

Game-Based Admission Control for Wireless Systems (Invited Paper)

Benjamin Yolken
yolken@stanford.edu

Nicholas Bambos
bambos@stanford.edu

Stanford University
Stanford, CA

ABSTRACT

Much previous work has examined the wireless power control problem using tools from game theory, an economic concept which describes the behavior of interdependent but non-cooperative users. In this paper, we expand these ideas to the antecedent process of deciding which users may participate in the network, i.e. the admission control problem. In particular, we propose three distinct *pricing schemes* for influencing users as they make their participation decisions. We fully characterize the equilibria induced by each and then test their performance in a simulated, wireless environment. Our preliminary results show that these schemes have the potential to produce high quality outcomes in an incentive-compatible way.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Economics, Performance, Theory

Keywords

admission control, game theory, Nash Equilibrium, wireless

1. INTRODUCTION

The last decade has seen an explosion in the diversity and ubiquity of wireless mobile devices. These are not just used for voice communication, but now, increasingly, as part of rich, but bandwidth-intensive, multimedia applications. This trend, in turn, has required the study of two inter-related issues in the allocation and use of wireless spectrum: (1) transmission *power control* and (2) *quality-of-service* (QoS) guarantees. The first is important because mobile devices, by design, have only limited battery capacity. The latter must be conserved so that the user is not

left “stranded” with a dead device before the opportunity to recharge. In addition, the choice of transmission power can have effects on other users. If some device is unnecessarily “blasting” at full-power, for instance, this may significantly increase the interference perceived by others nearby, forcing them to also increase their powers or drop out of the system.

The second concern, QoS, arises because of the applications mentioned previously. Users watching streaming video, for example, are very sensitive to image resolution, skipped frames, sound-image synchronization, etc. If the “quality” of the wireless link is not high enough, then these and other criteria may degrade, making users unhappy and significantly reducing the value they get from the underlying application. Although the environment of study has several plausible “quality” metrics, we focus in the sequel on the *signal-to-interference ratio* (SIR) in the associated wireless channel. This value is correlated with the channel’s theoretical information “capacity” and thus reflects “quality” in an application-agnostic way.

The previous two concerns are obviously not independent. As transmission power increases, SIR (and quality) also increase. This increase, however, comes at the expense of reduced battery life and increased interference experienced by neighboring users. Thus, the “optimal” choices for power and QoS expectations reflect *tradeoffs* between the two.

Two approaches have been developed in the literature for addressing the latter tradeoffs. The first, originally proposed by Foschini and Miljanic in [3], uses the concept of “hard,” SIR-targets. When users asynchronously adjust their powers, they can converge to a *Pareto-optimal* power point, i.e. one at which all targets are met with the minimal total power expenditure. The second, parallel approach, maps the power control problem into *game theory*, a branch of economics used to model strategic interactions among non-cooperating participants. The power-QoS tradeoff is explicitly encapsulated in user-specific *utility functions*. By selfishly seeking to maximize these functions, the users may converge to a *Nash Equilibrium*, i.e. a point at which no player has an incentive to unilaterally deviate. See [1, 4, 9], among others. A few works, such as [5, 10], have assumed utility functions with “soft” SIR targets, thus attempting to bridge the two approaches.

In this work, we consider the antecedent process of *admission control*, i.e. deciding who is allowed to use the wireless channel in the first place. This process, for the most part not emphasized in the previously cited works, plays a major role in the final system outcome. In the Foschini-Miljanic model, for instance, allowing everyone to participate may negate the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON’08 November 17-19, 2008, Maui, Hawaii, USA.
Copyright 2008 ICST 978-963-9799-36-3.

existence of the Pareto power point discussed above. In the game-oriented models, on the other hand, admission control is necessary for ensuring that the resulting Nash equilibria are desirable from the standpoint of efficiency, fairness, and other criteria.

This admission control process has obviously been examined in much previous work (see, for instance, [2, 6, 7]). Our original contribution here, however, is in considering this process within a framework that also includes SIR targets, user utilities, and economic incentives. In particular, we consider a variety of *payment schemes* which act as a “front door” to the overall system. These payments serve to keep out users who would have detrimental effects in the following power transmission game / update process. In addition, the payments could be a source of revenue for the wireless operator. As such, these schemes simultaneously address many important concerns in these wireless environments: power control, QoS, money, and others.

The remainder of this paper is organized as follows. In Section 2, we discuss the mathematical, wireless network model used for the analysis in the sequel. Section 3 adds to the latter the idea of *utility functions*, a framework useful for modeling user behavior and incentives in the previous environment. In Section 4, we explore various approaches to payment-based *admission control*, the main original contribution of this work. Finally, in Sections 5 and 6, we present some numerical examples, discuss directions for future research, and conclude the paper.

2. WIRELESS NETWORK MODEL

As in [10], consider a wireless network modeled as N interfering radio links. These links are assumed to occupy a single channel; orthogonal channels are not considered as they do not interfere with the former. Each user / link selects a transmission power, p_i , and experiences an SIR, $\gamma_i(\mathbf{p})$ as a function of its own power, the powers chosen by the other users, and the system environment. In particular, we assume

$$\gamma_i(\mathbf{p}) = \frac{G_{ii}p_i}{\sum_{i \neq j} G_{ij}p_j + \eta_i} \quad (1)$$

Using the notation commonly employed in the literature, $G_{ii} > 0$ represents the power-attenuation from the user i 's transmitter to its receiver, $G_{ij} > 0$ is the power-attenuation of user j 's signal from the transmitter of link j to the receiver of link i , and $\eta_i > 0$ represents the thermal noise experienced by user i . The exact values of these constants are exogenously determined by the system environment.

Assume that each user also has a “hard” target SIR denoted by $\bar{\gamma}_i > 0$. We seek to characterize the set of powers at which all of the latter are met or exceeded, i.e. $\gamma_i(\mathbf{p}) \geq \bar{\gamma}_i \forall i$. As done in [2], one can show that this set is equivalent to those power vectors, $\mathbf{p} > \mathbf{0}$, for which

$$(\mathbf{I} - \mathbf{F})\mathbf{p} \geq \mathbf{u}, \quad (2)$$

where

$$\begin{aligned} F_{ij} &= \begin{cases} 0 & \text{if } i = j \\ \frac{\bar{\gamma}_i G_{ij}}{G_{ii}} & \text{if } i \neq j \end{cases} \quad \forall i, j \\ u_i &= \frac{\bar{\gamma}_i \eta_i}{G_{ii}} \quad \forall i \end{aligned} \quad (3)$$

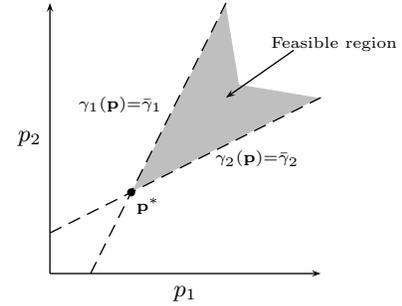


Figure 1: Illustration of \mathbf{p}^* geometry for $N = 2$ case. The set of all feasible powers form a “cone” with the former point at its tip. Any other point in this cone satisfies $\mathbf{p} \geq \mathbf{p}^*$ and thus requires greater power expenditure.

If $(\mathbf{I} - \mathbf{F})^{-1}$ is well-defined and componentwise positive, or, equivalently, if the spectral radius of \mathbf{F} , $\rho(\mathbf{F})$, is strictly less than 1, then the latter system has a solution given by

$$\mathbf{p}^* = (\mathbf{I} - \mathbf{F})^{-1}\mathbf{u} > \mathbf{0}, \quad (4)$$

This power vector is *Pareto optimal* in the sense that any other solution to the previous inequality system has $\mathbf{p} \geq \mathbf{p}^*$. Thus, the latter point is the “best” choice in that it satisfies all of the SIR targets in the most power-efficient way. See Figure 1 above for a geometric interpretation of this result. In addition, the system will converge to this point with the Foschini-Miljanic (FM), “best-response” power update scheme:

$$p_i(t+1) = p_i(t) \frac{\bar{\gamma}_i}{\gamma_i(t)} \quad \forall i \quad (5)$$

See [3] for more details.

If, on the other hand, $(\mathbf{I} - \mathbf{F})^{-1}$ is not well-defined ($\Leftrightarrow \rho(\mathbf{F}) \geq 1$), then there exists no point at which all SIR targets are simultaneously met. In fact, if users iteratively update their powers according to the FM algorithm above, the former will converge to infinity. As discussed in the introduction, *admission control* is needed to avoid this obviously undesirable outcome. We discuss this further in Section 4 below.

3. GAME THEORETIC MODEL

The previous model addresses the power-QoS tradeoff but does so in a way that ignores user incentives and unnecessarily restricts the space of possible outcomes. To address this, we impose user-specific *utility functions* and consider the participation and power level decisions within a *game-based* framework.

3.1 Utility and Value Functions

Many different utility function forms have been proposed in the literature, and many of these are indeed plausible, depending on the specific assumptions in the underlying environment. As in [10], however, we restrict our attention in the sequel to utilities that can be written as

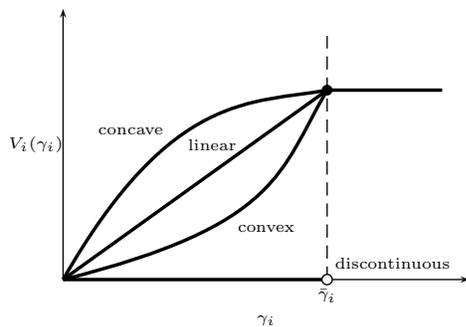


Figure 2: Some possible $V_i(\cdot)$ function forms for environments with “soft” target SIRs.

$$U_i(\mathbf{p}) = \begin{cases} V_i(\gamma_i(\mathbf{p})) - C_i(p_i) & \text{if } i \text{ participates} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

If the user participates in the system, then its utility is given by the *value*, $V_i(\cdot)$, it receives as a function of the perceived link quality / SIR minus the cost, $C_i(\cdot)$ that it has to pay. Although written as an explicit function of power, the latter costs could be fixed or dependent on other, exogenously determined parameters (e.g., $\bar{\gamma}_i$).

In addition, we assume that the former value functions are *target-based*. In particular, these values are all (weakly) increasing below the user targets but then *saturate* at $\bar{\gamma}_i$, implying that the users get no additional value for SIRs above and beyond the latter. This assumption is appropriate, for instance, in multimedia streaming applications, where there is some fixed, upper bound on user-perceived quality.

Of course, the description above allows for many possible functional forms on the intervals $[0, \bar{\gamma}_i)$. See Figure 2 above for some examples. For simplicity, however, we assume in the sequel that these have a binary, discontinuous structure, i.e.:

$$V_i(\gamma_i(\mathbf{p})) = \begin{cases} \hat{v}_i & \text{if } \gamma_i(\mathbf{p}) \geq \bar{\gamma}_i \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

for some constants $\hat{v}_i > 0$ for all i . The latter values may be private information, not known to the system operator. Many of our results below can be extended to the case of more general $V_i(\cdot)$ functions. This significantly complicates the analysis, however, and we thus leave this as an extension for future work.

3.2 Cost Functions

The value functions above are endogenous attributes of the participants and their preferences (e.g., SIR targets). We assume that the cost functions, however, are set by the operator. By adjusting the $C_i(\cdot)$, this operator has a “knob” by which it can direct the system towards specific, desired outcomes. In particular, these can be set to facilitate the *admission control* process described in the introduction.

To see this, note that, given the previously assumed utility function forms, user i will only participate if \hat{v}_i , the maximum value it can expect to obtain, is greater than or equal to its cost. Otherwise, it can do (weakly) better by sitting out. Thus, by setting the costs high or low, the operator

has the implicit power to keep some users out while allowing others through. Moreover, this process is done in an *incentive-compatible* way. No user is forced to leave the system against its will; rather, it chooses to not participate based on an open, observable charging policy.

These “admission price” choices and the outcomes they induce are discussed further in Section 4 below.

3.3 Admission and Power Control Game

Given the setup above, consider the *single-stage*, complete information *game* in which users simultaneously name their power choices, p_i , and then receive utilities according to the previously described form. Note that this game also includes the admission control problem, as users can choose $p_i = 0$ and thus signal their intent to not participate.

We then restrict our attention to those outcomes which lie in the set of *pure strategy Nash Equilibria*. More formally, this corresponds to those \mathbf{p} for which

$$U_i(p_i, \mathbf{p}_{-i}) \geq U_i(p'_i, \mathbf{p}_{-i}) \quad \forall i, p'_i \in [0, \infty) \quad (8)$$

where \mathbf{p}_{-i} represents the powers of all users other than i . At any Nash Equilibrium, therefore, no user has an incentive to unilaterally deviate. These points thus represent the “stable,” “expected” outcomes in a system with non-cooperating, selfish users.

4. ADMISSION CONTROL SCHEMES

In this section, we explore several possible $C_i(\cdot)$ choices and the admission control schemes they induce. For each, we describe the properties of the resulting Nash Equilibria (if they exist), and then discuss the advantages and disadvantages of the given scheme in more detail.

4.1 Fixed Admission Prices

The simplest approach for payment-based admission control is to set a *fixed, flat entrance price* that must be paid by all participants, irrespective of their SIR-targets, powers, etc. Within the above mathematical framework, this corresponds to having $C_i(\cdot) = C$ for some positive constant C , uniform across all users.

Before discussing the incentive-compatibility of this scheme, we need some additional notation. Given fixed power choices, \mathbf{p} , let \mathcal{P} represent the set of participants, i.e. the set of users for whom $p_i > 0$, and \mathcal{N} represent the negation of \mathcal{P} , i.e. the set of users not participating. Furthermore, let $\mathbf{F}_{\mathcal{P}}$ and $\mathbf{p}_{\mathcal{P}}^*$ denote, respectively, the interference matrix and resulting Pareto power point when Equations 3 and 4 are restricted to just those users in \mathcal{P} .

We now assert the following:

PROPOSITION 1. *Given some fixed admission price choice, C , the system has a Nash Equilibrium if and only if there exists a partition of the users such that:*

1. $\rho(\mathbf{F}_{\mathcal{P}}) < 1$
2. $\hat{v}_i \geq C \quad \forall i \in \mathcal{P}$
3. $\hat{v}_j \leq C \quad \forall j \in \mathcal{N}$

PROOF. Suppose a Nash Equilibrium, \mathbf{p} exists. Let \mathcal{P} and \mathcal{N} represent, respectively, the set of participants and non-participants implied by the powers in the latter vector. We

then, necessarily have that the three conditions above hold. If the first does not hold, then it is impossible for the users in \mathcal{P} to reach their targets, and thus they will constantly be increasing their powers. If the second or third do not hold, on the other hand, then at least one user has an incentive to unilaterally deviate, violating the assumptions on our equilibrium \mathbf{p} . Thus, \mathcal{P} and \mathcal{N} are, indeed, a partition of the desired form.

On the other hand, if the three conditions above hold, then there exists a Nash Equilibrium. In particular, we can set the powers of those users in \mathcal{P} according to $\mathbf{p}_{\mathcal{P}}$, with $p_i = 0$ for all non-participants. In this case, no users in either \mathcal{P} or \mathcal{N} have an incentive to deviate, giving us the desired result.

□

We thus have that there is a very simple Nash Equilibrium, provided that C is chosen properly. In particular, this admission fee must be “high enough” to ensure that the system is not “overloaded” by those users with $\hat{v}_i \geq C$. In any environment, a Nash Equilibrium with no participation can be induced if we set $C > \max_i \hat{v}_i$.

Note that if $C \neq \hat{v}_i \forall i$ and an equilibrium exists, then the $\mathcal{P} / \mathcal{N}$ partition implied by this equilibrium is unique. If C equals one or more user values, on the other hand, then there could exist multiple equilibria corresponding to different participation choices of those users “on the boundary.” The latter condition, however, is unlikely to be encountered in practice.

This scheme has the advantage of simplicity, but with at least two significant downsides. First, computing a “good” value of C may be difficult, particularly with limited information about the user \hat{v}_i values. If the operator sets this value too low, then an equilibrium may not exist, leading to continual power fluctuations and user unhappiness. On the other hand, if this value is set too high, then this may make admission unnecessarily restrictive, causing the system to be significantly underprovisioned.

Second, even with complete information, there is no guarantee that any C choice will lead to an *efficient* outcome. In other words, it may not be optimal from a “total user happiness” standpoint to just let in all those with $\hat{v}_i > C$. The reason is that these values do not necessarily incorporate the interference caused and experienced by each user; a high value user, may, for instance, “block” several lower valued ones, even if the combined values of the latter exceed the former. By ignoring these criteria, fixed pricing schemes can lead to significantly suboptimal allocations.

As a solution to the first concern, the operator could set C based on some kind of iterative, *price discovery* process. For instance, users could be asked to name their maximum “willingness to pay” in an *auction*, and then the final admission price could be set to the lowest value at which all participant targets can be met. This is a very promising approach, but one that requires a detailed study of truthfulness incentives and other concerns. We thus leave it as a topic for future research.

The second concern, on the other hand, cannot be addressed with the framework of flat admission prices. We thus need to consider other pricing schemes, ones that penalize users for the harm they cause to others. This is the subject of our next section.

4.2 Externality-Based Pricing

Much research in the field of economics has shown that overall system outcomes can be significantly improved when users are forced to internalize the effects of their actions on others (see, for instance, [8]). The easiest approach is often to incorporate these *externalities* directly into the observed system prices. Users who could, potentially, cause more harm to others are charged more and, as a result, weigh this result when deciding whether to participate.

This method can only work, however, if there is some way of quantifying these inter-user effects. In full generality, these could reflect the total drop experienced in “other user” utilities when each user decides to participate. This method, however, is computationally difficult and leads to non-uniform, potentially “unfair” pricing schemes.

A more palatable approach is to instead charge on the basis of some easily observed *proxy* for these externality effects. Before describing these, we need one more piece of notation. Let g_i represent the sum of the coefficients used to calculate user i 's interference effects on the other users. More formally

$$g_i = \sum_{j \neq i} G_{ji} \quad (9)$$

We describe the intuition behind this metric later in this section.

Given the above, we now consider the following two, non-flat pricing schemes:

1. Charge users in proportion to their transmission powers: $C_i(\cdot) = cp_i$
2. Charge users in proportion to the interference metric defined above: $C_i(\cdot) = cg_i$

where, in each case, $c > 0$ is some system-wide scaling term set by the operator.

Note that each of these captures the externalities described previously. The first does this by charging more as powers are increased. Since this increase in power also increases the interference experienced by the other users, the transmitting user is, in a sense, being penalized in proportion to this “harm.” The second approach increases payments in proportion to g_i , a metric that is increasing in the G_{ji} interference coefficients. Thus, i pays in proportion to the potential interference that it could cause by participating. Both approaches, therefore, force users to “pay” in a manner consistent with their impacts on the optimality of the overall system.

As in the previous section, it is also possible to characterize the equilibria resulting from these schemes. In particular, we have the following theorems:

PROPOSITION 2. *Suppose that prices are set as $C_i(\cdot) = cp_i$. Then, there exists a Nash Equilibrium if and only if there exists a partition of the users, \mathcal{P} and \mathcal{N} , such that:*

1. $\rho(\mathbf{F}_{\mathcal{P}}) < 1$
2. $\hat{v}_i \geq cp_{i,\mathcal{P}}^* \quad \forall i \in \mathcal{P}$
3. $\hat{v}_j \leq cp_{j,\mathcal{P}+j}^* \quad \forall j \in \mathcal{N}$

PROPOSITION 3. *Suppose that prices are set as $C_i(\cdot) = cg_i$. Then, there exists a Nash Equilibrium if and only if there exists a partition of the users, \mathcal{P} and \mathcal{N} , such that:*

1. $\rho(\mathbf{F}_{\mathcal{P}}) < 1$
2. $\hat{v}_i \geq cg_i \quad \forall i \in \mathcal{P}$
3. $\hat{v}_j \leq cg_j \quad \forall j \in \mathcal{N}$

The proofs are very similar to that for Theorem 1 above, so we omit the details. Note that the conditions for 2 are slightly more complicated than the ones in the other two theorems; the price in this case depends on power, which in turn varies based on the participants in \mathcal{P} .

As with flat, fixed prices, the schemes defined here are sensitive to operator-defined thresholds. If c is set too low, then an equilibrium may not exist because too many users will want to participate. As c is increased, the chances of getting an equilibrium increase. This, however, happens at the expense of setting this bar “too high” and keeping out users who could, potentially, be admitted. Setting this value optimally ultimately requires information about the user preferences, information that may not be available to the operator. As in the previous section, this problem could be addressed with an auction; we leave this as a topic for future work.

4.3 Convergence Algorithms

Given any of the pricing schemes above, a natural question to ask is how the system can reach an equilibrium, provided it exists. In the case of $C_i(\cdot) = C$ or $C_i(\cdot) = cg_i$, this convergence is straightforward; ignoring the case of “ties,” (i.e., $\hat{v}_i = C$ or $\hat{v}_i = cg_i$ for some i), all users with values strictly above their costs enter the system. These users can then use the Foschini-Miljanic update procedure (described previously in Section 2) to reach the corresponding Pareto power point, $\mathbf{p}_{\mathcal{P}}^*$.

The case of $C_i(\cdot) = cp_i$ is trickier because the final $\mathcal{P} / \mathcal{N}$ partition is hard to determine apriori. As is commonly done, we propose here using an iterative, “best response dynamics” algorithm to reach the equilibrium point. In particular, users “update” their powers at discrete intervals, in each instance setting these to maximize their utilities given the currently observed interference:

$$p_i = \operatorname{argmax}_{\bar{p}_i \geq 0} U_i(\bar{\mathbf{p}}_i, \mathbf{p}_{-i}) \quad (10)$$

The updates may be done either synchronously or asynchronously. These types of procedures have been proven to converge in other game-theoretic applications, and we have observed that the algorithm above seems to work well in our environment of study. Proving this rigorously, however, will be addressed in future research.

5. SIMULATED PERFORMANCE

The three pricing schemes above were studied in a simulated wireless network. In a manner similar to [10], this network was constructed by randomly placing transmitter-receiver pairs into a 1000×1000 grid. 25 such pairs were placed, each corresponding to a single system user. See Figure 4.3 above for the resulting network topology. The G_{ii} and G_{ij} parameters were computed assuming a power drop-off proportional to $\frac{1}{d^{\alpha}}$, and the thermal noise terms were set to $\eta_i = 10^{-8}$ for all i .

The user SIR targets were chosen uniformly at random from $[0, 1]$ intervals. Their values were then set according to:

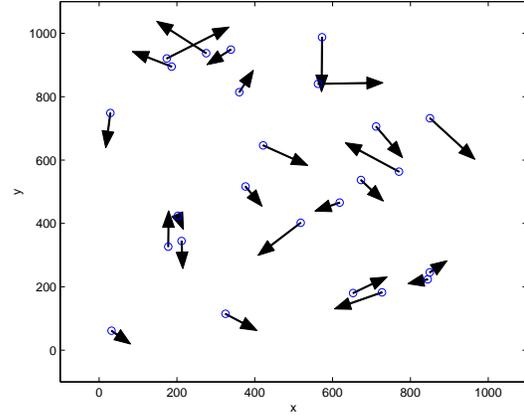


Figure 3: Randomly located links used in simulation. Arrows point from transmitters to receivers.

$$\hat{v}_i = \ln(1 + \bar{\gamma}_i) \quad (11)$$

for all i , thus reflecting the information theoretic capacities of these links at the targeted SIR levels.

Given this setup, we simulated the equilibrium power choices in each scheme as a function of the appropriate cost scaling term (i.e., either C or c). For the cases of flat prices and interference-based charging, this was done by first determining the $\mathcal{P} / \mathcal{N}$ partition as discussed in Section 4.3 above. We then ran the Foschini-Miljanic algorithm among those users in the former set. In the case of power-based pricing, we ran the synchronous, “best response dynamics” algorithm specified in Equation 10 above.

The results of these runs are shown in Figures 4 and 5 below. In cases where a Nash Equilibrium was found, we have plotted the total value of the associated participants, i.e. $\sum_{i \in \mathcal{P}} \hat{v}_i$. This is a good metric for the “quality” or “efficiency” of these outcomes. “Gaps” or undefined regions of these plots correspond to cost choices for which an equilibrium could not be found. For the cases of fixed and interference-based prices, we have also plotted the cardinality of the \mathcal{P} set, i.e. the number of users for whom either $\hat{v}_i > C$ or $\hat{v}_i > cg_i$.

For the first two schemes, we clearly see the “threshold” behavior discussed earlier. When costs are zero, all users want to participate, but no equilibrium exists because the system is too “overloaded.” As these costs are increased, however, users “drop out” in increasing numbers until, finally, a \mathbf{p}^* exists for those remaining. As costs are increased beyond this point, equilibria always exist, but the total value of those participating decreases. Eventually, these admission criteria become too strict for any user, and number of participants drops to 0. Note, however, that the peak total value is much lower in the fixed pricing scheme; by not incorporating power or interference criteria, it does not admit users in a capacity-aware manner.

The results from the power-based pricing scheme also show total value decreasing in cost. Here, however, this curve is not continuous; there are numerous “gaps” where our power control algorithm did not converge. We hypothesize that these correspond to places where, in fact, equilibria do not

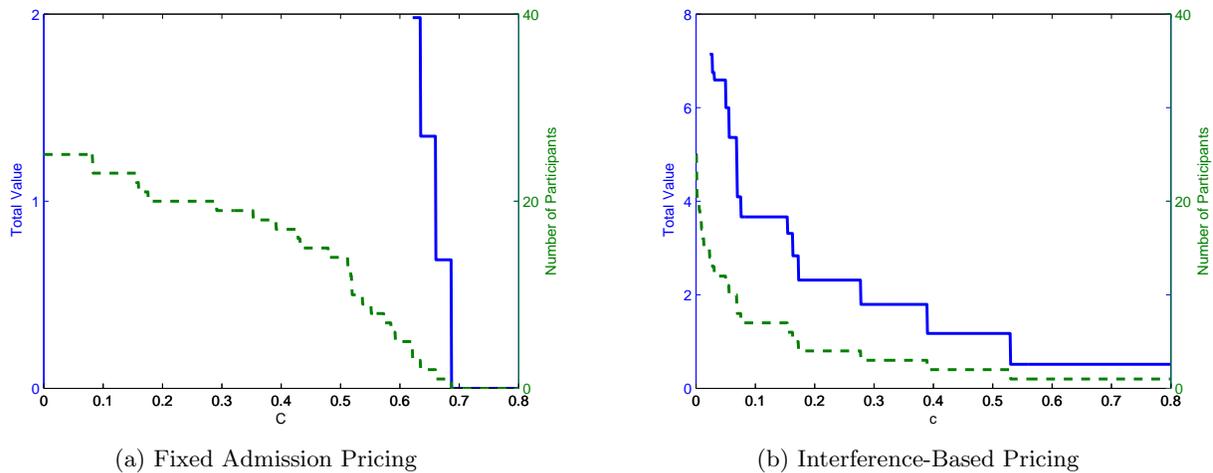


Figure 4: Total user value (solid lines) and number of participants (dotted lines) as functions of C and c scaling terms for fixed and interference-based pricing schemes. The former is plotted only for those costs at which an equilibrium exists; the latter is shown for all costs, even those where “too many” users want to participate and no feasible powers exist. We see that, as expected, the number of participants decreases until, eventually, an equilibrium exists. Beyond this point, the value of those users in equilibrium also decreases until eventually hitting zero.

exist. Unlike the previous two cases, power-based pricing does not admit any clear thresholds. Increasing c may exclude some users, but then allow in others by reducing the powers necessary to reach their targets. Because power is incorporated into the user costs, these participation criteria can be highly non-monotonic, thus leading to the discontinuities observed. Future research will explore this phenomenon further.

These simulations, therefore, show that our pricing schemes have the potential to achieve admission and power control in an incentive-compatible way. By adjusting the appropriate cost parameter, the operator can either raise or lower the bar for admission into the system, affecting both the existence and quality of the resulting equilibria points. Our results here are promising but obviously preliminary; future research will need to run these types of simulations with a wider variety of input parameters.

6. CONCLUSION

In this paper, we have thus looked at the wireless admission control problem from a game theoretic, payment-oriented perspective. Assuming binary utility valuations, we studied three different, price-based admission control schemes—flat pricing, power-based pricing, and interference-based pricing—and characterized the resulting Nash Equilibria for each. These schemes were tested in a simple 25 user, simulated wireless network. While somewhat limited, our simulations show that these approaches have the potential to achieve admission control in a way that is both incentive-compatible and value maximizing.

In future work, we will extend our results here to a wider set of utility function forms. In addition, we propose to more rigorously explore efficiency, convergence, fairness, operator revenue, and other properties of the resulting Nash Equilibria. We will also test our schemes more thoroughly

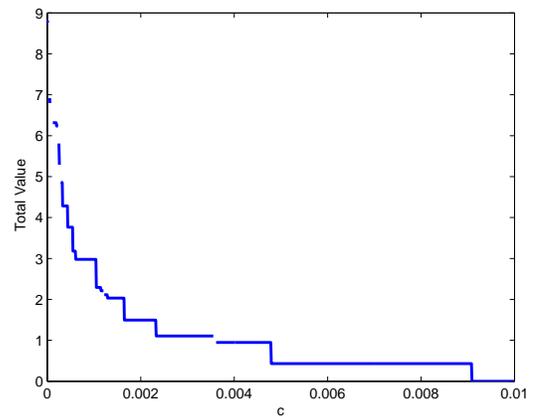


Figure 5: Total user value as a function of c scaling term for power-based pricing scheme. Gaps in the plot correspond to regions where the algorithm specified in Equation 10 did not converge. We hypothesize that the latter points correspond to regions in which no equilibria exist.

by running larger, more complex simulations. This is a very exciting topic of study, particularly as payment systems become increasingly integrated into consumer wireless systems.

7. REFERENCES

- [1] T. Alpcan, T. Basar, R. Srikant, and E. Altman. CDMA uplink power control as a noncooperative game. In *IEEE Conference on Decision and Control*, Dec. 2001.
- [2] N. Bambos, S. Chen, and C. Pottie. Channel access

- algorithms with active link protection for wireless communication networks with power control. *IEEE/ACM Transactions on Networking*, 8(5):583–597, Oct. 2000.
- [3] G. Foschini and Z. Miljanic. A simple distributed autonomous power control algorithm and its convergence. *IEEE Transactions on Vehicular Technology*, 42(4):641–646, Nov. 1993.
- [4] S. Gunturi and F. Paganini. Game theoretic approach to power control in cellular CDMA. In *IEEE Vehicular Technology Conference*, Oct. 2003.
- [5] S. Koskie and Z. Gajic. A Nash game algorithm for SIR-based power control in 3G wireless CDMA networks. In *IEEE/ACM Transactions on Networking*, Oct. 2005.
- [6] T. Kwon, Y. Choi, C. Bisdikian, and M. Naghsineh. Call admission control or adaptive multimedia in wireless/mobile networks. In *ACM International Workshop on Wireless Mobile Multimedia*, Oct. 1998.
- [7] M. Naghsineh and M. Schwartz. Distributed call admission control in mobile/wireless networks. *IEEE Journal on Selected Areas in Communications*, 14(4):711–717, May 1996.
- [8] N. Nisan. Algorithms for selfish agents. In *Symposium on Theoretical Aspects of Computer Science*, Mar. 1999.
- [9] C. Saraydar, M. Mandayam, and D. Goodman. Pricing and power control in a multicell wireless network. *IEEE Journal on Selected Areas in Communications*, 19(10):1883–1892, Oct. 2001.
- [10] B. Yolken and N. Bambos. Target-driven and incentive-aligned power control for wireless networks. In *IEEE Global Communications Conference*, Nov. 2007.