The effect of selfish behavior in mobile networks using CSMA/CA

Olav Queseth
Dept. of Signals, Sensors and Systems, Radio communication systems lab
KTH
Stockholm, Sweden
olav.queseth@radio.kth.se

Abstract— **It can be shown that in a communication system which uses CSMA/CA for medium access control greedy behavior can be discouraged by punishing the greedy users. However in a real system it is not always possible to punish a misbehaving user. Actually numerical experiments show that it is difficult to achieve enough punishment do discourage greediness since it takes time to detect that a user misbehaves and the time that a user can be punished is limited due to capture effects.**

Keywords: Game theory; CSMA/CA; unlicensed spectrum; etiquette rules; coexistence; radio resouce management

I. INTRODUCTION

Lately we have seen an increase in the use of unlicensed spectrum. The success of the IEEE 802.11 suite of standards and the success in the marketplace cannot be disputed. Today there are also numerous hotspot providers that provide high speed internet access in for example airports, cafés and fast food restaurants. The traditional telecom operators are also embracing this technology as a complement to their mobile service offerings which can be seen both from those who actually operates hotspots and the high interest in 3G and WLAN interworking.

In unlicensed spectrum the interference from other devices cannot be controlled by planning ahead and instead the devices have to be able to cope with the interference situation. There are many different methods for coping. One philosophy is to simply live with the interference. By introducing some redundancy the message will eventually get through even though some bits are lost. Slow frequency hopping, DS-CDMA and coding and interleaving techniques are examples of this approach.

The other philosophy is to avoid interference. For example a transmitter tries to find an empty frequency slot to transmit in. The automatic channel allocation methods are examples of this method. Or the transmitter tries to find an empty timeslot to transmit in. The CSMA/CA protocol can be viewed as an example of this method. However in order to avoid that all users start using the channel as soon as it becomes available there is a backoff mechanism added to the protocol that ensures some cooperation between the users. However this slightly violates the idea that users have to cope with the interference alone and introduces a degree of cooperation between the users. But by introducing this cooperation between users a possibility for individuals to be greedy also opens up. It is this possibility that is the focus of this paper.

Most of the interference handling techniques have been developed and/or evaluated in a setting where the aim is to avoid manual configuration or to add dynamic behavior to a system. A lot about this can be found in the literature, but they all assume that all users in the system share a common objective (maximizing system capacity) set by the system designer. What is unique with unlicensed spectrum is that there are many users of the spectrum that may not share the same objective. Many times the users are only interested in maximizing their own performance. The designers of the IEEE 802.11b standard were able to make all the users share the same common goal by making the standard in a specific way. For a single user deviating from this common goal is difficult. Modifying equipment is beyond the reach of most users and most products implement the resource sharing algorithm in the same way. But as the market matures it is increasingly important to show good performance and is easy to imagine that manufacturer X add proprietary “features” that enhance the performance at the expense of other equipment not made by manufacturer X. There already examples of this happening [1].

For the regulator it is important to ensure effective use of the spectrum and ensure economic growth. One way to achieve this may be to release spectrum for unlicensed use. The main benefits with unlicensed spectrum are that it is easy to deploy new services and that technology developments can quickly be implemented in new products. The drawback is that once spectrum has been released the regulator looses (almost) all possibilities to control the use. The main tool to control what is done in unlicensed spectrum is the rules that devices must follow. Determining which rules that result in efficient spectrum use is nontrivial. However there are some general observations. By creating a large set of rules that users must follow it is possible to make all the users behave like one system and thus maximize capacity. More rules usually mean that it is harder to control that they are followed. But maybe the most troublesome is that a lot of rules tend to limit the possibilities for innovation. At the extreme there can only be one specific service from one specific manufacturer.

Evaluating how users will behave under a specific set of rules is a little bit different than evaluating the performance of a distributed system where the users follow a specific algorithm. In a “traditional” radio system in licensed spectrum all users share a common goal (maximum revenue for the operator) but in unlicensed spectrum each user has his own objective and we can expect each user to try to maximize his own benefit.

II. PROBLEM

To model many users in a radio system with different objectives game theory has been used as an analytical tool. This has been used to study Aloha systems [2] as well as power control problems [3].

In this paper we model a radio system by letting the individual users be actors that try to maximize their utility, i.e. throughput. Each user can take a number of actions depending on the situation and what they know about the other users in the system. There are many possible algorithms that a user may employ to pick an action. This algorithm is known as a strategy. Depending on the actions that the various users take they will have a different payoff, e.g. average throughput.

Game theory makes a distinction between games that are played only once and games that are repeated many times. In games that are repeated the users have the possibility to decide their current actions on the outcome of previous rounds. Another difference is that in games that are played only once the users have no incentive to cooperate but in repeated games there is the possibility for a user to punish the other users. This property tends to encourage cooperation.

One of the assumptions usually done when analyzing repeated games is that the game is played infinitely. If the game is played a limited number of times the last game will be equivalent to a game played only once [4]. The reasoning is that it makes sense to not cooperate the last time. However the same argument can be recursively applied to all games and thus some of the results of the analysis become invalid.

M. Cagalj et al have studied the CSMA/CA protocol under the assumption that users are selfish and that they do not strictly follow the protocol [5]. It can be shown that in a CSMA/CA game that is played only once the best strategy is to start transmitting immediately, something that of course causes the system to break down completely.

However when the game is played repeatedly cooperating users can punish the misbehaving user and it is possible to reach an equilibrium point. Punishing a user is done by deliberately jamming a misbehaving user's transmissions. The responsibility for punishing misbehaving users can be distributed among the other users. The idea is that a misbehaving user can have his throughput reduced so much that he will stop misbehaving in order to not be punished any more.

In the aforementioned paper detection is done by measuring the throughput of the users and if a user has an average throughput that is sufficiently higher than the other users he is deemed to be misbehaving and thus punished.

In a radio system where users move around the time one user can punish another user is limited since they will eventually move so far apart that they cannot interfere with each other. This makes detection difficult since the time to measure is limited which makes averaging difficult. Due to the randomness inherent in the backoff procedure of the CSMA/CA protocol there will always be random variations that need to be averaged out. Another feature of radio systems is that if the received signal is sufficiently stronger than the interfering or jamming signals there is capture and the packet is received correctly anyway. The net result is that it is not possible to punish certain users.

This paper we investigate the possibility to achieve cooperation in a radio environment. We do this by determining the performance gains that can be achieved by acting greedily. We also have a look at the important components: detecting misbehaving users and punishing them.

III. MODELS AND ASSUMPTIONS

The relative successfulness of the various strategies have been evaluated using numerical experiments. The propagation and mobility parameters have been borrowed from macro-cellular systems. The reason is simply that these models were readily available and already implemented in MATLAB in the RUNE toolbox [6]. However the effects of the various strategies of the users should be similar in other environments.

The system consists of total 64 cells. We assume that there are 4 channels available which results in 16 cells per channel. To avoid border effects we use a wraparound technique to create a borderless surface to place the system on. The system is planned using a regular channel plan. The propagation is modeled using the Okumura-Hata model, i.e. the propagation loss can be described as: $L=28+35\log(R)+X$ [dB] where X is a random variable with mean 0 and standard deviation 8 dB. The cell radius is 1000m. The transmitter power is 1 W and the noise is set to result in a median SIR at the cell border of 25 dB. Thus the system is essentially interference limited.

There are on average three users per access point. The average mobile speed is 15 m/s (vehicular speeds) which results in a median staying time in each cell of 6.5 seconds. The mobility model is further explained in [6]. The handoff margin used is 3 dB and no handoff is performed when a user is transmitting a packet.

We study the full traffic case. I.e. all users always have packets to send. To be able to study how the strategies influence throughput we only have traffic from the users to the access point. It is not reasonable to assume that the access-points will have different strategies for different users.

For programming convenience the system is implemented as a slotted system. Each slot is 20 ms and each packet lasts 10 slots. To be correctly received the SIR for all slots in a packet have to be above 10 dB. A channel is considered to be free if the received signal level not more than 5 dB above the noise floor. Feedback information is assumed to be instantaneous and error free.

IV. STRATEGIES

Strategies are randomly assigned to the users. No user switches strategies during the numerical experiments. In the system there are three different strategies that a user can follow.

Strategy A is the greedy strategy. The user starts to transmit a packet as soon as it has finished the previous packet. If the other users are timid and listens before speaking, this strategy will essentially give the greedy user the full access to the bandwidth. However if all users implement the greedy strategy this is obviously not a good strategy since all users will interfere with each other.

Strategy B is to use the CSMA/CA protocol and follow the rules strictly, this is the timid strategy. The protocol implemented by these users is very similar to the protocol used by devices implementing IEEE 802.11 [7]. Whenever a user has finished transmitting a packet he draws a random number on the interval [0,FW]. The user then waits for this number of empty slots before transmitting the next packet. If there is someone else transmitting on the channel the counting down is suspended. If the packet is lost on the way the FW variable is doubled. If the packet is successfully transmitted the FW is reset to the initial value. The minimum FW value is set to 8 and the maximum is 256.

Strategy C is to follow the timid strategy (strategy B) initially. However if a user detects that another user is behaving in a greedy fashion he can punish that user. Detecting that a user is cheating is non-trivial. In the paper written by Cagalj et al [5] it is suggested that each user measure the throughput of all the other users and deem a user to be cheating when that user achieves a higher throughput than the rest of the users. In this paper we use a slightly simpler algorithm. The user that transmits 5 packets in a row without releasing the channel is considered to be greedy. Punishing a user is done by deliberately jamming a packet sent by the user. The user is punished for 5 packet times. We assume that this can be done since the packet header contains the address of the sender and thus it is possible to quickly determine which packets are sent by a specific user. If a user moves into another cell the punishment is stopped immediately. We also assume no communication between the cells thus it is not possible to punish a user if he is on another channel or in another cell.

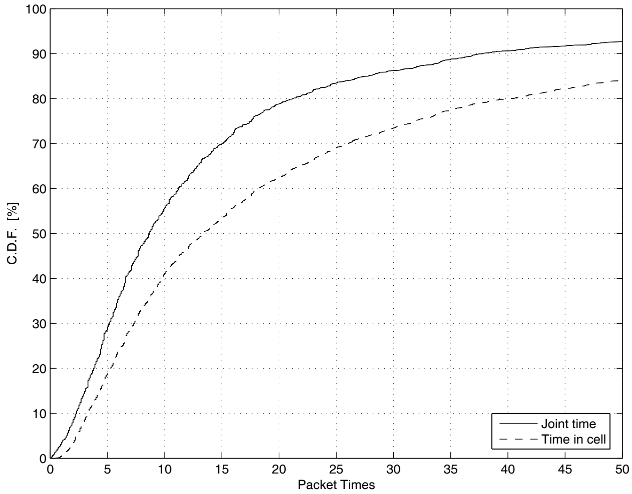

Figure 1. Distribution of time spent in a cell by a user and the “joint time” the time two users are colocated in a cell.

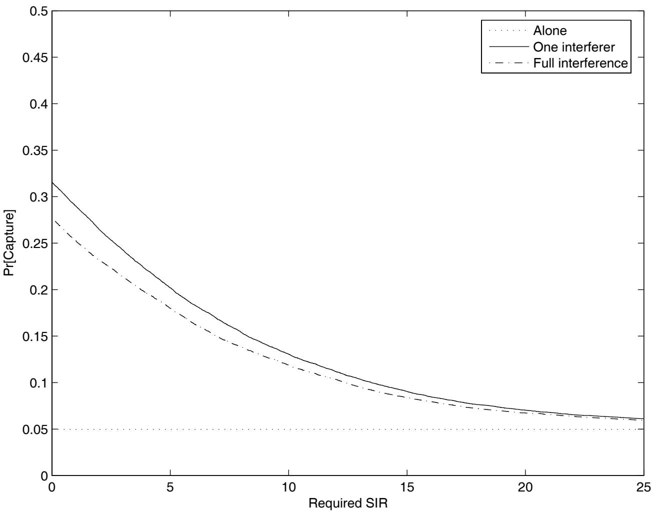

Figure 2. Probability of capture under the assumption that there is one interferer or that all users in a cell interfere with the transmission.

V. RESULTS

In all numerical experiment we simulate 400 seconds. The throughput per user is measured as the fraction of slots that a user sends useful data. In an ideal situation with only timid users they would achieve a throughput of roughly 33% each. But in our experiments it is slightly less. Handoffs, hidden terminals and contention times all reduce the throughput.

Obviously the characteristics of the mobility and how long users stay within a cell will influence the results. In Fig. 1 the time a user stays within a cell is plotted. There is also a plot of how long users spent together in a cell, the “joint time”. On average a user has time to send approximately 30 packets while in a cell.

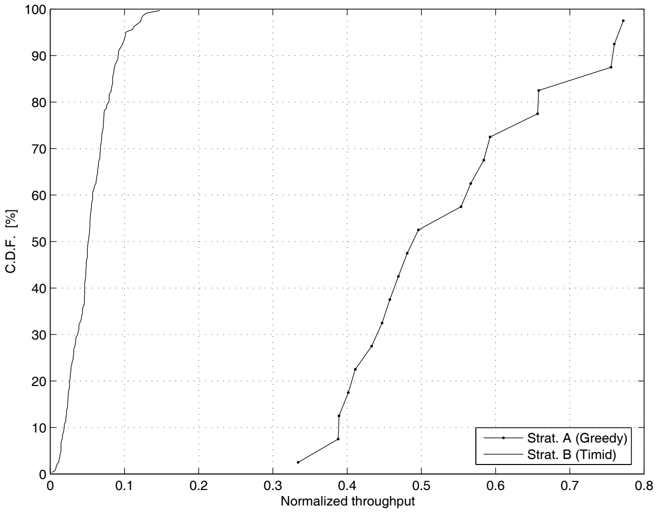

Figure 3. In a system with only timid users and only a few greedy (10%) users the performance improvement for the greedy users is substantial.

The propagation model as well as the user density has a large influence on the possibility of packet capture. Fig. 2 outlines the probability of capture for a given SIR requirement. It is worth noting that with the user load in the experiments a user has a 5% chance of being alone in a cell. It is also worth noting that at a SIR requirement of 10dB there is approximately a 10% chance that a packet will be captured in a cell, even though there is interference from the other users in the system.

In Fig. 3 we can see that being greedy indeed pays if there are only a few users that use strategy A. There are 10% users that follow strategy A and the rest follow strategy B (timid). Note that even the timid users get some data through. The reason is that in some instances there are no greedy users in a specific cell and thus they can get some packets through.

In Fig. 4 we increase the fraction of users that follow strategy A (greedy) to 80% we can see that the throughput drops for all users. One thing to note though is that there is actually some data that gets through for both categories of users. There are two effects that cause this. There is 5% chance that there is only one user in a cell and obviously a lonely user will always get his data through, even if he uses strategy B (timid). The second reason is the capture effect. As we have seen previously there is a 10% chance of packet capture.

We can clearly see that greedy users enjoy more throughput than the timid users, but does that still hold if the greedy users are punished? In the next experiment we introduce users that follow strategy C, i.e. they are able to punish users that hog the channel too long. In fig. 5 we can see the results of an experiment with 30% users following strategy C and 10% users following strategy A (greedy).

It can be seen that the strategy C users achieve approximately the same performance as the timid users. This is no surprise since they use the same channel access method as the timid users. The gain in punishing should result from the strategy A users being less greedy, but that is not implemented in these experiments. However we note that there is a drop in the performance for strategy A (greedy) users. But even though the performance drops there is still a significant gain in being greedy.

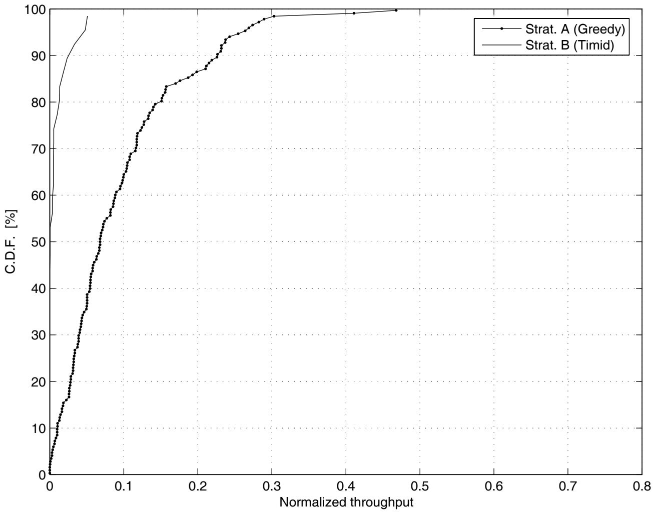

Figure 4. In a system with mostly greedy users only a very limited throughput is achieved.

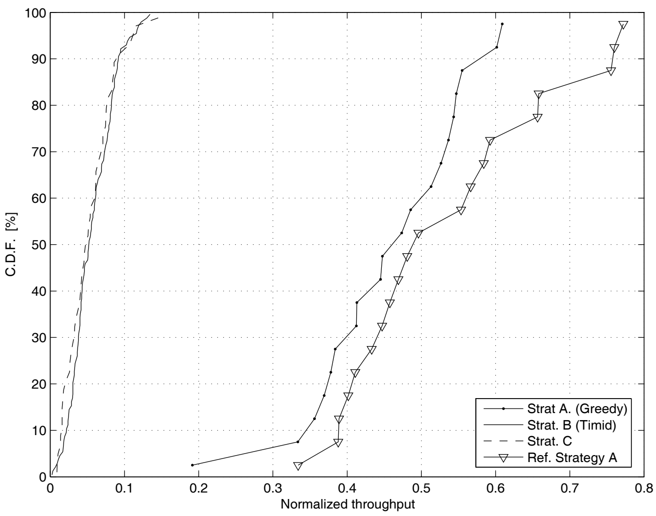

Figure 5. In a system with 30% users that ounish the other and 10% greedy users the greedy users just get a marginal degradation of their performance. The reference throughput is from fig. 3

In the next experiment we increase the portion of users that follow strategy C to 90%. We let 10% of the users be greedy. The result is shown in fig. 7. We can see that there is a slight degradation in performance for the users that are greedy. However there is still a large discrepancy in the performance of the strategy C users and the strategy A (greedy) users.

In the previous experiments the greedy users have only been punished for 5 packets. In the next experiment we let a user be punished indefinitely once he has been determined to be greedy. The result is shown in fig. 7. There is a performance loss compared to the previous cases. But there is still a large difference in throughput between greedy users and

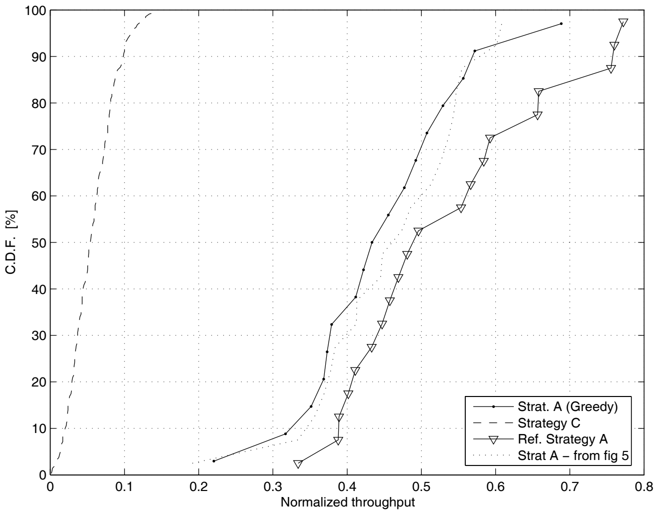

Figure 6. Throughput for a system with 90% strategy C users and 10% strategy A (greedy) users. In the graph the results for the greedy users from fig 3 and fig 5 are included for reference.

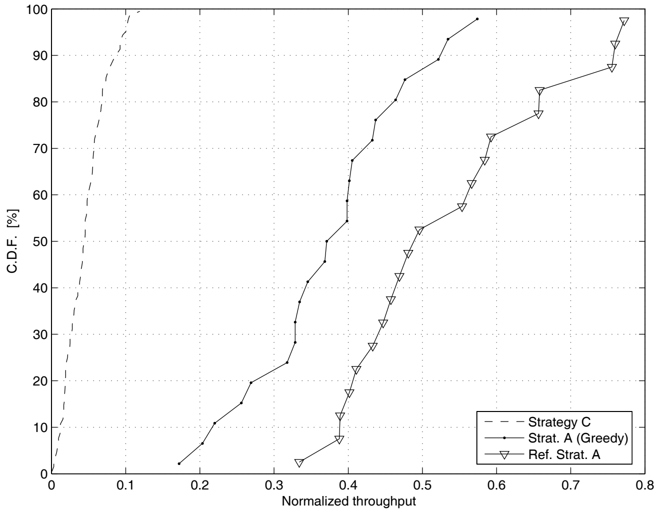

Figure 7. Normalised throughput for a system with 10% greedy users and where 90% of the users follow strategy C except that misbehaving users are punished indefinitely.

those that do not act greedily. Most of this is due to the difficulty in detecting misbehaving users. If we compare the throughput here with the one in fig 4 where all users effectively jam each other we can see the difference.

VI. DISCUSSION

The main outcome of this study is that we have shown that it is actually difficult to implement punishment in a radio system. One reason is the capture effects that let greedy users communicate 10% of the time event though they may be punished. Another reason is the difficulty in detecting misbehaving users. Detection time with the algorithms in this paper is on the order of 6-7 slots in the most favorable cases. If we compare that to how long users actually spend together in a cell we can see that in 50% of the cases there is not enough time to detect the misbehavior. Then to be effective the users must stay in the cell long enough so that the punishment can be experienced buy the misbehaving user.

If we compare the throughput for the greedy users, even if they are being punished, with what can be expected if everyone were well behaved (approx. 0.3 normalized throughput) we see that it still pays to be greedy. This is discouraging since to be effective it must be possible to punish a misbehaving user enough so that he realizes that it is better to be nice in the long run.

The problem seems to be to detect that a user is greedy. Here we have assumed that the information about who is nice and who is mean is not shared between users and not between access points. It would of course be possible to keep a "blacklist" and this would simplify the detection process.

The assumptions made here about user mobility may be slightly pessimistic. Currently unlicensed systems are mainly used where the users move at a pedestrian speed at most, on the other hand the cells are smaller. In a system where users stay longer in a cell it is easier to detects and punish a misbehaving user.

The mechanisms for detection and punishment used here are quite simple. With more sophisticated techniques it may be easier to detect when a user is greedy and administer punishment in a proportional fashion. On the other hand the greedy users in our experiments are easy to detect more complicated strategies for cheating may take longer to detect.

We have only studied the uplink communication in this paper. It is difficult to imagine that the accesspoint is greedy on behalf of only one user. This means that it is only in a fraction of the traffic that greediness really pays and thus may not be a large problem.

VII. CONCLUSIONS

It is beneficial for a user to act greedily, even if he may be punished. The resulting throughput is higher than for users that fully cooperate with each other. Thus it may be difficult to ensure cooperative behavior.

One problem is the capture effects. Another problem is properly detecting misbehaving users, more sophisticated techniques and communication between users may be a way to remedy this.

REFERENCES

[1] Atheros Communication inc. "Super G, Maximizing Wireless Performance" White paper: http://www.atheros.com/pt/atheros_superg_whitepaper.pdf (2004-05-25)

[2] A.B. MacKenzie and S.B. Wicker "Stability of Multipacket Slotted Aloha with Selfish Users and Perfect Information" IEEE Infocom 2003

[3] C.U. Saraydar N. B. Mandayam and D. J. Goodman "Efficient Power Control via Pricing in Wireless Data Networks" IEEE Trans. Communications Feb. 2002

[4] M. J. Osborne and A. Rubinstein "A course in game theory" MIT Press, Cambridge 1994

[5] M. Cagalj et al. "On cheating in CSMA/CD Ad Hoc Networks" EPFL-IC technical report No. IC/2004/27 http:// icwww.epfl.ch/publications/documents/ IC_TECH_REPORT_200427.pdf (2004-05-25)

[6] J. Zander and S.L. Kim "Radio Resource Management in Wireless Networks", Artech House 2001

[7] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications" (section 9) IEEE standard 802.11, 1999.