Multiple Access in Ad-Hoc Wireless LANs with Noncooperative Stations

Jerzy Konorski

Technical University of Gdansk
ul. Narutowicza 11/12, 80-952 Gdansk, Poland
jekon@pg.gda.pl

Abstract. A class of contention-type MAC protocols (e.g., CSMA/CA) relies on random deferment of packet transmission, and subsumes a deferment selection strategy and a scheduling policy that determines the winner of each contention cycle. This paper examines contention-type protocols in a noncooperative an ad-hoc wireless LAN setting, where a number of stations self-optimise their strategies to obtain a more-than-fair bandwidth share. Two scheduling policies, called RT/ECD and RT/ECD-1s, are evaluated via simulation It is concluded that a well-designed scheduling policy should invoke a noncooperative game whose outcome, in terms of the resulting bandwidth distribution, is fair to non-self-optimising stations.

1 Introduction

Consider N stations contending for a wireless channel in order to transmit packets. In a cooperative MAC setting, all stations adhere to a common contention strategy, C, which optimises the overall channel bandwidth utilisation, U: $C=\mathrm{argmax}U(x)$. In a noncooperative MAC setting, each station i self-optimises its own bandwidth share, U_i: $C_i^*= \mathrm{argmax}U_i(C_1^*,...,C_{i-1}^*,x,C_{i+1}^*,...,C_N^*)$. C_i^* is a *greedy* contention strategy and $(C_1^*,...,C_N^*)$ is a *Nash equilibrium* [3] i.e., an operating point from which no station has incentives to deviate unilaterally. Note that noncooperative behaviour thus described is *rational* in that a station intends to improve its own bandwidth share rather than damage the other stations'. This may result in unfair bandwidth shares for stations using C. For other noncooperative wireless settings, see [1,4].

In ad-hoc wireless LANs with a high degree of user anonymity (for security reasons or to minimise the administration overhead), noncooperative behaviour should be coped with by appropriate contention protocols. A suitable communication model is introduced in Sec. 2. The considered contention protocol under the name Random Token with Extraneous Collision Detection (RT/ECD) involves voluntary deferment of packet transmission. We point to the logical separation of a deferment selection strategy and a scheduling policy that determines the winner in a contention cycle. Sec.

E. Gregori et al. (Eds.): NETWORKING 2002, LNCS 2345, pp. 1141-1146, 2002.

3 outlines a framework for a noncooperative MAC setting. A scheduling policy called RT/ECD-1s is described in Sec. 4 and evaluated against the RT/ECD policy in terms of the bandwidth share guaranteed for a cooperative (c-) station (using C) in the presence of noncooperative (nc-) stations (using C_i*). Sec. 5 concludes the paper.

2 Noncooperative MAC Setting with RT/ECD

Our 'free-for-all' communication model consists of the following non-assumptions:

- neither N nor stations' identities need to be known or fixed,
- except for detecting carrier, a station need not interpret any packet of which it is not an intended (uni- or multicast) recipient.

This allows for full encryption and/or arbitrary encoding and formatting among any group of stations. To simplify and restrict the model we assume in addition

- single-hop transfer of packets with full hearability, and
- a global slotted time axis.

Any station is thus able to distinguish between v- and c-slots sensed (for 'void' and 'carrier'). An intended recipient of a successful transmission recognises also an s-slot (for 'success') and reads its contents. This sort of binary feedback allows for *extraneous collision detection* in the wireless channel as employed by the following RT/ECD protocol (Fig. 1). In a protocol cycle, a station defers its packet transmission for a number of slots from the range 0..D–1, next transmits a 1-slot *pilot* and senses the channel in the following slot. On sensing an s-slot containing a pilot, any intended recipient transmits a 1-slot *reaction* (a burst of non-interpretable carrier), while refraining from reaction if a v- or c-slot is sensed. A reaction designates the sender of a successful pilot as the winner and prompts it to transmit its packet in subsequent slots; a v-slot will mark the termination of this protocol cycle. If pilots collide, no reaction follows and the protocol cycle terminates with a no-winners outcome. In a full-hearability environment, RT/ECD operates much like CSMA/CA in the IEEE 802.11 Distributed Coordination Function [2], with the pilot/reaction mechanism resembling the RTS/CTS option. Note, however, that it is to provide ACK functionality rather than solve the hidden terminal problem; moreover, pilots only need to be interpreted by intended recipients, while reactions are non-interpretable.

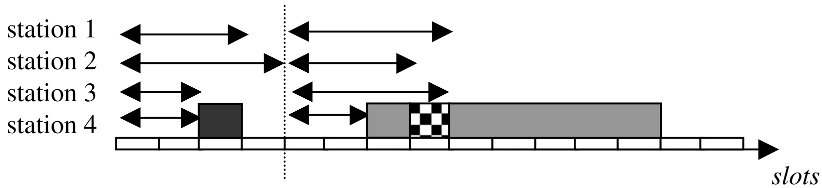

Fig. 1. RT/ECD, a no-winners protocol cycle followed by one where station 4 wins

To account for noncooperative behaviour, we assume that

- *NC* out of *N* stations are nc-stations that may use greedy deferment selection strategies (*NC* need not be known or fixed),
- the c-stations use a standard deferment selection strategy *S*, defined by the probabilities π_l of selecting a deferment of *l* slots ($l \in 0..D-1$), and
- all stations adhere to a common scheduling policy.

A simple greedy strategy might consist in introducing a downward $bias \in 0..D-1$ to the deferment distribution e.g., $\pi_0' = \pi_0 + ... + \pi_{bias}$ and $\pi_l' = \pi_{l+bias}$ for $l>0$. As shown in Sec. 4, this may leave the c-stations with a tiny fraction of the bandwidth share they would obtain in a cooperative setting (with $NC=0$).

3 Framework for a Noncooperative MAC Setting

Besides pursuing a greedy deferment selection strategy, an nc-station might try various 'profitable' departures from the protocol specification – for example, pretend to have transmitted a pilot and sensed a reaction. In RT/ECD-like protocols, however, such cheating must involve making false claims as to the presence or absence of carrier on the channel, which is easily verifiable. Therefore it suffices to design a scheduling policy so as to minimise the benefits of any conceivable greedy strategy vis-a-vis *S*. A greedy strategy can be expected to be

- *isolated* i.e., not relying on collusion with other nc-stations, and
- *rational*, meaning that deferment selection rules observed to increase own bandwidth share are more likely to be applied in the future, however, to stay responsive to a variable environment, no rules are entirely abandoned [3].

A reasonable scheduling policy is constrained to be

- *nontrivial*, in that no deferments should be known a priori to render other deferments non-winning (note that RT/ECD is a counterexample, deferment of length 0 being 'fail-safe'), and
- *incentive compatible*, in that channel feedback up to any moment in the deferment phase should not discourage further pilots (as a counterexample, imagine a scheduling policy whereby a second-shortest deferment wins).

Let $U_c(NC)$ be the bandwidth share obtained by a generic c-station in the presence of *NC* nc-stations. A fair and efficient scheduling policy is one that ensures $U_c(NC)$ '≥' $U_c(0)$ '≥' $U_c(0)|_{RT/ECD}$ for any *NC* and any greedy strategy, where '≥' reads 'not less or at least tolerably less than.' This means that the presence of nc-stations should not decrease a generic c-station's bandwidth share by an amount that its user would not tolerate. The latter 'inequality' implies that protection against nc-stations should not cost too much bandwidth in a cooperative setting, RT/ECD being a reference policy supposed, by analogy with IEEE 802.11, to perform well in a cooperative setting.

4 Evaluation of the RT/ECD-1s Scheduling Policy

While RT/ECD prevents any station from winning if a collision of pilots occurs, in RT/ECD-1s the first successful pilot wins no matter how many collisions precede it. A protocol cycle is illustrated in Fig. 2. A slot occupied by a pilot (or a collision of pilots) is paired with a following one, reserved for reactions. Stations whose pilots were not reacted to back off until the next protocol cycle. The lack of a second chance to transmit a pilot in the same protocol cycle creates a desirable 'conflict of interest' for an nc-station selecting its deferment. RT/ECD-1s is arguably nontrivial and incentive compatible. (A family of similar policies can be devised whereby the n^{th} successful pilot wins, or the last one if there are less than n; of these, RT/ECD-1s yields the best winner outcome vs. scheduling penalty tradeoff.)

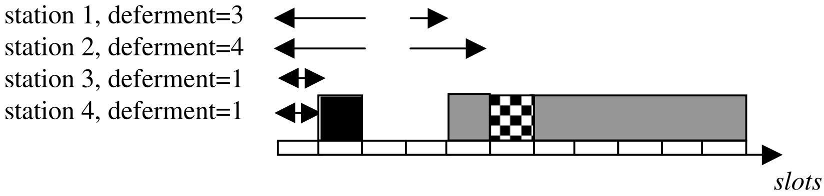

Fig. 2. RT/ECD-1S protocol cycle: stations 3, 4 back off when no reaction follows; station 1's first successful pilot wins (deferments are frozen during reaction slots)

In a series of simulation experiments, simple models of c- and nc-stations were executed to evaluate RT/ECD-1s against the backdrop of RT/ECD. In each simulation run, D=12, N=10 and $NC \in 0..N-1$ were fixed and packet size was 50 slots. Symmetric heavy traffic load was applied with one packet arrival per station per protocol cycle. The strategy S at the c-stations used a truncated geometric probability distribution over $0..D-1$ i.e., $\pi_l = const. \cdot q^l$ with parameter q=0.5, 1 or 2 (referred to symbolically as 'aggressive,' 'moderate' and 'gentle'). Two isolated and rational greedy strategies were experimented at nc-stations: Biased Randomiser (BR) and Responsive Learner (RL). The former introduced a downward bias as explained in Sec. 2; the *bias* value was optimised on the fly at each nc-station and occasionally wandered off the optimum to keep the strategy responsive to possible changes in other stations' strategies. The latter mimicked so-called *fictitious play* [5] by selecting deferments at random based on their winning chances against recently observed other stations' deferments. Once selected, a deferment was repeated consistently throughout the next update period of UP=20 protocol cycles. For simplicity, strategies were configured uniformly within all stations of either status, producing two noncooperative game scenarios: S vs. BR and S vs. RL.

Ideally, $U_c(NC) \equiv 1/N$ of the channel bandwidth. Scheduling penalties cause this figure to drop even in a cooperative MAC setting (at NC=0), whereas nc-stations may bring about a further decrease. For the S vs. BR scenario, Fig. 3 (*left*) plots $U_c(NC)$ (normalised with respect to $1/N$) as measured after the nc-stations have

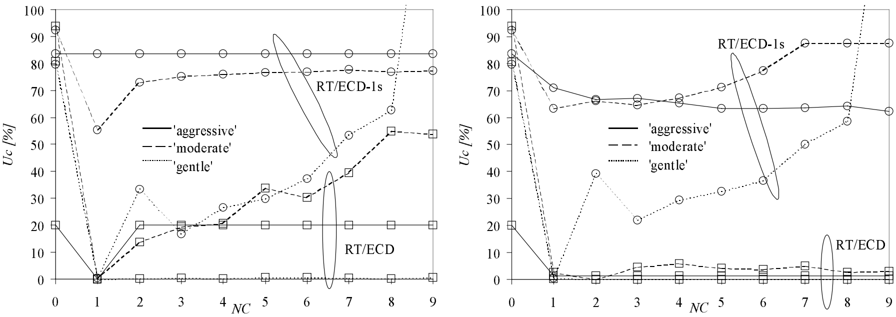

Fig. 3. C-station bandwidth share as a function of *NC*, *left*: S vs. BR, *right*: *S* vs. RL

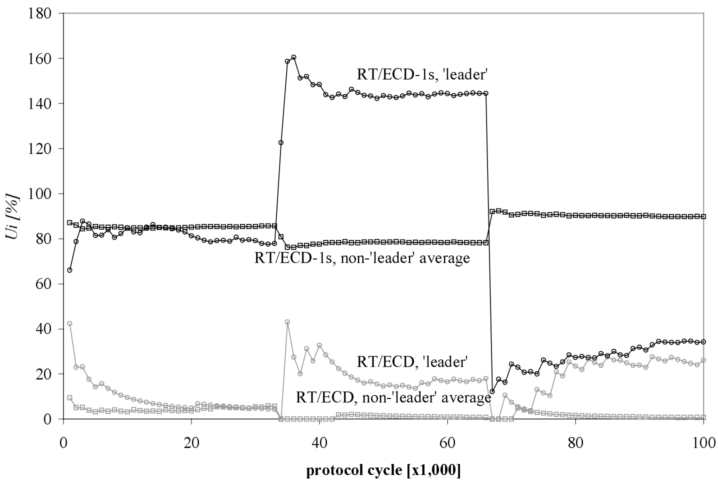

Fig. 4. RL vs. RL: Stackelberg 'leader' scenario

reached a Nash equilibrium with respect to *bias*. Note that while RT/ECD-1s is generally superior to RT/ECD, much depends on the parameter *q*: the 'gentle' value is not recommended, especially for a small *N*, while for the 'aggressive' value, the nc-stations detect that the optimum *bias* is 0, hence $U_c(NC)$ remains constant. Also, RT/ECD-1s has difficulty coping with *NC*=1. Fig. 3 (*right*) presents similar results for the *S* vs. RL scenario. Observe that under RT/ECD-1s, nc-stations' increased intelligence does not worsen $U_c(NC)$ significantly, which it does under RT/ECD. Again, much depends on *q*: although the 'moderate' value pays off in a cooperative setting, the 'aggressive' value offers more uniform guarantees for $U_c(NC)$ across various *N*. Lose-lose situations (with both the c- and nc-station bandwidth shares below $U_c(0)$) were observed under RT/ECD owing to this policy not being nontrivial.

Fig. 4 presents an RL vs. RL scenario where, after a third of the simulation run, one nc-station captures more bandwidth by lengthening its update period tenfold whenever a deferment of length 0 is selected. In doing so, it becomes a so-called Stackelberg 'leader' [5]. A form of protection, switched on after another third of the simulation run, is for a c-station to monitor its own and other stations' win counts over the last update period. If the former is zero and the latter nonzero, the station temporarily resorts to *S* with the 'aggressive' *q*. Under RT/ECD-1s, this quickly results in the 'leader' obtaining a less-than-fair bandwidth share. Under RT/ECD the protection is ineffective; moreover, the overall bandwidth utilisation remains poor.

5 Conclusion

Ad-hoc wireless LAN systems, with their preferences to user anonymity and a lack of tight administration, potentially constitute a noncooperative MAC setting. For a class of contention protocols relying on random deferment of packet transmission, c-stations are vulnerable to unfair treatment by nc-stations, which use greedy deferment selection strategies. The design of a scheduling policy has been shown to be quite sensitive in this respect. A framework for a reasonable scheduling policy and greedy strategies that might be expected from nc-stations has been outlined. A slotted-time scheduling policy called RT/ECD, analogous to CSMA/CA with the RTS/CTS option, and an improved variant thereof called RT/ECD-1s have been evaluated under heavy load to find that the latter guarantees the c-stations a substantially higher bandwidth share. This it does assuming that nc-stations behave rationally and seek a Nash equilibrium. In the experiments, RT/ECD-1s coped well with nc-stations using a randomisation bias or a fictitious play-type greedy strategy.

Several directions can be suggested for future work in this area:

- a game-theoretic study of RT/ECD-like scheduling policies aimed at establishing the mathematical properties of the underlying noncooperative games,
- model extensions to include multihop wireless LAN topologies (in particular, dealing with the problem of hidden stations); development of a suitable extension of RT/ECD-1s is under way, and
- access delay analysis to investigate the issues of QoS support.

References

1. Heikkinen, T.: On Learning and the Quality of Service in a Wireless Network. In: Proc. Networking 2000, Springer-Verlag LNCS 1815 (2000), 679-688
2. IEEE 802.11 Standard (1999)
3. Kalai, E., Lehrer, E.: Rational Learning Leads to Nash Equilibrium, Econometrica 61 (1993), 1019-1045
4. MacKenzie, A.B. and Wicker, S.B.: Game Theory and the Design of Self-Configuring, Adaptive Wireless Networks, IEEE Comm. Magazine, 39 (2001), 126-131
5. Milgrom, P., Roberts, J.: Adaptive and Sophisticated Learning in Normal Form Games, Games and Economic Behaviour 3 (1991), 82-100