Enabling Spectrum Sharing in Secondary Market Auctions

Citation

Published Version

Permanent link

Terms of Use

Share Your Story

Enabling Spectrum Sharing in Secondary Market Auctions

Ian A. Kash
School of Engineering and
Applied Sciences
Harvard University
Cambridge, MA 02138, USA
kash@seas.harvard.edu

Rohan Murty
School of Engineering and
Applied Sciences
Harvard University
Cambridge, MA 02138, USA
rohan@eecs.harvard.edu

David C. Parkes
School of Engineering and
Applied Sciences
Harvard University
Cambridge, MA 02138, USA
parkes@eecs.harvard.edu

ABSTRACT

Wireless spectrum is a scare resource, but in practice much of it is under-used by current owners. To enable better use of this spectrum, we propose an auction approach to dynamically allocate the spectrum in a secondary market. Unlike previous auction approaches, we seek to take advantage of the ability to share spectrum among some bidders while respecting the needs of others for exclusive use. Thus, unlike unlicensed spectrum (e.g. Wi-Fi), which can be shared by any device, and exclusive-use licensed spectrum, where sharing is precluded, we enable efficient allocation by supporting sharing alongside quality-of-service protections. We present SATYA (Sanskrit for "truth"), a strategyproof and scalable spectrum auction algorithm whose primary contribution is in the allocation of a right to contend for spectrum to *both* sharers and exclusive-use bidders. Achieving strategyproofness in our setting requires appropriate handling of the externalities created by sharing. We demonstrate SATYA's ability to handle heterogeneous agent types involving different transmit powers and spectrum needs through extensive simulations.

1. INTRODUCTION

Spectrum is a limited and expensive resource. For example, the 2006 Federal Communications Commission (FCC) auctions for 700 - 800 MHz raised almost $19 billion dollars. Hence, the barrier to entry for potential spectrum buyers is high. One can either buy a lease on spectrum covering a large area at a high price or use the limited spectral bands classified as unlicensed (e.g. Wi-Fi). Such unlicensed bands are subject to a "tragedy of the commons" where, since they are free to use, they are over-used and performance suffers [8]. Efforts such as the recent FCC ruling on white spaces are attempting to free additional spectrum by permitting opportunistic access [4]. However, such efforts are being met with opposition by incumbents (such as TV broadcasters and wireless microphones manufacturers) who have no incentive to permit their spectrum to be shared.

Motivated by these observations, many researchers and companies (*e.g.,* [7, 34, 19]) have proposed allowing spectrum owners and spectrum users to participate in a secondary market for spectrum where users are allocated the use of spectrum in a small area on a dynamic basis based on their short- or medium-term needs. This approach is beneficial for two reasons. First, it allows flexible approaches to determining how best to allocate spectrum rather than relying on the decision making of regulators like the FCC in the United States. Second, it provides an incentive for spectrum that is currently owned but unused or under-used (such as the television spectrum) to be made available by its owners. Note, by secondary market we mean, one in which the owner of a chunk of spectrum leases different frequencies to *other* users who bid for the spectrum. The FCC has also recognized the potential use of a secondary spectrum market and has begun encouraging spectrum owners in certain bands to sublease the spectrum [18].

Prior work has proposed a number of auction designs to support such a market. However, the possibility of **sharing** in such markets has not been sufficiently explored. Most auctions provide exclusive access: the allocation is such that no winners interfere. However, this may not be the most efficient use of spectrum. For example, devices like wireless microphones are only used occasionally, so even if they require exclusive access while in use some other device may be able to use the same spectrum on a secondary basis when they are not. This *heterogeneity* of devices and demands is a source of opportunities for sharing. Further, many devices are capable of using a medium access controller (MAC) to share bandwidth when given the right to contend.

Rather than full sharing, as in the Wi-Fi model, using an auction has two key advantages. First, it provides revenue and thus incentives for primary spectrum owners to open up spectrum to other uses. Second, it provides incentives for different potential users to describe (through bids) their distinct needs for spectrum access, be it exclusive or with sharing. With Wi-Fi, if too many people try to use the same access point, service degrades and may become unacceptable for all of them, and no one has an incentive to consider the (negative) externality their use imposes on others.

Current proposals for secondary-market spectrum auctions are either unable to support the externalities created by two interfering devices sharing the same channel, or do not scale to realistic problem sizes and interference graph topologies. Solutions that ignore the possibility of sharing rely on bid-

ders caring only about whether or not they are allocated a channel. With sharing, bidders also care about *with whom* they share the channel.

We present SATYA a scalable, strategyproof auction algorithm that permits different classes of spectrum users (sharing and exclusive) to co-exist and share the spectrum, while appropriately accounting for the resulting externalities. Allowing bidders to report arbitrary externalities would yield an intractable allocation problem. But the externalities in our setting have significant structure, and SATYA uses a simple, yet expressive, language to allow bidders to express their value for different allocations given probabilistic activation patterns, interference, and different requirements for shared vs exclusive-access spectrum. Using this language, we can quantify the utility of a bidder for an allocation in terms of the fraction of the bidder's demand that is satisfied. Bidders only interfere with other nearby bidders and, given a model for resolving contention by devices allocated shared spectrum and both simultaneously active, we can quantify the resulting externality.

Even without sharing, finding an optimal channel assignment involves solving a graph coloring problem and is NP-hard [20], and we are unable to find an optimal allocation fast enough to be reasonable for network deployment. Despite focusing on a single-parameter mechanism design problem by assuming that the components of a bidder's type that reflect its interference, usage patterns, and penalty for being blocked from accessing allocated (shared) spectrum by an exclusive-use device on the network are known, a key difficulty remains: *unlike in settings without externalities a straightforward greedy allocation approach that still allows sharing fails to be monotonic.*

An allocation algorithm in our domain is monotonic if, submitting a larger bid for access to some fraction of a channel when active, always leads to an allocation in which the bidder receives a (weakly) larger fractional share of channel capacity, in expectation with respect to the activation patterns of other devices, and whatever the bids of other devices. Monotonicity is well known to be necessary and sufficient for strategyproofness in single-parameter domains [29]. Strategyproofness, a property that guarantees that it is optimal for each bidder to report his true value regardless of the actions of other bidders, is desirable for two main reasons. First, it provides strategic simplicity for bidders: they do not need to perform any sophisticated reasoning about the actions of others in order to determine how to participate. Second, it greatly simplifies evaluating a potential auction algorithm. If the algorithm is strategyproof, we can simply assume bidders report their true valuations. Otherwise, we would need to analyze the auction to determine the structure of equilibrium bids; naively assuming that bidders would report truthfully can overestimate both the efficiency of the allocation that results and the revenue that is raised.

In providing monotonicity, SATYA therefore employs a novel combination of bucketing bids into intervals wherein they are treated equally within an interval (this idea was employed in Ghosh and Mahdian [15]), and a computational ironing procedure that is used to validate the monotonicity of an allocation and perturb the outcome as necessary to ensure monotonicity (this idea was introduced by Parkes and Duong [30]). These techniques prevent cases where, if a bidder raises his bid, the greedy algorithm selects a different allocation that, at the time, is as good or better, but ends up being worse.

To evaluate SATYA we use real world data sources to determine participants in the auction, along with the sophisticated Longley-Rice propagation model [3] and high resolution terrain information to generate conflict graphs. We compare the performance of SATYA against other auction algorithms and baseline computations. Our results show that, when spectrum is scarce, allowing sharing using SATYA increases social welfare by 40% over previous approaches.

In summary, this paper makes the following contributions:

- The first strategyproof, scalable auction design for dynamic spectrum access that allows *sharing* and exclusive access by appropriately dealing with the externalities this creates.
- An approach that accommodates different classes of wireless users, each with a different transmit power, spectrum access, and activation patterns.
- The use of sophisticated propagation models and real world data to demonstrate in simulation the efficacy of SATYA.

1.1 Related Work

There has been significant work on spectrum auctions where a regulatory agency, such as the FCC in the United States, sells rights to spectrum across large areas (see, e.g. [10, 11]). However, we focus on secondary-market auctions, where the existing owner of spectrum wishes to resell it to a large number of smaller users subject to interference constraints.

Most secondary-market spectrum auction algorithms do not allow auction participants that interfere to share a channel [6, 13, 17, 32, 34]. Among these designs, VERITAS [34] was the first spectrum auction algorithm based on a monotone allocation rule, and thus strategyproof. However, VERITAS does not support sharing. Zhou *et al.* [33] proposed TRUST, which uses a double auction for cases when multiple owners are selling channels. In terms of supporting technologies, the use of a spectrum database in facilitating secondary market auctions has been proposed [19].

Turning to sharing, Jia *et al.* [24] envision spectrum owners auctioning off spectrum rights to a secondary user when it is not otherwise being used by the owner, and investigate how revenue can be maximized in this setting. While winners share with the spectrum owner in this way, there is no sharing among participants and no interference is tolerated.

Gandhi *et al.* [14] use an approach that allocates many small channels, and effectively enables sharing. However,

their algorithm allows sharing only among bidders who want only a portion of a channel. Thus, it cannot take advantage of bidders who are not always active (e.g. wireless microphones). Their approach is not strategyproof, and they perform no equilibrium analysis. Kasbekar and Sarkar [25] use a strategyproof auction and allow bidders to express *arbitrary* externalities, but this lack of structure makes their approach intractable except in a very simple case.

The issue of externalities in auctions has been considered more generally. Jehiel *et al.* [22, 23] consider situations, such as the sale of nuclear weapons, where bidders care not just about winning but about who else wins. However, the settings studied are without the computational challenges or need for expressiveness of our domain. A number of papers have considered externalities in online advertising [9, 15, 16, 27, 31]. However, this work (and similarly that of Krysta *et al.* [28] on the problem of externalities in general combinatorial auctions) is not directly relevant, as the externalities in spectrum auctions have a spectral structure, of which SATYA takes advantage.

2. AGENT MODEL

In order to find opportunities to share among heterogeneous agents (e.g., a user with a wireless device, or a TV station), we need a language to describe the requirements of each *type* of agent.

Our model uses discrete intervals of time (called *epochs*), with auctions clearing periodically and granting the right to agents to contend for access to particular channels over multiple epochs. The ultimate allocation of spectrum arises through random activation patterns of agents and interference effects, and depends on specifics of the medium-access control (MAC) contention protocol. The effect of this MAC protocol is also modeled within SATYA.

The interference between agents and their associated devices is modeled through a *conflict graph*, $G = (V, E)$, such that each agent i is associated with a vertex ($i \in V$) and an edge, $e = (i, j) \in E$, exists whenever agents i, j would interfere with each other if they are both active in the same epoch and on the same channel.

Recall that we allow for exclusive-use and "willing to share" types of agents, where the former must receive access to a channel without contention from interfering devices whenever they are active, while the latter can still obtain value through contending for a fraction of the channel with other interfering devices. We say that a channel is *free*, from the perspective of agent i in a particular epoch, if no exclusive-use agent j, who interferes with i and is assigned the right to the same channel as i, is active in the epoch.

Formally, we denote the set of agent types T. Each type $t_i \in T$ is a tuple $T_i = (x_i, a_i, d_i, p_i, C_i, v_i)$, where

- $x_i \in \{0, 1\}$ denotes whether the agent requires *exclusive use* of a channel in order to make use of it ($x_i = 1$) or willing to share with another agent while both are active on the channel ($x_i = 0$).

- $a_i \in (0, 1]$ denotes the *activation probability* of the agent: the probability that the agent will want to use the channel, and be active, in an epoch. Our model assumes that the activations of agents in each epoch are independent of other epochs and other agents.

- $d_i \in (0, 1]$ is the *fractional demand* of the channel that an agent who is willing to share access requires in order to achieve full value when active. If $x_i = 1$ the agent cannot share, so $d_i = 1$.

- $p_i \in \mathbb{R}^+$ denotes the per-epoch penalty incurred by the agent when it is active, but the assigned channel is not free. Both exclusive use and non exclusive use agents can have a penalty.

- $C_i \subseteq \mathcal{C} = \{1, 2, \ldots\}$, where $\mathcal{C}$ is the set of channels to allocate, each corresponding to a particular spectrum frequency, denotes the channels that agent i is able to use (the agent is indifferent across any such channel.) An agent receives an allocation of $\perp$ when not assigned a channel.

- $v_i \geq 0$ denotes the per-epoch value received by the agent in an epoch in which it is active, the channel is free, and in the case of non exclusive-use types, the agent receives at least a share d_i of the available spectrum.

In this model, each agent demands a single channel. An extension to multiple channels is discussed in our technical report [26]. As the following examples show, this type space provides a rich language for agents to describe their intended usage pattern. In order to fully describe the auction we also require a model of what fraction of an agent's potential value is realized when it shares a channel with other agents. We defer a full description of this model to our technical report.

Examples

An agent who wishes to run a low-power (local) TV station on a channel would be unable to share it with others when active ($x_i = 1$), would be constantly broadcasting ($a_i = 1$), and would have a very large penalty p_i since it is unacceptable for the broadcast to be interrupted by someone turning on another (exclusive use) device. Another agent might want to use a device like a wireless microphone that also cannot share a channel when active ($x_i = 1$), but might be used only occasionally ($a_i = 0.05$) and might have a smaller value of p_i since it may be acceptable if it is occasionally unable to be used because there is another exclusive agent also trying to use the channel. For example, it might make sense to have several such devices share a channel if they interfere with each other sufficiently rarely.

There are also classes of agents capable of using a MAC and thus sharing a channel ($x_i = 0$). For example, someone who wants to run a wireless network could have constant traffic ($a_i = 1$) that consumes a large portion of the channel ($d_i = 0.9$), and might have a large penalty similar to a TV station because completely disconnecting users is unacceptable. However, such an agent is willing to share the channel

with other non-exclusive types, and pay proportionately less for a smaller fraction of the bandwidth. There might also be opportunistic data users, for example a delay tolerant network [21], who occasionally ($a_i = 0.2$) would like to send a small amount information ($d_i = 0.4$) if the channel is available. Such users might have a low or even no penalty as their use is opportunistic.

3. DESIGN OF SATYA

Turning to the design of SATYA, we assume that the only component of an agent's type that can be misreported is v_i, with some bid $b_i \neq v_i$ possible. This makes our auction an example of what Blum and Hartline [5] termed an *attribute auction*, where, in addition to the bid, the auctioneer knows some additional characteristics about each bidder. This is a reasonable assumption to make in practice. Most other characteristics, such as how often the agent makes use of the channel, how much of the channel he uses, whether his devices can use a MAC, and on what channels they can legally broadcast, can be observed by the auctioneer. Assuming he knows an agent's penalty is a somewhat stronger assumption, but we expect that the auctioneer will have at least a broad idea of how well different applications tolerate preemption. In addition, we assume that the auctioneer knows the structure of the conflict graph.

Even if no agents are permitted to share channels, finding the efficient allocation is NP-Hard [20], as assigning bidders to channels such that no two neighbors have the same channel is a graph coloring problem. Therefore we adopt the same approach as previous strategyproof algorithms for channel allocation, and seek to assign agents to channels greedily. But in our setting, the effect of externalities and sharing is that a straightforward greedy algorithm will fail to satisfy the key property of *monotonicity*, which is necessary and sufficient in providing for strategyproofness.

3.1 Externalities and Monotonicity

Let us first define the property of monotonicity in our setting. Given a joint bid vector $b = (b_1, \ldots, b_n)$ received from agents (with $b_j \geq 0$ for all j) and a joint type vector $t = (t_1, \ldots, t_n)$, an allocation algorithm $\mathcal{A}$ produces an allocation $\mathcal{A}(b, t)$. Each agent i has some utility $U_i(\mathcal{A}(b, t))$ for this allocation.

Fixing the bids $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ of other agents, an allocation algorithm is *monotone* if,

$$U_i(\mathcal{A}((b_i', b_{-i}), t)) \geq U_i(\mathcal{A}(b, t)),$$

for all bids $b_i' \geq b_i$. This insists that the expected share of a channel available to an agent, and thus its expected utility, (weakly) increases as the agent's bid increases.

With no sharing (and no externalities), a greedy allocation algorithm is monotone. However, Figure 1 shows how monotonicity can fail for simple greedy algorithms in the presence of sharing and externalities. The greedy algorithm considers each agent in (decreasing) order of bids and al-
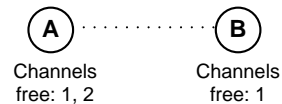

Figure 1: A potential violation of monotonicity. Nodes A and B are in contention range. At node A's location channels 1 and 2 are free; at B only channel 1 is free.

locates him to the best available channel in terms of social welfare (or no channel if that is better). If there is a tie, it uses some tie-breaking rule, such as the lowest channel number. If agent A has a lower bid than agent B, the algorithm assigns agent B to channel 1, then agent A to channel 2, and both are fully satisfied. If agent A raises his bid so that it is higher than agent B's bid, then the algorithm greedily assigns him to channel 1. Then, assuming sharing is better than not assigning B a channel, it has no other option than to assign agent B to channel 1, so the agents share the channel and are less well off.

3.2 The SATYA Algorithm

To prevent violations of monotonicity from happening while still using a greedy allocation rule, SATYA brings to bear a novel combination of (a) forbidding some shared allocations during the process (using a *bucketing* approach), and (b) canceling some shared allocations in a post-processing step (using an *ironing* approach). SATYA treats all agents in a bucket as if they bid the same amount, so changes of bid that do not change the bucket to which the agent is assigned have no effect on the allocation and thus do not violate monotonicity. Furthermore, agents in different bucket are allowed share only in a limited way, which prevents the greedy assignment from introducing externalities, and thus monotonicity violations, in such cases.

SATYA begins by assigning each agent i to a bucket K_i based on his bid b_i. There many ways this can be done as long as it is monotone in the agent's bid. For example, agent i with a bid b_i in the range $[2^\ell, 2^{\ell+1})$ could be assigned to bucket $K_i = \ell$. In general, we assume that this is done according to some function $\beta(k)$, such that bucket k contains all agents with bids b_i in the range $[\beta(k), \beta(k+1))$.

The agents are assigned channels greedily, in descending order of buckets, with the order of assignment for agents within the same bucket determined randomly. A channel c is considered available to allocate to agent i at some step in the algorithm, and given the intermediate allocation A, if,

- the channel c is in C_i;
- assigning i would not cause an externality to a neighbor from a higher bucket (combined demands of each neighbor, that neighbor's neighbors, and i's are less than 1);
- and the combined demands of i and his neighbors from higher buckets assigned to c are less than 1.

Note that this does not preclude allocations where some agent does not have his demand fully satisfied. It simply

requires that, in such cases, the agent is sharing with others in his own bucket.

For the next agent to be allocated, i, SATYA finds the channel that will have the maximum marginal effect on the total value of all currently allocated agents and agent i itself. To do so, for every channel c that is available to the agent, and including $\perp$ and thus not allocating any spectrum to the agent, SATYA estimates the expected value to each agent j after assigning i to c. This estimate differs from the agent's actual bid by assuming that each agent in a given bucket shares the same value. This is important for achieving monotonicity as we need to ensure the way an agent is treated depends only on his bucket.

Finally, agent i is greedily assigned to the channel that maximizes the sum of the expected bid values of each agent (already allocated, and itself) without leaving any agent with a negative utility. The decision could be to allocate $\perp$ and thus no spectrum to the agent. In the event of a tie, the agent is assigned to the lowest numbered among the tied channels (including preferring $\perp$, all else equal).

After all agents in a bucket are assigned channels, there is an *ironing* step in which monotonicity of the allocation is verified, and the allocation perturbed if this fails. First, the allocation procedure is re-run for each agent, to determine what would happen had he been in a lower bucket. These counterfactuals are used to determine if the agent might have been able to be allocated a channel in a lower bucket. If so, this might cause a monotonicity violation where an agent bids more but ends up less well off, and so the provisional allocation is modified by changing the assignments of the neighbors with whom he shared a channel to $\perp$. This is the ironing step: removing failure of monotonicity. In Figure 1, this would mean that rather than A and B sharing a channel, ironing would assign B to $\perp$ to ensure that monotonicity is not violated for A.

A full description of the allocation and pricing algorithms and a proof that this approach ensures monotonicity are available in the our technical report [26].

4. EVALUATION

In this section we compare the performance of SATYA to VERITAS. Since VERITAS does not permit sharing, we modify it slightly and implement VERITAS-S, which permits sharing as long as there are no externalities imposed (i.e. sharing is permitted only when the combined demands of sharing agents do not exceed the capacity of the channel). We also implement GREEDY, a version of SATYA without bucketing and ironing that provides higher overall efficiency. *GREEDY is neither strategyproof nor monotone.* Thus, agents' bids need not match their true values. However, to set as high a bar as possible, we assume they do so. Since it gets to act on the same information but has fewer constraints than SATYA, GREEDY serves as an upper bound for our experiments.

Parameters: As shown in Table 1, all our experiments use

| Agent Type	Act. Prob.	Value	Penalty	Demand
Exclusive-Continuous	1	[0, 1000]	10000	1
Exclusive-Periodic	[0.05, 0.15]	[0, 1000]	5000	1
Sharing-High	1	[0, 1000]	10000	[0.3, 1]
Sharing-Low	[0, 1]	[0, 1000]	5000	[0.3, 1]

Table 1: Mix of agents used in the evaluation

four classes of agents bidding for spectrum. Note that, in the table, we we have normalized the values so the table reflects the range of $a_i v_i$ rather than the range of v_i. Each class represents different applications. For example, a TV station serving a local community is an agent who wants exclusive access for a long period of time. A wireless microphone is an example of an agent who wants exclusive access but for short periods of time. A low-cost rural ISP is an example of a *Sharing-High* agent who expects to actively use the spectrum but can potentially tolerate sharing, and a regular home user is an example of a *Sharing-Low* class agent whose spectrum access pattern varies. Note, each class of agents may have different transmit powers and coverage areas than the others. Since our goal is to evaluate the efficacy of SATYA in exploiting opportunities for sharing, we assign 5% of the total agents as exclusive-continuous, 15% exclusive-shared, 30% Sharing-High, and the remaining 50% Sharing-Low.

Methodology: Each auction algorithm takes as input a conflict graph for the agents. To generate this conflict graph in a realistic manner, we implement and use the popular Longley-Rice [1] propagation model in conjunction with high resolution terrain information from NASA [2]. This sophisticated model estimates signal propagation between any two points on the earth's surface factoring in terrain information, curvature of the earth, and climactic conditions. We use this model to predict the signal attenuation between agents, and consequently the conflict graph between the bidding agents.

We use the FCC's publicly available CDBS [12] database to model the transmit power, location, and coverage area of Exclusive-Continuous users. Note, this information as well as the signal propagation predictions are sensitive to geographic areas. We model the presence of all other types of agents using population density information. Agents are scattered across a 25 mile x 25 mile urban area in a random fashion that factors in population density. Since each class of agent has a different coverage area, we determine that a pair of nodes conflicts if the propagation model predicts signal reception higher than a specified threshold.

We repeat each run of the experiment 10 times and present averaged numbers across runs. In these experiments, the number of channels is 5. In tuning SATYA, we experimented with a variety of methods for determining to which bucket to assign an agent. We omit these results for space reasons, but based on them use buckets of size 500 ($\beta(k) = 500k$).

In our experiments, we evaluate our approach using two metrics. The first, social welfare, is the sum of the valuations of winning agents (includes externalities). This measures how happy the participants are with the resulting allocation.

5
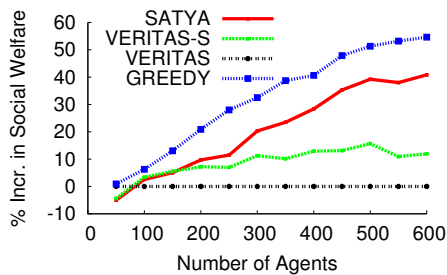
(a) Social Welfare
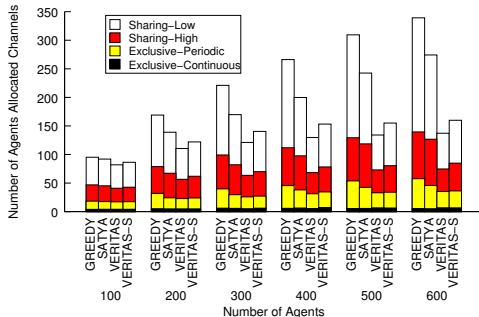
(b) Distribution of Allocated Agents

Figure 2: Quality of Allocation

Results for other measures of the efficiency are similar and are reported in our technical report [26]. The second is revenue, the sum of agents' payments. This measures the incentive spectrum owners have to make their spectrum available to the auctioneer.

4.1 Varying the Number of Agents

Figure 2 shows the performance of various algorithms as a function of the number of agents participating in the auction. As seen in Figure 2(a), the social welfare (relative to the baseline of VERITAS) attained by SATYA increases with an increase in the number of agents. This is a direct consequence of assigning channels to more agents capable of sharing the spectrum. This shows that, despite externalities from sharing, the additional agents allocated consider it valuable. At 600 bidders, SATYA realizes a gain of 25% over VERITAS-S and 40% over VERITAS in the total social welfare of the network. With fewer agents, all three algorithms demonstrate similar performance because almost all agents can either be allocated a channel of their own or are impossible to satisfy.

Overall, VERITAS-S and VERITAS do not make the best use of agents that can share. This is demonstrated in Figure 2(b), which is the distribution of different classes of agents assigned channels by each algorithm. As the number of agents increases, VERITAS-S and VERITAS significantly reduce the fraction of agents capable of sharing who are assigned channels (relative to SATYA). However, all algorithms demonstrate a similar performance in the fraction of exclusive bidders who are assigned channels. Hence, SATYA is capable of taking advantage of sharing by allocating channels to more of such users. As expected, GREEDY outper-
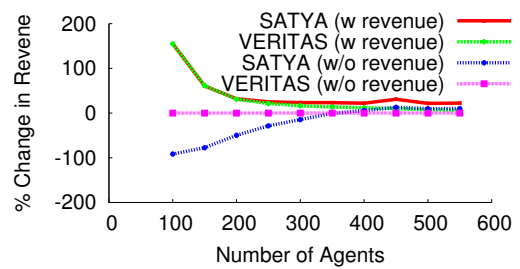

Figure 3: Revenue as a function of number of agents

forms all strategyproof auctions and is able to assign more sharing agents. The difference in performance between SATYA and GREEDY is primarily due to bucketing. Ironing does occur but has only a minor effect. Hence, the main takeaway is that SATYA *increases the number of allocated agents and thus social welfare.*

4.2 Measuring Revenue

We consider social welfare the most important measure of performance: a market that finds success in the long run will allocate resources to those that find the most value. However, in our setting revenue may also be important to provide an incentive for current spectrum owners to participate in the secondary market. First, we measure the total revenue obtained as a function of the number of agents bidding for spectrum without reserve prices. We do not include GREEDY in this analysis because it is not strategyproof and it is not clear what agents will bid and thus what the actual revenue would be. As seen in Figure 3, the revenue obtained by SATYA and is much lower than VERITAS for smaller numbers of agents. We omit VERITAS-S from the figure for readability, but its performance also suffers. This is a consequence of sharing making it easier to accommodate agents.

To improve revenue, we institute reserve prices, minimum bids agents must make to participate in the auction. VERITAS explored a similar opportunity to increase revenue by limiting the number of channels sold. Using a reserve price of 400, we experiment to measure revenue by varying the number of bidders. As Figure 3 shows, this increases revenue for the auctioneer significantly for all algorithms. The increase is most pronounced with 50 agents (not shown because the improvement is so large) where revenue goes from essentially zero to approximately ten thousand. SATYA, which without a reserve price lost revenue by being too efficient in allocating agents, benefits slightly more than VERITAS. With a large number of agents, the reserve price is essentially irrelevant because of the amount of competition; with 550 agents the gain is below 12%. Hence, the the main takeaway is that *with appropriate reserve prices, the increase in efficiency enabled by SATYA does not have to come at the the cost of revenue*. Additional results about reserve prices are available our technical report.

REFERENCES

[1] The ITS Irregular Terrain Model Algorithm, NTIA, Department of Commerce.

[2] Shuttle Radar Topograph Mission (SRTM), http://www2.jpl.nasa.gov/srtm/.

[3] OET BULLETIN No. 69, Longley-Rice Methodology for Evaluating TV Coverage and Interference. February 2004. URL http://www.ieee.or.com/Archive/uwb.pdf.

[4] FCC press release, Evaluation of the Performance of Prototype TV-Band White Space Devices. November 2008. URL http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-08-2243A3.pdf.

[5] A. Blum and J. D. Hartline. Near-optimal online auctions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1156–1163, 2005.

[6] M. M. Buddhikot and K. Ryan. Spectrum management in coordinated dynamic spectrum access networks. In *Proc. of IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks(DySPAN)*, 2005.

[7] M. M. Buddhikot, P. Kolodzy, S. Miller, K. Ryan, and J. Evans. DIMSUMNet: new directions in wireless networking using coordinated dynamic spectrum access. In *2005 International Conference on a World of Wireless, Mobile and Multimedia Networks (WOWMOM)*, pages 78–85, 2005.

[8] M. M. Bykowsky, M. A. Olson, and W. W. Sharkey. Modeling the efficiency of spectrum designated to licensed service and unlicensed operations. Technical report, FCC OSP Working Paper Series, 2008.

[9] F. Constantin, M. Rao, C.-C. Huang, and D. C. Parkes. On Expressing Value Externalities in Position Auctions. In *Proceedings of the 6th Ad Auctions Workshop*, 2010.

[10] P. Cramton. In M. Cave, S. Majumdar, and I. Vogelsang, editors, *Handbook of Telecommunications Economics*, chapter Spectrum Auctions, pages 605–639. Elsevier Science B.V., 2002.

[11] P. Cramton, Y. Shoham, and R. Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006. ISBN 0262033429.

[12] FCC Media Bureau. TVQ TV Database, http://www.fcc.gov/mb/video/tvq.html.

[13] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri. A general framework for wireless spectrum auctions. In *Proc. of IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks(DySPAN)*, 2007.

[14] S. Gandhi, C. Buragohain, L. Cao, H. Zheng, and S. Suri. Towards real-time dynamic spectrum auctions. *Computer Networks*, 52(4):879–897, 2008.

[15] A. Ghosh and M. Mahdian. Externalities in online advertising. In *17th International World Wide Web Conference (WWW)*, 2008.

[16] R. Gomes, N. Immorlica, and E. Markakis. Externalities in keyword auctions: An empirical and theoretical assessment. In *Internet and Network Economics, 5th International Workshop (WINE)*, pages 172–183, 2009.

[17] J. Huang, R. A. Berry, and M. L. Honig. Auction mechanisms for distributed spectrum sharing. In *Proc. of 42nd Allerton Conference*, 2004.

[18] S. B. Inc. Analysis of fcc secondary spectrum markets positions, policies and comments. .

[19] S. B. Inc. The secondary spectrum market: A licensing & leasing primer. .

[20] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu. Impact of interference on multi-hop wireless network performance. *Wireless Networks*, 11(4):471–487, 2005.

[21] S. Jain, K. Fill, and R. Patra. Routing in a Delay Tolerant Network. In *SIGCOMM*, 2004.

[22] P. Jehiel, B. Moldovanu, and E. Stacchetti. How (not) to sell nuclear weapons. *The American Economic Review*, 84(6):814–829, Sep. 1996.

[23] P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory*, 85:258–293, 1999.

[24] J. Jia, Q. Zhang, Q. Zhang, and M. Liu. Revenue generation for truthful spectrum auction in dynamic spectrum access. In *Proceedings of the 10th ACM Interational Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 3–12, 2009.

[25] G. S. Kasbekar and S. Sarkar. Spectrum auction framework for access allocation in cognitive radio networks. In *Proceedings of the 10th ACM Interational Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 13–22, 2009.

[26] I. Kash, R. Murty, and D. Parkes. Enabling spectrum sharing in secondary market auctions. Technical Report TR-07-10, Harvard EECS, 2010.

[27] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *Internet and Network Economics, 4th International Workshop (WINE)*, pages 585–596, 2008.

[28] P. Krysta, T. Michalak, T. Sandholm, and M. Wooldridge. Combinatorial auctions with externalities. In *Ninth International Conference on Autonomous Agents And Multiagent Systems (AAMAS)*, 2010.

[29] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.

[30] D. C. Parkes and Q. Duong. An ironing-based approach to adaptive online mechanism design in single-valued

domains. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*, pages 94–101, 2007.

[31] D. Reiley, S.-M. Li, and R. Lewis. Northern exposure: A field experiment measuring externalities between search advertisements. In *Proc. of 11th ACM Conference on Electronic Commerce (EC)*, 2010.

[32] A. P. Subramanian, M. Al-Ayyoub, H. Gupta, S. R. Das, and M. M. Buddhikot. Near optimal dynamic spectrum allocation in cellular networks. In *Proc. of 3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks(DySPAN)*, 2008.

[33] X. Zhou and H. Zheng. TRUST: A general framework for truthful double spectrum auctions. In *INFOCOM 2009. 28th IEEE International Conference on Computer Communications*, pages 999–1007, 2009.

[34] X. Zhou, S. Gandhi, S. Suri, and H. Zheng. ebay in the sky: strategy-proof wireless spectrum auctions. In *Proceedings of the 14th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 2–13, 2008.