On Spectrum Sharing Games

Magnús M. Halldórsson
Dept. of Computer Science
University of Iceland
IS-107 Reykjavik
mmh@hi.is

Joseph Y. Halpern[†]
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
halpern@cs.cornell.edu

Li (Erran) Li[*]
Center for Networking Research
Lucent Bell Labs
Holmdel, NJ 07733
erranlli@dnrc.bell-labs.com

Vahab S. Mirrokni
Computer Science and Artificial Intelligence
Laboratories (CSAIL)
MIT
Cambridge, MA 02139
mirrokni@theory.lcs.mit.edu

ABSTRACT

Each access point (AP) in a WiFi network must be assigned a channel for it to service users. There are only finitely many possible channels that can be assigned. Moreover, neighboring access points must use different channels so as to avoid interference. Currently these channels are assigned by administrators who carefully consider channel conflicts and network loads. Channel conflicts among APs operated by different entities are currently resolved in an ad hoc manner or not resolved at all. We view the channel assignment problem as a game, where the players are the service providers and APs are acquired sequentially. We consider the price of anarchy of this game, which is the ratio between the total coverage of the APs in the worst Nash equilibrium of the game and what the total coverage of the APs would be if the channel assignment were done by a central authority. We provide bounds on the price of anarchy depending on assumptions on the underlying network and the type of bargaining allowed between service providers. The key tool in the analysis is the identification of the Nash equilibria with the solutions to a maximal coloring problem in an appropriate graph. We relate the price of anarchy of these games to the approximation factor of local optimization algorithms for the maximum k-colorable subgraph problem. We also study the speed of convergence in these games.

[*]Research partially supported by NSF ANI-0335244.

[†]Work supported in part by NSF under grant CTC-0208535, by ONR under grants N00014-00-1-03-41 and N00014-01-10-511, by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grant N00014-01-1-0795, and by AFOSR under grant F49620-02-1-0101.

PODC'04, July 25–28, 2004, St. Johns, Newfoundland, Canada.

Categories and Subject Descriptors

F.m [**Theory of Computation, Miscellaneous**]

General Terms

Economics, Theory

Keywords

game theory, Nash equilibrium, price of anarchy, graph coloring, approximation algorithm, unit disk graph

1. INTRODUCTION

802.11 Wireless LANs, commonly known as WiFi, provide wireless network access to subscribers. They have been deployed in public hotspots, ranging from airports to hotels to Starbucks coffee shops. Fueled by the growing usage, service providers have been planning to provide wireless network access that covers larger areas. For example, Verizon has deployed WiFi for hundreds of hotspots in New York City [21] and MeshNetworks Corporation has been deploying city-wide WiFi networks to facilitate law enforcement and emergency response in cities such as Medford, Oregon [16] .

To understand the issues of interest to us, we need to briefly review some relevant details of 802.11 wireless LANs. (See [3] for further details.) An 802.11 network consists of a set of access points (APs). Each AP must be configured with a fixed transmission power. There is a constant number of possible transmission powers to choose from. Users can then access the Internet by communicating with their provider's APs using the 802.11 air interface. Each AP must be assigned a channel (i.e., a frequency that it transmits on) for it to service users. There are a small number of non-interfering channels; for example, 802.11b and 802.11g each have 3 such channels and 802.11a has 11. A user within an AP's coverage area then uses this channel to communicate with the AP. Channel access between users of the same AP is arbitrated by a media access control protocol (MAC). For example, in the DCF model of the 802.11 MAC protocol, if a user determines that the media is free, it sends a request to send (RTS) message; the AP replies with a clear to send (CTS) message; users that receive this message will defer media access on that channel for long enough to

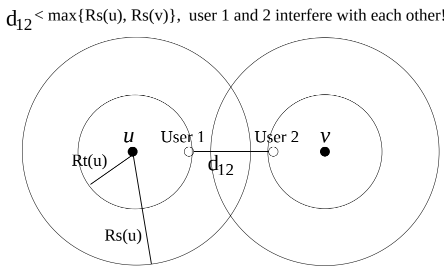

Figure 1: Potential interference between two APs.

guarantee that there is no interference with the user's message. To avoid interference in the wireless media between nearby APs and their users, nearby APs must use different channels.

In more detail, we can associate with each AP u two circular regions around it (see Figure 1). The smaller circle, denoted $R_t(u)$, represents u's *transmission range*. All messages sent by u can be correctly received by users in $R_t(u)$. The larger circle, denoted $R_s(u)$, represents u's *sensing range*. In practice, the radius of $R_s(u)$ is about twice that of $R_t(u)$. The actual size of $R_s(u)$ and $R_t(u)$ depends on the transmission power used by u. AP u's transmissions will interfere with the transmission of APs or users within $R_s(u)$ if they share the same channel. To avoid such interference, the distance $d(u,v)$ between u and v has to be greater than $R_t(u)+R_t(v)+\max\{R_s(u),R_s(v)\}$. That is, if APs u and v are greater than $R_t(u)+R_t(v)+\max\{R_s(u),R_s(v)\}$ apart, they can transmit using the same channel, since then no user in v's transmission range will be able to sense a message from u or its users, and vice versa. In Figure 1, user 1 of AP u is within the sensing range of user 2 of AP v. When AP u transmits to user 1, user 2, which is outside of the u's sensing range, may think the media is free. If user 2 then transmits on the same channel as u, its transmission will prevent user 1 from correctly decoding u's transmission.

Currently, APs are statically configured with a channel by administrators who carefully consider channel conflicts and network loads. Channel conflicts among different entities are resolved in an ad hoc manner or not resolved at all. We model the channel assignment problem as a game, where the players are the service providers. APs are set up or acquired by service providers sequentially. When an AP is set up or acquired, a channel must be chosen that does not interfere with the channels chosen for APs that were previously set up; if there is no such channel, the AP cannot be used.[1] The order that the APs are set up is determined exogenously (that is, by some agency outside the game, not the service providers themselves) and is arbitrary. For example, service provider 1 might set up 5 APs before service provider 2 sets up any. We assume that when a service provider sets up an AP, it knows about the APs that have already been set up and might interfere with it, but we do not make any assumptions about what the service providers know about other APs that have already been set up. For simplicity, we also assume that when a service provider sets up an AP, it does not know what APs will become available in the future. The only information it has is the APs that currently exist. For example, suppose that there is only one service provider and there are 3 APs, v_1, v_2, and v_3, which are placed so that v_2 interferes with each of v_1 and v_3, but v_1 and v_3 do not interfere with each other. If there is only one color and v_2 is acquired first, then the provider should choose not to give it a channel, since that will prevent it from giving a channel to v_1 and v_3. But, for simplicity, we assume that the service provider must choose a channel for v_2 as soon as it is set up.

The utilities of the service providers depend on how many users they can serve. We assume that there is a commonly known distribution of users. The utility to a service provider of setting up an AP u that is assigned a channel is the expected number of users in $R_t(u)$; if AP u is not assigned a channel, then its utility to the service provider is 0. The utility of a provider at the end of a game is just the sum of the utilities of the APs that it sets up.

A *socially optimal* assignment is one where the maximum number of users can be served. We would expect that a central authority would assign channels in a way that leads to a socially optimal assignment. Of course, there is no reason to believe that the socially optimal assignment is the one that arises in this game. Our interest is in seeing how far away we are from this assignment. In the language of Koutsoupias and Papadimitriou [14], we are interested in investigating the *price of anarchy*.

We can represent the game using a labeled graph $G=(V,E)$, where the vertices in V are the APs, and two vertices u and v are joined by an edge if they potentially interfere, i.e., if $d(u,v) \leq R_t(u)+R_t(v)+\max\{R_s(u),R_s(v)\}$. Each vertex also has a label, which represents the utility of the AP associated with that vertex being assigned a channel. G is called the *interference graph* induced by the game.

There is obviously a close connection between assigning channels to vertices and coloring the induced interference graph. Suppose that there are k channels. Consider a k-coloring of the graph, i.e., an association of some vertices to colors such that two adjacent edges are labeled with different colors. Clearly this corresponds to a feasible assignment of channels to APs. All APs that are assigned colors can safely communicate on the channel associated to the color without interference. The APs that are assigned a channel in a Nash equilibrium of the game correspond to a maximal subset of vertices that has been colored with k colors. A *maximal k-colored subset* of the induced graph is defined to be a subset of nodes with a specific coloring such that no additional nodes can be colored.[2] If there are any other vertices in the graph that can be colored, then the corresponding AP should have been assigned a channel. Conversely, given a maximal k-colored subset of the interference graph, there is a Nash equilibrium of the game where these are precisely the APs that are assigned a channel. In particular, this will be the case if the APs in the maximal set are set up before any other APs are set up. Thus, there is a 1-1 correspondence between maximal k-colored subsets of the graph and Nash equilibria of the game. Moreover, a socially optimal assignment corresponds to a maximal k-colored subset of maximum weight (that is, a maximal k-coloring where the sum of the weights of the vertices that are colored is maximum). Thus, the price of anarchy is simply the ratio of the total weight of a maximal k-coloring of minimum weight to the k-coloring of maximum weight. Given the close connection between k-colorings and Nash equilibria, we speak in the rest of

[1] We are implicitly assuming here providers follow the "social rule" of not assigning a channel to an AP if it interferes with the channel already assigned to another AP. This social rule can easily be imposed and enforced by a government agency such as the FCC. If we do not make this assumption, then we must define the utility to a service provider of owning an AP whose communications interfere with those of another AP, and discuss how providers bargain over spectrum allocation in this case.

[2] Note the distinction between maximal k-colored subgraph and maximal k-colorable subgraph. A maximal k-colorable subgraph is a subset of nodes that can be colored with k colors such that the subset is not a proper subset of another k-colorable subgraph. When $k=1$, the two definitions are identical.

the paper often of "vertices" and "colors" rather than "APs" and "channels".

It is not hard to show by example that, in the general case, the coverage of the APs in the network that results after channels are chosen can be arbitrarily far from socially optimal; that is, the price of anarchy is unbounded. However, we can do better if we assume that users are uniformly distributed and all APs must use the same transmission power. The interference graph G is then a *unit disk* graph: two vertices u and v are joined by an edge iff $d(u, v) \leq 2R_t + R_s$, where R_t and R_s are the common transmission and sensing ranges, respectively, of all APs. (We are implicitly taking $2R_t + R_s$ to be the "unit" here.) (We remark that the interference graph for 802.11 wireless networks is often modeled as a unit disk graph [3].) Moreover, the utility of a provider is proportional to the number of APs it sets up which are assigned a channel. In this case, we can show that the price of anarchy is at least 5 and at most $5 + \max(0, (k - 5)/k)$, where k is the number of channels. In particular, it follows that if there are at most 5 channels, then the price of anarchy is 5.

Because providers are forced to assign a channel to an AP as soon as it is set up, a provider may be able to do better just by changing the assignment to APs it controls. (This is already clear from the example given above where one provider controls v_1, v_2, and v_3.) It certainly seems reasonable to allow providers to change the channel assignments of APs that it controls. We assume that this is always possible in the remainder of the paper. Service providers may also be able to negotiate changes to assignments of vertices they control that are mutually advantageous. We are particularly interested in what we call *k-buyer–m-seller bargains*, where the m sellers uncolor certain vertices in exchange for payment from k buyers. (Some of the buyers may also be sellers.) The types of bargains we consider improve the total weight of the colored vertices (otherwise the changes will not be of benefit to both buyers and sellers). We are interested in the question of how close bargaining can bring us to a socially optimal assignment. In general, we expect it to be hard to negotiate arrangements involving many buyers and sellers. (By way of analogy, in sports, player trades typically involve two teams; trades involving more than two teams are quite rare.) We start an investigation of this issue in this paper by examining two limited types of bargaining situations, that we expect can be implemented relatively easily in practice.

- The first is a generalization of the situation described initially with APs v_1, v_2, and v_3. If v_2 is colored (i.e., has a channel assigned), v_1 and v_3 are not, v_1 and v_3 could be colored if v_2 were not colored, and the sum of the weights of v_1 and v_3 is greater than the weight of v_2, then we assume that the providers that own APs v_1 and v_3 can always offer the owner of v_2 sufficient utility (in terms of, say, money) so that v_2 is uncolored, while still themselves coming out ahead.[3] We do not go into the details of exactly what the offers are. All that matters is that, in equilibrium, the exchange will be made (i.e., v_2 will be uncolored and v_1 and v_3 colored). We call this a *local 2-buyer–1-seller bargain*.

- The second occurs if an AP is uncolored but its weight is greater than the sum of weights of all its neighbors of a particular color. In this case, we assume that the owner of that AP can offer the owners of the interfering APs sufficient utility so that the interfering APs will be uncolored. Again, we do not go into the details of exactly what the offers are. We call this *local 1-buyer–multiple-seller bargain*. Note that although many sellers may be involved, this really is a collection of 2-way arrangement, since the buyer can negotiate separately with each of the sellers.

[3]For this to be true, we must assume that the utility of money is the same for all players.

By allowing such bargains, we are effectively excluding certain Nash equilibria; thus, the price of anarchy may go down. We show that if local 2-buyer–1-seller bargains are allowed, then in the case that users are uniformly distributed and the power of all APs is the same, the price of anarchy is at most $3 + \max(0, 1 - 3/k)$ and at least 3. Moreover, if users are not uniformly distributed and the transmission power of APs is the same, then if a 1 buyer – multiple seller bargains are allowed, the price of anarchy is at most $5 + \max(0, (k - 5)/k)$ and at least 5.

In all these results, we assumed that the power with which an AP transmits is fixed, and not under the control of the service provider. If the service provider can choose the transmission power from among a finite set of possible transmission powers, we know that the price of anarchy is still unbounded, but if we allow local 1-buyer–m-seller bargains, the price of anarchy is at most 9 and at least $7/(1 + \epsilon)$ if users are distributed uniformly. However, the price of anarchy is still unbounded for local 1-buyer–m-seller bargains if users are not distributed uniformly. Interestingly, bargains covering constant-sized geometric regions do give us a bounded price of anarchy.

We also consider the speed of convergence to a Nash equilibrium in different variants of spectrum sharing games. We prove that in some special cases players converge to a Nash equilibrium after polynomial number of steps. But in the general case, we show that there exists an exponentially long path of improvements to a Nash equilibrium.

The rest of the paper is organized as follows: In Section 2, we define formally the relevant k-colorability problems and apply some standard results from the literature on k-colorability to spectrum sharing. In Section 3, we prove our main results on the price of anarchy. In Section 4, we examine how long it can take to converge to a Nash equilibrium. We discuss related work in Section 5, and conclude in Section 6.

2. GRAPH-THEORETIC PRELIMINARIES

As we observed in the introduction, there is a close connection between our spectrum-sharing game and k-colorability. We formally define the k-colorability problem here, review some standard results on the problem, and show how they apply to the spectrum-sharing game.

DEFINITION 2.1. *Given graph $G = (V, E)$, the* maximum induced k-colorable subgraph **Max k-CIS** problem *is that of finding a k-colorable subgraph of G with maximum number of vertices. (Recall that a graph is k-colorable if it is possible to color the nodes in such a way that no two adjacent vertices are colored with the same color.) If for each vertex of the graph G, we are given a weight $w(v)$, then **weighted Max k-CIS** is the problem of finding a k-colorable induced subgraph whose total weight is maximum.*

It is well known that deciding if a graph is k-colorable is NP-complete [8]. It follows that there is unlikely to be a polynomial time algorithm for finding an optimal channel assignment, even if one player owns all the APs. Indeed, it is hard to even approximate an optimal solution to the **Max k-CIS** problem. More precisely, recall that the *maximum independent set problem* (**Max-IS**) is that of finding a set of vertices of maximum cardinality which are pairwise nonadjacent. It is known that the problem of approximating the **Max-IS** to within a factor better than $\Omega(n^{1-\epsilon})$ (that is, the problem of finding a set of vertices in a graph with n vertices which are independent and whose cardinality is within a factor better than $\Omega(n^{1-\epsilon})$ of a maximum independent set in the graph) is NP-complete [11], and that the problem of finding an approximation to **Max** k**-CIS** is just as hard as that of finding an approximation to the maximum independent set, for any fixed k [12].

The situation is somewhat better for unit disk graphs (which correspond to the situation where all APs transmit with the same transmission power and users are uniformly distributed).

THEOREM 2.2. *There is a* 1.582*-approximation for* ***Max*** k***-CIS*** *in unit disk graphs.*

The proof follows from two facts: (1) we can reduce the problem of approximating **Max** k**-CIS** to that of approximating **Max-IS** and (2) there is a polynomial-time approximation scheme for **MIS** in unit disks [18]. Thus, even if some entity could assign channels, and was trying to do so in a way that maximizes potential usage, the best we can hope for even in the special case that all APs transmit with the same power and users are uniformly distributed is to get an assignment that is within a factor of roughly 1.5 of optimal.

3. THE PRICE OF ANARCHY

In this section, we prove that the price of anarchy is unbounded in the basic spectrum-sharing game for arbitrary graphs, even if players are computationally unbounded. We then show that the price of anarchy is bounded in unit disk graphs. Finally, we consider the extent to which allowing bargaining helps improve the price of anarchy.

Before proving these results, we prove a general result that allows us to reduce to games where there is only a single color/channel available. This allows us to simplify a number of arguments. Moreover, since a 1-coloring is just an independent set, this allows to apply results about the **Max-IS** problem. This result applies for all types of bargaining we consider. The key observation here is that the various types of bargains allowed impose constraints on the structure of an optimal coloring. For example, if we consider the weighted case and allow 1-buyer–multiple-seller bargains, then we do not allow solutions where a vertex is uncolored but has greater weight than all of its neighbors of a given color.

THEOREM 3.1. *Suppose the price of anarchy if there is only one channel for a spectrum-sharing game that allows a certain type of bargaining is* ρ*. Then, for all* k*, the price of anarchy for the same game with* k *channels is at most* $\rho+\max(0, 1-\rho/k)$ *and at least* ρ*.*

PROOF. For the lower bound, suppose that G is an interference graph where the price of anarchy is ρ'. Thus, there are maximal independent subsets X and Y of G such that $w(X) = \rho' w(Y)$. We construct a graph where, even with k colors, the price of anarchy is ρ'. The idea is to replace each vertex in G with k copies of that vertex at the same location. Then there is a Nash equilibrium that involves coloring each of the k vertices that corresponds to a vertex in X a different color, and similarly for Y. Thus, the price of anarchy in the game with k colors is still ρ'. Of course, we cannot set up k APs on top of each other, but we can achieve the same effect as follows. Suppose that we have a setting of APs that results in the interference graph G. Note that there must be an $\epsilon > 0$ such that if all distances are contracted by a factor of $(1-\epsilon)$, G would still be the graph corresponding to the resulting placement of APs. Now replace each AP by a cluster of k APs on the circumference of a circle of radius $\epsilon/2$ around the original AP, and we get the required graph.

For the upper bound, let X consist of the colored vertices in a Nash equilibrium to the given spectrum-sharing game, with color classes $X_1, \ldots, X_k$. Let Y be the vertices in a socially-optimal solution, with color classes $Y_1, \ldots, Y_k$. Let $C = X \cap Y$, let $Y' = Y \setminus C$, and let $Y_i' = Y \cap Y_i$, for $i = 1, \ldots, k$.

Observe that, for all i and j, the $X_j \setminus Y_i$ must a maximal independent subset in the subgraph of G induced by the vertices $(X_j \setminus Y_i) \cup Y_i'$. Clearly the vertices in X_j are independent, since they are all in the same color class. And if there are no edges from some vertex $y \in Y_i'$ to all the vertices in $X_j \setminus Y_i$, then there are no edges from y to any vertex in X_j (since there are no vertices from y to any vertex in Y_i, since they are all in the same color class). Then we can add y to X and color it with the same color as the vertices in X_j, contradicting the maximality of X. It follows that

$$|Y_i'| \ \le\ \rho|X_j \setminus Y_i|,$$

for otherwise the price of anarchy in the graph $Y_i \cup X_j$ with one color is greater than ρ, a contradiction. Summing up over j, it follows that $k|Y_i'| \le \rho|X \setminus Y_i|$. Summing up over i,

$$kw(Y') \ \le\ \sum_i \rho w(X \setminus Y_i) \ =\ \rho(kw(X) - w(C)).$$

Hence,

$$w(Y) \ =\ w(Y') + w(C) \ \le\ \rho w(X) + (1-\rho/k)w(C)$$
$$\le\ (\rho + \max(0, 1-\rho/k))w(X),$$

as desired. □

We start by considering the basic spectrum game, without bargaining. Our first result shows that, in general, the price of anarchy in this game is unbounded, even if all vertices have equal weight.

PROPOSITION 3.2. *The price of anarchy is unbounded in the basic spectrum-sharing game, no matter how many channels or players there are, even if all vertices have equal weight.*

PROOF. First suppose that $k = 1$. Consider a star graph, where the center vertex is connected to n other vertices. If the center vertex appears first, then none of the other vertices can be colored. In the optimal assignment, all the vertices other than the center vertex are colored. Thus, the price of anarchy is n. The fact that the price of anarchy is unbounded with k colors follows immediately from Theorem 3.1. □

Note that the star graph in Proposition 3.2 can be realized by assuming that the center vertex v transmits with high power, while the remaining vertices transmit with low power. We can think of the remaining APs as being placed on the circumference of a circle with center v. It is then easy to distribute the users so that all vertices have an equal number of users, and hence equal weight.

We can construct similar examples even if APs transmit with the same power (although in that case we must look at the weighted coloring problem, since APs have different utilities). However, we cannot construct such an example if APs all transmit with the same power and users are uniformly distributed.

To prove the result, recall that a graph is $\tau+1$*-claw free* if no vertex has more than τ neighbors in the graph, none of which are connected to each other. Unit disk graphs are known to be 6-claw free (see e.g. [15]).

THEOREM 3.3. *If all APs transmit with the same power and users are uniformly distributed, then the price of anarchy of the spectrum-sharing game is at most* $5 + \max(0, 1-5/k)$ *and at least 5.*

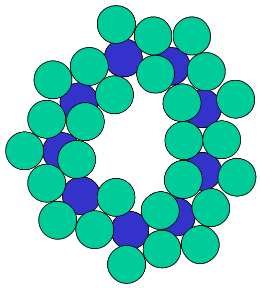

Figure 2: An interference graph for which the price of anarchy is 3 with 2-buyer–1-seller bargaining.

PROOF. By Theorem 3.1, it suffices to show that the price of anarchy is 5 if there is one channel. This follows from the well-known observation that the size of a maximal independent set in a 6-claw free graph is no more than a factor of 5 from the size of the largest maximal independent set in a 6-claw free graph, and a simple example showing that it can be 5, namely, a 6-claw free graph consisting of 6 vertices: a central vertex connected to 5 other vertices. The central vertex by itself is a maximal independent set, as are the other 5 vertices. □

It follows from Theorem 3.3 that for $k \leq 5$, we have a tight bound of 5 on the price of anarchy in unit disk graphs.

We now consider what happens if we allow bargaining. Our first result shows that if all APs transmit with uniform power, users are uniformly distributed, and we allow 2-buyer–1-seller bargaining, then the price of anarchy drops to at most 4.

THEOREM 3.4. *If all APs transmit with the same power, users are uniformly distributed, and 2 buyer–1 seller bargains are allowed, then the price of anarchy of the spectrum-sharing game is at most* $3 + \max(0, 1 - 3/k)$ *and at least* 3.

PROOF. For the upper bound, it suffices by Theorem 3.1 to show that the price of anarchy is 3 if there is one channel. This follows from the analysis of local optimization for independent set in 6-claw free graphs of [13] (see also [10]).

The lower bound is attained by the construction of Fig. 2. The 9 dark circles correspond to vertices in a maximal independent set, while the 27 light circles correspond to the optimal solution. Note that the set of 9 dark circles is stable with respect to 2-buyer–1-seller bargains, since two vertices have to be removed from the maximal independent set (i.e., uncolored) before any new vertex can be added. □

When users are not uniformly distributed, then the price of anarchy for the simple spectrum sharing game is not bounded, even if all APs transmit with the same power and we allow 2-buyer–1-seller bargaining. Indeed, we can show that the price of anarchy is unbounded unless bargains involve at least $\min(p, \tau)$ sellers, where p is the number of players and the interference graph is $(\tau + 1)$-claw free.

PROPOSITION 3.5. *If APs transmit with the same power but users may not be uniformly distributed, then the price of anarchy is unbounded unless bargains involve at least* $\min(p, \tau)$ *sellers, where* p *is the number of players and the interference graph is* $(\tau + 1)$*-claw free.*

PROOF. By Theorem 3.1, we can assume without loss of generality that there is only one channel. Consider a star with a center vertex of large weight and τ leaves of small weight that occupy a channel. The large-weight vertex cannot be bought unless all τ non-adjacent neighbors will be sold. The bound follows if we assume that each of the vertices is controlled by a different player. □

However, as we now show, if we allow 1-buyer–multiple-seller bargains, then the price of anarchy is bounded, even in the weighted case, provided that APs transmit with the same power.

THEOREM 3.6. *If APs transmit with the same power and 1-buyer–multiple-seller bargains are allowed, then the price of anarchy of the spectrum-sharing game is at most* $5 + \max(0, 1 - 5/k)$ *and at least* 5.

PROOF. This follows from Theorem 3.1, using a bound on simple weighted local improvements for maximum independent sets in 6-claw free graphs [2]. The key point is that the "local improvements" of [2] correspond to the result of a 1-player–multiple-seller bargain. □

The requirement that APs transmit with the same power is critical in Theorem 3.6.

PROPOSITION 3.7. *In the general case (where APs transmit with different powers and users are not uniformly distributed), then the price of anarchy of the spectrum sharing game is unbounded, even if multiple-buyer–multiple-seller bargains are allowed.*

PROOF. Consider a spectrum-sharing game with bargains involving up to t buyers and arbitrary number of sellers. Let p be a large value. Consider a graph consisting of vertex v of weight $p \cdot t$ and transmission power corresponding to a circle of radius p; vertices $v_1, \ldots, v_{5p}$ each of weight p and transmission power corresponding to a circle of radius 1; and vertices $v'_1, \ldots, v'_{5p}$, each of weight 1 corresponding to a circle of radius 1. The situation is illustrated in Figure 3. Note that the set $\mathrm{LOPT} = \{v, v'_1, \ldots, v'_{5p}\}$ is a maximal independent set, and has weight $5p + pt$. On the other hand, the set $\mathrm{OPT} = \{v_1, \ldots, v_{5p}\}$ is an independent set of weight $5p^2$. Further note that LOPT cannot be improved with any bargain involving less that t sellers. Hence, the price of anarchy for this instance is $p/(1 + t/5)$. By setting p large enough, we get an unbounded ratio. Furthermore, this example involves only two different weights and two different transmission powers. □

What happens if we allow more general bargains? As we now show, with sufficiently general bargains, we drive the price of anarchy arbitrarily close to 1. However, the bargains may involve arbitrarily many players, which would makes the coordination complexity unreasonable.

For a set X of vertices, let $w(X)$ denote the sum $\sum_{u \in X} w(u)$ of the weights of the vertices in X.

THEOREM 3.8. *Suppose that distances have been normalized so that, for any pair of nodes* u, v, *we have* $R_t(u) + R_t(v) + \max\{R_s(u), R_s(v)\} \leq 1$. *Thus, two vertices* u, v *such that* $d(u, v) > 1$ *do not have an edge between them in the interference graph. Further suppose that bargains involving arbitrary sets of vertices within distance* $\sqrt{2}d$ *are allowed. Then the price of anarchy in the spectrum-sharing game is at most* $d^2/(d-1)^2$.

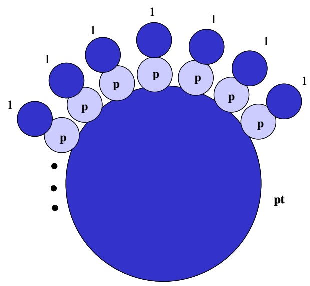

Figure 3: A family of interference graphs with unbounded price of anarchy.

PROOF. Consider a network with induced interference graph G. Suppose without loss of generality that there is one channel. Let LOPT consist of the vertices in a maximal independent subset of G (recall that a maximal 1-colored set is just a maximal independent set) after generalized bargaining, and let OPT be the vertices in a socially optimal independent set.

Consider a $d \times d$ rectangle R with integer coordinates that is half-closed in both directions (i.e., it contains all vertices in its interior and on its right or top boundaries, but not the vertices on its bottom and left boundaries). All vertices within R are of distance at most $\sqrt{2}d$. Let R' be the inner $(d-1) \times (d-1)$ rectangle obtained by removing a unit-length strip from each side of R. This separation ensures that no node within R' interferes with nodes outside R. Since no generalized bargains are possible,

$$w(\mathrm{OPT} \cap R') \;\leq\; w(\mathrm{LOPT} \cap R);$$

otherwise, it would be profitable to buy $(\mathrm{OPT} \cap R') \setminus (\mathrm{LOPT} \cap R)$ and sell $(\mathrm{LOPT} \cap R) \setminus (\mathrm{OPT} \cap R')$. If we sum over all possible d-by-d rectangles with integer coordinates, we count each node in OPT exactly $(d-1)^2$ times but each node in LOPT exactly d^2 times. Thus,

$$(d-1)^2 w(\mathrm{OPT}) \;\leq\; d^2 w(\mathrm{LOPT}),$$

as desired. □

If we assume that ownership is relatively local, so that all the APs within a distance d of each other are owned by a relatively small set of APs, this says that we can get a relatively small price of anarchy. Obviously, as d gets larger, the number of players likely to be involved will increase.

We next consider what happens if users are allowed to choose the transmission power of an AP. That is, when an AP becomes available, a user chooses a channel for it (if one is available) and a transmission power, subject to not interfering with other channels. We then allow the same bargaining procedures (and, as usual, allow players to make arbitrary changes among the APs that they control). Essentially the same example as in the proof of Proposition 3.5 shows that we need to allow multiple sellers in order to get a bounded price of anarchy in this case.

PROPOSITION 3.9. *Even if users are distributed uniformly, in the spectrum-sharing game with power control, the price of anarchy is unbounded unless bargains involve at least* $\min(p, \tau)$ *sellers, where* p *is the number of players and the interference graph is* $(\tau+1)$*-claw free.*

Our next result shows that if we allow multiple sellers, then we do in fact get a bounded price of anarchy.

THEOREM 3.10. *If users are distributed uniformly and 1-buyer–multiple seller bargains are allowed, then the price of anarchy of the spectrum-sharing game with power control is at most 9 and at least* $7-\epsilon$*, for any* $\epsilon > 0$*.*

PROOF. As usual, we can assume without loss of generality that there is one channel. Given a network with induced interference graph G, let LOPT consist of the vertices in a maximal independent subset of G after 1-buyer–multiple-seller bargaining, and let OPT be the vertices in an independent set of greatest weight after bargaining. We divide the vertices in OPT into two groups. A vertex in OPT is *small* if it interferes with at least one vertex of greater weight in LOPT; otherwise it is *large*. These two groups are denoted as $S(OPT)$ and $L(OPT)$ respectively.

We prove this theorem using the following geometric lemma.

LEMMA 3.11. *Let* u *be an AP in LOPT with transmission range circle* C*. Let* N_S *(*N_L*) be the set of neighbors of* u *in* $S(OPT)$ *(*$L(OPT)$*). Then the sum of the weights of nodes in* N_S *is at most* $9 - |N_L|$ *times the weight of* u*.*

PROOF. Let β be the ratio between the sensing and transmission range radii, and let $\phi = 1 + \beta/2$. We construct a set S of circles. Around each small node v in N_S, draw a circle of radius $R_t(v) + R_s(v)/2 = \phi R_t(v)$. For each large node w in N_L, draw a circle of radius $\phi R_t(u)$ with center at distance $2R_t(u)+R_s(u) = 2\phi R_t(u)$ from u along the line from w to u.

We claim that none of these circles in S intersect. Suppose the circles corresponding to nodes v and w intersect. We consider here the case when v is small and w is large; the other cases are similar. Then, the distance from v to w is bounded by

$$\begin{aligned} d(v,w) &< \phi R_t(v) + \phi R_t(u) + (d(w,u) - 2\phi R_t(u)) \\ &= d(w,u) + \phi(R_t(v) - R_t(u)) \\ &= (R_t(w) + R_s(w) + R_t(u)) + \phi(R_t(v) - R_t(u)) \\ &= R_t(w) + R_s(w) + R_t(v) + \beta/2[R_t(v) - R_t(u)] \\ &< R_t(w) + R_s(w) + R_t(v). \end{aligned}$$

Then, v and w interfere, and cannot both be contained in OPT, a contradiction.

All centers of circles in S are within distance $2\phi R_t(u)$ from u, and all the circles are therefore contained within a circle centered at u of radius $3\phi R_t(u)$. This circle is $(3\phi)^2$ times the area of C. The fraction of the area used by transmission range circles is at most $1/\phi^2$, or at most $(3\phi/\phi)^2 = 9$ times the area of C. Of that, the circles in S that derive from nodes in N_L contribute $|N_L|$ to the factor. Finally, recall that our assumption is that weight of a node corresponds to the area of its transmission range circle. □

A large node u in OPT is larger than all the circles it intersects in LOPT. ¿From the local optimality of LOPT, for any u in OPT, $w(u) \leq w(N(u))$. Let $N_L(u)$ be the set of large neighbors in OPT of node u. Thus,

$$w(L(\mathrm{OPT})) \;\leq\; \sum_{u \in L(\mathrm{OPT})} \sum_{v \in N(u)} w(v) \;\leq\; \sum_{v \in \mathrm{LOPT}} |N_L(v)| w(v).$$

For a small node u in OPT, let $v = B(u)$ be some larger circle in LOPT that interferes with C. From Lemma 3.11,

$$\sum_{u \in \mathrm{OPT},\; B(u)=v} w(u) \;\leq\; (9 - |N_L(v)|) w(v).$$

Thus,

$$w(S(\mathrm{OPT}) \ \le \ \sum_{v \in \mathrm{LOPT}} (9 - |N_L(v)|)w(v).$$

Adding together the two inequalities, and summing up over all nodes in LOPT, we have $w(\mathrm{OPT}) \le 9w(\mathrm{LOPT})$.

For the lower bound, we sketch an example where the ratio is arbitrarily close to 7. Arrange two concentric circles C_0, C_L of radius 1 and $1+\epsilon$, respectively. Around C_0, arrange 6 circles $C_1, C_2, \ldots, C_6$ of radius 1. The circles $C_0, \ldots, C_6$ do not overlap, but C_L intersects them all. Hence, the price of anarchy for this instance is $7/(1+\epsilon)$.

4. CONVERGENCE TO NASH EQUILIBRIA

We have assumed that the order that APs are set is determined exogenously. Clearly, if there are n APs altogether, there will be at most n steps before they are all set up. But now suppose that bargaining moves are interspersed with the setting up of APs. How many steps will it take before all the APs are set up and we reach a local optimum, so that no further bargaining can improve the situation? In this section, we address that question. For the purposes of this section, we call a bargaining move or coloring of a new vertex a *local improvement*, since it improves the payoff for at least one agent and does not make any other agent worse off.

This question is particularly easy to answer in the unweighted case, where all APs transmit with the same power and users are uniformly distributed. In that case, each local improvement increases the number of colored vertices by at least one, thus after at most n local improvements, the resulting color assignment is a Nash equilibrium. In the weighted case, the same argument shows that the number of local improvements is finite. In fact, we can easily prove the following:

PROPOSITION 4.1. *In the weighted spectrum sharing game, players will converge to a local optimum after finitely many local improvements, no matter what kind of bargains are allowed. Furthermore, if all weights are integers bounded by a polynomial in the number of vertices, then players will converge to a local optimum after a polynomial number of local improvements.*

PROOF. Each local improvement increases the value of the color assignment, and there are only finitely many color assignments, thus the number of possible improvements is finite. If all weights are integers and are polynomial in the number of vertices, then the total weight of colored vertices is also polynomial in the number of vertices. After each local improvement, the total weight increases by at least one. thus players converge to a local optimum after polynomial number of local improvements. □

The assumption that weights are integers bounded by a polynomial in the number of vertices is critical in Proposition 4.1. If we allow arbitrary weights, then we show that there is always an order of local improvements that reaches a local optimum in a linear number of steps, but in some graphs, there may also be orders that take exponentially many steps.

THEOREM 4.2. *Suppose that local improvements are of two kinds: coloring a new vertex and changing the coloring via a 1-buyer multiple-seller bargain. In the weighted spectrum sharing game on unit disk graphs, it may take exponentially many local improvements to converge to a Nash equilibrium.*

PROOF. Assume $k = 1$, i.e., the number of available colors is one. Consider the graph $G = (V, E)$, where $V = V_1 \cap V_2$, $V_i = \{v_{i,0}, \ldots, v_{i,n}\}$, $w(v_{1,j}) = 2^j$, and $w(v_{2,j}) = 2^j - \epsilon$. Vertex $v_{2,j}$ is connected to $v_{1,j-2}, v_{1,j-1}, v_{1,j}$ As we show in the full paper, this graph is a unit disk graph. We start from the empty coloring. The set of improvements is as follows. We start with an empty coloring:

1. Color vertices v_{10} and v_{11}.
2. Color v_{22} and uncolor v_{10} and v_{11}.
3. Color v_{12} and uncolor v_{22}.
4. Color v_{23} and uncolor v_{12}.
5. Color v_{13} and uncolor v_{23}.
6. Do items 1,2,3 again.
7. Color v_{24} and uncolor v_{12} and v_{13}.
8. Color v_{14} and uncolor v_{24}.
9. Color v_{25} and uncolor v_{14}.
10. Color v_{15} and uncolor v_{25}.
11. Do items 1 to 7 again.

Note that each of the above steps corresponds to coloring a new vertex or a 1-buyer multiple-sellers bargaining. We continue the above sequence by a similar pattern. Using induction, we can easily show that the number of local improvements is at least 2^{n-1} for graph G. We leave the details to the full paper. □

Although, the number of local improvements to a Nash equilibrium can be exponential, it is worth noting that from an empty coloring we can find a path of length at most n of local improvements to a Nash equilibrium. This can be done first ordering the vertices in decreasing order by weight, and then coloring the vertices in a greedy way, starting with the one of highest weight. After the coloring is completed, it is easy to see that no 1-buyer–multiple-seller bargain can improve the situation.

5. RELATED WORK

There are two bodies of work related to ours. The first is work on the price of anarchy in other contexts. Large distributed systems such as the Internet often involve many economic agents. Game theory suggests that, if they follow their own selfish interests in a noncooperative manner, they will end up playing a Nash equilibrium. Koutsoupias and Papadimitriou [14] first proposed investigating the price of anarchy, that is, how far a Nash equilibrium can be from the socially optimal solution to the problem. They studied the price of anarchy of a scheduling problem on parallel machines with selfish jobs. Since their work, much progress has been made in understanding the price of anarchy in many situations. See [4, 6, 9, 19, 22] for a representative sample of the papers on this topic.

Our results are based on relating the Nash equilibrium of the spectrum-sharing game and local optimization algorithms for maximum k-colorable subgraphs. Even-dar et al. [5] studied the convergence time to Nash equilibria of a scheduling game by relating that game to local optimization algorithms for the scheduling problem. See [7, 9, 17] for other work in this spirit.

The second body of relevant work is on spectrum-sharing mechanisms. Efficient spectrum-sharing mechanisms have been proposed by Aftab [1] and Satapathy and Peha [20]. Satapathy and Peha proposed a spectrum-sharing etiquette for devices accessing the free frequency band. Aftab [1] presented an artificial economy approach to the problem. Each vertex is assigned an artificial budget. Nodes use this wealth intelligently to bid dynamically for the right to transmit. These papers consider dynamic channel access. To the best of our knowledge, we are the first to study spectrum sharing in the static case, where a vertex (AP) holds a channel indefinitely unless it releases the channel voluntarily. Our model seems more appropriate for the large 802.11 networks that are being set up by service providers.

6. CONCLUSIONS AND FUTURE WORK

Spectrum sharing is an inherently distributed problem, with no central authority to coordinate and arbitrate channel allocation. It is important that spectrum sharing be efficient, allowing as many users as possible to use the network. With this in mind, we have modeled spectrum sharing as a game between providers, and analyzed the price of anarchy. We show that if we assume that providers are able to use easily implementable bargaining procedures, the price of anarchy is bounded by a constant if users are distributed uniformly or every AP uses the same transmission power.

There are many avenues for future research. We intend to tighten our upper and lower bounds on the price of anarchy. The convergence issues are not completely resolved. In particular, we would like to find useful conditions that guarantee polynomial-time convergence to a Nash equilibrium. In addition, we would like to further investigate the general weighted power-control game where the weight is not just a function of the area within transmission range. We are also interested in investigating the effect on the price of anarchy of allowing different types of bargaining procedures.

Acknowledgments

We thank Mohammad Taghi Hajiaghayi for helpful comments and discussions especially on Section 4.

7. REFERENCES

[1] O. Aftab. Economic mechanisms for efficient wireless coexistence. MIT Technical Report MIT-LCS-TR-876, 2002.

[2] E. M. Arkin and R. Hassin. On local search for weighted k-set packing. *Mathematics of Operations Research*, 23:640–648, 1998.

[3] Y. Bejerano and R. Bhatia. MiFi: A framework for fairness and QoS assurance in current IEEE 802.11 networks with multiple access points. In *Proc. INFOCOM*, 2004.

[4] A. Czumaj, P. Krysta, and B. Vocking. Selfish traffic allocation for server farms. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 287–296, 2002.

[5] E. Even-dar, A. Kesselman, and Y. Mansour. Convergence time to Nash equilibria. In *Proc. 30th International Colloq. on Automata, Languages, and Programming (ICALP)*, pages 502–513, 2003.

[6] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. On a network creation game. In *Proc. 22nd Annual Symp. Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[7] A. Fabrikant, C. Papadimitriou, and K. Talwar. On the complexity of pure equilibria. In *Proc. 37th Annuals Symp. on Theory of Computing (STOC)*, 2004.

[8] M. R. Garey and D. S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, CA, 1979.

[9] M. Goemans, L. Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad hoc networks. In *Proc. MOBIHOC*, 2004.

[10] M. M. Halldórsson. Approximating discrete collections via local improvements. In *Proc. Sixth ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 160–169, Jan. 1995.

[11] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.

[12] D. S. Hochbaum. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Company, Boston, MA, 1997.

[13] C. A. J. Hurkens and A. Schrijver. On the size of systems of sets every t of which have an SDR, with an application to the worst-case ratio of heuristics for packing problems. *SIAM Journal of Discrete Mathematics*, 2(1):68–72, Feb. 1989.

[14] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proc. 16th Annual Conf. Theoretical Aspects of Computer Science*, Lecture Notes in Computer Science, Volume 1563, pages 404–413. Springer-Verlag, Berlin, 1999.

[15] M. V. Marathe, H. Breu, H. B. Hunt III, and S. S. R. a nd D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.

[16] Meshnetworks, Inc. Medford Mobilizes Wireless Workforce. http://www.meshnetworks.com/pdf/cs_medford_or.pdf.

[17] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Random-Approx, to appear*, 2004.

[18] T. Nieberg, J. L. Hurink, and W. Kern. A new PTAS for maximum independent sets in unit disk graphs. Department of Applied Mathmematics, University of Twente, Technical Memorandum No. 1688, 2003.

[19] T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[20] D. P. Satapathy and J. M. Peha. A novel co-existence algorithm for unlicensed fixed power devices. In *Proc. IEEE Wireless Communication and Networking Conference (WCNC)*, volume 3, pages 1403–1409, 2000.

[21] Verizon. Verizon WiFi Networks. http://www.verizon.net/wifi/.

[22] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proc. 43rd Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.