Noname manuscript No.
(will be inserted by the editor)

On Spectrum Sharing Games*

Magnús M. Halldórsson · Li (Erran) Li · Joseph Y. Halpern · Vahab S. Mirrokni

February 19, 2010

Abstract Efficient spectrum-sharing mechanisms are crucial to alleviate the bandwidth limitation in wireless networks. In this paper, we consider the following question: can free spectrum be shared efficiently? We study this problem in the context of 802.11 or WiFi networks. 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 (i.e., not in a coordinated way) 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 optimally 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

* A preliminary version of this paper appeared in ACM Symposium on Principles of Distributed Computing(PODC), 2004

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.

Magnús M. Halldórsson
School of Computer Science, Reykjavik University, 101 Reykjavik, Iceland.
E-mail: mmh@ru.is

Joseph Y. Halpern
Department of Computer Science, Cornell University, Ithaca, NY 14853.
E-mail: halpern@cs.cornell.edu

Li Li, Center for Networking Research, Lucent Bell Labs, Murray Hill, NJ 07974.
E-mail: erranlli@research.bell-labs.com · Vahab S. Mirrokni
Theory Group, Google Research, New York, NY
E-mail: mirrokni@theory.csail.mit.edu

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.

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

1 Introduction

The Federal Communications Commission (FCC) currently allocates two types of spectrums: dedicated and free. Dedicated spectrum is allocated exclusively to the spectrum owner wheas free spectrum can be used by anyone. For dedicated spectrum, recent measurements by Shared Spectrum Company [27] shows that the maximal total spectrum occupancy in New York City is only 13%. These findings support the ongoing effort by FCC and various interested parties to propose a more efficient spectrum-allocation model. For example, one model proposes that the spectrum owner should allow other users access as long as they keep the interference temperature under some threshold. In this paper, we consider the efficiency of the other extreme, where the spectrum is given to various providers, according to some policy. For definiteness, we consider the issues that arise in the context of 802.11 Wireless LANs, commonly known as WiFi.

WiFi provides wireless network access to subscribers. It has been deployed in public hotspots, ranging from airports to hotels to 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 [28] and MeshNetworks Corporation has been deploying city-wide WiFi networks to facilitate law enforcement and emergency response in cities such as Medford, Oregon [19].

To understand the issues of interest to us, we need to briefly review some relevant details of WiFi (see [5] 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 three such channels and 802.11a has twelve. 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 Distributed Coordination Function (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 guarantee that there is no interference with the user's message. To avoid interference in the wireless media between nearby APs and their users, those nearby APs must use different channels.

1.1 Network model

We can associate with each AP u two circular regions around it (see Figure 1). The smaller circle, with radius denoted by $R_t(u)$, represents u's *transmission range*. All messages sent by u can be correctly received by users in $R_t(u)$. The larger circle, with radius denoted by $R_s(u)$,

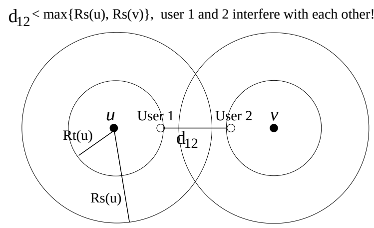

Fig. 1 Potential interference between two APs.

represents u's *sensing range*. The sensing range of a transmitter is the maximum distance at which a receiver can sense the transmitter's signal. The receiver knows if the channel is busy, but it may or may not be able to decode the signal. In practice, $R_s(u)$ is about twice $R_t(u)$. The APs form communication links with users within their transmission range. Note that, each user also has its own transmission range and interference range. Two links can interfere if a party to one link (a user, or an AP) is within sensing range of a party to the other link. Since users can show up anywhere within an AP's transmission range, to avoid interference among users of neighboring APs, as shown in [5], the Euclidean distance $d(u,v)$ between AP u and AP v has to be more than $R_t(u) + R_t(v) + \max\{R_s(u), R_s(v)\}$. That is, if APs u and v are more 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 effect, each AP u has a *potential interference range* $R_i(u)$ of $R_t(u) + R_s(u)$, and users of other APs that are within that distance of u can potentially be affected by u's communication. To be conservative, we therefore require APs u and v using the same channel to be separated by a distance of $R_I = \max(R_i(u) + R_t(v), R_i(v) + R_t(u)) = R_t(u) + R_t(v) + \max\{R_s(u), R_s(v)\}$. In Figure 1, user 1 of AP u is within the sensing range of user 2 of AP v. While AP u transmits to user 1, user 2, which is outside of 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. Thus, if APs with the same channel are separated by a distance smaller than R_I, there is potential for interference among clients in overlapping regions. These clients are subject to starvation and persistent collisions. The network performance will degrade. Therefore, we only allow APs to use the same channel if they are separated by more than R_I.

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 (i.e., not in a coordinated way) or not resolved at all. This ad hoc method of channel allocation can be quite inefficient [2]. Solutions assuming that all APs use the same software have been proposed [2,21], but they do not consider the possibly selfish behavior of different entities operating the APs, and work therefore only in a cooperative environment.

2 Modeling the game

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 so as not to interfere with channels chosen earlier by other APs; if there is no such channel, the AP cannot be used.

We are implicitly assuming here that providers follow the "social rule" of not assigning a channel to an AP if it interferes with the channel already assigned to another AP. Without rules about interference in place, there will be continual interference between users of nearby APs, leading to serious complaints about service. Although there is currently no such rule for the unlicensed band of 802.11 networks, this social rule can easily be imposed and enforced by a government agency such as the FCC. Such rule enforcement is, however, outside the scope of the game that we consider.

The order in which APs are set up is determined exogenously (that is, by some process outside the scope of the game) 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 assume that service providers do not know anything about what APs will be set up in the future. 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 channel and v_2 is acquired first, then, if the provider knew that it would acquire v_1 and v_3 (and channel assignments cannot be changed), it would not assign a channel to v_2, since it would be better off assigning a channel to v_1 and v_3 instead. However, since we assume that channel providers do not know the future, we assume that they assign channels to APs as they acquire them. (We later also allow providers to reassign channels, so that if v_1 and v_3 are acquired, then the channel assigned to v_2 can be reassigned to v_1 and v_3.)

The utilities of the service providers depend on how many users they can serve. We assume that there is a fixed and 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. This is a simplification of the true utility, which is difficult to model precisely and to reason about; our definition of utility captures, for example, the peak utilization characteristics of the operation.

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 when the players do the assignment on their own. In the language of Koutsoupias and Papadimitriou [16], we are interested in investigating the *price of anarchy*.

We can represent the game using a labeled graph $G = (V,E)$, where the APs are vertices in V, and two vertices u and v are joined by an edge if the APs 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. A move of a player in this game corresponds to (re)assigning channels to some or all of the APs it controls.

3 Our results

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 (see Proposition 1). 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 remark that the interference graph for 802.11 wireless networks is often modeled as a unit disk graph [5], by taking $R_t + R_s/2$ to be the "unit".) Moreover, the utility of a provider is proportional to the number of APs it sets up that 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 (Theorem 2). 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 *n-buyer–m-seller bargains*, where the m sellers unassign channels from certain APs in exchange for payment from n buyers. (Note that some of the buyers may also be sellers.) Of course, we are interested in bargains that benefit both buyers and sellers. As long as both buyers and sellers place the same monetary value on access to users (which we assume), then these will be bargains that result in more users having access to the network. For if unassigning a channel so that others can be assigned results in more users having access, then there will be a price that one provider can pay another in order to unassign the channel that will make both providers better off.

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 by examining two limited types of bargaining situations that we expect can be implemented relatively easily in practice.

– We call the first type of bargain a *(local) 2-buyer–1-seller* bargain. By local, we mean buyers and sellers form a star topology in our labeled graph G. This occurs when there are three APs v_1, v_2, and v_3 as in our earlier example, where v_2 is currently assigned a channel which interferes with v_1 and v_3 and, by assigning the channel to v_1 and v_3 instead, more users would have access to the system. Thus, the owners of APs v_1 and v_3 can always offer the owner of v_2 sufficient money to unassign the channel, while still coming out ahead themselves. (The owners of v_1, v_2, and v_3 do not all have to be different for such a bargain to be struck.) We do not go into the details of exactly what the offers are. All that matters is that, in equilibrium, the channel assignments will be changed appropriately.
– We call the second type of bargain a *(local) 1-buyer–multiple-seller* bargain. This occurs when an AP v cannot be assigned a channel because all the channels have already been assigned to its neighbors, $v_1,\ldots,v_m$, while the expected number of users of AP v is higher than the total expected number of users of $v_1,\ldots,v_m$. Here again, the owner of v should be able to offer the owners of $v_1,\ldots,v_m$ enough money for them to unassign the channels to $v_1,\ldots,v_m$, although we do not go into the details of the actual bargaining. Note that although many sellers may be involved, this really is a collection of 2-way arrangements, since the buyer can negotiate separately with each of the sellers.

By allowing such bargains, we are effectively changing the game in a way that may reduce the set of Nash equilibria; thus, the price of anarchy may go down. We show that if 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 (Theorem 3), where k is the number of channels available. Moreover, if users are not uniformly distributed and the transmission power of the APs is the same, then if 1-buyer–multiple-seller bargains are allowed, the price of anarchy is at most $5+\max(0,(k-5)/k)$ and at least 5 (Theorem 4).

In all these results, we have 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 7.441 and at least $7/(1+\varepsilon)$, for any $\varepsilon > 0$, if users are distributed uniformly (Theorem 6). However, the price of anarchy is still unbounded for local 1-buyer–*m*-seller bargains if users are not distributed uniformly (Proposition 1). Interestingly, bargains covering constant-sized geometric regions (of diameter $\sqrt{2d}$, for a given parameter d) do give us a bounded price of anarchy.

The results on the price of anarchy are summarized in Table 3; in the table, we use r^* to denote a lower bound of r and an upper bound of $r+\max(0,1-r/k)$, where k is the number of channels available.

Scenario →	Uniform user distribution		Arbitrary user distribution	
Game ↓	Unit power	Different power	Unit power	Different power
No-bargaining game	5^*	∞	∞	∞
2-buyer–1-seller	3^*	∞	∞	∞
1-buyer–multiple-sellers	5^*	$7 \le r \le 7.441$	5^*	∞
Dist-$\sqrt{2}d$ bargains	$1 \le r \le d^2/(d-1)^2$	$1 \le r \le d^2/(d-1)^2$	$1 \le r \le d^2/(d-1)^2$	$1 \le r \le d^2/(d-1)^2$

Table 1 Summary of results on the price of anarchy r.

We also consider the speed of convergence to a Nash equilibrium in several variants of spectrum-sharing games. We prove that in some special cases players converge to a Nash equilibrium after polynomially-many 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 4, we define the games formally as graph problems. In Section 5, we prove our main results on the price of anarchy. In Section 6, we examine how long it can take to converge to a Nash equilibrium. In Section 7, we present approximation algorithms for the optimal channel allocation problem. Each player needs to solve this problem to compute its best response function. We discuss related work in Section 8, and conclude in Section 9.

4 Game Formulations

Representation. We represent APs as points in the Euclidean plane. Each point v has two associated concentric circles: a transmission region of radius $R_t(v)$, and a sensing region of radius $R_s(v)$. These radii are determined by the transmission power of the AP. We assume that there is a fixed ratio between the transmission and sensing range. There is also a non-negative weight $w(v)$ attached to each vertex v, representing the utility of assigning that AP a channel. Given a set S, let $w(S) = \sum_{v\in S} w(v)$ be the total weight of the set. We assume the utility to be linear in the number of users that can be serviced, and the same for all providers. Our results can be extrapolated to more general metrics, to allow for imperfect signal propagation, antennas, or other environmental features. For this purpose, we state results in terms of the independence property of the metric, or the claw-freeness of the intersection graph.

Given power settings, we have a graph $G = (V,E)$, where the vertices V are the APs and vertices u and v are adjacent if $d(u,v) \leq R_t(u) + R_t(v) + \max\{R_s(u), R_s(v)\}$, i.e. if transmissions involving the corresponding APs could potentially interfere. The graph will be *labeled*, or *colored*, with a subset of the vertices receiving one of k colors. We view this partial coloring as a function mapping V to the set $\{1,2,\ldots,k\} \cup \{\perp\}$, with $\perp$ representing an empty label. The coloring represents the assignment of channels to some of the APs, with the social rule guaranteeing that adjacent vertices receive different colors.

Spectrum sharing games. A *spectrum-sharing game* is characterized by a graph and a set of legal moves. We can think of the nodes of the graph as the set of APs that will ultimately be set up. The vertices of the graph are partitioned between the *players* of the game. Intuitively, player i's vertices are the APs that i will ultimately set up. At any point, player i can color (or recolor) a subset of the vertices that he has set up; intuitively, these are the APs that i has currently set up. Thus, a move at time m in the game consists of the players assigning colors to or changing the colors of the APs they have currently set up and possibly transferring money. We consider different games depending on the constraints on money transfer. For example, if there is no money transfer, then we are modeling the game without bargaining. To model n-buyer–m-seller bargains, we allow money transfer between groups of up to $n+m$ player, where up to n players can decrease their money total and up to m can increase their money total. The utility to a player at the end of the game is the sum of the weights of the nodes he has colored (i.e., the sum of the utilities of the AP that it sets up) together with the utility of its money it has. For simplicity, we assume that utility is linear in money, and that all the players value money the same way. Thus, the total utility of a game position can be identified with the sum of the weights of the colored vertices. Clearly, a Nash equilibrium of a spectrum-sharing game must be a *locally-optimal coloring*—no additional vertices can be colored.

A game position corresponds to a *k-colored subgraph*. A socially-optimal assignment of channels corresponds to a maximum weighted (induced) k-colorable subgraph. The *price of anarchy* is the maximum ratio between the utilities of a socially-optimal solution and of a locally-optimal solution, i.e., the ratio between the best and the worst Nash equilibrium. This is identical to the *performance ratio* of local search algorithms that find the local optima with respect to the given game operations. We consider the following games.

Basic game. The *basic spectrum-sharing game* is one with no cooperation or bargaining between players. A player's move in this game is to color some of its nodes, in a manner consistent with the social rules. A player may change the coloring of APs it owns (or uncolor them) if that results in better utility. This in turn can cause other players to change. It is easy to see that a Nash equilibrim of this game must be a a *maximal k-colored subgraph*, i.e. one where no further vertices can be colored, given the current coloring. Indeed, there is clearly a one-to-one correspondence between maximal k-colored subgraphs and Nash equilibria in a game without bargaining. (Observe the subtle difference between a *maximal k-colorable subgraph* and a maximal k-colored graph. The latter is one that admits no local improvements—that is, no additional vertices can be assigned a color; the former is one where any additional vertex would make the subgraph non-k-colorable. Consider a path with four vertices, where the first and third vertices have been colored by colors 1 and 2, respectively. If there are only $k = 2$ colors, then this coloring gives a maximal k-colored subgraph. But the only maximal k-colorable subgraph is the whole graph.) Given the correspondence between maximal k-colored subgraphs and Nash equilibria, in the rest of the paper, we use the following terminology interchangeably: vertices and APs, colors and channels, k-colored subgraphs and solutions, and locally optimal solutions and Nash equilibria.

Games with bargaining. As we said earlier, we also consider *spectrum-sharing games with bargaining*, where players can bargain with other players, and pay them for uncoloring their vertices. In a *n-buyer m-seller bargaining game*, each local improvement can involve up to n buyers seeking to add vertices to their colored set, and m sellers that are willing to uncolor their vertices in exchange for a payment. A move in this game takes a colored graph S, uncolors up to m vertices and colors up to k vertices in such a way that if S' is the resulting colored graph, then $w(S') > w(S)$. As we said earlier, we are implicitly assuming that all players have the same utility for money as well as for users, so that such a trade will be rational.

Games with power control. We consider both the situation when all APs transmit with the same power, and when the power can be variable. In the former case, the interference graph is, without loss of generality, a *unit disk graph*, since after scaling distances by $2R_t + R_s$, two vertices are adjacent iff their distance is at most 1. In the latter case, the games involve *power control* operations, where users increase or decrease the utility of their APs by adjusting their transmission power and thus transmission and sensing radii. This can also involve a number of buyers and/or sellers. As before, the decrease in the utility of APs that decrement their power must be more than offset by the increased utility of APs that can increment their transmission radii.

A power-control game can also be (implicitly) represented as a coloring game on the graph with vertices for all possible power settings of each AP. In practice, there are only a fixed number of power settings for each AP. Our upper bounds are independent of available settings, while the lower bounds focus on the asymptotic case.

User distribution. We also consider two extreme cases for the distribution of the users. In the uniform case, the users are distributed uniformly and homogeneously over the plane. Then the utility of an AP is proportional to the area of its transmission region, which again is proportional to the square of the transmission radius. If, in addition, power is uniform, this amounts to counting the number of nodes colored. In the second scenarios, we assume an arbitrary, possibly worst-case, distribution of the users.

Computing an optimal solution is in general a hard problem; indeed, even finding an approximately optimal solution may be hard. However, as we discuss in Section 7, there are effective approximations possible in the planar setting (i.e., when G is a unit disk graph).

5 The Price of Anarchy

In this section, we prove that the price of anarchy for arbitrary graphs in the basic spectrum-sharing game (without bargaining) is unbounded, 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.

We first 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 us to apply results about the *maximum independent set* (**Max-IS**) problem. This result applies to all the 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 1 *Consider a spectrum-sharing game with a given local-optimality criterion. Suppose that the price of anarchy in the case of a single channel is ρ and, more specifically, that for any locally-optimal independent set X and any independent set Y we have that*

$$w(Y \setminus X) \leq \rho \cdot w(X \setminus Y).$$

Then, for any 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 upper bound, given a spectrum-sharing game on a graph $G = (V,E)$, 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 $Y_i' = Y' \cap Y_i = Y_i \setminus C$, for $i = 1, \ldots, k$.

Let H_j be the graph induced by $(V - X) \cup X_j = V - \cup_{j' \neq j} X_{j'}$. X_j is locally optimal in H_j, since otherwise we could have extended X_j without affecting the rest of our solution. In particular, $Y_i' \cup (Y_i \cap X_j)$ is contained in H_j, so we have, by assumption, that

$$w(Y_i') \ \leq \ \rho \cdot w(X_j \setminus Y_i).$$

Summing up over j, it follows that $k \cdot w(Y_i') \leq \rho \cdot w(X \setminus Y_i)$, and summing over i,

$$\begin{aligned} k \cdot w(Y') = \sum_i k \cdot w(Y_i') \ &\leq \ \sum_i \rho \cdot w(X \setminus Y_i) \\ &= \ \rho(k \cdot w(X) - w(C)). \end{aligned}$$

Hence,

$$\begin{aligned} w(Y) \ = \ w(Y') + w(C) \ &\leq \ \rho w(X) + (1-\rho/k) w(C) \\ &\leq \ (\rho + \max(0, 1-\rho/k)) w(X), \end{aligned}$$

as desired.

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 assigning a different color to each of the k vertices that correspond to a vertex in Y, and similarly for Y. Thus, the price of anarchy in the game with k colors is still ρ'. Even if we cannot set up k APs on top of each other, 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 $\varepsilon > 0$ such that if all distances are contracted by a factor of $(1-\varepsilon)$, 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 $\varepsilon/2$ around the original AP, and we get the required graph.

We start by considering the spectrum-sharing 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 1 *The price of anarchy is unbounded in the spectrum-sharing game without bargaining, 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 has been selected, 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 1.

Note that the star graph in Proposition 1 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.

Recall that a graph is *$(\tau+1)$-claw free* if each vertex has at most τ mutually non-adjacent vertices. Unit disk graphs are known to be 6-claw free (see e.g. [18]). This follows from the observation that if two non-overlapping unit circles centered at y_1 and y_2 overlap a unit circle centered at x, then the angle $\angle y_1 x y_2$ must be greater than 60°.

Observation 1 *A unit disk graph is 6-claw free.*

Theorem 2 *If all APs transmit with the same power and users are uniformly distributed, then the price of anarchy of the basic spectrum-sharing game is at most* $5 + \max(0, 1 - 5/k)$ *and at least 5.*

Proof Consider a maximal independent set X and an optimal independent set Y. Each node in $Y \setminus X$ can dominate at most 5 vertices in $X \setminus Y$, since G is 6-claw free. Thus, the size of $Y \setminus X$ is at most 5 times that of $X \setminus Y$, and the upper bound follows by Theorem 1.

A simple example shows that the price of anarchy 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 2 that for $k \leq 5$, we have a tight bound of 5 on the price of anarchy in unit disk graphs.

5.1 Spectrum sharing games with bargaining

We now consider what happens if we allow bargaining.

5.1.1 2-buyer–1-seller 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 *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*.*

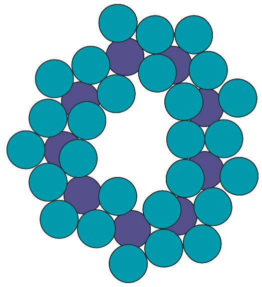

Fig. 2 An interference graph for which the price of anarchy is 3 with 2-buyer–1-seller bargaining.

Proof For the upper bound, it suffices by Theorem 1 to show that the price of anarchy is 3 if there is one channel. This follows from the analysis of 2-opt local optimization for unweighted independent sets in 6-claw free graphs of [15] (see also [13]). We repeat it here for completeness.

Let X be a solution locally optimal under 2-opt, i.e. local improvements involving adding up to two vertices and removing up to one. Let Y be any other independent set. Let $C = X \cap Y$, $X' = X \setminus C$, Y_1 be the set of vertices in $Y \setminus C$ with one neighbor in X and Y_2 those in $Y \setminus C$ with two or more neighbors in X. Observe that C, Y_1 and Y_2 partition Y. Then, no two nodes in X are neighbors of the same node in Y_1, due to local optimality. Hence, $|Y_1| \leq |X'|$. Also, each node in X' has at most 5 neighbors in Y. Hence, counting edges between nodes in X' and Y, we have $2|Y_2| + |Y_1| \leq 5|X'|$. Adding the two inequalities gives $2|Y \setminus C| \leq 6|X \setminus C|$. Hence, $|Y| = |Y \setminus C| + |C| \leq 3|X \setminus C| + |C| \leq 3|X|$.

The lower bound is attained by the construction of Fig. 2. The 9 dark circles correspond to vertices in a maximal independent set X, while the 27 light circles correspond to the optimal solution Y. Each node in X is adjacent to four nodes in Y_2 and one in Y_1. With only one seller, there can be only one node that can be added (namely, the sole adjacent node in Y_1), which gives no benefit. Hence, X is locally optimal under any multiple-buyer–1-seller bargains.

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 2 *Suppose 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 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. Suppose the channel is currently occupied by the τ leaves. Then, the large-weight vertex cannot be bought unless all its τ non-adjacent neighbors are sold. The bound now follows if we assume that each of the vertices is controlled by a different player.

5.1.2 1-buyer–multiple-seller bargaining

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 4 *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 By Theorem 1, it suffices as before to consider only the case of a single channel. Arkin and Hassin [3] and Bafna, Narayan and Ravi [4] showed that the performance ratio of a simple local improvement heuristic for the maximum weight independent set problem in $\tau+1$-claw free graphs equals τ. Their heuristic corresponds to the addition of a single vertex with the removal of as many as needed (but at most τ). This corresponds to bargains involving one buyer and multiple sellers. In general, the result of [3] shows that the performance ratio involving n buyers is $4+1/n$.

The requirement that APs transmit with the same power is critical in Theorem 4.

Proposition 3 *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 t sellers. Let z be a value larger than t. Let β denote the ratio between the sensing and transmission range. Let $q = 5z/\beta$. Consider a graph consisting of vertex v of weight $z \cdot t$ and transmission power corresponding to an interference radius of z; vertices $v_1, \ldots, v_q$ each of weight z and transmission power corresponding to a transmission radius of 1; and vertices $v'_1, \ldots, v'_q$, each of weight 1 with power corresponding to a transmission radius of 1. The situation is illustrated in Figure 3. Observe that the q circles $v_1, \ldots, v_q$ can be distributed so that they stay mutually non-overlapping while their transmission ranges overlap with the interference range of v. Note that the set $\mathrm{LOPT} = \{v, v'_1, \ldots, v'_q\}$ is a maximal independent set of weight $zt+q$. On the other hand, the set $\mathrm{OPT} = \{v_1, \ldots, v_q\}$ is an independent set of weight $qz = 5z^2/\beta$. Any exchange involving only t buyers offers a gain of only zt at the cost of giving up the heavy node, for a total cost of $zt+t$, resulting in a negative improvement. The same holds for bargains involving $t+1$ sellers or fewer. Hence, the price of anarchy for this instance is

$$\frac{qz}{q+zt} = \frac{z}{1+zt/q} = \frac{z}{1+5t/\beta}\ .$$

This grows linearly with z, resulting in an unbounded ratio. Note that this example involves only two different weights and two different transmission powers.

5.1.3 Dist-$\sqrt{2}d$ bargaining

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

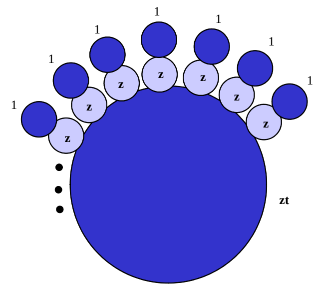

Fig. 3 A family of interference graphs with unbounded price of anarchy.

Theorem 5 *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, any pair of vertices u, v with $d(u, v) > 1$ are not adjcent 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 $d^2/(d-1)^2$.*

Proof Consider a network with induced interference graph G. Suppose without loss of generality that there is one channel. Let LOPT be 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 from the top-left, i.e., containing 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 mutual distance at most $\sqrt{2}d$. Let R' be the inner $(d-1) \times (d-1)$ rectangle obtained by removing unit-length strips from the bottom and the right edges 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.

5.2 Power control games

We next consider what happens if users are allowed to choose the transmission power of APs. 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 2 shows that we need to allow multiple sellers in order to get a bounded price of anarchy in this case.

Proposition 4 *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. We prove this theorem using two geometric lemmas.

Lemma 1 *Let x be a point in the plane and let S be a set of disks of radius at most 1 whose centers are located within distance at most* 2 *from x. Then, the total area of the disks in S is at most* 7.441π*. If, additionally, the circles in S differ in diameter by a factor of at most* 1.10157*, then the area of the circles is at most* 7π*.*

Proof Let C_2, C_3 be the circles centered at x of radius 2 and 3, respectively. All circles in S are contained inside C_3. This gives a quick upper bound of 9π on the area of S.

Consider the annulus A_2 formed by the region between C_2 and C_3, and let X be any disc in S that overlaps A_2. Let $\theta = \angle pxp'$, where p and p' are the points of intersection of X with C_2, i.e., θ is the angle of the sector induced by the overlap of X with A_2. We examine the area $A(X,A_2)$ of overlap of X and A_2 relative to the angle θ. It is easily observed that $A(X,A_2)/\theta$ is maximized when X is largest possible and centered as far from x as possible. In our case, it is maximized when X is a unit circle with a center on C_2. In this case, the sector of overlap between X and A_2 spans $\theta = 4 \cdot arcsin(0.25) = 57.91°$. The amount $A(X,A_2)$ of overlap of such a circle with A_2 can be computed to be $\pi - 1.40307 \approx 1.73852$. The maximum amount of coverage of A_2 by circles in S is therefore at most $360/57.91 \cdot 1.73852 \approx 10.8076 \leq 3.441\pi$. Adding the area inside C_2, we obtain a maximum total area of overlap of circles in S with C_3 to be 7.441π.

When the circles in S are of similar diameter, we can apply known results about the packing of circles inside circles. The smallest radius of a circle that can contain 8 circles of radius r is $3.30476r$ (see [11]). Observe that $3.30476/1.10157 > 3$. Thus, if the difference in the radius is at most a factor of 1.10, we can shrink the circles in S to the size of the smallest one and still at most 7 of them can fit inside C_3. The desired result follows.

Lemma 2 *Let u be an AP and let* $N(u)$ *be a set of neighbors of u that are themselves mutually non-interfering. Assume the weight of an AP is proportional to the area of its transmission region. Let* $N_S(u)$ ($N_B(u)$) *be the APs in* $N(u)$ *that are smaller or equal (bigger) than u, respectively. Then,*

$$w(N_S(u)) \leq (7.441 - |N_B(u)|)w(u).$$

If the transmission radii of APs in N_S *differ from that of u by a factor of at most 1.10, then* 7.441 *can be replaced by 7.*

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 disks. Around each small node v in N_S, form a disk in S of radius $R_t(v)+R_s(v)/2 = \phi R_t(v)$. For each big node w in N_B, form a disk 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. Let w' denote the new position of big node w. Let R_x denote $\phi R_t(x)$, for a node x. Note that $R_{w'} = R_u$ and $d(w',w) = d(u,w) - 2R_u$.

We claim that none of the disks in S intersect. Suppose the disks of nodes v and w' intersect. That is,

$$d(v,w') < R_v + R_{w'} = R_v + R_u.$$

We consider here the case when v is small and w is big; the other cases are similar. Since u interferes with w and u is smaller than w, we have by definition that

$$d(w,u) < R_t(w) + R_s(w) + R_t(u) = R_w + R_u/\phi.$$

Then, using the triangular inequality, we have that

$$\begin{aligned} d(v,w) &\leq d(v,w') + d(w',w) \\ &< (R_v + R_u) + (d(w,u) - 2R_u) \\ &< (R_w + R_u/\phi) + (R_v - R_u) \\ &= (R_w + R_v/\phi) + (1 - 1/\phi)(R_v - R_u). \end{aligned}$$

Since v is smaller than u, this is less than $R_w + R_v/\phi = R_t(w) + R_s(w) + R_t(v)$, and thus v and w interfere, which is a contradiction. Hence, the disks in S are non-intersecting. Observe also that the centers of all disks in S are of distance at most $2R_u$ from u.

We now apply Lemma 1, scaling the radii by R_u. By our assumption, the area covered corresponds to the weights of the node. Hence, $w(S) \leq 7.441 \cdot w(u)$. Of that, the circles in S that come from nodes in N_B contribute a $|N_B|$ term. The circles in S that come from nodes in N_S are unchanged in size. Hence, $w(N_S(u)) \leq (7.441 - |N_B(u)|)w(u)$.

Theorem 6 *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 7.441 and at least 7. If power assignments differ by at most a factor of 7, the price of anarchy is exactly 7.*

Proof As before, 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 a socially-optimal independent set. 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 *big*. These two groups are denoted as OPT_S and OPT_B respectively. A big node u in OPT is larger than all those it intersects in LOPT. From the local optimality of LOPT, it holds for any u in OPT that $w(u) \leq w(N(u))$. Thus,

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

Let $f : \mathrm{OPT}_S \rightarrow \mathrm{LOPT}$ be a mapping of each small node u in OPT to some larger node in LOPT that interferes with u. Define the inverse mapping, $g : \mathrm{LOPT} \rightarrow 2^{\mathrm{OPT}_S}$, such that $g(v) = f^{-1}(v)$ contains the nodes in OPT_S that map to node $v \in \mathrm{LOPT}$. Note that g forms a partition of OPT_S. From Lemma 2, for each v in LOPT we have that

$$\sum_{u \in g(v)} w(u) \ \leq \ (7.441 - |N_B(v)|) w(v).$$

Summing over nodes in LOPT, we get that

$$w(OPT_S) \ \leq \ \sum_{v \in \mathrm{LOPT}} (7.441 - |N_B(v)|) w(v).$$

Adding together the two inequalities gives $w(\mathrm{OPT}) \leq 7.441w(\mathrm{LOPT})$.

When power assignments are nearly identical, Lemma 2 yields an improved bound of 7, resulting in the same upper bound for the price of anarchy.

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 $1+\varepsilon$, respectively. Around C_0, arrange 6 circles $C_1, C_2, \ldots, C_6$ of radius 1. Note that these circles correspond to transmission radii of $1/\phi$ and $(1+\varepsilon)/\phi$, respectively. 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 $7/(1+\varepsilon)$. Since ε can be made arbitrarily small, we obtain a lower bound of 7 on the price of anarchy.

We conjecture that the true price of anarchy is 7. In lieu of formal evidence, we have verified this computationally by formulating a nonlinear optimization problem. When solved using the LGO global optimization suite [17], we obtain an optimal solution of 7, achieved only when packing 7 circles.

Claim The price of anarchy of the spectrum-sharing game with power control is 7.

6 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 5 *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 5. 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.

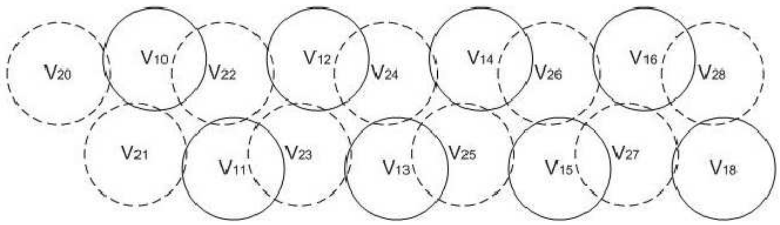

Fig. 4 The unit disk graph used in Theorem 7 with $n = 8$.

Theorem 7 *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 $k = 1$, that is, that there is only one color available. 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 $w(v_{2,j}) = 2^j - \varepsilon$. Vertex $v_{2,j}$ is connected to $v_{1,j-2}, v_{1,j-1}, v_{1,j}$. The graph for $n = 8$ is depicted in Figure 4. It is not hard to show that 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 can continue the sequence above using a similar pattern. Using induction, it is straightforward to show that the number of local improvements is at least 2^{n-1} for graph G.

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 by 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.

7 Approximation Algorithms for Optimal Channel Allocation

Consider the situation of a service provider a after all the APs have been acquired (and there may have been perhaps some bargaining). At that point, a owns a collection of APs, and would like to color them in an optimal way. This is the problem we already considered. However, now there is an additional complication. Service provider a may be able to pay other providers to uncolor some of their APs. Which ones should it pay? That depends in part, of course, on the cost. Suppose, for simplicity, that all the vertices of the other service providers have already been allocated colors and have a publically-known (non-negotiable) price of uncoloring. Then we can associate with each vertex v owned by provider a and color c the cost $p(v,c)$ of coloring v with c. This is just the sum of the prices required to uncolor all the neighboring nodes owned by other service providers that are colored by c. (The price is 0 if no neighbors of v are colored by c.) The following definition captures the essence of the problem in which we are interested.

Definition 1 Let $G=(V,E)$ be a graph and let $\{1,\ldots,k\}$ be the available colors. With each vertex v are associated $k+1$ numbers: $w(v)$, $p(v,1)$, $\ldots$, $p(v,k)$, where $w(v)$ is the utility of vertex v and $p(v,c)$ is the cost of coloring vertex v with color $c \in C$. The **Max weighted** ***k*****-LIS** problem is the problem of finding a proper coloring $C: V \rightarrow \{1,\ldots,k\} \cup \{\perp\}$ that maximizes the total weight of colored vertices minus the total price of assigning these colors to these vertices, i.e., $\sum_{v\in V(G),C(v)\neq\perp} w(v) - \sum_{v\in V(G),C(v)\neq\perp} p(v,C(v))$.

Max-LIS is a generalization of the maximum induced k-colorable subgraph problem. Thus, it is not approximable for general graphs. However, for unit disk graphs we can design efficient approximation algorithms for this problem. If we treat k, the number of colors, as a constant, then we can in fact design a polynomial time approximation scheme (PTAS) for the problem, i.e, a $1+\varepsilon$-approximation for any constant $\varepsilon > 0$.

Theorem 8 *There is a PTAS for the unweighted* ***Max-LIS*** *problem for any constant k on unit disk graphs.*

Proof The algorithm is based on that of Nieberg, Hurink, and Kern [23] for finding maximum independent sets in unit disk graphs. Let G be a unit disk graph. Assume that there are k colors available. Let $N^r(v)$ consist of all vertices u such that $d(u,v) \le r$, and let $L_r(v)$ be the maximum number of vertices in $N^r(v)$ that can be colored (with the k colors). Observe that $N^r(v)$ is contained within a circle $C_r(v)$ of radius r centered at v. The number of vertices in $N^r(v)$ that can be colored with one color is then at most r^2, as each has area π compared with the area of $C_r(v)$ of πr^2. Thus, $L_r(v) \le kr^2$.

Let $\rho = 1+\varepsilon$. Let r' be such that $\rho^{r'} \ge kr'^2$. Since $\rho > 1$, it is always possible to find such an r'. We claim that there exists a constant r such that $L_{r+1}(v) \le \rho L_r(v)$. Otherwise, if $L_{r+1}(v) \ge \rho L_r(v)$ for all $r \le r'$, then it holds $L_r(v) \ge \rho^r > kr^2$, which contradicts our previous bound on $L_r(v)$.

We claim that one can find a $1/\rho$-approximation to an optimal **Max-LIS** set in time $n^{O(kr^2)}$, where $n=|G|$. We prove this by induction on n. We proceed as follows. Choose a vertex $v \in G$ and $r' \le r$ such that $L^{r'}(v) < \rho L^{r'+1}(v)$. Note that we can find r' in time $n_1^{O(kr^2)}$, where $n_1 = |N^{r'}(v)|$ by exhaustive search of all the colorings. In this time, we can also find an optimal **Max-LIS** set of $N^{r'}(v)$. Since $L^{r'}(v) < \rho L^{r'+1}(v)$, the optimal **Max-*****k*****-LIS** set of $N^{r'}(v)$ is a $1/\rho$-approximation to the optimal **Max-*****k*****-LIS** set for $N^{r'+1}(v)$. Now consider the graph $G' = G - N^{r'+1}(v)$. By the induction hypothesis, we can find a $1/\rho$ approximation to an optimal **Max-LIS** set for G' in time $n_2^{O(kr^2)}$, where $n_2 = |G'|$. Note that all vertices of G' and $N^r(v)$ are nonadjacent, so we can combine the optimal solution for $N^{r'}(v)$ with the solution for G' to get a $1/\rho$-approximation to the optimal solution for G.

The algorithm above does not solve the problem for the weighted case. In the weighted case, we can design a PTAS if the graph is given as a set of disks in the plane.

Theorem 9 *There is a PTAS for the* **weighted Max-*k*-LIS** *problem (where k is a constant) on unit disk graphs, when the graph is given as a set of disks in the plane.*

Proof We give a randomized algorithm, that is easily derandomized. Let $\varepsilon > 0$ be given and let $d = 2 + 2/\varepsilon$. Pick values α, ψ uniformly at random from $\{0, 1, \ldots, d-1\}$. Construct a graph G' containing all the disks from G except those with centers (x, y) such that $\lfloor x \rfloor \equiv \alpha \pmod{d}$ or $\lfloor x \rfloor \equiv \beta \pmod{d}$. The graph G' now consists of disjoint subgraphs. Each subgraph C has independence number at most $(d-1)^2$, and thus the maximum number of vertices in a k-colorable subgraph in C is at most kd^2. There are at most $\binom{|C|}{kd^2}$ such sets to search through. Hence, by exhaustive search, we can find a optimal solution to **Max-LIS** on G' in time n^{kd^2}. Let OPT (OPT') be the weight of the maximum weight **Max-LIS** solution on G (G'). By the linearity of expectation, $E[OPT'] = \frac{d^2-2d+1}{d^2} OPT$. Thus, the expected performance ratio is bounded by

$$\frac{d^2}{d^2 - 2d + 1} \leq \frac{d}{d-2} = 1 + \varepsilon.$$

The algorithms in Theorem 8 and 9 are not polynomial time if k, the number of colors, is part of input. Instead, if k is a part of input, the same local optimization algorithm for **weighted Max-CS** gives a 6-approximation for **Max weighted *k*-LIS**. The proof is similar to that of Theorem 1, as it involves only a comparison between optimal and locally-optimal solutions.

Theorem 10 *There is a polynomial time 6-approximation algorithm for the* ***Max weighted k-LIS*** *problem in unit disk graphs.*

8 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 in a Nash equilibrium. Koutsoupias and Papadimitriou [16] 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 [6, 8, 12, 25, 29] for a representative sample of the papers on this topic.

Our results are based on relating the Nash equilibrium of the spectrum game and local optimization algorithms for maximum k-colored subgraphs. Even-Dar et al. [7] studied the convergence time to Nash equilibria of a scheduling game by relating that game to local optimization algorithms for the scheduling problem. See [9, 12, 20] for other work in this spirit. Subsequent to the conference version of our paper, Panagopoulou and Spirakis [24] considered a graph-coloring game similar in spirit to those considered here. In thier game, the set of players is the set of vertices, and the action space of each player is the set of colors. The payoff that a vertex v receives is the total number of vertices that have chosen the same color as v, unless a neighbor of v has also chosen the same color. Thus, their payoff function is quite different from ours.

The second body of relevant work is on spectrum-sharing mechanisms. Efficient spectrum-sharing mechanisms have been proposed in several works [1,26,14]. Satapathy and Peha [26] 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. Huang et al. [14] present an auction-based centralized mechanism to allocate transmitting power to a group of users such that the interference at any user is below a threshold. They consider spread spectrum system where there is no fixed channelization. Thus, these papers either consider dynamic channel access or power allocation in spread spectrum system. 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. Subsequent to our work, Felegyhazi and Hubaux [10] study a power-control game among cellular providers. The key difference is that, they assume each base station broadcasts a pilot signal using a pilot channel, and users communicate with the base station with the strongest received pilot signal. The utility of a base station is the coverage area where it has the strongest pilot signal.

9 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 provided that users are distributed uniformly or every AP uses the same transmission power.

There are many avenues for future research. Some of the bounds on price of anarchy can be tightened, and the convergence issues are not completely resolved. In particular, useful conditions that guarantee polynomial-time convergence to a Nash equilibrium would be valuable. Further investigations on general weighted power-control game where the weight is not just a function of the area within transmission range, as well as the effect of other types of bargaining procedures on the price of anarchy, would be interesting.

Finally, spectrum sharing games for other systems are an open topic. We have disallowed an AP a channel if it interferes with existing APs. Alternatively, providers can agree on a social rule where the usage of a channel is shared in some fair manner. We have also assumed that channel width is fixed, while one could consider spectrum sharing games where channel width can be adapted. Moscibroda et al. [22] have considered the centralized channel assignment problem.

Acknowledgements We thank MohammadTaghi Hajiaghayi for helpful comments and discussions especially on Section 6.

References

1. O. Aftab. Economic mechanisms for efficient wireless coexistence. MIT Technical Report MIT-LCS-TR-876, 2002.
2. A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self-management in chaotic wireless deployments. In *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 185–199, 2005.
3. E. M. Arkin and R. Hassin. On local search for weighted k-set packing. *Mathematics of Operations Research*, 23:640–648, 1998.
4. V. Bafna, N. Narayan, and R. Ravi. Nonoverlapping local alignments (weighted independent sets of axis-parallel rectangles). *Disc. Appl. Math.*, 71:41–53, 1996.
5. 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*, pages 1229–1240, 2004.
6. 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.
7. 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.
8. 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.
9. A. Fabrikant, C. Papadimitriou, and K. Talwar. On the complexity of pure equilibria. In *Proc. 37th Annuals Symp. on Theory of Computing (STOC)*, pages 604–612, 2004.
10. M. Felegyhazi and J.-P. Hubaux. Wireless operators in a shared spectrum. In *Proc. INFOCOM*, pages 1–11, April 2006.
11. E. Friedman. `www.stetson.edu/~efriedma/packing.html`. Attributed to Braaksma, 1963.
12. M. Goemans, L. Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad hoc networks. In *Proc. MOBIHOC*, pages 55–66, 2004.
13. M. M. Halldórsson. Approximating discrete collections via local improvements. In *Proc. Sixth ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 160–169, Jan. 1995.
14. J. Huang, R. Berry, and M. L. Honig. Auction-based spectrum sharing. *ACM/Springer Journal of Mobile Networks and Applications (MONET)*, 11(3):405–418, 2006.
15. 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.
16. 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.
17. LGO. Pintér Consulting Services. `http://www.pinterconsulting.com/`.
18. 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.
19. Meshnetworks, Inc. Medford Mobilizes Wireless Workforce. `http://www.meshnetworks.com/pdf/cs_medford_or.pdf`.
20. V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proc. 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), LNCS 3122*, pages 183 – 194, 2004.
21. A. Mishra, D. Agrawal, V. Shrivastava, S. Banerjee, and S. Ganguly. Distributed channel management in uncoordinated wireless environments. In *Proceedings of the 9th Annual International Conference on Mobile Computing and Networking (MOBICOM)*, pages 170–181, 2006.
22. T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan. Load-aware spectrum distribution in wireless lans. In *Proc. 16th IEEE International Conference on Network Protocols (ICNP)*, pages 137–136, 2008.
23. T. Nieberg, J. L. Hurink, and W. Kern. A robust PTAS for maximum independent sets in unit disk graphs. In *Proc. 30th Workshop on Graph-Theoretic Concepts in Computer Science (WG). Springer LNCS 3353*, pages 214–221, 2004.
24. P. N. Panagopoulou and P. G. Spirakis. A game theoretic approach for efficient graph coloring. In *ISAAC '08: Proceedings of the 19th International Symposium on Algorithms and Computation*, pages 183–195, Berlin, Heidelberg, 2008. Springer-Verlag.
25. T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
26. 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.

27. Shared Spectrum Company. Comprehensive spectrum occupancy measurements over six different locations, August 2005. `http://www.sharedspectrum.com/?section=nsf_summary`.
28. Verizon. Verizon WiFi Networks. `http://www.verizon.net/wifi/`.
29. 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.