

Packet Scheduling in Wireless LANs – A Framework for a Noncooperative Paradigm

Jerzy Konorski

Technical University of Gdansk, Poland

Key words: wireless LAN, EY-NPMA, Random Token, noncooperative stations

Abstract: Contention-based packet scheduling policies incorporated into MAC protocols in wireless networks attempt to schedule one packet transmission per protocol cycle and are optimised to reduce the scheduling penalty while distributing the bandwidth fairly among the network stations. The paper points out the possibility of there being some noncooperative stations that, instead of adhering to a common-goal optimum policy, try to maximise their individual service rates to the detriment of the cooperative stations. A framework for noncooperative scheduling policies is postulated with analogies drawn from the auction paradigm; desirable features include verifiability at various levels, fairness and low performance cost. Upon discussing possible noncooperative station behaviour, four single-cycle noncooperative scheduling policies are proposed as modifications of the well-known EY-NPMA and Random Token policies. Results of a preliminary simulation study are presented to demonstrate that these modifications do prevent or at least discourage noncooperative stations from stealing bandwidth from cooperative ones.

1. INTRODUCTION

Scheduling packet transmissions in wireless LANs has been intellectually challenging mainly because of the lack of 'natural' scheduling facilities present in satellite or wired LANs such as immediate network-wide feedback, physical ordering of stations, collision detection by sensing the channel while transmitting, on-the-fly modification of other stations' packets etc. The challenge is aggravated by user mobility, which makes it difficult to track actions of a particular station, and changing or non-existent station id's,

as in ad-hoc networks. A large variety of distributed scheduling policies have been devised to be incorporated in the wireless MAC protocols (see [7], [8] for a survey and details of some representative policies). All of them impose a scheduling penalty arising from contention, polling, reservation, distributed election and similar mechanisms that consume a certain portion of the channel bandwidth in the form of control packet transmission, synchronisation to time slots or various collision stand-offs. Thus apart from purely random access protocols, a generic *protocol cycle* consists of a scheduling phase and a packet transmission phase (the former may itself be composed of some sub-phases). An important distinction can be made between *single-cycle* and *multiple-cycle* policies that span, respectively, only one protocol cycle (as in ETSI's HIPERLAN [2] or IEEE 802.11's CSMA/CA [10]) and many consecutive protocol cycles (as in PRMA [5] and token-passing protocols). Single-cycle policies can be further divided into *contention-based* and *reservation-based*; the former schedule one packet transmission per cycle, whereas the latter usually attempt to schedule more (an example is the CRT protocol [3]).

Hereafter we shall focus on single-cycle, contention-based policies. Their main goals are:

- high bandwidth utilisation (minimum scheduling penalty) and
- fairness (equal service rates as perceived by all network stations).

While the contradictory nature of these goals has been recognised since the introduction of high-speed wired LANs [6], rapid development and deployment of MAC technologies add a new dimension to the fairness problem. Namely, fairness has so far been striven for with the silent assumption that all the network stations will adhere to the designed policy and co-operate for the common goal. However, this goal will probably be ruined if noncooperative stations are present and pursue their own goals (i.e., greedy maximisation of individual service rates). Consider a CSMA/CD station that refuses to back off after a collision; a HIPERLAN station that always generates long elimination bursts; a CRT station that sends multiple reservation packets per cycle while only one is allowed – clearly, such stations have a competitive edge on the other, cooperative stations. Suitably modified MAC interfaces (like bogus parts in the aircraft industry) can be easily manufactured and distributed, if half-legally, on a commercial basis. This calls for a new generation of *noncooperative scheduling policies* that prevent noncooperative stations from stealing the bandwidth from cooperative ones. Note that a similar evolution has been observed in WANs and internets: from globally optimised mechanisms (e.g., routing) to architecting independently administered, noncooperative networks [4], often based on game theory and seeking a Nash equilibrium point. This paper is meant as a stimulus to a discussion along these lines.

The material is presented as follows. In Sec. 2, the network model is described along with a simplified specification of two exemplary scheduling policies: EY-NPMA of ETSI's HIPERLAN and Random Token (RT) first proposed in [1]. Sec. 3 discusses possible noncooperative station behaviour and outlines the framework of noncooperative scheduling policies. In Sec. 4, modifications of EY-NPMA and RT are proposed to improve fairness in the presence of noncooperative stations. Sec. 5 contains the results of a simulation study.

2. NETWORK MODEL AND SCHEDULING POLICIES

The following assumptions and non-assumptions will govern our network model:

- A1. The network uses a single-frequency radio channel shared on a time-division basis.
- A2. There are N network stations of which NC are noncooperative; the stations never go down or switch off.
- A3. All stations hear one another (a single-hop configuration).
- A4. The maximum station-to-station propagation delay constitutes a common *time slot*; all stations synchronise to the boundaries of successive time slots.
- A5. The transceiver at a station, when in the receive mode, is able to sense the channel at various, policy-specific levels of accuracy, distinguishing idle/busy states, idle/single transmission/collision states or measuring the total carrier power on the channel.
- A6. Stations need not reveal their id's, or have any permanent id's at all; noncooperative stations do not present themselves as such (thus N and NC remain unknown to all stations).
- A7. The employed cooperative scheduling policy (i.e., adhered to by the cooperative stations) is known to all stations.

As examples of cooperative scheduling policies we take EY-NPMA and RTCA (RT with Collision Avoidance, a slightly improved version of RT) specified below in a simplified version. Both are single-cycle, contention-based policies whose viability has been confirmed by numerous implementations (of the former) and analysis (of the latter).

EY-NPMA: at the beginning of a protocol cycle each station examines the contents of its packet buffer and, if nonempty, considers itself active, enters the *elimination phase* and transmits an *elimination burst* of random length (1.. E_{max} time slots). Thereupon, if the channel is sensed idle, the

station enters the *yield phase* in the next time slot; if the channel is sensed busy, the station backs off (i.e., defers its packet transmission to a future protocol cycle). In the yield phase, a packet is transmitted after a random *yield delay* ($1..Y_{max}$ time slots) unless the channel is sensed busy before that, indicating a packet transmission by another station participating in the yield phase. *Figure 1* illustrates the protocol cycle (note that the declaration of priorities and resynchronisation issues of the original EY-NPMA have been left out).

RTCA: in the elimination phase, an active station delays its packet transmission for a random *elimination timeout* ($1..E_{max}$ time slots), then transmits the packet unless sensing the channel busy with another station's packet before that. Again, resynchronisation issues are left out. However, to allow direct comparisons with EY-NPMA, we shall improve RT somewhat in the spirit of CSMA/CA; namely, instead of a whole packet, a station transmits a 1-slot pilot packet to discourage those active stations who have chosen longer elimination timeouts. Next, it enters the yield phase similarly as in EY-NPMA, along with the other stations that transmitted their pilot packets in the same time slot. The protocol cycle is illustrated in *Figure 2*.




Figure 1. EY-NPMA simplified protocol cycle




Figure 2. RTCA simplified protocol cycle

It is interesting to view the elimination phase as a kind of an auction where the stations bid their elimination bursts or timeouts. EY-NPMA corresponds to a first-price sealed-bid auction, and RTCA – to a Dutch auction (where the item being auctioned is awarded to the earliest bidder who matches the progressively descending price) [9]. In our scheduling context, these auctions have a few peculiar features:

- the maximum price of the item i.e., the right to transmit a packet, is bounded (by E_{max}),
- some bidders are willing to buy the item at any price and therefore can cheat by choosing their bids deterministically rather than at random, and
- the auctioneer's goal is to make the item equally affordable to all bidders rather than maximise the selling profit.

3. FRAMEWORK FOR NONCOOPERATIVE SCHEDULING

In EY-NPMA, a simple strategy of a greedy noncooperative station could consist in transmitting its elimination burst and subsequently, upon sensing the channel still busy, resuming the elimination burst with a good chance of outbidding the other active stations (which would turn the sealed-bid auction into a classical English one). In RTCA, a noncooperative strategy could consist in joining in the yield phase despite not having transmitted a pilot packet. This raises the issue of *verifiability*.

Suppose a dedicated network station has at its disposal a monitoring device equipped with a small-aperture antenna able to track the activity of any other station it locks upon. It is easy to see that a noncooperative station using any of the above strategies can thus be identified and neutralised by imposing whatever sanctions one might think of (e.g., heavy penalties on the user or just jamming the ensuing packet transmission). A hypothetical noncooperative scheduling policy forbidding the above strategies would then be *small-aperture verifiable*. A stronger verifiability constraint would involve a monitoring device of less channel sensing accuracy e.g., of the idle/busy (carrier detection) type. Now noncooperative stations could use the above strategies and get away with it, making the corresponding noncooperative scheduling policy *carrier unverifiable*. However, another noncooperative strategy could be devised for EY-NPMA whereby a station starts transmitting its packet right after the longest elimination burst has terminated. Similarly, in RTCA a noncooperative station could start transmitting its packet without prior transmission of a pilot packet i.e., not having subjected itself to the elimination phase. A hypothetical noncooperative scheduling policy forbidding these strategies would be

carrier verifiable. In fact, a taxonomy of noncooperative scheduling policies could be proposed based on the strength of the verifiability constraints they meet. In our study we shall require that a noncooperative scheduling policy be at least small-aperture verifiable. When deciding on its bidding (elimination phase) strategy, a greedy noncooperative station is thus left with small-aperture unverifiable options, given that it is not prepared to risk sanctions imposed by the monitoring device.

One can imagine that a noncooperative strategy will in general consist in bidding longer elimination bursts or shorter elimination timeouts than resulting from the cooperative randomisation mechanism (assume for simplicity that the noncooperative behaviour does not extend to the yield phase). Such a policy can also be soundly assumed to be

- *stationary* – in particular, a noncooperative station will remain so all the time, without ever reforming to become cooperative again,
- *inconspicuous* – attempting to mimic cooperative behaviour in order to conceal its greediness e.g., by randomising bids in successive protocol cycles rather than insistent bidding extreme values in each cycle,
- *rational* – based on the knowledge of the cooperative bidding strategy and aiming to maximise individual service rates rather than to damage the other stations', and
- *isolated* – only aware of a possible presence of other noncooperative stations and not their number, location etc., in particular unable to collude with any other noncooperative station.

In response, a noncooperative scheduling policy should ensure that a noncooperative station does not outperform a cooperative one significantly in the long run, at least for a wide range of NC/N ratios. One might allow significant unfairness for small NC/N on the premise that the few noncooperative stations will not steal enough bandwidth to worry about and besides, small NC/N will not persist too long as more and more cooperative stations, lured by the prospect of stealing some extra bandwidth, turn noncooperative. Likewise, at high NC/N unfairness is less harmful since there are few cooperative stations to steal bandwidth from.

Note in passing that detection of noncooperative stations based on monitoring the source statistics of successful packet transmission is out of the question because

- it would span multiple protocol cycles – in fact, a large enough number to maintain statistical credibility,
- inconspicuous noncooperative strategies would make statistical inference even harder, and
- source information is not necessarily deducible from packets (assumption A6 of Sec. 2).

4. NONCOOPERATIVE SCHEDULING POLICIES

Any verifiable noncooperative scheduling policy should prevent forceful bidding (insistent bidding long elimination bursts or short elimination timeouts), at the same time not deterring cooperative stations from bidding at all. Therefore, some straightforward policies have to be ruled out e.g., imposing penalties upon stations whose elimination bursts were longest or exclusion from the yield phase of stations whose elimination timeouts were shortest. Here we propose four modifications: two of EY-NPMA, called EY-NPMA/(a,b) and EY-NPMA/2ndMAX, and two of RTCA, called RTCA/1stCOLL and RTCA/1stSINGLE. They differ in the channel sensing accuracy required at each station: idle/busy for EY-NPMA/(a,b), total carrier power for EY-NPMA/2ndMAX and idle/single transmission/collision for the RTCA modifications. See *Figure 3* and *Figure 4* for illustration.

EY-NPMA/(a,b): upon termination of its elimination burst, a station counts the number of time slots when the channel is still sensed busy (with longer bursts). If and only if that number is not greater than a and greater than b , the station is allowed to join in the yield phase (a and b are integer parameters, $E_{max} \geq a > b \geq 0$). Thus only bidders within a predetermined range from the highest bid are the winners.

EY-NPMA/2ndMAX: upon termination of its elimination burst, a station that senses the channel still busy counts the number of power drops around the boundaries of successive time slots (reflecting subsequent termination of longer elimination bursts). If and only if that number is 1, the station joins in the yield phase. Thus only second-highest bids win.

RTCA/1stCOLL: upon transmitting its pilot packet, a station waits for the other stations' reaction in the next slot. If the channel is sensed idle, it joins in the yield phase; otherwise it backs off and all other active stations start the elimination phase anew, with E_{max} decreased by 1. An active station reacts to a single pilot packet by transmitting a 1-slot burst and refrains from reacting if a collision of pilot packets was sensed. Thus in terms of the Dutch auction, an earliest bidder wins only if his bid is not unique.

RTCA/1stSINGLE: similarly as in RTCA/1stCOLL except that active stations react to collisions of pilot packets and refrain from reacting upon single pilot packets. Having transmitted a pilot packet and sensed the channel idle in the next slot, a station starts transmitting its packet immediately. Thus an earliest bidder wins only if his bid is unique.

Commenting from the auction perspective, the above modifications encourage the bidders to outbid each other while ensuring that the item is also affordable to the less wealthy bidders. EY-NPMA/(a,b) and EY-NPMA/2ndMAX resemble the 'reverted' Vickrey auction where the second-highest bid wins, but the winner pays the highest bid.




Figure 3. EY-NPMA/(2,0) (top) and EY-NPMA/2ndMAX (bottom) protocol cycles




Figure 4. RTCA/1stCOLL (top) and RTCA/1stSINGLE (bottom) protocol cycles

5. PERFORMANCE STUDY

A simulation experiment has been set up to investigate the performance of the four noncooperative scheduling policies of Sec. 4 and compare them to the basic EY-NPMA and RTCA policies. The network and scheduling policy parameter setting is shown in *Table 1*.

Table 1. Network parameter setting for simulation

number of stations	$N=10$
number of noncooperative stations	$NC=1..10$
elimination burst/timeout at cooperative stations	$1..E_{max}=15$ time slots uniform distribution
elimination burst/timeout at noncooperative stations	$1..E_{max}=15$ time slots shifted distribution (see below)
yield delay	$1..Y_{max}=3$ time slots uniform distribution

The noncooperative bidding strategy is assumed to be identical at all noncooperative stations and consist in shifting the distribution of the elimination burst/timeout by an integer constant $1 \leq m \leq E_{max}$. That is, $elimination_burst := \langle E+m \rangle_{1..E_{max}}$ time slots and $elimination_timeout := E_{max} - \langle E+m \rangle_{1..E_{max}} + 1$ time slots, where E is uniformly distributed over $1..E_{max}$ and $\langle \cdot \rangle_{I..J}$ denotes clipping of a random variable so that it falls into the interval $I..J$.

As a principal performance measure we take the per-cycle average station service rate $P_{succ} = Prob\{\text{station transmits packet successfully} \mid \text{station active in the present cycle}\}$. Ideally, $P_{succ} = 1/N$, although in practice $P_{succ} < 1/N$ on account of imperfect collision resolution. Note that P_{succ} is indicative, through queuing theory relationships, both of the average packet throughput and packet delay; thus focusing on P_{succ} relieves us from setting specific characteristics of offered load, packet sizes and buffer capacities. However, the choice of E_{max} and Y_{max} is sensitive since one has to weigh the growing efficiency of collision resolution against the degradation of throughput as they increase; the above values have been optimised experimentally.

In *Figure 5* and *Figure 6*, P_{succ} is plotted against NC and m for the basic EY-NPMA and RTCA policies, separately for cooperative and noncooperative stations. As expected, the plots show a high degree of unfairness, with the service rates at cooperative stations quickly dropping to a tiny fraction of the all-cooperative values (at $NC = 0$) as NC and m increase. On the other hand, noncooperative stations steal a large portion of the bandwidth until they become too many and so have few cooperative stations to steal from.

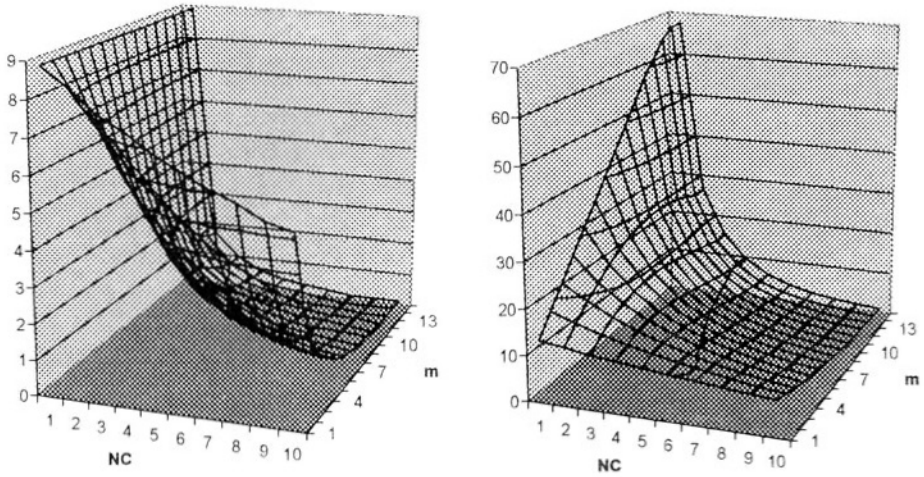


Figure 5. EY-NPMA, P_{succ} in % for cooperative (left) and noncooperative (right) stations

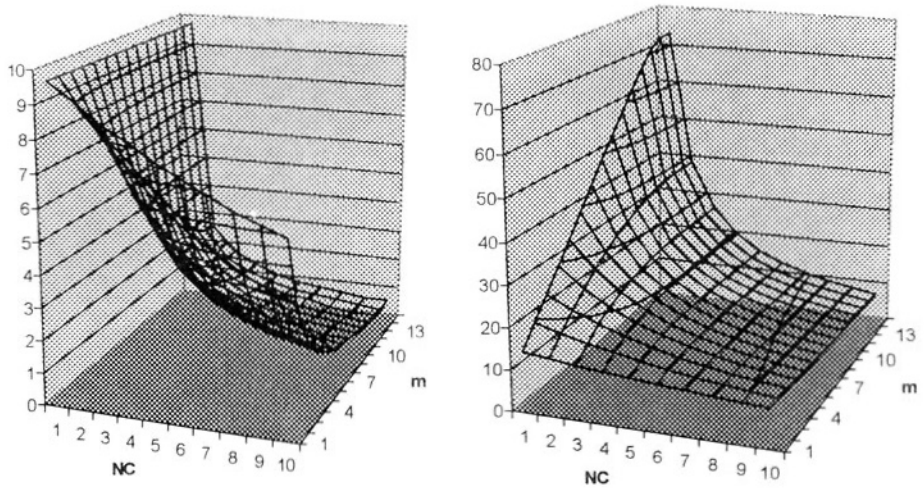


Figure 6. RTCA, P_{succ} in % for cooperative (left) and noncooperative (right) stations




Figure 7. EY-NPMA/(2,0), P_{succ} in % for cooperative (left) and noncooperative (right) stations

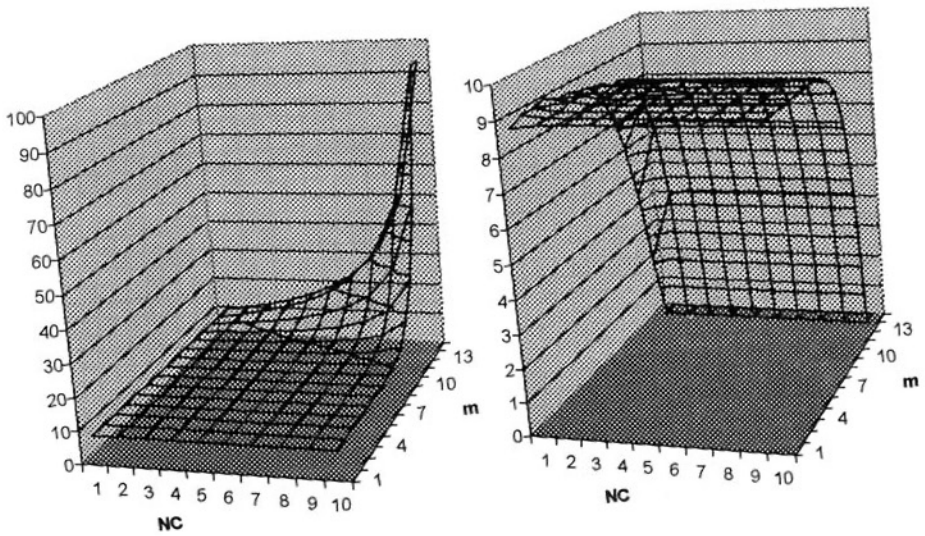


Figure 8. EY-NPMA/2ndMAX, P_{succ} in % for cooperative (left) and noncooperative (right) stations




Figure 9. RTCA/1stCOLL, P_{succ} in % for cooperative (left) and noncooperative (right) stations




Figure 10. RTCA/1stSINGLE, P_{succ} in % for cooperative (left) and noncooperative (right) stations

The situation for EY-NPMA is remedied by EY-NPMA/(2,0) (Figure 7) and EY-NPMA/2ndMAX (Figure 8), which appear an ideal disincentive for a station to become noncooperative: P_{succ} is hardly greater for noncooperative than for cooperative stations; moreover, cooperative stations even gain an advantage when they are few. This comes at a price, however – the all-cooperative P_{succ} is smaller than in Figure 5 – and the price is higher for EY-NPMA/(2,0) than for EY-NPMA/2ndMAX (a drop from 8.9% to 6.1% vs. 7.9%, respectively). The latter observation stems from the fact that while in EY-NPMA/2ndMAX there are winners of the elimination phase in each protocol cycle, in EY-NPMA/(2,0) there may be none.

RTCA/1stCOLL and RTCA/1stSINGLE (Figure 9 and Figure 10) paint a bit less ideal picture. Although the unfairness is distinctly reduced (cf. Figure 6), it still persists at some NC and m (as another illustration of the difficulty of the Dutch auction design). However, RTCA/1stSINGLE looks quite attractive in that it guarantees reasonable service rates for cooperative stations over almost the whole parameter range, while disfavouring noncooperative stations except for very small (thus not persistent) NC/N . (Paradoxically, this policy looks more suited for forceful bidding than RTCA/1stCOLL.) The all-cooperative P_{succ} drop (refer again to Figure 6) is insignificant for RTCA/1stSINGLE and acceptable for RTCA/1stCOLL – from 9.7% to 7.9%.

6. CONCLUSION

We have pointed out the need for noncooperative scheduling policies in contention-based MAC protocols for single-channel wireless LANs. Like in WAN internets, where the noncooperative management paradigm has already set in, one should account for the presence of noncooperative entities that will pursue their individual goals instead of adhering to a common-goal optimum policy. Building on two well-known protocols, EY-NPMA and RT, and drawing on the auction paradigm, we have proposed four noncooperative scheduling policies and investigated their performance via simulation to confirm that they prevent, or at least discourage, noncooperative stations from stealing the channel bandwidth from cooperative ones. This comes at the cost of a modest drop of the all-cooperative station service rate. Some questions for future research are:

- How can RTCA/1stCOLL and RTCA/1stSINGLE be further improved?
- What would a noncooperative scheduling policy look like in response to a general noncooperative strategy (rather than based specifically on a distribution shift, m)?
- Do effective carrier verifiable noncooperative scheduling policies exist?

- How to construct noncooperative counterparts of reservation-based and priority scheduling policies – is there a systematic way of seeking a Nash equilibrium analogue?

REFERENCES

- [1] I. Chlamtac and A. Ganz, *Evaluation of the Random Token Protocol for High-Speed and Radio Networks*, IEEE J. Select. Areas Commun., **SAC-5**, July 1987
- [2] ETSI TC RES, Radio Equipment and Systems; *High Performance Radio Local Area Network (HIPERLAN); Services and Facilities; Version 1.1*, RES 10, Jan. 1995
- [3] J. Konorski, K. Nowicki and J. Wozniak, *The Cooperative Random Token Protocol for High-Speed Radio LANs*, Proc. GLOBECOM'92, Orlando FL, Dec.1992
- [4] Y.A. Korilis, A.A. Lazar and A. Orda, *Architecting Noncooperative Networks*, IEEE J. Selected Areas Commun., **13**, Sept. 1995
- [5] S. Nanda, D.J. Goodman and U. Timor, *Performance of PRMA: A Packet Voice Protocol for Cellular Systems*, IEEE Trans. Veh. Technol., **40**, Aug. 1991
- [6] *Special Issue on Media Access Techniques for High-Speed LANs and MANs*, Computer Networks and ISDN Syst., **26**, March 1994
- [7] *Special Issue on Mobile Radio Communications*, IEEE J. Selected Areas Commun., **SAC-2**, July 1984
- [8] *Special Issue on Packet Radio Networks*, Proc. IEE, **75**, Jan. 1987
- [9] R. Wilson, *Auction Theory*, [in] J. Eatwell, M. Milgate and P. Newman (eds), The New Palgrave, MacMillan, London 1987.
- [10] *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Draft Standard IEEE 802.11, P802.11/D1