Game Theory in Communications: Motivation, Explanation, and Application to Power Control

Allen B. MacKenzie, Stephen B. Wicker
Cornell University
School of Electrical and Computer Engineering
Ithaca, NY 14583

Abstract— **Game theory is a set of tools developed to model interactions between agents with conflicting interests, and is thus well-suited to address some problems in communications systems. In this paper we present some of the basic concepts of game theory and show why it is an appropriate tool for analyzing some communication problems and providing insights into how communication systems should be designed. We then provided a detailed example in which game theory is applied to the power control problem in a CDMA-like system.**

I. Motivation

Game theory is a tool for analyzing the interaction of decision makers with conflicting objectives. Economists have long used it as a tool for examining the actions of economic agents such as firms in a market. In recent years, it has seen some application by computer scientists to problems such as flow control and routing (e.g. [1] and [2]), but we believe that it can be applied fruitfully to a much broader class of problems in communications systems.

Modern day communications systems are often built around standards. Some such standards are open, such as the TCP/IP standard on which the internet is based. Other standards, such as IS-95 (CDMA), contain intellectual property which must be licensed by the developer. In most cases, though, devices to access these systems are being built by a variety of different manufacturers. In many cases, these manufacturers may have an incentive to develop products which behave "selfishly" by seeking a performance advantage over other network users at the cost of overall network performance. In other cases, end users may have the capability to alter products to behave in a "selfish" manner. Given our reliance on standards, it seems that we should design and build systems that are prepared to cope with users who behave selfishly. If possible, such systems should make selfish behavior unprofitable, so that users will prefer to behave in a manner which is optimal for the system as a whole. When this is not possible, designers should at least be aware of the impact that selfish users would have on the operation of the specified system.

Note that while specifications can establish the "rules of interaction," it is difficult for a specification to enforce a specific algorithm for the end user to execute. For instance, consider a slotted Aloha system. While the system designer can specify that users use slotted Aloha for access to a given channel, it may be impossible for the designer to ensure that every user uses the Pseudo-Bayesian algorithm to estimate the number of backlogged users and choose a retransmit probability. In this case, it is impossible for a central controller to know what retransmit probability the end user is using, making it difficult to enforce such a choice.

Another reason that game theory is an appropriate tool in the setting of communications networks is that game theory deals primarily with distributed optimization — individual users, who are selfish, make their own decisions instead of being controlled by a central authority. Many of the problems which must be solved in a communications system are known to be NP-hard; as a result solving these optimization problems centrally becomes computationally infeasible as network size increases. Because game theory focuses on distributed solutions to system problems, we expect systems designed with game-theoretic concerns in mind to be highly scalable. (For examples of communications problems which are NP-hard, see [3] and [4].)

In some sense, game theory is better suited to solving communication problems — where the "agents" are likely to be computers — than to solving economics problems. One of the main obstacles faced by economists applying game theory to the study of human beings and human institutions is game theory's strong rationality assumption. Game theory typically assumes that all players seek to maximize their utility functions in a manner which is perfectly rational. Obviously, human players are seldom perfectly rational. When the players of a game become computerized agents, though, it is reasonable to assume that the device will be programmed to maximize the expected value of some utility function (at least in so far as this maximization is computationally feasible). Thus the strong rationality assumption seems more reasonable for machines than for people.

II. Explanation

Having shown that game theory may be an appropriate tool to solve some problems in communications systems, we will now present a brief overview of some of the most important concepts of game theory and some particular concepts which will be important in our application.

A *game* has three components: a set of players, a set of possible actions for each player, and a set of utility functions mapping action profiles into the real numbers. We denote the set of players as $\mathcal{I}$ and usually take $\mathcal{I}$ to be finite with $\mathcal{I} = \{1, 2, 3, \ldots, I\}$. For each player $i \in \mathcal{I}$ we denote by A_i the set of possible actions that player i can take, and we let $A = A_1 \times A_2 \times \cdots \times A_I$ denote the space of all action profiles. Finally, for each player $i \in \mathcal{I}$ we let $u_i : A \rightarrow \mathbf{R}$ denote player i's utility function. (At this point, one often defines mixed strategies and utility functions of mixed strategy profiles. Unfortunately, there are some technical problems which must be dealt with when the action sets are uncountable, as they will be in our example. Rather than delve into those issues here, we restrict our attention to pure strategies.) One more notational issue before we move on: Suppose that $a \in A$ is a strategy profile and $i \in \mathcal{I}$ is a player; we let $a_i \in A_i$ denote player i's action in a and a_{-i} denote the actions of the other $I-1$ players.

The most important equilibrium concept in game theory is the concept of Nash Equilibrium. A *Nash equilibrium* is an action profile at which no user may gain by unilaterally deviating. So a Nash equilibrium is a stable operating point

This material is based upon work supported by a National Science Foundation Graduate Research Fellowship, by National Science Foundation Grant NCR-9725251, and by the Multidisciplinary University Research Initiative (MURI) under the Office of Naval Research Contract N00014-00-1-0564.

0-7803-7206-9/01/$17.00 © 2001 IEEE

Fig. 1. An example game in matrix form.

		Player 2		
		L	M	R
Player 1	U	3,2	3,0	1,1
	D	1,0	1,0	2,1

because no user has any incentive to change strategy. More formally, a *Nash equilibrium* is a strategy profile a such that for all $\tilde{a}_i \in A_i$,

$$u(a_i, a_{-i}) \geq u(\tilde{a}_i, a_{-i})[5].$$

Pareto efficiency is another important concept for our application of game theory. An action profile $a \in A$ is said to be Pareto efficient if there is no action profile $\tilde{a} \in A$ such that for all i,

$$u_i(\tilde{a}) \geq u_i(a).$$

with strict inequality for at least one i. In other words, an action profile is said to be Pareto efficient if it is impossible to improve the utility of any player without harming another player.

Simple two-player games are often represented in the form of a matrix. For an example game matrix which we will use to illustrate several of our definitions, see figure 1. In this case, we assume that players choose their moves simultaneously. Player 1 chooses a row, and player 2 chooses a column. The ordered pair in each box represents the payoff which each player receives when that "strategy profile" (choice of row and column) is realized; player 1's payoff is listed first in the ordered pair.

In the game of figure 1, there are two pure-strategy Nash equilibria. The first is (U,L) and the second is (D,R). In addition, note that (U,L) is the only Pareto efficient point in the matrix. Hence, not every Nash equilibrium is Pareto efficient. In some games, in fact, none of the Nash equilibria are Pareto efficient. (See, for instance, the vast literature on the Prisoner's dilemma.)

Repeated games with observable actions are a class of games which has been studied extensively. The basic idea in a repeated game with observable actions is that a simple game, the stage game, is played repeatedly by the same set of players. After each play of the stage game, all of the players learn what strategies were chosen by their opponents in the last round. As a result, players can condition their choice of strategies on past actions of their opponents. This gives rise to an enlarged strategy space; a strategy, s_i for player i is now a map from the set of possible histories, H, to the set of actions for player i, A_i.

Consider a repeated version of the game in figure 1. In this case, the following set of strategies form a Nash equilibrium: Player 1 always plays U. Player 2 plays L unless player 1 has played D in a previous period, in which case player 2 plays M. The reason that this is a Nash equilibrium is simple. If player 1 always plays U, then playing L except after player 1 has played D is an optimal response for player 2. Similarly, if player 2 will always play L unless player 1 has played D in the past, then always playing U is an optimal choice of strategy for player 1. Nonetheless, something seems wrong here. What if player 1 "accidentally" plays D in one stage? (Perhaps player 1's action must be communicated through a noisy channel, for instance.) Then player 2 will play M forever after even though she would be better off playing L or R. Player 2's strategy seems foolish off the equilibrium path. Thus, we need a stronger equilibrium concept for repeated games.

After each possible history $h \in H$, the players in essence start a new game, called the *subgame starting at h*. Like any other game, the concept of Nash equilibrium applies to the subgame starting at h. A *subgame perfect equilibrium* is a strategy profile such that for every $h \in H$ the players will play a Nash equilibrium in the subgame starting at h. The game in figure 1 has infinitely many subgame perfect equilibria (as do most repeated games). One example of a subgame perfect equilibrium is the strategy profile in which player 1 always plays U and player 2 always plays L.

The concept of subgame perfection can be extended to a much wider class of games than repeated games with observed actions. For a more thorough treatment of these and other topics in game theory, a simple introductory game theory textbook at an undergraduate level is [6]. An excellent intermediate text is [5]. For those with a lifetime to devote to the study of game theory, we suggest [7].

III. Application

The power control problem in a CDMA-based system is one example of a problem in communication networks that is appropriate for the use of game-theoretic tools. In the power control problem, each user's utility is increasing in her signal-to-interference-and-noise ratio (SINR) and decreasing in her power level. If all other users' power levels were fixed, then increasing one's power would increase one's SINR. In this case, a user could simply trade power for SINR. In a real system, though, raising one's power has other consequences; when a user raises her transmission power, this action increases the interference seen by other users, driving their SINRs down, inducing them to increase their own power levels. Game theory is a good tool for analyzing this situation.

The power control problem for data users in a CDMA-like system was first framed as a game theory problem in [8], [9]. This work was further expanded in [10], [11], [12]. In all of these papers, very similar utility functions are developed and utilized. In this paper we will utilize the same utility function used in [12]. We note, however, that the issue of the proper utility function for data users on a wireless network deserves further research.

In a wireless system, suppose that users transmit information at the rate R bits/second in L bit packets over a spread-spectrum bandwidth of W (Hz). Let p_j be the power transmitted by user j; we assume that users choose their power levels from the set of non-negative real numbers, $p_j \in [0, \infty)$. We define the signal-to-interference-and-noise ratio of user j to be

$$\mathrm{SINR}_j = \gamma_j = \frac{W}{R} \frac{h_j p_j}{\sum_{\forall i \neq j} h_i p_i + \sigma^2}$$

where h_j is the path gain from user j to the base station and σ^2 is the power of the background noise at the receiver. We assume that the background noise is additive white Gaussian noise (AWGN).

Given these preliminaries, the utility function of user j has the unit of bits/J and is expressed by

$$u_j(p_j, \gamma_j) = \frac{R}{p_j}(1 - 2\mathrm{BER}(\gamma_j))^L \quad [12]$$

where $\mathrm{BER}(\gamma_j)$ is the bit error rate achieved by a given transmission scheme.

Fig. 2. An example of a user's utility function.

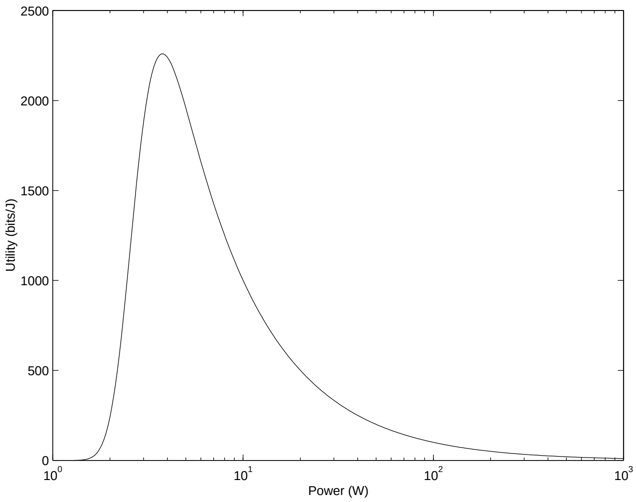

If our transmission scheme is non-coherent frequency shift keying (FSK) in an AWGN channel, then we have

$$u_j(p_j, \gamma_j) = \frac{R}{p_j}(1 - e^{-0.5\gamma_j})^L \quad [12].$$

Observe that $\lim_{p_i \downarrow 0} u_i = 0$ [8]. For our purposes, we define $u_i(0, \gamma_j) = 0$.

An example of one user's utility function, assuming that all other users' transmit powers are fixed, is shown in figure 2. In this figure, the horizontal axis shows the power at which the user transmits, and the vertical axis shows the user's utility. If the user's transmit power is too low, then the user's received power at the base station will be lower than the received powers of other users. This will cause the user's SINR to be low, and hurt the user's performance. This is reflected by the drop in the user's utility function as $p_i \downarrow 0$. If the user's transmit power is too high, then she is squandering precious battery power while having little impact on her bit error rate. This is reflected by the drop in the user's utility function as $p_i \to \infty$.

To this point, we have defined our user's objective functions as well as the space from which they choose their transmit powers (the non-negative real numbers). In the vocabulary of game theory, we have defined the utility functions and the strategy space. Suppose that we let each user unilaterally decide how much power to transmit. The outcome for each user is a function of that user's own decisions as well as the decisions of the other users. What will the users decide to do? The users will attempt to make the best possible choices, taking into account that the other users are doing the same thing. By assumption, our users have complete information about each other. We further assume that our users are completely rational. Then, according to game theory, our users will choose an operating point which is a Nash Equilibrium.

Implicitly assuming a one-shot game, Shah, Mandayam, and Goodman prove that the power control game as described here has a unique Nash Equilibrium [8]. This Nash Equilibrium has the property that all users have the same received power at the base station, and hence all users have the same SINR [8]. In addition to its intuitively appealing "fairness," this property is optimal for despreading the received signals in a CDMA system [13].

Another desirable characteristic of the outcome of a game (or any optimization problem involving several different objective functions) is Pareto efficiency. It is easy to see that if the power control problem were centralized, then the centralized controller would never want to choose an outcome that was Pareto inefficient — a centralized controller would always want to improve the outcome for a user if such an improvement could be made without harming the rest of the users. The Nash Equilibrium of the power control game is shown to be Pareto inefficient in [8].

Here, we will look at two alternative power control games. The first game, which we call the Refereed Game, allows the base station to "referee" the game by punishing users who attempt to "cheat." The second game is a repeated game, in which we assume that the players are not myopic, but consider the impact of their current actions on future play. A third alternative is the pricing game as proposed in [12]; in this scenario, users are charged for the interference that they cause to other users.

A. The Refereed Game

The problem with the Nash Equilibrium solution to the basic power control game is that it is Pareto inefficient. If all users decrease their transmit powers by a factor $\alpha < 1$, then all users can obtain a higher utility [8].

We will show that if the base station acts as a "referee," then it is possible to achieve a solution which is a Pareto improvement over the Nash Equilibrium of the simple pricing game. Furthermore, it is possible to construct a scheme where actual power control decisions are still left to the discretion of individual users. In these cases, the base station's only power control function is to "enforce" the operating point chosen by the users. Numerical simulation suggests that the outcome obtained in the Refereed Game is a Pareto improvement over the best outcome obtained by linear pricing.

The number of ways in which such a system could be designed are vast, so we will only consider a typical example. Suppose that we wish for our system to be "fair" in the sense that all users will have the same received power (and hence the same signal to interference ratio) — a property possessed by the Nash Equilibrium operating point. In this case, we want the system to seek a received power level which optimizes some measure of overall system performance.

Consider the level of received power ($\tilde{p}_i = p_i \cdot h_i$) that an individual user would choose given that all other users will choose their power levels so that they have the same received power at the base station. In this case, if N is the number of users currently in the cell then user j's signal to noise ratio becomes

$$\gamma_j = \frac{W}{R}\frac{\tilde{p}_j}{(N-1)\tilde{p}_j + \sigma^2}.$$

User j's utility function is then

$$u_j(\tilde{p}_j) = \frac{ERh_j}{\tilde{p}_j}(1 - \exp(\frac{-W}{2R}\frac{\tilde{p}_j}{(N-1)\tilde{p}_j + \sigma^2}))^L.$$

Given that all users have the same received power and that $L > 1$, we can show that $u_j(\tilde{p}_j)$ has a single local maximum on $\tilde{p}_j \in (0, \infty)$ and that this local maximum is the global maximum on the domain (see [14]). Furthermore, this maximum is at the same level of received power for **all** users, regardless of their path gain. This is obvious from the fact that h_j appears in $u_j(\tilde{p}_j)$ only as a multiplicative factor. Thus, the same value of $\tilde{p}_j$ will maximize $u_j(\tilde{p}_j)$ regardless of the value of h_j. In other words, all users desire the same operating point given that the system operation is such that all users will have the same received power. Thus, our system should seek this operating point. We will denote the value of $\tilde{p}_j$ which maximizes the above utility function by $\tilde{p}*$.

We believe that the operating point at which all users' transmissions have received power $\tilde{p}*$ is Pareto efficient.

It is difficult to find a closed-form expression for $\tilde{p}*$. It is, however, relatively easy to determine the value of $\tilde{p}*$ numerically. Furthermore, $\tilde{p}*$ depends only on the parameters of the system (the spreading factor, W/R, and the packet length, L), the number of users in the cell, N, and the level of AWGN, σ^2; since these parameters are readily available to the base station, the base station can compute the desired received power and communicate it to the mobile users. Alternatively, the individual users may compute the desired received power themselves, using information provided by the base station.

Once the most desirable operating point has been determined, the system must "punish" any user whose received power is higher than the chosen threshold. Inevitably, some terminals will exceed the threshold unintentionally; for instance, a terminal coming out of a deep fade may underestimate its path gain and transmit with excessive power. The system must handle this situation gracefully; in other words, the punishment should be adequate to ensure that the user would be better off operating at the socially desirable operating point, but not so severe that a user's performance is seriously damaged by a transmission with slightly too much power.

Suppose that user j's transmission has a received power $\tilde{p}_j$ which is x W higher than the target received power $\tilde{p}_t$. This user's signal to noise ratio will increase by

$$\Delta\gamma_j = \frac{W}{R}\frac{x}{(N-1)\tilde{p}_t + \sigma^2}.$$

Recall that our user's bit error rate is $0.5\exp(-0.5\gamma_j)$. Thus, this increment in γ_j will improve user j's BER over the BER of the chosen operating point, multiplying it by a factor of $\exp(-0.5\Delta\gamma_j)$.

In order to "punish" this user for improving her BER to the detriment of other users, the base station can simply increase the errant user's BER by randomly inverting data bits in the user's packet with a certain probability, q_{bi}. Giving the user a BER equal to the desired operating BER is adequate punishment. In this case, the user will have expended extra power in order to obtain the same BER — resulting in a lower utility than would have been obtained at the socially desirable operating point. To obtain this result, the base station should invert user data bits with probability

$$q_{bi} = \frac{\exp(0.5\Delta\gamma_j) - 1}{2(1-\exp(-0.5\gamma_j))}\exp(-0.5\gamma_j).$$

In order for this scheme to be successful, some information must be communicated from the base station to the mobile terminals. First, the base station must inform all users of the target received power at each instant (or must provide them with current system parameters so that they can calculate the target received power themselves). Second, the base station must inform all users of the power level at which their transmissions were received. This is similar to the amount of control information required in both the simple power control game and the power control game with pricing.

B. The Repeated Power Control Game

From a game theory perspective and for the purpose of simplifying base station operations, it is desirable to remove the base station from its role as a referee. In order to do this while maintaining the efficiency of the refereed game, the users will have to discipline themselves. Unlike the base station in the previous section, however, users cannot punish power violations instantaneously. Some period of time will be required for the users to recognize (or be notified) that another user is "cheating" and respond.

Although previous work in [8] has proposed an iterative algorithm to search out the desired operating point, the game-theoretic analysis has assumed that the power control game is a one-shot game. In other words, users are myopic; their only concern is the current value of the utility function. These users are unable to consider the consequences of their actions on future iterations. In order to allow cooperation to develop between users, it is important to model the power control game as a repeated game. By modeling the situation as a repeated game, a user who "cheats" in the current time slot may be punished by the other users in future time slots.

In addition to modeling the situation as a repeated game, we will also require that the power control game have an infinite horizon. In other words, a user must always expect to transmit again in the next period. As we have stated above, punishment in the repeated game will not be instantaneous. If a user knew when her last transmission was coming, then she could exploit this information to "cheat;" her immediate withdrawal from the game would then allow her to go unpunished. The infinite horizon assumption seems reasonable, however. In general, there are two ways that a user might drop out of the game — she can either leave the cell or she can stop actively using the network. Provided that handoff decisions are made by a central decision maker, the user will be unable to predict with precision when she will pass into a new cell. If the network is being used interactively — for example, by surfing the web — then it will be impossible for the mobile terminal to predict conclusively when the user will be done with her task.

We will assume a discrete time model. In each time slot, every user transmits one packet. Furthermore, we assume every user knows the received power of all transmissions in the previous time slots.

Each transmission of a packet gives rise to some utility, which is calculated via the same utility function which is used for the one-shot game. The user values the repeated game by taking a discounted sum of the utilities earned in the transmission of each packet. The discount rate, $\delta \in (0,1)$, is a measure of the value that the user places on the future. If we let $u_{n}{}_{n=0}^{\infty}$ be the sequence of utilities achieved in repeated rounds of the one-shot game (by the transmission of a stream of packets), then the user wishes to maximize the utility for the repeated game, given by

$$u_{rep} = \sum_{n=0}^{\infty} u_n \delta^n.$$

We assume that δ is very close to one. Since packets in a wireless network will come in such quick succession, it seems unlikely that the user's valuation of the current packet will be drastically different than her valuation of the next packet.

Now consider a particular choice of strategies and the resulting system operation. We presume that after each packet is received the system is able to instantaneously announce to all users the number of users in the cell and the power with which each user's transmission was received. We further assume that all of the "cooperating" users are striving for the same operating point as was sought in the refereed game. The desired received power may be either announced by the base station or calculated at the mobile terminal.

As long as no user exceeds the desired received power, the system operates normally. If a user has exceeded the desired receive power, however, then during the next packet period, the rest of the users will punish the wayward user by increasing their powers to the Nash equilibrium of the one-shot game. Once adequate punishment has been dispensed, the system returns to normal. According to our simulations, punishment generally lasts only for the duration of one packet transmission.

As long as no one cheats, the operating point for our repeated game system will be the same as the operating point for the refereed game. In the numerical results of the next section, the same results are shown for both the refereed game and the repeated game. In reality, however, the repeated game system will occasionally have to punish a user, resulting in decreased system performance for everyone. In the repeated game, though, the base station is able to be passive regarding power control — it simply reports to all users the power with which transmissions were received.

Finally, we observe that users who play this strategy are playing a subgame perfect equilibrium.

C. Comparison of Power Control Schemes

We have chosen a static situation in which to compare outcomes produced by our systems with those produced by both the simple power control game and the power control game with pricing, described in [12]. We suppose that 9 users are at the expected distances from the center of a circular cell given that the users are uniformly distributed around the cell of radius 5 km. The expected location of the closest user is 1.42 km from the center of the cell; the expected location of the farthest user is 4.74 km from the center of the cell.

We use the PCS Extension to the Hata Model to predict large-scale propagation loss [15]. We assume a center frequency of 1.9 GHz, an effective mobile height of 2 m, and an effective base station height of 80m. We also assume a packet length of 64 bits ($L = 64$), a transmission rate of 10 kbps ($R = 10000$), and a bandwidth expansion factor of 100 ($W/R = 100$). Finally, we assume the power of the AWGN at the base station is 5×10^{-15} W ($\sigma^2 = 5 \times 10^{-15}$).

Figures 3 and 4 show the results for our scenario. Figure 3 shows the utility versus distance from the base station for this case. From this graph, it is easy to see that the refereed game and the repeated game provide Pareto improvement over the best Pareto improvement which can be achieved with linear pricing. The logarithmic scale allows us to see what happens to the user(s) located farthest from the base station.

IV. Conclusion

We have shown that game theory is an appropriate tool for analyzing a variety of problems encountered in the design and analysis of a communications network. After presenting some basic game-theoretic concepts, we have presented an application to the power control problem in CDMA. Using game-theoretic insights we have demonstrated two power control schemes which obtain outcomes which we believe to be globally optimal. Hence, by using game-theory we have eliminated the threat which a user selfishly controlling her power might pose to the equilibrium operation of a CDMA system with distributed control.

Fig. 3. Comparison of user utility.

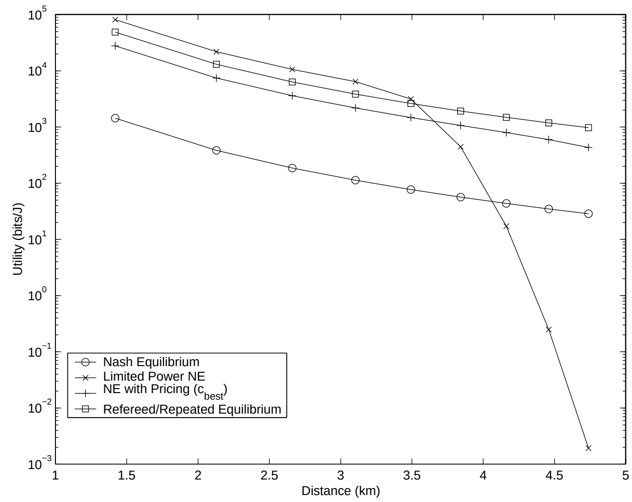

Fig. 4. Comparison of user transmit power.

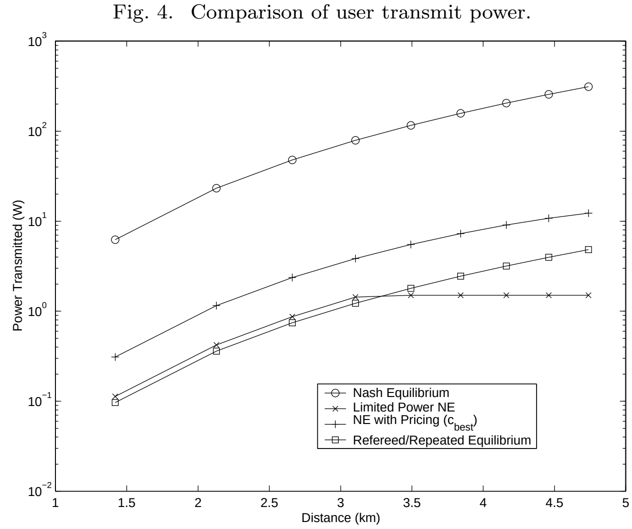

References

[1] C. Douligeris and R. Mazumdar, "A game theoretic approach to flow control in an integrated environment with two classes of users," in *Proceedings of the Computer Network Symposium*, Apr. 1988, pp. 214–221.

[2] T. Roughgarden and E. Tardos, "How bad is selfish routing?," in *Proceedings of the 41st Annual IEEE Symposium on the Foundations of Computer Science*, 2000.

[3] W. K. Hale, "Frequency assignment: Theory and application," *Proceedings of the IEEE*, vol. 68, no. 12, pp. 1498–1573, 1980.

[4] S. Verdu, "Computational complexity of optimum multiuser detection," *Algorithmica*, vol. 4, no. 3, pp. 303–312, 1989.

[5] Drew Fudenberg and Jean Tirole, *Game Theory*, The MIT Press, Cambridge, 1991.

[6] Ken Binmore, *Fun and Games*, D.C. Heath and Company, Toronto, 1992.

[7] J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, revised edition, 1980.

[8] V. Shah, N.B. Mandayam, and D.J. Goodman, "Power control for wireless data based on utility and pricing," in *Nineth IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications*, Sept. 1998, vol. 3, pp. 1427–1432.

[9] Viral M. Shah, "Power control for wireless data services based on

utility and pricing," M.S. thesis, Rutgers, The State University of New Jersey, May 1998.

[10] D. Goodman and M. Mandayam, "Power control for wireless data," in *IEEE International Workshop on Mobile Multimedia Communications*, 1999, pp. 55–63.

[11] D. Goodman and N. Mandayam, "Power control for wireless data," *IEEE Personal Communications*, vol. 7, no. 2, pp. 48–54, Apr. 2000.

[12] C.U. Saraydar, N.B. Mandayam, and D.J. Goodman, "Pareto efficiency of pricing-based power control in wireless data networks," in *1999 IEEE Wireless Communications and Networking Conference*, Sept. 1999, vol. 1, pp. 231–235.

[13] Andrew J. Viterbi, *CDMA: Principles of Spread Spectrum Communication*, Addison-Wesley Publishing Company, New York, 1995.

[14] A. B. MacKenzie and S. B. Wicker, "Game-theoretic approaches to distributed power control in CDMA wireless data networks," unpublished.

[15] Theodore S. Rappaport, Rias Muhamed, and Varun Kapoor, "Propagation models," in *The Mobile Communications Handbook*, Jerry D. Gibson, Ed., chapter 22, pp. 355–369. CRC Press, 1996.