Stochastic Loss Aversion for Random Medium Access

George Kesidis⋆ and Youngmi Jin⋆⋆

[1] CS&E and EE Depts, Penn State University
gik2@psu.edu
[2] EE Dept, KAIST, South Korea
youngmi_jin@kaist.ac.kr

Abstract. We consider a slotted-ALOHA LAN with loss-averse, non-cooperative greedy users. To avoid non-Pareto equilibria, particularly deadlock, we assume probabilistic loss-averse behavior. This behavior is modeled as a modulated white noise term, in addition to the greedy term, creating a diffusion process modeling the game. We observe that when player's modulate with their throughput, a more efficient exploration of play-space (by Gibbs sampling) results, and so finding a Pareto equilibrium is more likely over a given interval of time.

Keywords: ALOHA MAC, Pareto equilibria, diffusion machine.

1 Introduction

The "by rule" window flow control mechanisms of, *e.g.*, TCP and CSMA, have elements of both proactive and reactive communal congestion control suitable for distributed/information-limited high-speed networking scenarios. Over the past ten years, game theoretic models for medium access and flow control have been extensively explored in order to consider the effects of even a single end-user/player who greedily departs from such prescribed/standard behaviors [1,6, 9,13–16,23–25,28]. Greedy end-users may have a dramatic effect on the overall "fairness" of the communication network under consideration. So, if even one end-user acts in a greedy way, it may be prudent for all of them to do so. However, even end-users with an noncooperative disposition may temporarily not practice greedy behavior in order to escape from sub-optimal (non-Pareto) Nash equilibria. In more general game theoretic contexts, the reluctance of an end-user to act in a non-greedy fashion is called loss aversion [7].

In this note, we focus on simple slotted-ALOHA MAC for a LAN. We begin with a noncooperative model of end-user behavior. Despite the presence of a stable interior Nash equilibrium, this system was shown in [13,14] to have a large domain of attraction to deadlock where all players' transmission probability is one and so obviously all players' throughput is zero (here assuming feasible demands and throughput based costs). To avoid non-Pareto Nash equilibria, particularly those

⋆ G. Kesidis was supported by American NSF CISE/CNS grant 0916179.
⋆⋆ Y. Jin was supported by the Korean NRF grant number 2010-0006611.

V. Krishnamurthy et al. (Eds.): GameNets 2012, LNICST 105, pp. 236–247, 2012.

involving zero throughput for some or all users, we assume that end-users will *probabilistically* engage in non-greedy behavior. That is, a stochastic model of loss aversion, a behavior whose aim is long term communal betterment.

We may be able to model a play that reduces net-utility using a single "temperature" parameter T in the manner of simulated annealing (*e.g.*, [12]); *i.e.*, plays that increase net utility are always accepted and plays that reduce net utility are (sometimes) accepted with probability decreasing in T, so the players are (collectively) less loss averse with larger T. Though our model of probabilistic loss aversion is related that of simulated annealing by diffusions [10, 29], even with a free meta-parameter (η or ηw below) possibly interpretable as temperature, our modeling aim is not centralized annealing (temperature cooling) rather decentralized exploration of play-space by noncooperative users.

We herein do not model how the end-users will keep track of the best (Pareto) equilibria previously played/discovered[1]. Because the global extrema of the global objective functions (Gibbs exponents) we derive do not necessarily correspond to Pareto equilibria, we do not advocate collective slow "cooling" (annealing) of the equivalent temperature parameters. Also, we do not model how end-user throughput demands may be time-varying, a scenario which would motivate the "continual search" aspect of the following framework.

The following stochastic approach to distributed play-space search is also related to "aspiration" of repeated games [3, 8, 18], where a play resulting in suboptimal utility may be accepted when the utility is less than a threshold, say according to a "mutation" probability [17, 26]. This type of "bounded rational" behavior been proposed to find Pareto equilibria, in particular for distributed settings where players act with limited information [26]. Clearly, given a global objective L whose global maxima correspond to Pareto equilibria, these ideas are similar to the use of simulated annealing to find the global maxima of L while avoiding suboptimal local maxima.

This paper is organized as follows. In Section 2, we formulate the basic ALOHA noncooperative game under consideration. Our stochastic framework (a diffusion) for loss aversion is given in Section 3; for two different modulating terms of the white-noise process, the invariant distribution in the collective play-space is derived. A two-player numerical example is used to illustrate the performance of these two approaches in Section 4. We conclude in Section 5 with a discussion of future work.

2 A Distributed Slotted-ALOHA Game for LAN MAC

Consider an idealized[2] ALOHA LAN where each user/player $i \in \{1, 2, ..., n\}$ has (potentially different) transmission probability v_i. For the collective "play" $\underline{v} = (v_1, v_2, ..., v_n)$, the net utility of player i is

[1] The players could, *e.g.*, alternate between (loss averse) greedy behavior to discover Nash equilibrium points, and the play dynamics modeled herein for breadth of search (to escape non-Pareto equilibria).

[2] We herein do not consider physical layer channel phenomena such as shadowing and fading as in, *e.g.*, [16, 25].

$$V_i(\underline{v}) = U_i(\theta_i(\underline{v})) - M\theta_i(\underline{v}), \tag{1}$$

where the strictly convex and increasing utility U_i of steady-state throughput

$$\theta_i := v_i \prod_{j \neq i}(1 - v_j)$$

is such that $U_i(0) = 0$, and the throughput-based price is M. So, the throughput-demand of the i^{th} player is

$$y_i := (U')^{-1}(M).$$

This is a quasi-stationary game wherein future action is based on the outcome of the current collective play $\underline{v}$ observed in steady-state [5].

The corresponding continuous Jacobi iteration of the better response dynamics is [13, 14, 27]: for all i

$$\frac{\mathrm{d}}{\mathrm{d}t}v_i = \frac{y_i}{\prod_{j \neq i}(1 - v_j)} - v_i \quad =: \quad - E_i(\underline{v}), \tag{2}$$

cf. (6). Note that we define $-E_i$, instead of E_i, to be consistent with the notation of [29], which seeks to minimize a global objective, though we want to maximize such objectives in the following.

Such dynamics generally exhibit multiple Nash equilibria, including non-Pareto equilibria with significant domains of attraction. Our ALOHA context has a stable deadlock equilibrium point where all players always transmit, *i.e.*, $\underline{v} = \underline{1} := (1, 1, ..., 1)$ [13, 14].

3 A Diffusion Model of Loss Aversion

Generally in the following, we consider *differently* loss-averse players. Both examples considered are arguably *distributed* (information limited) games wherein every player's choice of transmission probability is based on information knowable to them only through their channel observations, so that consultation among users is not required. In particular, players are not directly aware of each other's demands (y).

3.1 Model Overview

We now model stochastic perturbation of the Jacobi dynamics (2), allowing for suboptimal plays despite loss aversion, together with a sigmoid mapping g to ensure plays (transmission probabilities) $\underline{v}$ remain in a feasible hyper-rectangle $D \subset [0, 1]^n$ (*i.e.*, the feasible play-space for $\underline{v}$): for all i,

$$\mathrm{d}u_i = -E_i(\underline{v})\mathrm{d}t + \sigma_i(v_i)\mathrm{d}W_i \tag{3}$$

$$v_i = g_i(u_i) \tag{4}$$

where W_i are independent standard Brownian motions. An example sigmoid is

$$g(u) := \gamma(\tanh(u/w) + \delta), \tag{5}$$

where $1 \leq \delta < 2$ and $0 < \gamma \leq 1/(1+\delta)$. Thus, $\inf_u g(u) = \inf v = \gamma(-1+\delta) \geq 0$ and $\sup_u g(u) = \sup v = \gamma(1+\delta) \leq 1$. Again, to escape from the domains of attraction of non-Pareto equilibria, the deterministic Jacobi dynamics (*i.e.*, $-E_i(\underline{v})\mathrm{d}t$ in (3)) have been perturbed by white noise ($\mathrm{d}W_i$) here modulated by a diffusion term of the form:

$$\sigma_i(v_i) = \sqrt{\frac{2h_i(\underline{v})}{f_i(v_i)}},$$

where

$$f_i(v_i) := g_i'(g_i^{-1}(v_i)).$$

For the example sigmoid (5),

$$f(v) = \frac{\gamma}{w}\left(1 - \left(\frac{v}{\gamma} - \delta\right)^2\right).$$

In the following, we will consider different functions h_i leading to Gibbs invariant distributions for $\underline{v}$.

Note that the discrete-time (k) version of this game model would be

$$\begin{aligned} u_i(k+1) - u(k) &= -E_i(\underline{v}(k))\varepsilon + \sigma_i(\underline{v}(k))N_i(k) \\ v_i(k+1) &= g_i(u_i(k+1)), \end{aligned} \tag{6}$$

where the $N_i(k)$ are all i.i.d. normal $\mathsf{N}(0, \varepsilon)$ random variables.

The system just described is a variation of E. Wong's diffusion machine [29], the difference being the introduction of the term h instead of a temperature meta-parameter T. Also, the diffusion function σ_i is player-i dependent at least through h_i. Finally, under the slotted-ALOHA dynamics, there is no function $E(\underline{v})$ such that $\partial E/\partial v_i = E_i$, so we will select the diffusion factors h_i to achieve a tractable Gibbs stationary distribution of $\underline{v}$, and interpret them in terms of player loss aversion.

Note that in the diffusion machine, a common temperature parameter T may be slowly reduced to zero to find the minimum of a global potential function (the exponent of the Gibbs stationary distribution of $\underline{v}$) [20, 21], in the manner of simulated annealing. Again, the effective temperature parameter here (η or ηw) will be constant.

3.2 Example Diffusion Term h_i Decreasing in v_i

In this subsection, we analyze the model when, for all i,

$$h_i(v_i) := \eta y_i (1 - v_i)^2. \tag{7}$$

with $\eta > 0$ a free meta-parameter (assumed common to all players). So, a greedier player i (larger y_i) will generally tend to be less loss averse (larger h_i), except when their current retransmission play v_i is large.

Theorem 1. *The stationary probability density function of $\underline{v} \in D \subset [0,1]^n$, defined by (4) and (3), is*

$$p(\underline{v}) = \frac{1}{Z} \exp\left(\frac{\Lambda(\underline{v})}{\eta Y} - \log H(\underline{v})\right), \tag{8}$$

where: the normalizing term

$$\begin{aligned} Z &:= \int_D \exp\left(\frac{\Lambda(\underline{v})}{\eta Y} - \log H(\underline{v})\right) d\underline{v}, \\ D &:= \prod_{i=1}^n (\gamma_i(-1+\delta_i),\ \gamma_i(1+\delta_i)) \\ \Lambda(\underline{v}) &:= \prod_{i=1}^n \frac{y_i}{1-v_i} - \sum_{j=1}^N \left(\frac{v_j}{1-v_j} + \log(1-v_j)\right) \prod_{i \neq j} y_i \\ H(\underline{v}) &:= \prod_{j=1}^n (1-v_i)^2, \ \textit{and} \\ Y &:= \prod_{j=1}^n y_j. \end{aligned}$$

Remark: Λ is a Lyapunov function of the deterministic ($\sigma_i \equiv 0$ for all i) Jacobi iteration [13, 14].

Proof. Applying Ito's lemma [19, 29] to (3) and (4) gives

$$\begin{aligned} \mathrm{d}v_i &= g_i'(u_i)\mathrm{d}u_i + \frac{1}{2} g_i''(u_i)\sigma_i^2(\underline{v})\mathrm{d}t \\ &= [-f_i(v_i)E_i(\underline{v}) + \frac{1}{2} g_i''(g_i^{-1}(v_i))\sigma_i^2(\underline{v})]\mathrm{d}t \\ &\quad + f_i(v_i)\sigma_i(\underline{v})\mathrm{d}W_i, \end{aligned}$$

where the derivative operator $z' := \frac{\mathrm{d}}{\mathrm{d}v_i} z(v_i)$ and we have just substituted (3) for the second equality. From the Fokker-Planck (Kolmogorov forward) equation for this diffusion [19, 29], we get the following equation for the time-invariant (stationary) distribution p of $\underline{v}$: for all i,

$$0 = \frac{1}{2}\partial_i(f_i^2\sigma_i^2 p) - [-f_i E_i + \frac{1}{2}(g_i'' \circ g_i^{-1})\sigma_i^2]p,$$

where the operator $\partial_i := \frac{\partial}{\partial v_i}$.

Now note that

$$\begin{aligned} f_i^2(v_i)\sigma_i^2(\underline{v}) &= 2h_i(v_i)f_i(v_i) \quad \text{and} \\ g_i''(g_i^{-1}(v_i))\sigma_i^2(v_i) &= 2h_i(v_i)g_i''(g_i^{-1}(v_i))/f_i(v_i) \\ &= 2h_i(v_i)f_i'(v_i). \end{aligned}$$

So, the previous display reduces to

$$\begin{aligned} 0 &= \partial_i(h_i f_i p) - (-E_i f_i + h_i f_i')p \\ &= (h_i \partial_i p + h_i' p + E_i p) f_i, \end{aligned}$$

where the second equality is due to cancellation of the $h_i f_i' p$ terms. For all i, since $f_i > 0$,

$$\begin{aligned} \frac{\partial_i p(\underline{v})}{p(\underline{v})} = \partial_i \log p(\underline{v}) &= -\frac{E_i(\underline{v})}{h_i(v_i)} - \frac{h_i'(v_i)}{h_i(v_i)} \\ &= \frac{1}{\eta Y}\partial_i \Lambda(\underline{v}) + \frac{2}{1-v_i}. \end{aligned} \tag{9}$$

Finally, (8) follows by direct integration. □

Unfortunately, the exponent of p under (7),

$$\tilde{\Lambda}(\underline{v}) := \frac{\Lambda(\underline{v})}{\eta Y} - \log H(\underline{v}), \tag{10}$$

and both its component terms Λ and $-\log H$, remain maximal in the deadlock region near $\underline{1}$. Under first-order necessary conditions for optimality, $\nabla \tilde{\Lambda} = \underline{0}$, demand is less than achieved throughput for every user i:

$$y_i = \frac{v_i \prod_{j \neq i}(1-v_j)}{1 + 2\eta \prod_j (1-v_j)}. \tag{11}$$

3.3 Example Diffusion Term h_i Increasing in v_i

The following alternative diffusion term h_i is an example which is instead increasing in v_i, but decreasing in the channel *idle time* from player i's point-of-view [2, 11],

$$h_i(\underline{v}) := \frac{\eta v_i}{\prod_{j \neq i}(1-v_j)}. \tag{12}$$

That a user would be less loss averse (higher h) when the channel was perceived to be more idle may be a reflection of a "dynamic" altruism [2] (*i.e.*, a player is more courteous as s/he perceives that others are). The particular form of (12) also leads to another tractable Gibbs distribution for $\underline{v}$.

Theorem 2. *Using (12), the stationary probability density function of the diffusion* $\underline{v}$ *on* $[0, 2\gamma]^n$ *is*

$$p(\underline{v}) = \frac{1}{W} \exp(\Delta(\underline{v})) \tag{13}$$

where

$$\Delta(\underline{v}) = \sum_{i=1}^{n} \left(\frac{y_i}{\eta} - 1 \right) \log v_i + \frac{1}{\eta} \prod_{i=1}^{n} (1 - v_i), \tag{14}$$

and W *is the normalizing term.*

Proof. Following the proof of Theorem 1, the invariant here satisfies also satisfies (9):

$$\begin{aligned} \partial_i \log p(\underline{v}) &= -\frac{E_i(\underline{v})}{h_i(\underline{v})} - \partial_i \log h_i(\underline{v}) \\ &= \frac{y_i}{\eta v_i} - \frac{1}{\eta} \prod_{j \neq i} (1 - v_j) - \frac{1}{v_i}. \end{aligned}$$

Substituting (12) gives:

$$\partial_i \log p(\underline{v}) = \left(\frac{y_i}{\eta} - 1 \right) \frac{1}{v_i} - \frac{1}{\eta} \prod_{j \neq i} (1 - v_j).$$

So, we obtain (14) by direct integration. □

3.4 Discussion

Note that if $\eta > \max_i y_i$, then Δ is strictly decreasing in v_i for all i, and so will be minimal in the deadlock region (unlike $\tilde{\Lambda}$). So the stationary probability in the region of deadlock will be low. However, large η may result in the stationary probability close to $\underline{0}$ being very high. So, we see that the meta-parameter η (or ηw) here plays a more significant role (though the parameters δ and γ in g play a more significant role in the former objective $\tilde{\Lambda}$ owing to its global extremum at $\underline{1}$).

For small $\eta < \min_i y_i$, note that $\Delta(\underline{1}) = 0$, *i.e.*, it is not a maximal *singularity* at $\underline{1}$ as $\tilde{\Lambda}$. Also, the differences in role played by η in the two Gibbs distributions (8) and (13) is apparent from the first-order necessary conditions for optimality of their potentials:

$$\begin{aligned} \nabla \Lambda(\underline{v}) &= \underline{0} \Leftrightarrow y_i &= v_i \prod_{j \neq i} (1 - v_j) \\ \nabla \Delta(\underline{v}) &= \underline{0} \Leftrightarrow y_i - \eta &= v_i \prod_{j \neq i} (1 - v_j), \end{aligned}$$

so that here demand is more than achieved throughput. Thus, under the potential Δ, if $0 < \eta < \min_i y_i$, then the Gibbs distribution is maximal at points $\underline{v}$ where the throughputs $\underline{\theta} = \underline{y} - \eta\underline{1}$, *i.e.*, all users' achieved throughputs are less than their demands by the same constant amount η. So, the meta-parameter η may be used to deal with the problem of excessive total demand $\sum_i y_i$.

Finally note that the Hessian of Δ has all off-diagonal entries $1/\eta$ and i^{th} diagonal entry $-(y_i - \eta)/(\eta v_i^2)$. Assume that the reduced demands $\underline{y} - \eta\underline{1}$ are feasible and achieved at $\underline{v}$. If $y_i - \eta > (n-1)v_i^2$ for all users i (again where n is the number of users), then by diagonal dominance, Δ's Hessian is negative definite at $\underline{v}$ and hence is a local maximum there. The sufficient condition of diagonal dominance is achieved in the special case when $v_i < 1/(2n)$ for all i because for all i:

$$y_i - \eta \;=\; v_i \prod_{j\neq i}(1 - v_j) \approx v_i(1 - \sum_{j\neq i} v_j),$$

where the approximation is accurate since $\sum_j v_j < 1/2$ by assumption, and

$$(n-1)v_i + \sum_{j\neq i} v_j < 0.5 + 0.5 = 1,$$

$$\textit{i.e.,} \quad \frac{y_i - \eta}{\eta v_i^2} \;\approx\; \frac{v_i(1 - \sum_{j\neq i} v_j)}{\eta v_i^2} > (n-1)\frac{1}{\eta}.$$

This special case obviously does not include the classical, static choice for slotted ALOHA of $v_i = 1/n$ for all i, which leads to optimal total throughput (for the identical users case) of $1/e$ when n is large.

4 Numerical Example

For an $n = 2$ player example with demands $\underline{y} = (8/15, 1/15)$ and $\eta = 1$, the two interior Nash equilibria are the locally stable (under deterministic dynamics) at $\underline{v}_a^* = (2/3, 1/5)$ and the (unstable) saddle point at $\underline{v}_b^* = (4/5, 1/3)$ (both with corresponding throughputs $\underline{\theta} = \underline{y}$) [13, 14]. Again, $\underline{1}$ is a stable deadlock boundary equilibrium which is naturally to be avoided if possible as both players' throughputs are zero there, $\underline{\theta} = \underline{0}$. Under the deterministic dynamics of (2), the deadlock equilibrium $\underline{1}$ had a significant domain of attraction including a neighborhood of the saddle point $\underline{v}_b^*$.

The exponent of p (potential of the Gibbs distribution), $\tilde{\Lambda}$, for this example is depicted in Figure 1. $\tilde{\Lambda}$ has a shape similar to that of the Lyapunov function Λ, but without the same interior local extrema or saddle points by (11). The extreme mode at $\underline{1}$ is clearly evident.

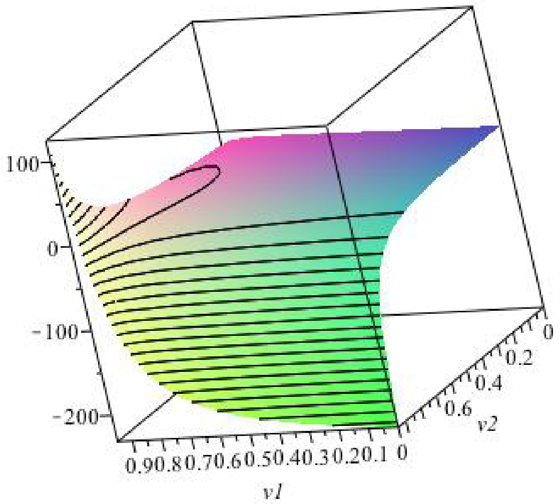

Fig. 1. The potential/exponent (10) of the Gibbs distribution (8) for $n = 2$ players with demands $\underline{y} = (8/15, 1/15)$

4.1 Small η

For the case where $0 < \eta < \min\{y_1, y_2\}$, we took $\eta = 0.01$ for the example above. The potential Δ of the Gibbs distribution (13) is depicted in Figure 2. Compared to $\tilde{\Lambda}$ in Figure 1, $\underline{v} = \underline{1}$ is not a local extremum under Δ (and does

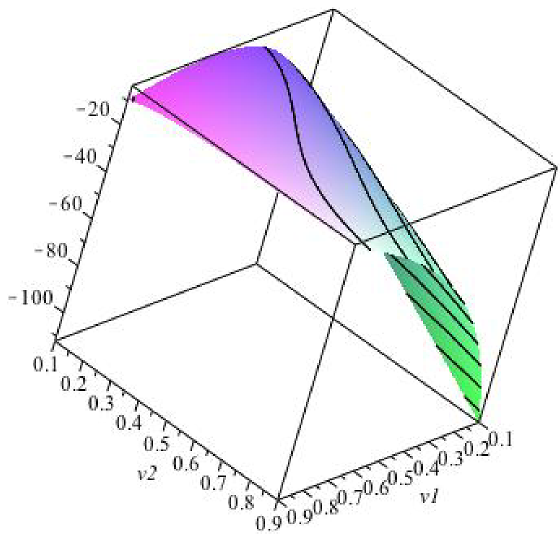

Fig. 2. The potential Δ of (13) for $n = 2$ players with demands $\underline{y} = (8/15, 1/15)$ under (12) with $\eta = 0.01$

not have a domain of attraction). The function Λ under demands $\underline{y} - .01 \cdot \underline{1}$, denoted Λ^* (recall the discussion at the end of Section 3.4), is depicted in Figure 3 and, again, is similar to that depicted in Figure 1. For purposes of reference in these figures, the following table compares these quantities at the points $\underline{v}^*$ that achieve the demands $\underline{y}$ under Λ:

v_1^*, v_2^*	Λ	Δ	Λ^*
$\frac{4}{5}, \frac{1}{3}$	.059	−4.6	.037
$\frac{2}{3}, \frac{1}{5}$	.057	−3.7	.046

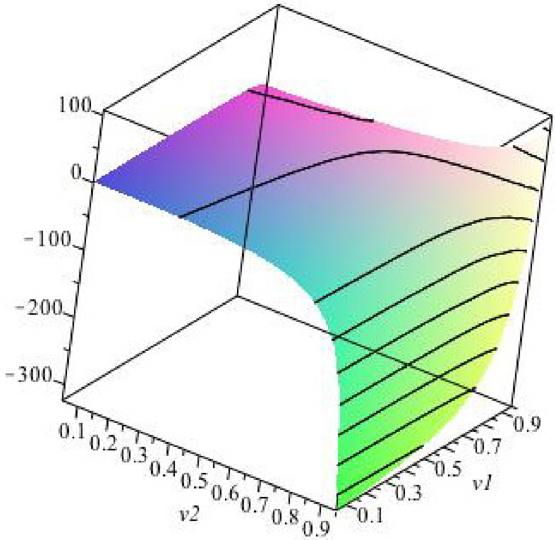

Fig. 3. The component Λ of the potential of (8) for $n = 2$ players with demands $\underline{y} = (8/15, 1/15) - 0.01 \cdot \underline{1}$

4.2 Large η

See [22] for a numerical example of this case, where we illustrate how the use of (12) results in dramatically less sensitivity to the choice of the parameters δ and γ governing the range of the play-space D.

5 Conclusions and Future Work

The diffusion term (12) was clearly more effective than (7) at exploring the play-space, but the interior local maxima of the Gibbs distribution are at points where achieved throughput is less than demand by the "temperature" meta-parameter η. In future work, we plan to explore other diffusion factors h and consider a model with power based costs, *i.e.*, Mv instead of $M\theta$ in the net utility (1). Also, we will study the effects of asynchronous and/or multirate play among the users [2, 4, 15].

References

1. Altman, E., Boulogne, T., El-Azouzi, R., Jiménez, T., Wynter, L.: A survey on networking games in telecommunications. Comput. Oper. Res. 33(2), 286–311 (2006)
2. Antoniadis, P., Fdida, S., Griffin, C., Jin, Y., Kesidis, G.: CSMA Local Area Networking under Dynamic Altruism (December 2011) (submitted)
3. Bendor, J., Mookherjee, D., Ray, B.: Aspiration-based reinforcement learning in repeated interaction games: an overview. International Game Theory Review 3(2&3), 159–174 (2001)
4. Bertsekas, D.P., Tsitsiklis, J.N.: Convergence rate and termination of asynchronous iterative algorithms. In: Proc. 3rd International Conference on Supercomputing (1989)
5. Brown, G.W.: Iterative solutions of games with fictitious play. In: Koopmans, T.C. (ed.) Activity Analysis of Production and Allocation. Wiley, New York (1951)
6. Cagalj, M., Ganeriwal, S., Aad, I., Hubaux, J.P.: On Selfish Behavior in CSMA/CA networks. In: Proc. IEEE INFOCOM (2005)
7. Camerer, C.F., Loewenstein, G.: Behavioral Economics: Past, Present, Future. In: Camerer, C.F., Loewenstein, G., Rabin, M. (eds.) Advances in Behavioral Economics. Princeton Univ. Press (2003)
8. Cho, I.-K., Matsui, A.: Learning aspiration in repeated games. Journal of Economic Theory 124, 171–201 (2005)
9. Cui, T., Chen, L., Low, S.H.: A Game-Theoretic Framework for Medium Access Control. IEEE Journal on Selected Areas in Communications 26(7) (September 2008)
10. Gidas, B.: Global optimization via the Langevin equation. In: Proc. IEEE CDC, Ft. Lauderdale, FL (December 1985)
11. Heusse, M., Rousseau, F., Guillier, R., Dula, A.: Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless LANs. In: Proc. ACM SIGCOMM (2005)
12. Holley, R., Stroock, D.: Simulated Annealing via Sobolev Inequalities. Communications in Mathematical Physics 115(4) (September 1988)
13. Jin, Y., Kesidis, G.: A pricing strategy for an ALOHA network of heterogeneous users with inelastic bandwidth requirements. In: Proc. CISS, Princeton (March 2002)
14. Jin, Y., Kesidis, G.: Equilibria of a noncooperative game for heterogeneous users of an ALOHA network. IEEE Communications Letters 6(7), 282–284 (2002)
15. Jin, Y., Kesidis, G.: Dynamics of usage-priced communication networks: the case of a single bottleneck resource. IEEE/ACM Trans. Networking (October 2005)
16. Jin, Y., Kesidis, G.: A channel-aware MAC protocol in an ALOHA network with selfish users. IEEE JSAC Special Issue on Game Theory in Wireless Communications (January 2012)
17. Kandori, M., Mailath, G., Rob, R.: Learning, mutation, and long run equilibria in games. Econometrica 61(1), 29–56 (1993)
18. Karnadikar, R., Mookherjee, D., Ray, D., Vega-Redondo, F.: Evolving aspirations and cooperation. Journal of Economic Theory 80, 292–331 (1998)

19. Karatzas, I., Shreve, S.E.: Brownian Motion and Stochastic Calculus. Springer (1991)
20. Kesidis, G.: Analog Optimization with Wong's Stochastic Hopfield Network. IEEE Trans. Neural Networks 6(1) (January 1995)
21. Kesidis, G.: A quantum diffusion network. Technical Report 0908.1597 (2009), http://arxiv.org/abs/0908.1597
22. Kesidis, G., Jin, Y.: Stochastic loss aversion for random medium access. Technical report (January 9, 2012), http://arxiv.org/abs/1201.1776
23. Lee, J.W., Chiang, M., Calderbank, R.A.: Utility-optimal random-access protocol. IEEE Transactions on Wireless Communications 6(7) (July 2007)
24. Ma, R.T.B., Misra, V., Rubenstein, D.: An Analysis of Generalized Slotted-Aloha Protocols. IEEE/ACM Transactions on Networking 17(3) (June 2009)
25. Menache, I., Shimkin, N.: Fixed-Rate Equilibrium in Wireless Collision Channels. In: Chahed, T., Tuffin, B. (eds.) NET-COOP 2007. LNCS, vol. 4465, pp. 23–32. Springer, Heidelberg (2007)
26. Montanari, A., Saberi, A.: Convergence to equilibrium in local interaction games. In: FOCS (2009)
27. Shamma, J.S., Arslan, G.: Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria. IEEE Trans. Auto. Contr. 50(3), 312–327 (2005)
28. Wicker, S.B., MacKenzie, A.B.: Stability of Multipacket Slotted Aloha with Selfish Users and Perfect Information. In: Proc. IEEE INFOCOM (2003)
29. Wong, E.: Stochastic Neural Networks. Algorithmica 6 (1991)