Game Theory and the Design of Self-Configuring, Adaptive Wireless Networks

Allen B. MacKenzie and Stephen B. Wicker, Cornell University

This material is based on 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.

ABSTRACT

Game theory provides a wealth of tools that can be applied to the design and operation of communications systems. In this article, we provide a brief introduction to game theory. We then present applications of game theory to problems in random access and power control. In the case of random access, we examine the behavior of selfish users in a simplified Aloha system; surprisingly, rational selfish users do not implement the "always transmit" strategy that one might expect. In the case of power control, we show that game theoretic techniques can yield an optimal operating point without the intervention of an external controller.

INTRODUCTION

The design and deployment of a cellular network is a tedious process. The propagation characteristics of the area are estimated, base station sites are carefully selected, and the equipment is installed. Once the network is operational, adjustments must often be made, up to and including the placement of additional cell towers and repeaters to eliminate dead spots. Additional adjustments may be necessary as usage patterns evolve. All in all, this is a time-consuming and manpower-intensive series of tasks. If endless funding is available, this may not present a problem, but as the cellular approach gives way to modern, flexible, multiservice wireless networks, financial considerations may require that some of the configuration and adaptation involved in network layout and management be handled automatically. Furthermore, it may not even be possible to determine ideal operating configurations, leaving human intervention not only expensive, but ineffective.

In this article we propose a new approach to wireless network configuration and management that places the decision-making burden on the individual terminals. Control is thus distributed and local, and network scalability is enhanced. In this scenario the network consists of a community of local agents. Design and operational decisions are made without explicit representations of the global environment or even of the other users. Global network configuration and performance are emergent phenomena that are determined solely by the decisions of individual agents. Although our examples here present simple problems at the physical layer (power control) and the medium access control sublayer (random access), we believe that this approach is applicable to a wide range of problems at all layers of a communications network.

The first step in designing such a system is determining the extent of knowledge and range of actions available to the agents. At one extreme we have the tools made available from the study of self-organization in the insect world. *Stigmergy* is the term used for seemingly complicated behavior that emerges from extremely simple local rules. Termites can build a nest, for example, by individually following rules that, should we choose to replicate them in code, would amount to only a few lines of assembly language [1]. At the other extreme we have the tools of game theory, in which individual agents are generally assumed to be completely "rational." In this article we will focus on the game-theoretic approach.

We begin with a brief introduction to game theory. We then apply game-theoretic techniques to two standard problems in wireless networks: random access and power control. We will show that the tools of game theory lead to strategies in which optimal behavior emerges "naturally" from the selfish interests of the agents and the rules of the games.

GAME THEORY

Some concepts of game theory date back centuries, but modern game theory began in the mid-20th century. One of its earliest modern

0163-6804/01/$10.00 © 2001 IEEE

applications was as a tool for modeling decision making by aggressive superpowers. A more enduring application has been as a powerful array of techniques for modeling economic behavior. The basic unit of game theory is, of course, the game. A game has three basic elements:

- A description of strategic interaction between players
- A set of constraints on the actions the players can take
- A specification of the interests of the players

Games are usually represented in one of two forms: the *normal form* and the *extensive form*. The normal form game for two players is represented as a bimatrix, as shown in Fig. 1. An extensive form game is depicted as a tree, where each node represents a decision point for one of the players. The normal form is easier to analyze, but the extensive form captures the structure of a real game in time.

Figure 1 shows our normal-form version of a coordination game known as "The Battle of the Sexes." In this game, Abelard and Eloise would like to attend a concert together. Unfortunately they have different tastes in music: Abelard would prefer to attend a Rolling Stones concert, while Eloise would prefer an opera by Mozart. Both, however, would rather go to either performance together than attend their favorite alone. In the normal form version of this game, the rows represent Abelard's choice of strategies, while the columns represent Eloise's choices. In this case the same strategies are available to both, although this need not be the case. Given strategy selections by both players, we go to the corresponding bimatrix element and read off the payoffs for the two players, with the row player (Abelard) getting the first number and the column player (Eloise) the second. The higher number represents the greater payoff.

The game proceeds by having each player simultaneously announce their choices. In this simple game, we assume that each player is stuck with whatever choice he or she makes. Suppose that Abelard and Eloise both choose to go to the Rolling Stones concert. Note that even if one of them *could* change their choice of strategy at this point, neither would. The (Rolling Stones, Rolling Stones) strategy pair, like the (Mozart, Mozart) pair, is thus a *Nash equilibrium* — a selection of strategies such that neither player can improve his or her payoff by changing strategies while the other players' strategies remain fixed. The Nash equilibrium is thus, in a sense, a stable operating point for a system defined by a game.

We need not limit Abelard and Eloise to the choice of one strategy or the other. Such a choice is called a *pure strategy*, in contrast to a *mixed strategy* that is a probability distribution on a player's available pure strategies. For example, Eloise may decide that she will attend each concert with probability 0.5. To obtain the payoff when one or both players chooses a mixed strategy, we simply compute the expected value of each player's payoff. As an example, suppose that both players choose a 50/50 mixed strategy.

		Eloise	
		Rolling Stones	Mozart
Abelard	Rolling Stones	2, 1	0, 0
	Mozart	0, 0	1, 2

■ Figure 1. *Coordination game with conflicting interests (Battle of the Sexes).*

		Senator Column	
		Don't confess	Confess
Senator Row	Don't confess	1 yr, 1yr	10 yrs, 0 yrs
	Confess	0 yrs, 10 yrs	3 yrs, 3 yrs

■ Figure 2. *The Prisoners' Dilemma.*

An expected value analysis shows that each player can expect a payoff of 0.75.

"The Battle of the Sexes" does have a Nash equilibrium in mixed strategies. In the mixed strategy Nash equilibrium, each player chooses his or her preferred concert with probability 2/3 and chooses the other concert with probability 1/3. This equilibrium gives each player an expected payoff of 2/3.

Note that Nash equilibria do not always entail the same payoffs. The three equilibria we have identified in this game offer three different payoffs to each of the players. The concept of Pareto efficiency can be used to compare different outcomes. An outcome is said to be *Pareto efficient* if it is impossible to increase the payoff of any player without decreasing the payoff of another player. In our example, the pure strategy Nash equilibria are Pareto efficient; the mixed strategy equilibrium is not.

The economist and Nobel Laureate John Nash showed that if each player in an n-player game has a finite number of pure strategies, then the game has a Nash equilibrium in pure or mixed strategies [2, 3]. Nash equilibria are often associated with "rationality." In other words, it would be irrational for players with complete knowledge of the game to choose any combination of strategies that does not constitute a Nash equilibrium. In the realm of economics, people often select strategies that are not rational. Fortunately we need not worry about this issue, for we are interested in programmed agents that will always do what we tell them to do.

Figure 2 shows a two-player normal form game that has achieved great fame as the *Prisoners' Dilemma*. We will give a slightly varied version of the setting. Two senators have been caught accepting bribes. There is not sufficient evidence for a full conviction of both senators, so the Justice Department offers each a deal. Each is told that they may either confess and

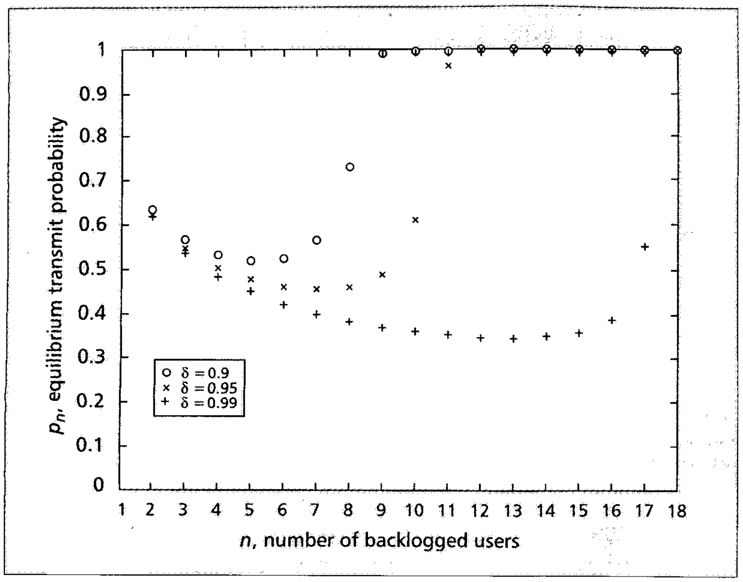

■ **Figure 3.** *Symmetric Nash equilibrium retransmit probabilities for* G(n).

testify against the other (the *defection* strategy) or remain silent and suffer the potential consequences (the *cooperation* strategy). The resulting payoffs are shown in the figure in the form of the resulting prison sentence; obviously, players in this game would prefer lower payoffs.

The original version of the Prisoners' Dilemma is due to Merrill Flood and Melvin Dresher. They developed the game at the RAND Corporation while trying to model the interaction of the then Soviet Union and the United States in a nuclear standoff. Note that this game has only one Nash equilibrium: (confess, confess). Fortunately neither side in the Cold War opted for the Nash equilibrium strategy.

The Prisoners' Dilemma becomes even more interesting if we choose to repeat the game. First consider the case in which the players are told that they will repeat the game a fixed, finite number of times. We assume that their sentences for the entire game are the sum of their sentences at each stage. Such a repeated game is said to have a *finite horizon*. You might imagine that the prisoners would want to cooperate with each other — if they didn't, their counterpart would punish them in the next stage. Due to the finite horizon, this is not the case. In the last game, there is no possibility of punishment in a subsequent game, so the rational players defect. They know this, so at the penultimate game they also defect, and so on back to the first game.

Now suppose that the players do not know when the game is going to end. In this *infinite horizon* game there is always a possibility of punishment, and one Nash equilibrium strategy profile is for each player to always cooperate unless his or her opponent has defected in the past, in which case he or she always defects.

The *utility function* is another concept that is important to applications of game theory. We refer the reader to [2] for a development of von Neumann and Morgenstern utility.

We now have sufficient tools to pursue simple self-configuration and adaptation games in wireless networks. The reader who would like to learn more about game theory should consider Binmore's *Fun and Games* [2] as a starting point. Fudenberg and Tirole's *Game Theory* [3] is an intermediate text that provides a detailed treatment of the field.

Random Access Games

As a first example of a situation for which game theory is an appropriate analysis tool, we consider random access to a communications channel. Users who wish to transmit typically wish to do so as soon as possible. If multiple users try to transmit simultaneously, though, all accesses fail; in addition, unsuccessful attempts to transmit may be costly. The users trying to transmit have conflicting objectives; game theory gives us some insight into this situation. Specifically, we will examine slotted Aloha, one of the best known random access protocols in existence [4, 5].

In slotted Aloha, time is divided into *slots* and via some method of synchronization, all users are presumed to know where the slot boundaries are located. When a user wishes to access the shared channel the user waits until the next slot boundary and then begins attempting to transmit. If two or more users try to transmit in the same slot, the users become "backlogged" and must attempt to transmit again in a future slot. Obviously, if the users use the same deterministic algorithm for determining the slot in which they will retry, they will collide again and again. So the users typically use a random retransmission algorithm.

Most studies of Aloha presume that the system designer dictates the retransmission strategy users will use, but we examine the situation in which users selfishly determine their own retransmission strategies. For our simplified model of Aloha, we will be using an extensive form game we call the "collision game." For now, we assume that players know the number of backlogged users, n. Let $G(n)$ be the game in which there are currently n users backlogged. In each stage of $G(n)$ each of the n backlogged players must decide whether to transmit (T) or wait (W). If one player decides to transmit and the rest decide to wait, the player who transmits will receive a payoff of 1, and each of the other $(n - 1)$ players will play $G(n - 1)$ in the next period. If either no users transmit or more than one user transmits, all players will play $G(n)$ again in the next period. Players place a lower value on payoffs in later stages than on current payoffs. This is represented by a per period discount factor $\delta < 1$. Let $u_{i,n}$ represent user i's utility from playing $G(n)$ and let K be a random variable denoting the number of other users who transmit in a given slot. For $n = 1$ the player should transmit and achieve utility $u_{i,1} = 1$ and for $n > 1$ we express $u_{i,n}$ as a function of player i's action (T) or wait (W)) recursively:

$$u_{i,n}(\mathrm{T}) = P[K = 0] + \delta u_{i,n} P[K > 0]$$

$$u_{i,n}(\mathrm{W}) = \delta u_{i,n-1} P[K = 1] + \delta u_{i,n} P[K \neq 1]$$

That is, for $n > 1$:

That is, for $n > 1$:

$$u_{i,n}(T) = \frac{P[K-0]}{1-\delta \cdot P[K>0]}$$

$$u_{i,n}(W) = \frac{\delta \cdot P[K=1]}{1-\delta \cdot P[K \neq 1]} u_{i,n-1}$$

This game has some very simple asymmetric Nash equilibrium strategies. For instance, if n users are backlogged, user 1 can transmit in period 1, user 2 in period 2, and so on until user n transmits in period n. Such an asymmetric equilibrium is not applicable to Aloha, however, because there is no way for us to distinguish between users. So, from now on we will confine ourselves to searching for symmetric equilibria, in which all users play the same strategy.

In addition, due to the stationary nature of our game (the players face the same situation repeatedly until one user successfully transmits), we will confine our players to the following strategies: each user i selects a vector of transmit probabilities $p_i = (p_{i,1}, p_{i,2}, \ldots)$ where $p_{i,n}$ represents the probability that player i will transmit in a period where $G(n)$ is being played. Note that each player's decision of whether or not to transmit is independent of all other players' decisions.

Even such a simple model produces some interesting insights. Figure 3 shows the transmit probability that produces a symmetric Nash equilibrium for different levels of backlog. This value seems to be unique provided that users at lower levels of backlog are also playing the symmetric equilibrium. Interestingly, while the transmit probability initially decreases with increasing n, it ultimately increases and approaches 1 asymptotically. This is due to the fact that as the number of backlogged users increases, the probability of exactly one user transmitting in a slot goes to zero. Thus, the probability of a collision goes to one, and since transmissions are costless in $G(n)$, users opt to transmit because as the utility of the current game goes to zero, transmitting becomes marginally better than waiting.

So far, our game accounts only for the fact that players wish to transmit sooner rather than later. Certainly transmission, whether or not it is successful, will have a cost if players are wireless users with finite battery life. It is easy to augment our model with a fixed transmission cost, c. We call this new game $\tilde{G}(n)$ where there are n backlogged users. This gives us a new utility function:

$$u_{i,n}(T) = \frac{P[K=0]-c}{1-\delta \cdot P[K>0]}$$

Figure 4 shows that including the transmission cost, c, in our model produces the desirable effect that users decrease their transmission probability as n increases. Preliminary results indicate that a selfish Aloha system achieves stability for sufficiently low attempt rates.

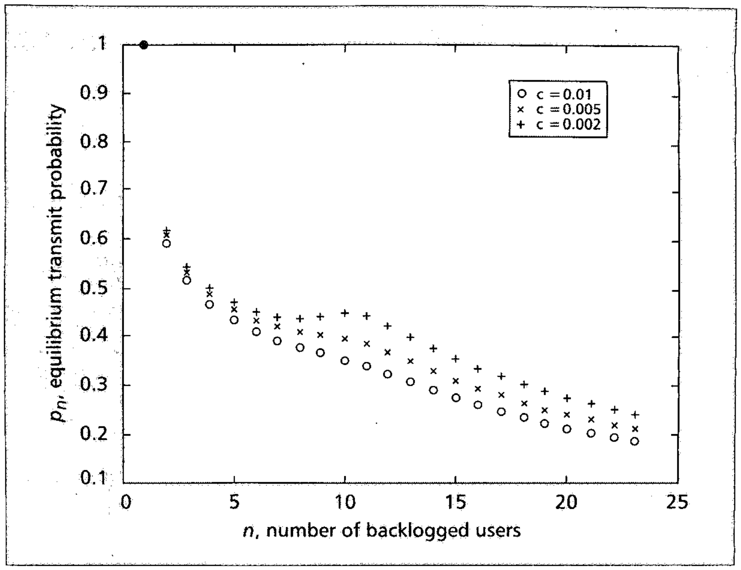

■ **Figure 4.** *Symmetric Nash equilibrium retransmit probabilities for $\tilde{G}(n)$ with transmission cost c. Discount rate fixed at $\delta = 0.95$.*

Power Control Games

The power control problem in a code-division multiple access (CDMA)-like system is a second 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 signal-to-interference-and-noise ratio (SINR) and decreasing in power level. We will model this trade-off with a utility function. If all other users' power levels were fixed, increasing one's power would increase one's SINR. Raising one's power has other consequences, though; when a user raises his or 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 [6, 7]. This work was further expanded in [8–10]. In all of these papers, very similar utility functions are developed and utilized. In this article we will utilize the same utility function used in [10]. We note, however, that the issue of the proper utility function for data users on a wireless network deserves further research.

Suppose that users in a wireless system transmit information at the rate R b/s in L bit packets. Let p_j be the power transmitted by user j; we assume that users choose their power levels from the set of nonnegative real numbers, $p_j \in [0, \infty)$. Finally, let γ_j be the SINR of user j. (Note that user j's SINR is a function of his/her transmitted power, the power transmitted by other users, the amount of background noise, and the path gain between each user and the base station.) If our transmission scheme is noncoherent frequency shift keying (FSK) in an additive white Gaussian noise (AWGN) channel, then the following has been shown to be a useful utility function:

$$u_j(p_j, \gamma_j) = \frac{R}{p_j}(1-e^{-0.5\gamma_j})^L \quad [10].$$

To this point, we have defined our user's utility functions and their strategy spaces. Suppose

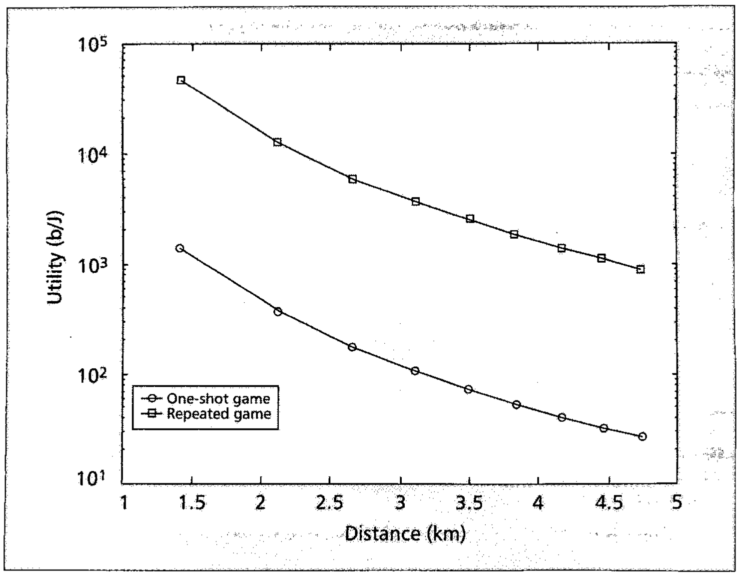

■ **Figure 5.** *A comparison of user utility.*

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? In game-theoretic terms, we have defined a power control game. 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. Then, according to game theory, rational users will choose an operating point that 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 [6]. 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 [6]. In addition to its intuitively appealing "fairness," this property is optimal for despreading the received signals in a CDMA system [11].

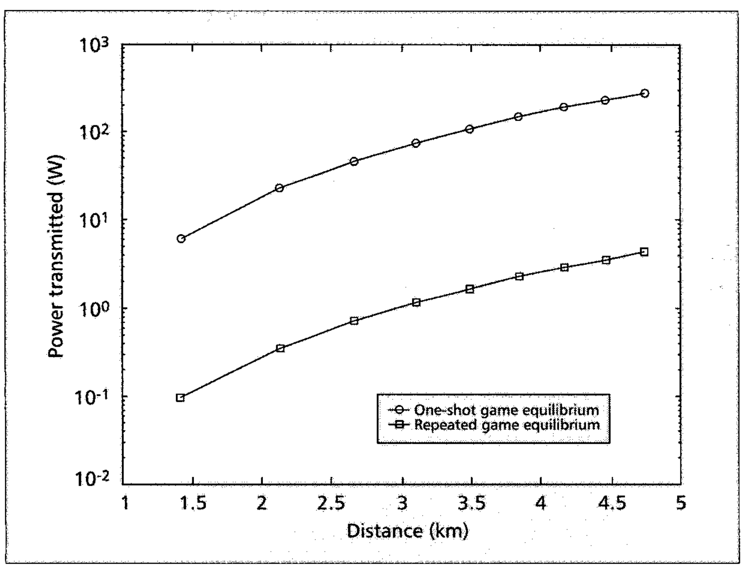

■ **Figure 6.** *A comparison of user transmit power.*

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, 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 [6].

Here, we will look at an alternative power control game. We will model the game as a repeated game, in which we assume that the players are not myopic, but consider the impact of their current actions on future play.

When the power control game is analyzed as a one-shot game, users are myopic; their only concern is the current value of the utility function. By modeling the power control game as a repeated game, we create users who can consider the consequences of their actions. A user who "cheats" in the current time slot may be punished by other users in future time slots.

We will require that our repeated power control game have an infinite horizon. In other words, a user must always expect to transmit again in the next period. Punishment in the repeated game will not be instantaneous. If a user knew when his last transmission was coming, he could exploit this information to cheat; his immediate withdrawal from the game would then allow him to go unpunished. The infinite horizon assumption seems reasonable, however, since a mobile terminal would rarely know when a transmission might end. 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; this information must be broadcast by the base station.

A general strategy for a repeated game specifies the player's (user's) action for each possible game history. Implementing an arbitrary strategy is extremely difficult, though. The usual restriction, then, is that each player's strategy be implementable with a finite-state machine. Each state specifies the strategy that will be played. After each repetition of the constituent game, the outcome of the game determines the transition between states.

Each transmission of a packet gives rise to some utility, which is calculated via the same utility function 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 the user places on the future. 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 the valuation of the next packet.

The repeated game has an infinite number of Nash equilibria. In fact, a well known theorem in game theory says that any feasible set of payoffs in a repeated game can be achieved by a Nash equilibrium of that game provided that δ is close enough to 1 [3]. As system designers, we can choose a "desirable" equilibrium. By the definition of a Nash equilibrium, no individual user will have any incentive to deviate from our chosen equilibrium.

We select our desired equilibrium operating point based on two properties: fairness and Pareto efficiency. We seek an operating point which is fair in the sense that all users will have the same received power at the base station. As noted earlier, this is optimal for despreading [11]. Along this continuum, we choose the received power that gives the users' the highest utility. It turns out that all users' utility functions peak at the same received power level.

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. This is a Nash equilibrium strategy when played against other users using the same strategy. Hence, the best strategy for a user entering a system in which everyone is playing fair and administering punishment to cheaters is to do likewise.

We have chosen a static situation with users distributed uniformly within a 5 km circular cell to compare the operating point of a repeated game with that of a one-shot game. Figures 5 and 6 show the results for our scenario. Figure 5 shows the utility vs. distance from the base station for this case. From this graph, it is easy to see that the repeated game provides Pareto improvement over the one-shot Nash equilibrium. Similarly, Fig. 6 shows that users utilize much lower transmit powers in the repeated game.

Conclusion

We have shown that the use of simple tools from game theory can lead to self-coordinating behavior in relatively complex networks. In particular, we have shown that desirable behavior in power control and random access protocols can be obtained from autonomous selfish agents. In our continuing work, we plan to apply this approach from the physical to the application layer in an effort to develop truly adaptive self-configuring systems. For further details on work mentioned herein, see [12, 13].

References

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, New York: Oxford Univ. Press, 1999.

[2] K. Binmore, *Fun and Games*, Toronto: D. C. Heath and Co., 1992.

[3] D. Fudenberg and Jean Tirole, *Game Theory*, Cambridge, MA: MIT Press, 1991.

[4] N. Abramson, "The Aloha System — Another Alternative for Computer Communications," *AFIPS Conf. Proc.*, 1970, vol. 36, pp. 295–98.

[5] L. G. Roberts, "Aloha Packet System with and without Slots and Capture," Tech. Rep. ASS Note 8, Stanford Res. Inst., ARPA, Network Info Ctr., 1972.

[6] V. Shah, N.B. Mandayam, and D. J. Goodman, "Power Control for Wireless Data Based on Utility and Pricing," *9th IEEE PIMRC*, Sept. 1998, vol. 3, pp. 1427–32.

[7] V. M. Shah, "Power Control for Wireless Data Services Based on Utility and Pricing," M.S. thesis, Rutgers Univ., May 1998.

[8] D. Goodman and M. Mandayam, "Power Control for Wireless Data," *IEEE Int'l. Wksp. Mobile Multimedia Commun.*, 1999, pp. 55–63.

[9] D. Goodman and N. Mandayam, "Power Control for Wireless Data," *IEEE Pers. Commun.*, vol. 7, no. 2, Apr. 2000, pp. 48–54.

[10] C. U. Saraydar, N. B. Mandayam, and D. J. Goodman, "Pareto Efficiency of Pricing-Based Power Control in Wireless Data Networks," *1999 IEEE Wireless Commun. and Networking Conf.*, Sept. 1999, vol. 1, pp. 231–35.

[11] A. J. Viterbi, *CDMA: Principles of Spread Spectrum Communication*, New York: Addison-Wesley, 1995.

[12] A. B. MacKenzie and S. B. Wicker, "Selfish Users in Aloha: A Game Theoretic Approach," to appear, *IEEE VTC Proc.*, Fall 2001.

[13] A. B. MacKenzie and S. B. Wicker, "Game Theory in Communications: Motivation, Explanation, and Application to Power Control," to appear, *IEEE GLOBECOM 2001 Proc.*

Biographies

ALLEN B. MACKENZIE (macken@ece.cornell.edu) received his B.Eng. in electrical engineering from Vanderbilt University in 1999. He is currently a Ph.D. student of Prof. S. Wicker at Cornell University. His research focuses on applications of game theory to communications networks.

STEPHEN B. WICKER (wicker@ece.cornell.edu) received his B.S.E.E. with High Honors from the University of Virginia in 1982. He received his M.S.E.E. from Purdue University in 1983 and a Ph.D. degree in electrical engineering from the University of Southern California in 1987. He is professor and associate director of the School of Electrical and Computer Engineering at Cornell University. He is a co-author of Turbo Coding (Kluwer, 1999) and author of Error Control Systems for Digital Communication and Storage (Prentice Hall, 1995). He is also co-editor of Reed-Solomon Codes and Their Applications (IEEE Press, 1994), and a recent associate editor for Coding Theory and Techniques for IEEE Transactions on Communications. He is currently a member of the Board of Governors of the IEEE Information Society and the DARPA Information Science and Technology study group. He teaches and conducts research in self-configuring wireless networks, distributed artificial intelligence, and error control coding. In 1998 he was awarded the Cornell College of Engineering Michael Tien Teaching Award. In 2001 he was awarded the Ruth and Joel Spira Excellence in Teaching Award for the School of Electrical Engineering.

We have shown that the use of simple tools from game theory can lead to self-coordinating behavior in relatively complex networks. In particular, we have shown that desirable behavior in power control and random access protocols can be obtained from autonomous, selfish agents.