A Game for Ad Hoc Network Connectivity in the Presence of Malicious Users

George Theodorakopoulos and John S. Baras
Department of Electrical and Computer Engineering
Institute for Systems Research
University of Maryland
College Park, Maryland 20742
Email: {gtheodor, baras}@isr.umd.edu

***Abstract*— Ad hoc network users are resource constrained: Before transmitting data, they have to take into account the energy expenditure involved. Even if a user is, in principle, willing to spend energy to improve network connectivity, his actual decision will be heavily influenced by the decisions of his neighboring users, since they act as relay nodes for him. Moreover, some of the neighbors may not be as benign as he is; in fact, they could be outright malicious.**

We are abstracting the tradeoff between spending energy and increasing connectivity by modeling the interaction of a user with his one-hop neighbors in a game theoretic fashion. Two types of users exist: Good users willingly trade energy for connectivity, but only if they expect their neighbors to do the same; Bad users try to destroy connectivity, but also lure the Good users to waste energy. Within our model for user behavior and sophistication, we explore outcomes that can arise in this graphical game.

I. Introduction

In ad hoc wireless networks, nodes are facing a tradeoff between preserving their limited energy and promoting the connectivity of the network by forwarding other users' data. Ideally, each node would like to forward only its own traffic, and discard everybody else's, as this would maximize its lifetime and throughput. Unfortunately, if everybody does this, then the network throughput will decrease almost to zero: Only traffic between one hop neighbors will be delivered. So, cooperation to some degree is necessary, if the network is to exist at all.

A growing body of literature, a comprehensive overview of which is in [1], deals with circumstances under which the cooperation between nodes can be sustained. A model used often in this literature (but also more generally in power control in ad hoc networks [2]) is a game theoretical representation of the users in the network (an exception is [3]). The players in game theory attempt to maximize an objective function which takes the form of a payoff. Users make choices and each user's payoff depends not only on his own choice, but also on those of the other users. Hence, in the wireless network context, a user's payoff depends not only on whether he decides to cooperate (by transmitting other users' data) or not, but also on whether his neighbors will decide to cooperate.

Our main contribution is that we distinguish between two types of users: Good and Bad. Good users are cooperative, as long as their neighbors are cooperative. Bad users aim to disrupt the network and waste the energy of the Good users. In other words, we are considering malicious users with interests that are diametrically opposite to those of the Good users.

The literature is almost exclusively considering selfish users, and there is no degree of selfishness that can approximate the payoffs of our Bad users. For example, Félegyházi, Hubaux and Buttyán [1] assume that the payoff function of a user is non-decreasing in the throughput experienced by the user. Our Bad users do not care about their data being transmitted. For the same reason, the model proposed by Urpi, Bonuccelli, and Giordano [4] does not apply (as the authors themselves point out). The only work, to our knowledge, that mentions malicious users in a game theoretic framework is DaSilva and Srivastava [5], and that is not the main focus of their work. They only analyze a simple example in which the Bad users always cooperate and the Good users incur a negative payoff if they cooperate when a Bad user exists in the network.

Another feature of our work is that we take into account the topology of the network, in the sense that a user's actions are influenced only by the actions of his neighbors, and not the whole network's, as is done, for example, in [5] or [6]. The work by Félegyházi, Hubaux and Buttyán [1] is taking into account the actual communication routes (shortest paths) between sources and destinations, so the topology matters even more in their case. This is the next natural step for our work: to see how things change when Bad users influence not just the actions of their immediate neighbors, but also the payoffs and actions of the users that are further away.

The game is played repeatedly in rounds, in which users choose whether to cooperate or not. A minor contribution is that Good users follow the fictitious play process [7]: They estimate the future probability of a neighbor's cooperation by the past frequency of his cooperations (the percentage of rounds in which that neighbor has cooperated so far). Bad users, on the other hand, are assumed to know everything that has ever happened in the network, and are also aware of the decision algorithm that the Good use (i.e. they can predict what the Good will play).

The basic idea behind fictitious play (a player taking into account the frequency of another player's actions so far) is simple, intuitive, and uses only locally available information, so similar things have appeared in the literature. However, the

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

main ingredient of fictitious play is that this frequency is used to predict the future behavior of the other player, an ingredient that does not exist in the related work.

Ref. [1] uses strategies that depend only on what happened in the previous time instant (time slot, according to their terminology). Ref. [6] is using a modified version of Generous Tit For Tat, but they have no notion of topology and, consequently, of neighborhoods. In their setting, each user is comparing his own frequency of cooperation to the aggregate frequency of cooperation of the rest of the network. Altman, Kherani, Michiardi, and Molva [8] proposed a scheme for punishing users whose frequency of cooperation is below the one dictated by a certain Nash equilibrium.

The importance of making the explicit connection to the concept of fictitious play is that many results could be transferred from that area. For example, if all the players have the same utility function, then the fictitious play process always converges to a Nash equilibrium [9].

This paper is organized as follows: We describe our system model in section II. In section III we describe the fictitious play strategy that the Good follow; we analyze and compare two fundamental strategies for the Bad; we find the optimal solution in two example graphs. Section IV concludes.

II. System model

We consider a wireless network of Good and Bad users, modeled as an undirected graph $G = (V, E)$, where V_G is the set of Good users, V_B is the set of Bad users, $V_G \cap V_B = \emptyset$, and $V = V_G \cup V_B$. An edge (i, j) means that i and j are within communication range, but not necessarily that the link is being used for transmission. Each user has a choice between cooperating (C) to activate the adjacent links, or defecting (D), which effectively kills all adjacent links. A link (i, j) is considered to be active if and only if both users i and j choose to cooperate (C). Each user chooses only one action (C or D) for all links, not one action per link.

We now describe what Good and Bad users want to achieve, and how that translates to game theoretic payoffs. The objective of each Good user is to activate the wireless links to as many of his neighbors as possible. So, for each adjacent link that is active, a Good user receives a payoff of N. However, choosing C costs energy, regardless of what the neighbors choose. So, a C incurs a negative payoff of $-E$ per adjacent link. It holds that $N \geq E \geq 0$, otherwise the Good would not have an incentive to play C, so everyone would always play D. In the real operation of the network, a C would correspond to sending/forwarding and receiving traffic; a D would effectively correspond to shutting down communication: neither send, nor receive.

The objective of each Bad user is the exact opposite of the Good users'. They want to disrupt connectivity by keeping links inactive, and they also want the Good users to waste their energy. So, the Bad users incur a negative payoff of $-N$ per active adjacent link, but gain E for each Good neighbor of theirs that plays C. Moreover, they do not spend energy when cooperating. We assume that there are no links between

Good \ Bad	C	D
C	$N - E, E - N$	$-E, E$
D	$0, 0$	$0, 0$

Good \ Good	C	D
C	$N - E, N - E$	$-E, 0$
D	$0, -E$	$0, 0$

Fig. 1. The games on the links for Good versus Bad and Good versus Good neighbors.

Bad users, so they never gain anything in terms of increased connectivity by activating links.

All these considerations are incorporated in the two payoff matrices of Fig. 1. Note that there is no Bad vs. Bad game, since no Bad users are neighbors. Also note that the payoffs for a Good user depend only on the actions of himself and the other player, but not on the type of the other player. This comes from a modeling decision, namely that the Good users do not mind cooperating with Bad users: We look at cooperation only from the point of view of connectivity (traffic forwarding under energy constraints), which means that the Bad users do not have any incentive to cooperate in order to, e.g., get access to Good user communications.

The game is a repeated game with simultaneous moves. It is played in rounds, at every round each player chooses an action independently of all the others, and the action of each user is announced to all his neighbors. Then, the payoffs along each link are computed and each user receives the sum of the payoffs along his adjacent links. After that, the users decide on their action for the next round. The notation for the payoffs of user i will be: $R_i(a_i a_j|x)$, where a_i is i's action, a_j is j's (the neighbor's) action, $a_i, a_j \in \{C, D\}$, and x is i's type ($x \in \{\text{Good, Bad}\}$). If a_j is omitted, we will be referring to the payoff of user i when he plays action a_i. To illustrate, the total payoff of user i for a single round is (N_i is the set of i's neighbors)

$$R_i(a_i|x) = \sum_{j \in N_i} R_i(a_i a_j|x). \tag{1}$$

We are assuming that the Good users are constrained to play according to fictitious play, which means that a Good user's neighbor who has so far played c times C and d times D is assumed to play C with probability $\frac{c}{c+d}$. So, at each round each Good user is choosing the action that maximizes his expected payoff given the probability estimates for the actions of the neighbors. The motivation for considering fictitious play behavior is that we do not want to overburden the users with excessive computation requirements. Moreover, if all players are Good, the fictitious play process converges to a Nash equilibrium [9].

The Bad users, however, are not constrained to play according to fictitious play. They can play according to any strategy they want, and they also know everything that is happening in the network (the actions taken by any player at any previous

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

point in time). They also know the types of everyone else, and they can coordinate their actions with the other Bad users.

The payoff for the repeated game is the average of the per round payoffs. That is, if v_i^t is the payoff for user i at round $t = 0, 1, \ldots$, his payoff for the whole game is $\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} v_i^t$. The main thing we want to find is the maximum total payoff that the Bad can achieve, and the corresponding Good payoff. Remember that the Good maximize their payoff within the constraints of fictitious play, whereas the Bad can do anything they want. The payoff function being the long term average, we only care about the steady state behavior.

III. Analysis

We first look at what the Good players will be doing. The Good want to maximize their expected payoff according to fictitious play. The expected payoffs for each of the possible actions of a Good player are (N_i is the set of i's neighbors, $a_j^t \in \{C, D\}$ is the action of player j at time t):

$$\begin{aligned} R_i(C|G) &= \sum_{j\in N_i} \{\Pr(a_j^t = C) R_i(CC|G) \\ &\qquad + (1 - \Pr(a_j^t = C)) R_i(CD|G)\} \\ &= \sum_{j\in N_i} \left\{\Pr(a_j^t = C) N - E\right\} \\ &= N \cdot \sum_{j\in N_i} \Pr(a_j^t = C) - |N_i| \cdot E \\ R_i(D|G) &= 0. \end{aligned} \tag{2}$$

So, if $\sigma_i(a_i|x)$ is the probability that player i will play action a_i, given that he is of type x (Good or Bad), and $\mathbf{1}\{P\}$ is the indicator function for the predicate P ($\mathbf{1}\{P\} = 1$ iff P is true), the following holds:

$$\begin{aligned} \sigma_i(C|G) &= \mathbf{1}\{R_i(C|G) > R_i(D|G)\} \\ &= \mathbf{1}\left\{ N \cdot \sum_{j\in N_i} \Pr(a_j^t = C) - |N_i| \cdot E > 0 \right\} \\ &= \mathbf{1}\left\{ \sum_{j\in N_i} \Pr(a_j^t = C) > \frac{|N_i| \cdot E}{N} \right\} \end{aligned} \tag{3}$$

But the Good player i does not know the value $\Pr(a_j^t = C), j \in N_i$, so he will estimate it according to the rules of fictitious play. We denote by q_j^t the estimated probability that player j will play C in round $t+1$. For $t = 0, q_j^0 = p$ (prior estimate for all j). For $t \geq 1, q_j^t$ is the fraction of Cs in $a_j^{1\ldots t}$ weighting the prior with M rounds:

$$q_j^t = \frac{Mp + \sum_{n=1}^{t} \mathbf{1}\left\{a_j^n = C\right\}}{M + t}, \tag{4}$$

and the vector of frequencies of the Good (Bad) players at time t is denoted by $\vec{q}_G^{\,t}$ ($\vec{q}_B^{\,t}$). So, a Good player's action a_i^t is determined as follows ($\alpha_i^t = 1$ for C, 0 for D):

$$a_i^t = \mathbf{1}\left\{ \sum_{j\in N_i} q_j^{t-1} > \frac{|N_i| \cdot E}{N} \right\}, \tag{5}$$

and the vector of actions of the Good (Bad) players at time t is denoted by $\vec{a}_G^{\,t}$ ($\vec{a}_B^{\,t}$).

We now turn to what the Bad players will be doing. In general, a strategy for the Bad is a choice of $\vec{q}_B$. This encompasses all possible strategies of the Bad, since the Good follow fictitious play, and therefore only respond to different frequencies: nothing more elaborate will be detected by the Good. The outcome of a particular strategy (i.e., the actions that the Good will choose and the sum of the payoffs of all the Bad users) can be found by setting all the Good to playing C, and iteratively looking for Good whose payoff is negative, setting their action to D, and continuing until there are no changes.

Looking at the payoffs in Fig. 1, we see that in a Good versus Bad game, playing D is a weakly dominating action for the Bad: Regardless of what the Good plays, the Bad never has a reason to prefer C over D. So, a strategy that would make sense for the Bad would be "Always D". However, this may cause the Good to play D themselves in order to save energy, in which case the Bad users' payoff will drop to zero. So, it seems it might pay off for the Bad to play C once in a while so as to keep the Good playing C. We first analyze the "Always D" strategy and then the "Keep the Good at C".

A. Strategy #1: Always D

According to this strategy the Bad always play D, that is, for all t, $\vec{a}_B^{\,t} = \vec{0}$. In this case, we can prove the following theorem:

Theorem 1: If a Good player ever plays D, then he never plays C again.

Proof: Let $q_{N_i}^t = \sum_{j\in N_i} q_j^t$, and $\gamma_i = \frac{E}{N}|N_i|$. Then, (5) becomes $a_i^t = \mathbf{1}\{q_{N_i}^t > \gamma_i\}$. We will show that, for all i, $q_{N_i}^t$ is non-increasing in t, so if it ever goes below γ_i (i.e. user i plays D), then it will stay below γ_i.

We denote by $C_{N_i}^t$ the number of neighbors of user i that choose action C at time t, i.e. $C_{N_i}^t = \sum_{j\in N_i} \mathbf{1}\{a_j^t = C\}$. For simplicity, we take the weight of the prior to be $M = 1$.

Here is how the quantity $q_{N_i}^t$ evolves:

t	$q_{N_i}^t$
0	pN_i
1	$\frac{1}{2}(pN_i + C_{N_i}^1)$
2	$\frac{1}{3}(pN_i + C_{N_i}^1 + C_{N_i}^2)$
$\vdots$	$\vdots$
n	$\frac{1}{n+1}(pN_i + C_{N_i}^1 + C_{N_i}^2 + \cdots + C_{N_i}^n)$
$n+1$	$\frac{1}{n+2}(pN_i + C_{N_i}^1 + C_{N_i}^2 + \cdots + C_{N_i}^n + C_{N_i}^{n+1})$

Assume that the first Good player i who plays D, does so at round n. This means that $q_{N_i}^n < \gamma_i$, and that his neighbors who played D are Bad. Since the Bad players will keep playing D, we have $C_{N_i}^{n+1} \leq C_{N_i}^n$, which implies that $q_{N_i}^{n+1} \leq q_{N_i}^n < \gamma_i$. So Player i will keep playing D forever, and as far as his behavior is concerned, it is as if he is Bad, too. We can then repeat the argument for the remaining Good nodes. ∎

Eventually, this process will converge to an equilibrium, since the number of "Good nodes who still play C" is non-increasing (by Theorem 1) and bounded below by zero.

This Bad user strategy is the one that causes as many Good as possible to play D. That is, "Always D" is the proper strategy for the Bad to destroy the connectivity as much as possible (deactivate as many links as they can). However, our Bad users have a double objective which also includes wasting the energy of the Good.

B. Strategy #2: Keep the Good at C

We now look at what happens when the Bad want to cause as many Good as possible to play C, so as to benefit as much as possible from wasting their energy. This strategy is at the other extreme of the previous one, and we will see that it does not always bring the highest payoff to the Bad either.

Given that the Good will always play C (i.e. $\vec{q}_G = \vec{1}$, which we will see how to guarantee), we want to find the vector $\vec{q}_B$ of frequencies for the Bad players (frequencies of playing C) that maximizes the sum of payoffs of the Bad. The payoff of a Bad player j is (on the average, when his frequency of Cs is q_j):

$$R_j(q_j) = \sum_{i \in N_j} \{q_j(E-N) + (1-q_j)E\} = (E - Nq_j)N_j \tag{6}$$

because we are assuming that a Bad player only has Good neighbors. So, the quantity to be maximized is

$$\sum_{j \in V_B} (E - Nq_j)N_j. \tag{7}$$

Since

$$\arg\max_{\vec{q}_B} \sum_{j \in V_B} (E - Nq_j)N_j = \arg\min_{\vec{q}_B} \sum_{j \in V_B} q_j N_j, \tag{8}$$

we can solve, equivalently,

$$\min_{\vec{q}_B} \sum_{j \in V_B} q_j N_j \tag{9}$$

subject to (condition for the Good to always play C):

$$\sum_{j \in N_i} q_j > \frac{|N_i| \cdot E}{N}, \text{ for all } i \in V_G, \tag{10}$$

and also subject to (condition for each Bad to have a positive payoff):

$$0 \leq \vec{q}_B \leq \frac{E}{N}. \tag{11}$$

The idea is that the Bad will adjust their play so as to keep their frequencies equal to the computed $\vec{q}_B$.

For the example graph in Fig. 2, the linear program for the strategy "Keep the Good at C" would be as follows

$$\min q_1 + q_7 + q_{10} \quad \text{s.t.}$$
$$\begin{array}{rcl} q_1 + 1 & > & 2\frac{E}{N} \\ q_7 + 2 & > & 3\frac{E}{N} \\ q_{10} + 1 & > & 2\frac{E}{N} \\ 0 \leq \; q_1, q_7, q_{10} & \leq & \frac{E}{N} \end{array}$$

and would lead to the all the Good nodes playing C, and the Bad nodes to playing C with the corresponding frequencies.

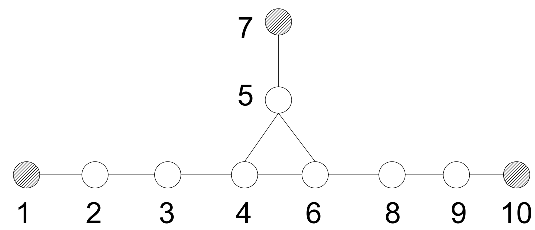

Fig. 2. Example Graph 1. The "Keep the Good at C" strategy is optimal. The Bad users are shaded.

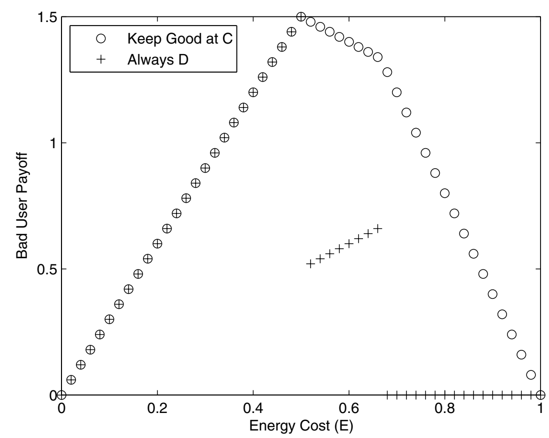

Fig. 3. Comparison of "Keep the Good at C" with "Always D", for the graph of Figure 2. Connectivity benefit N is normalized to 1.

The result depends on what the actual value of $\frac{E}{N}$ is. Namely (R_B and R_G are the total payoffs for Bad and Good nodes, respectively),

$\frac{E}{N}$	(q_1, q_7, q_{10})	R_B
$0 \leq \cdot \leq \frac{1}{2}$	$(0, 0, 0)$	$3E$
$\frac{1}{2} < \cdot \leq \frac{2}{3}$	$(2\frac{E}{N} - 1, 0, 2\frac{E}{N} - 1)$	$2N - E$
$\frac{2}{3} < \cdot \leq 1$	$(2\frac{E}{N} - 1, 3\frac{E}{N} - 2, 2\frac{E}{N} - 1)$	$4(N - E),$

while R_G would be $14(N - E) - R_B$ in each case.

The "Always D" strategy, on the other hand, would result in

$\frac{E}{N}$	(q_1, q_7, q_{10})	R_B	R_G
$0 \leq \cdot \leq \frac{1}{2}$	$(0, 0, 0)$	$3E$	$14(N - E) - R_B$
$\frac{1}{2} < \cdot \leq \frac{2}{3}$	$(0, 0, 0)$	E	$6(N - E) - R_B$
$\frac{2}{3} < \cdot \leq 1$	$(0, 0, 0)$	0	$0.$

So we can see (Fig. 3) that the "Keep the Good at C" strategy does at least as good as the "Always D" strategy, and for some values of $\frac{E}{N}$ it actually does better. Actually, the "Keep the Good at C" strategy is the optimal one for this graph, and for any value of $\frac{E}{N}$, as we can see in Figure 3.

C. Finding an optimal strategy

However, the optimal strategy of the Bad is not always to keep all the Good at C. We provide the following counterexample (Fig. 4). The idea is that in certain cases it is best for a Bad user to keep some but not all his Good neighbors at C.

In this case, the "Keep the Good at C" strategy would lead to the following trivial linear program

$$\begin{array}{rrcl} \min q & \text{s.t.} & & \\ & q+1 & > & 2\frac{E}{N} \\ & q+3 & > & 4\frac{E}{N} \\ 0 \leq & q & \leq & \frac{E}{N}, \end{array}$$

the solution of which would be

$\frac{E}{N}$	q	R_B	R_G
$0 \leq \cdot \leq \frac{1}{2}$	0	$4E$	$10(N-E)-R_B$
$\frac{1}{2} < \cdot \leq 1$	$2\frac{E}{N}-1$	$4(N-E)$	$10(N-E)-R_B.$

The "Always D" strategy, on the other hand, would lead to

$\frac{E}{N}$	q	R_B	R_G
$0 \leq \cdot \leq \frac{1}{2}$	0	$4E$	$10(N-E)-R_B$
$\frac{1}{2} < \cdot \leq \frac{3}{4}$	0	$3E$	$8(N-E)-R_B$
$\frac{3}{4} < \cdot \leq 1$	0	0	$0.$

So, in this case there is no clear winner, since each strategy outperforms the other for different values of $\frac{E}{N}$. They are equivalent when $\frac{E}{N} \in [0, \frac{1}{2}]$, "Always D" is better when $\frac{E}{N} \in [\frac{4}{7}, \frac{3}{4}]$ and worse when $\frac{E}{N} \in [\frac{1}{2}, \frac{4}{7}] \cup [\frac{4}{7}, 1]$.

In general, the best strategy depends heavily on the topology and the placement of the Bad users. However, it seems that the two strategies outlined above play a special role, and perhaps a combination of the two will prove to be optimal.

IV. Conclusion

Ad hoc networks depend on the cooperation of their members to operate successfully. However, since it costs energy to cooperate, users have a large incentive to be selfish. In this paper, we consider what happens if some of the network users are not selfish, but outright malicious. We find that some of the Good users may be forced to stop cooperating so that they preserve their energy. Depending on the relative value of energy cost versus network connectivity, the Bad users may actually increase their payoffs by being cooperative once in a while, thus wasting the energy of the Good ones. Unfortunately, we are unable to compare our results to the relevant literature, since the literature is not considering malicious users within a game theoretic framework to any significant extent.

In the future, we plan to elaborate on the exact connection between the topology and the optimal strategy that the Bad users can play against the simple fictitious play of the Good ones. We would like to see if fictitious play is sufficient, or a more complicated Good strategy is needed.

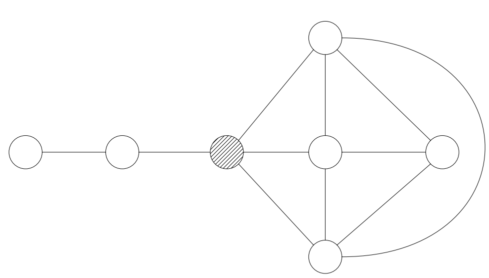

Fig. 4. Example Graph 2. The mixed strategy is optimal.

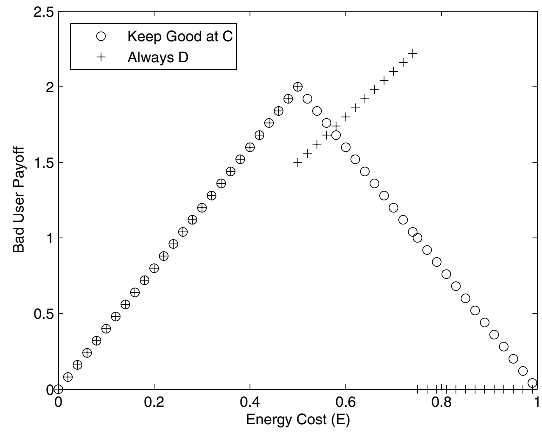

Fig. 5. Comparison of "Keep the Good at C" with "Always D", for the graph of Figure 4. Connectivity benefit N is normalized to 1.

References

[1] M. Félegyházi, J.-P. Hubaux, and L. Buttyán, "Nash Equilibria of Packet Forwarding Strategies in Wireless Ad Hoc Networks," *IEEE Transactions on Mobile Computing*, vol. 5, no. 5, pp. 463–476, May 2006.
[2] Y. Xing and R. Chandramouli, "Distributed discrete power control for bursty transmissions over wireless data networks," in *Proc. of ICC 2004*, Paris, France, June 2004, pp. 139–143.
[3] L. Buttyan and J.-P. Hubaux, "Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks," *ACM/Kluwer Mobile Networks and Applications*, vol. 8, no. 5, 2003.
[4] A. Urpi, M. Bonuccelli, and S. Giordano, "Modelling Cooperation in Mobile Ad Hoc Networks: A Formal Description of Selfishness," in *Proc. of WiOpt 2003: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, INRIA Sophia-Antipolis, France, Mar. 2003.
[5] L. A. DaSilva, and V. Srivastava, "Node Participation in Ad Hoc and Peer-to-Peer Networks: A Game-Theoretic Formulation", in *Workshop on Games and Emergent Behaviors in Distributed Computing Environments*, Birmingham, UK, Sep. 2004.
[6] V. Srinivasan, P. Nuggehalli, C.-F. Chiasserini, and R. Rao, "Cooperation in Wireless Ad Hoc Networks," in *Proc. of IEEE Infocom 2003*, San Francisco, March 30-April 3 2003.
[7] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT Press, 1998.
[8] E. Altman, A. Kherani, P. Michiardi, and R. Molva, "Non-cooperative Forwarding in Ad Hoc Networks," INRIA, Tech. Rep. RR-5116, 2004.
[9] D. Monderer and L. S. Shapley, "Fictitious Play Property for Games with Identical Interests," *Journal of Economic Theory*, vol. 68, no. 1, pp. 258–265, January 1996.

[0] Work prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. Research also supported by the U.S. Army Research Office under grant No DAAD19-01-1-0494. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsors.