N by lying, a seller i needs to change the point where the budget-balance is satisfied, which requires the seller to change the total selling price. Let $S(N)$ be the total selling price when N channels are sold and $R(N)$ be the total revenue. Let $U_i(x, N)$ be the utility of seller i when it asks x and N channels are sold. Since the seller-side design is truthful, we have $U_i(v_i, N) \geq U_i(x, N)$, where v_i is i 's true valuation and x can be any value. We consider the following cases.

- 1) i is a winning seller: i cannot change N as long as it still wins. If it changes N and now loses, its utility does not improve.
- 2) i is a losing seller. We consider the following cases.
 - Lie by asking higher. Obviously, N cannot increase in this case since $S(N)$ does not reduce for any N . i is a losing seller in the original case, so it is not within the top N . i cannot become the top N by asking higher.
 - Lie by asking lower. We only consider when i wins because its utility is still 0 if it still loses. We consider the following cases.
 - Case 1: N does not change: i 's utility cannot improve because $U_i(v_i, N) \geq U_i(x, N)$.
 - Case 2: N reduces to M : i is not in the top M when it bids truthfully, so $U_i(v_i, M) = 0$. Thus, i either receive 0 utility or negative utility because $U_i(v_i, M) \geq U_i(x, M)$.
 - Case 3: N increases to M . We consider the following cases:
 - i is in the top M when it bids truthfully, and the $(M + 1)$ th asking price does not change. Thus, $S(M)$ does not change. In the original case, only N channels are sold, so we know $S(M) > R(M)$. Hence, budget-balance is not satisfied, and this case cannot happen. i is not among the top M sellers when it asks truthfully. Hence, $U_i(v_i, M) = 0$. i 's utility cannot increase in this case because $U_i(v_i, M) \geq U_i(x, M)$. ■

Theorem 6: Our buyer-side design is truthful when N is fixed.

Proof: It is easy to see that the design is monotonic, *i.e.*, if a winner wins at v , it still wins if it bids $b > v$, as bidding higher does not reduce the group bid and all other decisions are bid-independent. Next, we verify that the price we charge is the critical value. Specifically, we charge a winner the first losing group bid in its subgraph divided by k , if k members of its group bid higher than that. We find the maximum k , so no more members can be admitted and still afford the fair share. To prove it is the critical value, we should show: 1) a buyer still wins if it bids higher; and 2) a buyer loses if it bids lower. 1) holds because the buyer's group bid (which is greater than k times of the critical value) is still higher than the first losing group and the merge process is bid-independent. To see 2), we consider three cases: i) k remains the same; ii) k increases; and iii) k decreases. For i), if a buyer bids lower than that value, it cannot afford its fair share for the same k even if its group still wins. ii) cannot happen because other members' bids do not change and no new members can afford the fair share. In iii), the fair share is even larger as k decreases, so the buyer does not win either. Thus, a buyer never wins if it bids lower than the price we charge. Therefore, the buyer side with bid-independent merge is truthful when N is fixed according to Theorem 2. ■

Theorem 7: Bid-independent merge-based allocation does not allow a buyer to unilaterally change N and gain.

Proof: The proof has the same structure as the proof to Theorem 4, but from a buyer side. ■

Theorem 8: In our extension to incorporate seller locality, the price we pay to a seller according to our nonuniform seller-side pricing is the critical value of that seller.

Proof: Suppose seller i becomes budget-balanced after dropping T sellers. To prove i is charged with its critical value, we show: 1) i wins if it asks lower; and 2) i loses if it asks higher.

To prove 1): since i is not dropped in the original case and it now asks lower, it cannot be dropped before T sellers are dropped. Thus, we consider two cases: i) If i becomes budget-balanced before T sellers are dropped, it wins. ii) If i does not become budget-balanced before T sellers are dropped, the same $T - 1$ sellers are dropped as in the original case. That is because the order of their asking prices does not change, regions without i do not change, and i 's regions are still not budget-balanced as in the original case. Thus, the same T th seller is dropped, and i becomes budget-balanced in the same region. Hence, i wins.

To prove 2): similarly we consider: i) If i is dropped before T sellers are dropped, i does not win. ii) If i is not dropped before T sellers are dropped, the same $T - 1$ sellers are dropped as in the original case because all regions that have i are still not budget-balanced and other regions are not affected by i 's asking price. Now it selects the highest asking seller that is not budget-balanced to drop. In the original case, the T th dropped seller, denoted as w , was selected. Now i asks higher than w , so i is dropped instead and loses. ■

REFERENCES

- [1] N. Ahmed and S. Keshav, "SMARTA: A self-managing architecture for thin access points," *Proc. CoNEXT*, 2006, Art. no. 9.
- [2] X. Feng, Y. Chen, J. Zhang, Q. Zhang, and B. Li, "TAHES: A truthful double auction mechanism for heterogeneous spectrums," *IEEE Trans. Wireless Commun.*, vol. 11, no. 11, pp. 4038–4047, Nov. 2012.

- [3] M. Gonen, R. Gonen, and E. Pavlov, "Generalized trade reduction mechanisms," in *Proc. 8th ACM EC*, New York, NY, USA, 2007, pp. 20–29.
- [4] "Handover," [Online]. Available: <http://xa.yimg.com/kq/groups/19564168/454394754/name/05-Handover.pdf>
- [5] I. Poole, "GSM handover or handoff," Handover tutorial [Online]. Available: http://www.radio-electronics.com/info/cellular-telecomms/gsm_technical/handover-handoff.php
- [6] S. S. Hong, J. Mehlman, and S. Katti, "Picasso: Flexible RF and spectrum slicing," in *Proc. ACM SIGCOMM*, 2012, pp. 37–48.
- [7] Q. Huang, Y. Tao, and F. Wu, "Spring: A strategy-proof and privacy preserving spectrum auction mechanism," in *Proc. IEEE INFOCOM*, 2013, pp. 827–835.
- [8] G. Iosifidis, L. Gao, J. Huang, and L. Tassiulas, "An iterative double auction for mobile data offloading," in *Proc. WiOpt*, 2013, pp. 154–161.
- [9] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *Proc. ACM MobiCom*, 2003, pp. 66–80.
- [10] A. Lapets and A. Levin, "Restricted truthful combinatorial auction mechanisms," 2007.
- [11] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. IT-28, no. 2, pp. 129–137, Mar. 1982.
- [12] U. V. Luxburg, "A tutorial on spectral clustering," *Statist. Comput.*, vol. 17, no. 4, pp. 395–416, 2007.
- [13] R. P. McAfee, "A dominant strategy double auction," *J. Econ. Theory*, vol. 56, no. 2, pp. 434–450, Apr. 1992.
- [14] M. Meila and J. Shi, "A random walk view of image segmentation," in *Proc. AISTATS*, 2001.
- [15] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *J. Econ. Theory*, vol. 29, no. 2, pp. 265–281, 1983.
- [16] A. Y. Ng, M. I. Jordan, and Y. Weiss, "On spectral clustering: Analysis and an algorithm," in *Proc. NIPS*, 2001, pp. 849–856.
- [17] T. Nieberg, J. Hurink, and W. Kern, "A robust PTAS for maximum weight independent sets in unit disk graphs," in *Proc. Int. Conf. Graph-Theoretic Concepts Comput. Sci.*, 2004, pp. 214–221.
- [18] J. Padhye *et al.*, "Estimation of link interference in static multi-hop wireless networks," in *Proc. ACM IMC*, 2005, p. 28.
- [19] V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki, "Pie in the sky: Online passive interference estimation for enterprise WLANs," in *Proc. NSDI*, 2011, pp. 337–350.
- [20] K. Tan, H. Shen, J. Zhang, and Y. Zhang, "Enable flexible spectrum access with spectrum virtualization," in *Proc. IEEE DYSpan*, Sep. 2012, pp. 47–58.
- [21] H. R. Varian, "Position auctions," *Int. J. Ind. Org.*, vol. 25, no. 6, pp. 1163–1178, Dec. 2007.
- [22] H. R. Varian, "Online ad auctions," *Amer. Econ. Rev.*, vol. 99, no. 2, pp. 430–434, Apr. 2009.
- [23] S. Wang *et al.*, "TODA: Truthful online double auction for spectrum allocation in wireless networks," in *Proc. New Frontiers Dynam. Spectrum*, Apr. 2010, pp. 1–10.
- [24] W. Wang, B. Li, and B. Liang, "District: Embracing local markets in truthful spectrum double auctions," in *Proc. IEEE SECON*, 2011, pp. 521–529.
- [25] F. Wu and N. Vaidya, "Small: A strategy-proof mechanism for radio spectrum allocation," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 81–85.
- [26] L. Xiang, G. Sun, J. Liu, X. Wang, and L. Li, "A discriminatory pricing double auction for spectrum allocation," in *Proc. IEEE WCNC*, Apr. 2012, pp. 1473–1477.
- [27] E. Yao, L. Lu, and W. Jiang, "An efficient truthful double spectrum auction design for dynamic spectrum access," in *Proc. CROWNCOM*, Jun. 2011, pp. 181–185.
- [28] S. Yun, D. Kim, and L. Qiu, "Fine-grained spectrum adaptation in wifi networks," in *Proc. ACM MobiCom*, Sep. 2013, pp. 327–338.
- [29] X. Zhou, S. Gandhi, S. Suri, and H. Zheng, "eBay in the sky: Strategy-proof wireless spectrum auctions," in *Proc. MobiCom*, 2008, pp. 2–13.
- [30] X. Zhou *et al.*, "Practical conflict graphs for dynamic spectrum distribution," in *Proc. ACM SIGMETRICS*, New York, NY, USA, 2013, pp. 5–16.
- [31] X. Zhou and H. Zheng, "TRUST: A general framework for truthful double spectrum auctions," in *Proc. IEEE INFOCOM*, 2009, pp. 999–1007.




Wei Dong received the Ph.D. degree in computer science from The University of Texas at Austin, Austin, TX, USA, in 2013, under advice of Prof. Yin Zhang and Prof. Lili Qiu.

He is a Senior Software Engineer with Google, San Francisco, CA, USA, working on mobile ads quality. His research interests include network economics and pricing, mobile systems and applications, security and privacy, and social networks.



Swati Rallapalli received the Ph.D. degree in computer science from The University of Texas at Austin, Austin, TX, USA, in 2014. Her advisor was Prof. Lili Qiu.

She is a Research Staff Member with the Cloud-based Networking Group, IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. Her research interests are in the areas of mobile computing and user data analytics, including localization, mobility prediction, physical analytics, cellular network economics, and mobile social networks.



Lili Qiu (S'96–M'01–SM'07) received the Ph.D. degree in computer science from Cornell University, Ithaca, NY, USA, in 2001.

She is a Professor with the Department of Computer Sciences, The University of Texas at Austin, Austin, TX, USA. She was a Researcher with Microsoft Research, Redmond, WA, USA, during 2001 to 2004. Her research interests are in Internet and wireless networking, with special focus on wireless network performance and management.

Prof. Qiu is an ACM Distinguished Scientist and a recipient of the NSF Career Award in 2006 and a Google Faculty Research Award.



K. K. Ramakrishnan (S'76–A'83–M'03–SM'04–F'05) received the M.S. degree in automation from the Indian Institute of Science, Bangalore, India, in 1978, and the M.S. and Ph.D. degrees in computer science from the University of Maryland, College Park, MD, USA, in 1981 and 1983, respectively.

He joined AT&T Bell Labs in 1994 and has been with AT&T Labs—Research, Florham Park, NJ, USA, since its inception in 1996. Prior to 1994, he was a Technical Director and Consulting Engineer in Networking with Digital Equipment Corporation, Littleton, MA, USA. Between 2000 and 2002, he was with TeraOptic Networks, Inc., Sunnyvale, CA, USA, as Founder and Vice President. At AT&T Labs—Research, he was involved in several technical and strategic activities in networking and information distribution. He has published nearly 200 papers and has over 100 patents issued in his name.

Dr. Ramakrishnan is an AT&T Fellow, recognized for his contributions to communications networks, including congestion control, traffic management, and VPN services. He has been on the Editorial Board of the *IEEE/ACM TRANSACTIONS ON NETWORKING* and *IEEE Network* and has been a member of the National Research Council Panel on Information Technology for NIST.



Yin Zhang (M'02–SM'07–F'15) received the Ph.D. degree in computer science from Cornell University, Ithaca, NY, USA, in 2001.

He is an Adjunct Professor with the Department of Computer Sciences, The University of Texas at Austin, Austin, TX, USA. Between 2001 and 2004, he was a member of the Networking Research Department, AT&T Labs—Research. His research interests span several areas of computer networks, including network measurement, network management, and network security.

Prof. Zhang received the NSF CAREER Award in 2006.