Stability of Slotted Aloha with Multipacket Reception and Selfish Users

Allen B. MacKenzie and Stephen B. Wicker

Abstract

Aloha is perhaps the simplest and most-studied medium access control protocol in existence. Only in the recent past, however, have researchers begun to study the performance of Aloha in the presence of selfish users. In this paper, we present a game-theoretic model of multipacket slotted Aloha with perfect information. We show that this model must have an equilibrium and we characterize this equilibrium. Using the tools of stochastic processes, we then establish the equilibrium stability region for a variety of well-known channel models.

I. Introduction

Aloha and its variants have been central to the understanding of communications network theory for many years. Simple to describe and straightforward to analyze, Aloha is probably the most widely studied medium access control protocol in existence. Furthermore, systems using variants of the Aloha protocol are widely deployed. [I need 2-3 examples of modern systems which use Aloha.]

Aloha was first proposed in [1]; the slotted variation was introduced in [2]. Early research sought methods to stabilize the protocol and provide for retransmission control to make use of available feedback information, for examples see [3], [4].

The operation of slotted Aloha is straightforward. All nodes accessing the medium are synchronized, and time is divided into slots. When a node has a packet to send, it may attempt to transmit it in any slot. Conventionally, a node with a newly arrived packet will attempt to transmit in the first slot after packet arrival; packets being retransmission because of a collision will be transmitted probabilistically. In this paper, however, we assume that newly arrived packets and packets awaiting retransmission are treated identically.

In conventional Aloha models, it was assumed that if exactly one packet was transmitted in a slot then that packet would be received without error; otherwise, all transmitted packets were destroyed. Obviously, this model is somewhat pessimistic as differences in received power, etc., may make it possible for a packet to be received even in the presence of one or more interferers. Furthermore, recent advances in receiver technology, as well as uses of multiple channels, etc., have made it possible for more than one packet to be successfully received simultaniously. This desire for more accurate channel modeling as well as technological progress have led to the development of multipacket reception (MPR) models for Aloha. The most widely-used MPR model was developped by Ghez, Verdu, and Schwartz in the 1980s [5], [6]. We adopt their MPR model in this work.

Despite the bounty of work invested in understanding Aloha, all of the studies of Aloha which we are aware of have ignored the performance of Aloha in the presence of selfish users. In an age of ubiquitous communications terminals and open communications standards, we must begin to treat user devices as selfish agents. Consider the case of the internet — while system designers can recommend various congestion reduction backoff algorithms, actually assuring that end users are running a selected algorithm is nearly impossible. Similar problems exist in most network protocols based on open standards.

The communications research community has failed to investigate the performance of medium access control protocols in the presence of selfish agents. This has lead to the deployment of protocols and systems which can be easily hijacked or manipulated. A better understanding will

1) help us understand our current protocols better and
2) lead to the development of better protocols in the future.

In the remainder of this paper, we will provide a brief introduction to game theory and then present a game-theoretic model for slotted Aloha with multipacket reception. We will then show that an equilibrium of this model must exist, and we will show how we can characterize the equilibria of these games; using this characterization, we will compute (analytically or numerically, depending on the complexity) the stability region for a slotted multipacket Aloha system with selfish users and perfect information.

II. Game Theory

The appropriate tool for the analysis of a system in which a group of users with conflicting interests interact is game theory. Developped principally by economists for the study of the interaction of agents in a market, game theory can also be readily applied to problems in communications systems when the users are thought of as intelligent agents who seek to maximize some measure of their own well-being.

The most basic setting of game theory is the normal form game. Three elements define a normal form game:

- a set of users $\mathcal{I}$ (usually taken to be finite),
- a set of actions for each user A_i, $i \in \mathcal{I}$ which together define a set of possible action profiles $A = \times_{i\in\mathcal{I}} A_i$, and
- a utility function for each user $u_i : A \to \mathbb{R}$.

When the game is played, each player i selects an action from his own set of actions A_i. These selections are made without any knowledge of the selections made by others. The selections of all players taken together define an action profile, $a \in A$, and each player i receives the payoff $u_i(a)$.

Ordinarily, we assume that a player is not limited to choosing actions directly from A_i. Instead, we allow players to choose "strategies" or mixed actions which are probability distributions over A_i. Let $\Sigma_i = \Delta(A_i)$ be the set of probability distributions over A_i. What does it mean to select a strategy which is a probability distribution over actions? We adopt the simple explanation that a player who selects a mixed strategy will use a random device (such as a series of coin flips) to determine which action in A_i that she will play. Now, when each player selects a strategy $\sigma_i \in \Sigma_i$, the action profile is $\sigma \in \Sigma = \times_{i\in\mathcal{I}} \Sigma_i$. What are player utilities if a mixed strategy profile is selected? The usual assumption is that players are expected utility maximizers. That is, we extend the simple definitions of u_i in the most straightforward way. If the A_i are finite, then we have (in a slight abuse of notation)

$$u_i(\sigma) = \sum_{a\in A} \sigma(a) u_i(a).$$

Once such a game has been defined, game theory defines a solution concept which attempts to specify what we should "expect" to occur if rational players play the game. The most widely known solution concept is the Nash Equilibrium. For convenience, we will sometimes write an action profile $a \in A$ as (a_i, a_{-i}) where a_i denotes the action chosen by player i and a_{-i} denotes the actions chosen by everyone else; we will use a similar notation for mixed strategy profiles $\sigma \in \Sigma$. An action profile $a \in A$ is said to be a Nash Equilibrium if for every player $i \in \mathcal{I}$

$$u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i}) \forall a'_i \in A_i.$$

That is, an action profile is a Nash Equilibrium if no player can gain by unilaterally deviating from the specified profile.

An identical definition holds for Nash Equilibria in mixed strategies. A mixed strategy profile $\sigma \in \Sigma$ is said to be a Nash Equilibrium if for every player $i \in \mathcal{I}$

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \forall \sigma'_i \in \Sigma_i.$$

Because of the assumption that players are expected utility maximizers, however, the inequality in this definition is equivalent to the more easily checked inequality

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(a'_i, \sigma_{-i}) \forall a'_i \in A_i.$$

The Nobel prize winning result of John Nash was that for all finite games (games with finite sets $\mathcal{I}$ and A_i), there exists a Nash equilibrium, possibly in mixed strategies.

The results in the remainder of this paper will require a more complex type of game, with a correspondingly more complex notion of an equilibrium. The basic notions of a game, a strategy, and an equilibrium will continue to hold, however.

III. Problem Model

We examine an Aloha system in which selfish users make transmission decisions in an effort to maximize their utility. The channel model we employ is taken from [7]; this model assumes that when the number of successes in a given slot depends only upon the number of transmissions. Specifically, the channel is defined by a MPR matrix

$$E = \begin{bmatrix} \epsilon_{10} & \epsilon_{11} & 0 & 0 & 0 & 0 \\ \epsilon_{20} & \epsilon_{21} & \epsilon_{22} & 0 & 0 & 0 \\ \multicolumn{6}{c}{\dots\dots\dots\dots\dots\dots} \\ \epsilon_{n0} & \epsilon_{n1} & \epsilon_{n2} & . & \epsilon_{nn} & 0 \\ \multicolumn{6}{c}{\dots\dots\dots\dots\dots\dots} \end{bmatrix}$$

where ϵ_{nk} is defined as the probability that k packets are successfully received in a slot where n packets are transmitted. This model is applicable to a wide variety of channels with capture and to some systems using CDMA.

We define the expected number of successes in a transmission of size n to be

$$e_n = \sum_{k=0}^{n} k\epsilon_{nk}.$$

We assume that all users who transmit in a given slot are equally likely to be successful. If n users transmit and k are successful, then the probability that any particular user's transmission is successful is k/n. More usefully, if n users transmit, the probability that a particular user's transmission is successful is given by

$$\sum_{k=0}^{n} \epsilon_{nk}\frac{k}{n} = \frac{e_n}{n}.$$

We assume that users interact as players in a game. A user enters the game when she has a packet to transmit; she leaves the game when her packet has been successfully transmitted. In each slot while she is in the game, the user can choose either to transmit or to wait. We assume that in each slot the users know how many users are currently participating in the game; in other words, it is a game of perfect information.

We assume that users enter the game according to an exogenous random process, and the number of arrivals in each slot are independent and identically distributed random variables with distribution $\alpha \in \Delta(\mathbb{Z}^+)$ where $\mathbb{Z}^+$ is the set of nonnegative integers. Let the expected number of arrivals per slot be denoted λ:

$$\lambda = \sum_{k=0}^{\infty} k\alpha(k).$$

For the moment, we allow for the possability that $\lambda = \infty$.

A user's immediate payoff is determined by whether or not she transmits and whether or not she is successful. We normalize the value of a successful transmission to 1, and we assume that the cost of transmitting is $c \in (0,1)$. So, the immediate payoff from a successful transmission is $1-c$; the payoff from an unsuccessful transmission is $-c$. The payoff from not transmitting in a particular slot is 0.

We further assume that the users share a common per-slot discount rate $\rho \in (0,1)$, and that the goal of each user is to maximize her discounted sum of payoffs. That is, when a user enters the game at time t_0, her goal is to maximize the expectation of $\sum_{t=t_0} \infty \rho^{t-t_0} u_t$ where u_t is the immediate payoff defined in the previous paragraph for each slot t where she is still in the game and is 0 for all slots after she successfully transmits. Finally, we assume that users are expected utility maximizers.

Since we assume perfect information, a strategy in this game is a mapping from the number of users currently in the game to a transmission probability. That is, a strategy is a function $\sigma : \mathbb{Z}^{++} \to [0,1]$ where $\mathbb{Z}^{++}$ is the set of positive integers. We are looking for an equilibrium strategy in the sense of Nash. A strategy σ_0 is said to be an equilibrium strategy in the sense of Nash if given that all other players in the game are playing σ_0, σ_0 is an optimal strategy to play. Note that we are requiring that in equilibrium all players play the same strategy. This reflects a feature of Aloha: the players are indistinguishable.

This game belongs to a general class of games known as Games of Population Transition, which we introduce in [8]. In that work, we provide a general existence proof for games of population transition. Here, we present a special case of the theorem with a proof appearing in the Appendix.

Theorem 1: For any MPR slotted Aloha system with MPR matrix E, any cost parameter $c \in [0,1]$, any arrival distribution α, and any discount rate $\rho \in [0,1)$, there exists an equilibrium strategy σ (not necessarily unique) of the MPR slotted Aloha game.

IV. Analytical Results

In this section we develop the primary results concerning the stability of a MPR Aloha system with selfish users. We begin by providing conditions on the MPR Aloha game under which the equilibrium Markov chain over the number of users in the system, $\{X_t\}_{t=0}^{\infty}$ is irreducible and aperiodic. We then characterize the equilibria of these games and show how to compute the stability region of such a game. In the following section we will show how these results apply to some particular systems.

The main result of the previous section was that an equilibrium exists for any MPR slotted Aloha system. In the current section we wish to characterize those equilibria, particularly with respect to their stability. Observe that the MPR matrix for a channel E, the arrival distribution α, and a particular equilibrium behavior σ define a Markov chain over the number of users currently contending for the channel. Given the current state $n \in \mathbb{Z}^+$, all of the current users transmit with probability $\sigma(n)$. The next state is then determined by the number of successful transmissions and the number of new arrivals. In order to apply many of the results of Markov chain theory, it is necessary that the given Markov chain be irreducible. We provide sufficient conditions on E and α which will guarantee that any Markov chain resulting from an equilbrium strategy σ will be irreducible and aperiodic. Similar conditions are presented in [5], [6]. In that case, however, a fixed retransmission probability is assumed, and specific restrictions are placed on that probability. We place no restrictions on the equilibrium strategy σ other than those which are implied by our conditions on E and α. Furthermore, in [5], [6] it is assumed that all arriving users will transmit in the next slot; we assume that arriving users will use the same retransmit probability as other users.

$$\alpha(0) \neq 0 \tag{1}$$
$$\exists N, \{e_n\}_{n\geq N} \text{is nonincreasing} \tag{2}$$
$$\epsilon_{10} < 1-c \tag{3}$$
$$\tag{4}$$

Condition **??** guarantees that if the Markov Chain is irreducible then it must be aperiodic because it ensures that state 0 is aperiodic. Condition **??** ensures that transmitting with certainty is not a best response in any state $n \geq N$ [check this and explain further]. Finally, condition **??** ensures that in every state it is possible to make ... In addition, condition **??** ensures that $\sigma(n) = 0$ is not a best response for any state n. Suppose not. That is, suppose that $\sigma(n) = 0$ for some n, then in state n a user can obtain a payoff of $\epsilon_{11} - c$ by transmitting. But this is greater than the payoff from waiting; hence, everybody waits is not a best response.

We have now found sufficient conditions on E and α to ensure that the Markov chain on the number of users induced by an equilibrium strategy σ must be irreducible and aperiodic. For the remainder of this section, we assume that those sufficient conditions hold.

For a given MPR matrix which meets the sufficient conditions, we know that there exists N such that $\{e_n\}_{n\geq N}$ is a nonincreasing sequence. Furthermore, we know that for an equilibrium strategy σ, $\forall n \geq N$, $0 < \sigma(n) < 1$. That is, for $n \geq N$, players are mixing between transmit and wait for large enough n. Because our players are expected utility maximizers, for an equilibrium strategy to use a nondegenerate mixing strategy, it must be the case that the payoffs from all strategies in the support of the mixture are the same. In this case, that means that if everyone else is playing σ, then the payoff from transmitting when $n \geq N$ must equal the payoff from waiting. If we write out expressions for these two quantities, then we have for $n \geq N$,

$$\sum_{k=0}^{n-1} \binom{n-1}{k} \sigma(n)^k (1-\sigma(n))^{n-1-k} \frac{e_{k+1} + (k - e_k)}{k+1} - c \tag{5}$$
$$\tag{6}$$

When we found conditions for the irreducibility and aperiodicity of the Markov Chain over the number of users contending for the channel, we showed that under these conditions the value function $v(n)$ must be non-increasing. Provided that the sequence $\{e_n\}$ is bounded, it is straightforward to show that we must also have $lim_{n\to\infty} v(n) = 0$. It follows that in the limit as $n \to \infty$, most of the terms in equation **??** go to zero. Let $\gamma = lim_{n\to\infty} n\sigma(n)$, supposing that the limit on the right-hand side exists. Then in the limit as $n \to \infty$, we can rewrite the first term on the left-hand side of equation **??** and apply the Poisson approximation to obtain

$$\sum_{k=1}^{\infty} \frac{e^{-\gamma}\gamma^{k-1} e_k}{k!} = c$$

For a particular channel model and a particular value of the cost parameter c, it is possible to solve this equation for γ; we call the solution $\hat{\gamma}$.

Once the value of $\hat{\gamma}$ is known, one can use this value to compute the throughput of the system as the number of users becomes large. Let D_n denote the expected drift of the Markov chain representing the number of users contending for the channel when there are n users contending. We can break this expected drift into two pieces: a positive component, representing new arrivals, and a negative component, representing departures. Hence we have the following expression for D_n:

$$D_n = \lambda - \sum_{k=0}^{n} \binom{n}{k} \sigma(n)^k (1-\sigma(n))^{n-k} e_k.$$

In the limit as $n \to \infty$, we can again apply the Poisson approximation to obtain

$$\lim_{n\to\infty} D_n = \lambda - \sum_{k=0}^{\infty} \frac{e^{-\hat{\gamma}}\hat{\gamma}^k}{k!} e_k.$$

Lemma 1: Given an irreducible, aperiodic Markov chain such that

1) $D_n < \infty$ for all n and
2) $\lim_{n\to\infty} < 0$

then the Markov chain is positive recurrent [9].

From this equation, using standard results of Markov Chain drift analysis, we can conclude that the Markov Chain will be positive recurrent if

$$\lambda < \sum_{k=1}^{\infty} \frac{e^{-\hat{\gamma}}\hat{\gamma}^k}{k!} e_k.$$

In this section we have shown how to calculate the stability region for a perfect-information MPR Aloha system. Although these equations can be solved analytically for only a few simple channel models, we have found that numerical solutions can be obtained relatively easily. In the next section we apply these results to a number of well-known channel models.

In conventional Aloha analyses, the perfect information analysis is used to provide bounds on performance for the case of imperfect information. Regrettably, this standard argument does not hold in the case of game-theoretic analysis. It is possible to construct simple games in which imperfect information actually improves players' payoffs. Although that seems unlikely in the current case, it is important to investigate the case of imperfect information carefully before attempting to draw conclusions from the bounds derived above for the perfect information case.

V. Example Systems

The capture model described in this paper (taken from [5], [6]) is quite general. In this section we present results for several different channels. Since these models appear elsewhere in the literature, our focus here will be on computing the maximum throughput which can be supported by such a channel if the contending users are selfish.

In all of these examples, we will assume that the arrival distribution α is Poisson with rate λ. We note, however, that the exact distribution is not important provided that the conditions for irreducibility are satisfied; for most of the models here, this means simply that $\exists n > 1$, such that $\alpha(n) > 0$.

A. General Capture

There are several models for capture in Aloha. These models are primarily based on either power discrimination or time descrimination. Our discussion here will focus on power discrimination, but the results for time discrimination are similar. As special cases of this result, we will include the conventional collision channel and the perfect capture channel.

Assume that the system can capture a maximum of one packet per slot, and that if any packet is captured it will be the packet with the highest received power; denote that power P_1. Finally, assume (as in [10], [11], [12]) that whether or not this packet is received depends solely on the second highest received power, denoted P_2. Specifically, assume that the highest powered packet will be received if and only if $P_1/P_2 > K$ where $K \geq 1$ is a system dependent constant. Note that $K = 1$ denotes perfect capture while $K = \infty$ denotes the conventional collision model.

If we make the conventional assumption of power law fading, then the received power will be $P = cr^{-\beta}$ where c and β are system dependent constants and r is the user's distance from the receiver. If r_1 is the distance of the closest user and r_2 is the distance of the second closest user, then the packet will be captured if $r_2 > br_1$ where $b = K^{1/\beta}$.

Assume that all transmitting users are distributed uniformly in a circle of radius 1, and that the positions of the transmitting users are independent from one slot to the next. If there are k transmitting users in a slot, what is the probability that a packet is captured? When the users are distributed uniformly over a circle of radius 1, their distances from the center are distributed as a random variable R with probability density function $f_R(r) = 2r$. The probability of capture can then be easily computed as 1 if $k = 1$ and $1/b^2$ otherwise.

Hence we have

$$E = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 1-1/b^2 & 1/b^2 & 0 & 0 & 0 & 0 \\ 1-1/b^2 & 1/b^2 & 0 & 0 & 0 & 0 \\ \multicolumn{6}{c}{\dotfill} \end{bmatrix}.$$

Hence we have $e_1 = 1$ and $e_k = 1/b^2$ for $k \neq 1$.

We substitute these values into equation IV in an attempt to solve for $\hat{\gamma}$. This equation does not yield to analytical solution methods for arbitrary values of b.

Note that the conventional collision channel, $b = \infty$, does not satisfy the sufficient conditions for irreducibility which we provided. It is straightforward to show, however, that for $n > 1$ any equilibrium must have $\sigma(n) < 1$. Hence the Markov chain will be irreducible and aperiodic, as desired. When $b = \infty$, we can solve equation IV to obtain $\hat{\gamma} = -\ln c$. Hence the system will be stable provided that $\lambda < -c \ln c$.

For the perfect capture channel, $b = 1$, solving equation IV is impossible using standard functions. Using Mathematica, however, it is possible to solve the equation to obtain $\hat{\gamma} = 1/c + ProductLog(-e^{-1/c}/c)$, where $ProductLog(x)$ is the principle solution for w of $x = we^w$. Hence the system will be stable for arrival rates $\lambda < 1 + cProductLog(-e^{-1/c}/c)$.

The curves in figure **??** show the maximum throughput bound as a function of c for several different values of b including the perfect capture ($b = 1$) and conventional collision ($b = \infty$) models. Not surprisingly, as $c \to 0$ the throughput bound goes to $1/b^2$ because $\sigma(n) \to 1$ as transmitting becomes costless. Also as expected, as $c \to 1$, the throughput bound goes to zero. When the cost of transmission approaches the value of a successful transmission, users will not transmit unless the probability of success is extremely high.

Another somewhat surprising result is that at $c = 1/e$ the conventional collision channel is stable for all $\lambda < 1/e$. Hence for this particular value of the parameter c the selfish Aloha channel supports the same throughput as Aloha with unselfish users.

The maximum throughput under optimal (total throughput maximizing) control with perfect information is computed to be $1/b^2 + (1 - 1/b^2)e^(- b^2/(b^2 - 1))$ in [7]. From figure **??**, it appears that for each value of b, there exists some value of c which obtains the same maximum throughput as the optimal controller.

This observation leads to a design suggestion. Suppose that c, the cost of transmission is fixed relative to the value of a successful transmission, is fixed. Then the designer should seek a technology with a capture parameter b such that c is the cost at which the throughput bound for b is maximized. Regretably, changes that decrease b (i.e. decrease the SINR at which a packet can be successfully decoded) usually increase the cost of transmission c. Nevertheless, this simple example shows that it may be possible for a designer to adjust system parameters in order to make a system perform better in the presence of selfish users.

B. q-frequency Hopping Model

The q-frequency hopping model assumes that q conventional collision channels exist in a system. After choosing to transmit, a user selects one of the q channels on which to transmit, all with equal probability. If exactly one user transmits on a given channel, then the user on that channel is successful; if more than one user transmits on a given channel, then no users on that channel are successful.

Not surprisingly, adding conventional collision channels increases the capacity of the system linearly.

C. CDMA Model

Although similar to the q-frequency hopping model, with the prevelance of direct sequence spread spectrum communications systems, it is worthwhile to consider a CDMA Aloha model. We assume that all transmissions are spread with a spreading factor of W [is this an appropriate symbol?] using one of N different available spreading codes. [Do we need more detail about the spreading codes?] When a user decides to transmit a packet, she selects one of the available spreading codes at random. The most interesting feature of this model is its soft failure. Using the results of [], we presume that the bit error rate falls off as the number of transmitters increases.

VI. Conclusions

In this paper we have found the stability region for a slotted Aloha system with multipacket reception and selfish users for the case of perfect information. We have computed this region for a variety of common channel models and have shown that while the stability region is dependent on the cost parameter, c, in some cases it may be as large as the stability region of a centrally controlled system.

Although we caution that these results do not necessarily provide bounds on stability for the case of imperfect information (e.g. ternary feedback), we believe that they are still useful in cases where the number of users contending for the channel can be reliably estimated.

In addition, we observed that the results for the general capture channel provide some insight for system designers attempting to design for selfish users. We expect that further study will reveal other ways to

The next research challenge in this area concerns the development of results for the case of practical interest — the case of imperfect information. An example of this is ternary feedback. Preliminary results suggest the existence of equilibria in such games, but the computation of equilibria is much more difficult than in the case of complete information.

We further contend that a similar approach could be profitable for investigating other medium access control protocols as well as the interaction of nodes at higher layers of the protocol stack.

References

[1] N. Abramson, "The Aloha system — another alternative for computer communications," in *AFIPS Conference Proceedings*, 1970, vol. 36, pp. 295–298.

[2] L.G. Roberts, "Aloha packet system with and without slots and capture," Tech. Rep. ASS Note 8, Stanford Research Institute, Advanced Research Projects Agency, Network Information Center, 1972.

[3] B. Hajek and T. Van Loon, "Decentralized dynamic control of a multiaccess broadcast channel," *IEEE Transactions on Automated Control*, vol. 27, pp. 559–569, 1982.

[4] R. L. Rivest, "Network control by bayessian broadcast," Tech. Rep. MIT/LCS/TM-285, MIT Laboratory for Computer Science, 1985.

[5] Sylvie Ghez, Sergio Verdu, and Stuart C. Schwartz, "Stability of multipacket aloha," in *Proceedings of the Twenty-first Annual Conference on Information Sciences and Systems*, Baltimore, 1987, pp. 262–267, The Johns Hopkins University.

[6] Sylvie Ghez, Sergio Verdu, and Stuart C. Schwartz, "Stability properties of slotted aloha with multipacket reception capability," *IEEE Tranactions on Automatic Control*, vol. 33, no. 7, pp. 640–649, July 1988.

[7] Sylvie Ghez, Sergio Verdu, and Stuart C. Schwartz, "Optimal decentralized control in the random access multipacket channel," *IEEE Transactions on Automatic Control*, vol. 34, no. 11, pp. 1153–1163, Nov. 1989.

[8] A. B. MacKenzie and S. B. Wicker, "Games of population transition and medium access control," in preparation.

[9] A.G. Pakes, "Some conditions for ergodicity and recurrence of markov chains," *Operations Research*, vol. 17, pp. 1058–1061, 1969.

[10] N. Abramson, "The throughput of packet broadcasting channels," *IEEE Transactions on Communications*, vol. COM-25, pp. 117–128, 1977.

[11] L. G. Roberts, "Aloha packet system with and without slots and capture," *Computer Communications Review*, , no. 5, pp. 28–42, 1975.

[12] C. Namislo, "Analysis of mobile radio slotted aloha networks," *IEEE Journal on Selected Areas in Communications*, vol. SAC-2, pp. 583–588, 1984.

APPENDIX

Proof: In this appendix we will use a standard fixed point argument to show that the MPR Aloha game must have an equilibrium. As a side-effect, we will also prove the existence of a value function, $v : \mathbb{Z}^{++} \to \mathbb{R}$ representing the expected discounted value to a player of being in the game when there are n players contending for the channel.

Suppose that all players in the game are playing a particular strategy σ. We begin by constructing a value function for this strategy, which we denote v_σ. Suppose that we begin with n players in the game; if all players transmit with probability $\sigma(n)$, then the expected immediate reward for the players is given by

$$\sum k = 0n \binom{n}{k} \sigma(n)^k (1 - \sigma(n))^{n-k} \frac{e_k}{k} - \sigma(n)c.$$

[I can simplify this argument a great deal, I believe. All I have is a Markov chain with a known bounded immediate reward function. I believe that there is a basic theorem from Markov chain theory that provides for the existence of a reward function.]

For convenience, we define the following two functions. $u_{\sigma,T}(n)$ is the expected reward from transmitting when there are n users contending for the channel and everyone else is playing strategy σ. Similarly, $u_{\sigma,W}(n)$ is the expected utility from waiting.

$$u_{\sigma,T}(n) = \tag{7}$$
$$u_{\sigma,W}(n) = \tag{8}$$

Now, given the value functions v_σ we will show how to compute a best response to σ. For all σ and all n let

$$\Phi\sigma(n) = \arg \max_{p \in [0,1]} p u_{\sigma,T}(n) + (1-p) u_{\sigma,W}(n).$$

Observe that due to the very simple linear nature of this maximization problem, the value of this correspondance will be either 0, 1, or the whole interval $[0, 1]$. Let $\mathcal{S}_0$ be the set of all possible strategies. In order to prove our result, we wish to show that the correspondance $\Phi : \mathcal{S}_0 \rightrightarrows \mathcal{S}_0$ satisfies the requirements of the Glicksburg-Fan Fixed Point Theorem.

Lemma 2: Given an upper semi-continuous point to convex set correspondance $\Phi : S \rightrightarrows S$ of a convex compact subset S of a convex Hausdorff linear topological space into itself there exists a fixed point $x \in \Phi(x)$.

Let $\mathcal{S}$ be the set of all functions from $\mathbb{Z}^{++}$ to $\mathbb{R}$ and endow this set with the sup norm, $||s||_{\mathcal{S}} = \sup_{n \in \mathbb{Z}^{++}} s(n)$. Then $\mathcal{S}$ is a convex Hausdorff topological vector space.

Now let $\mathcal{S}_0$ be the set of all functions from $\mathbb{Z}^{++}$ to $[0, 1]$. This is a convex compact subset of $\mathcal{S}$.

We have already seen that Φ is a point to convex set correspondance. All that remains to ensure the existence of a fixed point via Glicksberg-Fan is to show that Φ is upper semi-continuous. Suppose not. That is, suppose that there are sequences σ^k and $\tilde{\sigma}^k$, such that $\sigma^k \to \sigma$, $\tilde{\sigma}^k \to \tilde{\sigma}$, $\tilde{\sigma}^k \in \Phi\sigma^k$, but $\tilde{\sigma} \notin \Phi\sigma$. Then there exists n such that $\tilde{\sigma}(n)$ is not a best response to $\sigma(n)$ under valuation v_σ. But we know that some best response exists, hence there exists a $p \in [0, 1]$ and an $\epsilon > 0$ such that

$$p u_{\sigma,T}(n) + (1-p) u_{\sigma,W}(n) > \tilde{\sigma}(n) u_{\sigma,T}(n) + (1 - \tilde{\sigma}(n)) u_{\sigma,W}(n) + 2\epsilon$$

Now for k sufficiently large, we have ■