

# Throughput and Fairness in A Hybrid Channel Access Scheme for Ad Hoc Networks

Yu Wang J.J. Garcia-Luna-Aceves

Department of Computer Engineering

University of California at Santa Cruz

Santa Cruz, CA 95064, U.S.A.

{ywang,jj}@cse.ucsc.edu

**Abstract**—A novel hybrid channel access scheme that combines sender-initiated and receiver-initiated collision-avoidance handshakes is proposed for multi-hop ad hoc networks. The new scheme is compatible with the popular IEEE 802.11 MAC protocol and involves adding very simple queue management and book-keeping work mechanisms. Simulation experiments show that the new scheme can alleviate the fairness problems existing in applications running on either UDP or TCP with almost no degradation in throughput. More importantly, it is also shown that without explicit information exchange among nodes, the fairness problem cannot be solved conclusively.

## I. INTRODUCTION

Many contention-based channel access schemes have been proposed for multi-hop ad hoc networks in the literature [1–4]. Collision avoidance is very important in these schemes to combat the adverse effects of hidden terminals [2, 5] and can be largely divided into two categories, sender-initiated and receiver-initiated. In sender-initiated schemes, a node which has a packet to send initiates the collision avoidance handshake with an intended receiver. Usually, the handshake comprises the exchange of short ready-to-send (RTS) and clear-to-send (CTS) control packets between a sender-receiver pair followed by the transmissions of the actual data packet and optional acknowledgment packet. The RTS and CTS packets carry information about the duration of the handshake and serve as a channel reservation scheme to notify overhearing nodes to defer their access to the shared channel to avoid collisions. On the other hand, in receiver-initiated schemes, a node has to poll its neighbors actively to see if they have packets for itself.<sup>1</sup> The rationale behind receiver-initiated schemes is that, a receiver usually has better knowledge of the contention around itself and collision avoidance is more important at the receiver's side, given that the receiver needs to receive relatively long data packets successfully and such packets are more vulnerable to interference. It has been shown that, if the polled nodes always have packets for the polling node, receiver-initiated

schemes with proper collision avoidance procedures can outperform sender-initiated schemes due to reduced overhead of control packets [4]. Otherwise, the performance may degrade due to wasted transmissions of polling packets that poll inactive nodes with no packet for the polling node. The degradation in performance is more conspicuous in light to medium traffic load, unless a good traffic predictor is available at the polling node.

Despite the potential benefits of receiver-initiated schemes, they have not received wide acceptance. One reason is that sender-initiated schemes are more straightforward, because a sender has full knowledge of the packets in its queue and it can initiate the collision avoidance handshake only when necessary. On the other hand, for receiver-initiated schemes, a good traffic estimator and an appropriate polling discipline that can be adapted to the dynamic environments of ad hoc networks are mandatory and they have not been investigated sufficiently so far. Another reason is the prevalent acceptance of the IEEE 802.11 MAC protocol in the research community of ad hoc networks, which uses sender-initiated collision avoidance scheme. Many performance enhancements have been proposed and they are confined to the sender-initiated framework stipulated by the IEEE 802.11 MAC protocol.

Despite its popularity, the IEEE 802.11 MAC protocol can suffer severe fairness problems in multi-hop ad hoc networks where location-dependent contention is common. As is already pointed out in the research literature [6–9], some nodes in such networks are at a disadvantage in contending with other nodes due to their locations and may suffer severe degradation in throughput. Additionally, the commonly used binary exponential backoff (BEB) scheme, despite its robustness against repetitive collisions, can aggravate the fairness problem, because the node that succeeds in the last transmission period will gain access to the shared channel again with much higher probability while other nodes suffer starvation. Because it is difficult to provide quality-of-service (QoS) assurances without the fairness problem solved, some schemes have been proposed in the recent past to address the fairness problem. These schemes can be largely divided into two categories. In the first category, the goal is to achieve max-min fairness [6, 10, 11]. To

This work was supported in part by the US Air Force/OSR under Grant No. F49620-00-1-0330.

<sup>1</sup>Sometimes it is more appropriate to speak of *polling* and *polled* nodes rather than sender and receiver in receiver-initiated schemes.

be specific, these schemes try to reduce the ratio between maximum throughput and minimum throughput of flows, should it be at either a node's level or a flow's level. In the second category, the approach used in fair queueing for wireline networks is adapted to multi-hop ad hoc networks, taking into account the salient characteristics of such networks such as location-dependent contention, distributed coordination and possible spatial reuse [7–9, 12, 13]. These schemes usually exhibit some form of tradeoff between throughput and fairness. Nodes that are *leading* in channel access (in terms of throughput) will decrease their channel access activities, while nodes that are *lagging* will increase their channel access activities. In this way, nodes are encouraged to compete fairly but at the cost of increased contention, which may degrade the overall throughput. Despite the differences of backoff algorithms and information exchange among these schemes, the underlying channel access scheme remains largely the basic sender-initiated collision avoidance handshake, which can be less effective than a receiver-initiated scheme when a receiver has better knowledge of the contention around itself than the sender and can make use of such knowledge.

This motivates us to design an adaptive channel access scheme that makes use of both sender-initiated and receiver-initiated handshakes, because a receiver-initiated handshake is more desirable in some cases and a better tradeoff between throughput and fairness may be achieved. The new hybrid scheme should have the following desired properties. The scheme should fit within the IEEE 802.11 framework, even though it combines both sender-initiated and receiver-initiated handshake, and nodes implementing the new scheme should not break the existing network. The new scheme should still be simple and not introduce new types of control packets, as they may complicate implementation of the finite state machine of the protocol and degrade the overall network throughput unnecessarily when the basic sender-initiated scheme suffices.

The rest of the paper is organized as follows. In Section II, the new hybrid scheme is specified, which in fact is a very simple extension to the existing IEEE 802.11 MAC protocol and involves only some additional queue management and book-keeping mechanisms. In Section III, simulation experiments with the original IEEE 802.11 MAC protocol and the new hybrid scheme are presented for both UDP- and TCP-based traffic. It is shown that various degrees of the fairness problem exist in the original IEEE 802.11 MAC protocol even for simple network configurations with only two competing flows. Although the new hybrid scheme cannot solve the fairness problem conclusively, it can alleviate the fairness problem in some cases with almost no degradation in throughput. Section IV concludes this paper with directions for future work.

## II. THE NEW HYBRID CHANNEL ACCESS SCHEME

Talucci and Gerla [3] proposed MACA-BI (Multiple Access with Collision Avoidance - By Invitation) which was the first receiver-initiated MAC protocol. Garcia-Luna-Aceves and

Tzamaloukas [4] advanced that work and proposed several collision-free RIMA (receiver-initiated multiple access) protocols. Here collision-free means that, once a node sends a data packet, the data packet can be received by the receiver successfully, given that the channel is ideal without impairment and the only cause of failure to receive a packet is concurrent transmissions from multiple nodes. RIMA protocols achieve this collision-free property by introducing some additional types of short control packets and enforcing various collision-avoidance waiting periods. However, the receiver-initiated handshake in our proposed hybrid channel access scheme is simpler than that in the RIMA protocols. Firstly, it does not introduce new types control packets. Instead, only CTS packet is used as the polling packet. This is to maintain compatibility with the original IEEE 802.11 MAC protocol. Secondly, it does not include the various collision-avoidance waiting periods enforced in RIMA protocols. Instead, nodes defer access to the shared channel according to the network allocation vector (NAV) included in those overheard packets, which specifies the duration of the ensuing handshake. The reason is that the IEEE 802.11 MAC protocol itself cannot ensure collision-free data packet transmissions. We opt not to introduce additional collision-avoidance procedures and try to maintain compatibility with the existing protocols. Hence, the receiver-initiated collision avoidance handshake just includes three-way CTS-data-ACK exchange between polling and polled nodes. Though it is not expected that the hybrid scheme will improve throughput due to the lack of rigid collision-avoidance procedures, we do expect that it may still alleviate the fairness problem, because both a sender and a receiver can initiate a collision avoidance handshake alternately and the burden of contending for the shared channel is distributed to participating nodes according to the different degrees of contention they experience.

Our hybrid channel access scheme is built around the framework of the IEEE 802.11 MAC protocol. A node that implements this scheme operates alternately in two modes, sender-initiated (SI) and receive-initiated (RI). The SI mode is the default mode, which is in effect the same as the original IEEE 802.11 MAC protocol. The usual four-way RTS-CTS-data-ACK handshake is used in the SI mode. The RI mode is the new mode introduced in the hybrid scheme. The aforementioned receiver-initiated three-way collision avoidance handshake is used in the RI mode, which is triggered only when the SI mode does not perform well. In this mode, more cooperation between a pair of sending and receiving nodes is required, because both of them need to enter the RI mode before the receiver-initiated handshake can be initiated.

To facilitate our exposition, the states of both sending and receiving nodes are shown in Figure 1 and are explained separately.

A sender enters the *RI setup* mode when it sends the same RTS packet for more than one half of the times allowed in the IEEE 802.11 MAC protocol but gets no response from the

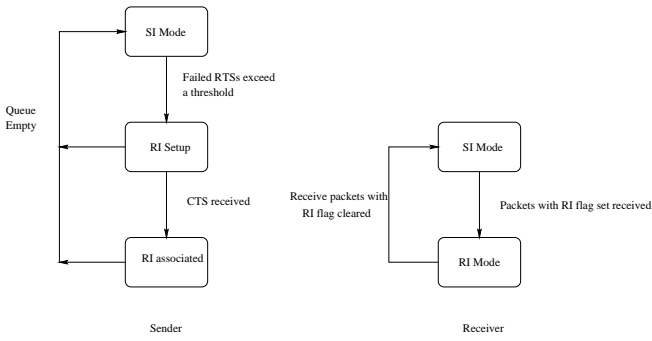


Fig. 1. State transition diagram of sending and receiving nodes

intended receiver. Failure to solicit a response from the intended receiver usually implies that contention around the receiver is so severe that the receiver is prevented from responding. Hence, it is more appropriate to let the receiver start the collision-avoidance handshake when this happens. After the sender enters the *RI setup* mode, it sets the RI flag in all the subsequent RTS packets and other packets that it sends out and requests the intended receiver to enter the RI mode as well. During this stage, the node keeps sending RTS packets following the usual collision-avoidance procedures, because it has not established an *association* with the intended receiver. There are two possible outcomes. One outcome is that the node never gets any CTS packet from the intended receiver. In this case, the sender may declare the receiver down after it has to drop a few packets. The other outcome is that it receives CTS packet from the intended receiver. In this case, the sender enters the *RI-associated* mode and will not send an RTS to the receiver thereafter. This helps reduce the contention around the receiver and also makes the sender available for accepting polling requests from the receiver. To keep the receiver in the RI mode, the sender keeps setting the RI flag in all the data packets that it sends out. The RI flag is cleared only when the sender's queue becomes empty.

The addition of the RI flag is the only necessary change to the frame structures in the IEEE 802.11 standard. Figure 2 illustrates the frame structure of the IEEE 802.11 RTS frame (ref. Fig. 13 in Page 35 and Fig. 16 in Page 41 of the IEEE 802.11 standard [1]). As the *More data* bit is not used in ad hoc mode according to the standard, it may be reused as the RI flag to indicate if the RI mode is on or not. Nodes that do not implement the hybrid channel access scheme can safely ignore this bit.

The receiver enters and stays in the RI mode when it receives RTS packets or data packets destined to it with the RI flag set. The receiver then generates RI-response packets (which are in fact self-initiated CTS packets) and multiplexes them with other data packets in its MAC queue. However, the receiver should not generate RI-response packets indiscriminately when it receives a packet with the RI flag on, lest serious fairness problem may occur. This can be explained as follows. When an RI-response packet becomes the head-of-line (HOL)

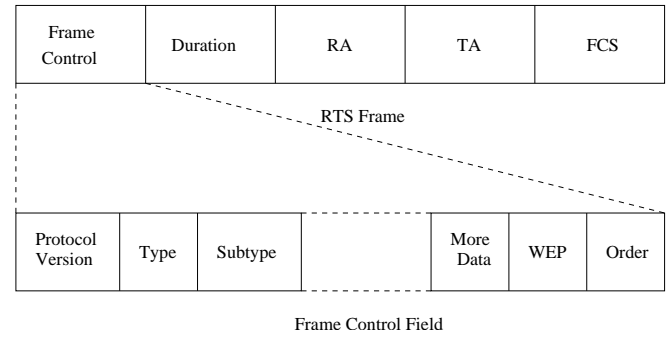


Fig. 2. Illustration of the IEEE 802.11 frame structure

packet of a receiver's queue, the node will send a self-initiated CTS to the sender, which in fact serves as the ready-to-receive (RTR) packet to poll the sender in the RIMA protocols [4]. If the sender replies with a data packet with the RI flag still on, which implies that there are more packets in its sending queue, the receiver will add another RI-response packet to the end of its queue. If there is no packet for other nodes intervened in the MAC queue, the receiver will be *locked into* the sender and will keep sending CTS packets to it. In this way, they may monopolize the shared channel for a long time, which obviously defeats the purpose of the hybrid scheme. Hence, when a node receives a packet with the RI flag on, it checks its HOL packet to see whether it is an RI-response packet for the node that just sent this packet. If so, the RI request is ignored; otherwise, it is added to the end of its MAC queue.

The RI-response packets are treated like RTS packets for normal data packets. That is, when they are served via a successful receiver-initiated CTS-data-ACK handshake or when they are transmitted more than the times allowed for RTS packets in the IEEE 802.11 standard, they will be removed from the MAC queue. Such precautions are necessary. One reason is to avoid excessive delay or deadlock when the sending node is down or moves out of range of the receiving nodes. Another reason is to ensure fairness so that neighboring nodes may still get chances to initiate handshake with the receiver or other nodes.

The above specification clearly shows that, with some additional queue management and book-keeping work, the existing IEEE 802.11 can be easily extended to support a receiver-initiated scheme while maintaining compatibility.

### III. SIMULATION RESULTS

In our simulation experiments, we focus on how two competing flows share the available channel resource in some simple network configurations. These configurations are shown in Figures 3 and 4, in which a dashed line means that two nodes can hear each other's transmissions and an arrow indicates an active flow between two nodes. Nodes without any line in-between are hidden from each other.

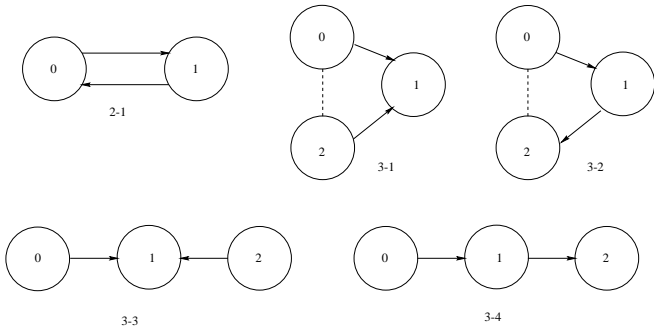


Fig. 3. Networks with 2 or 3 nodes

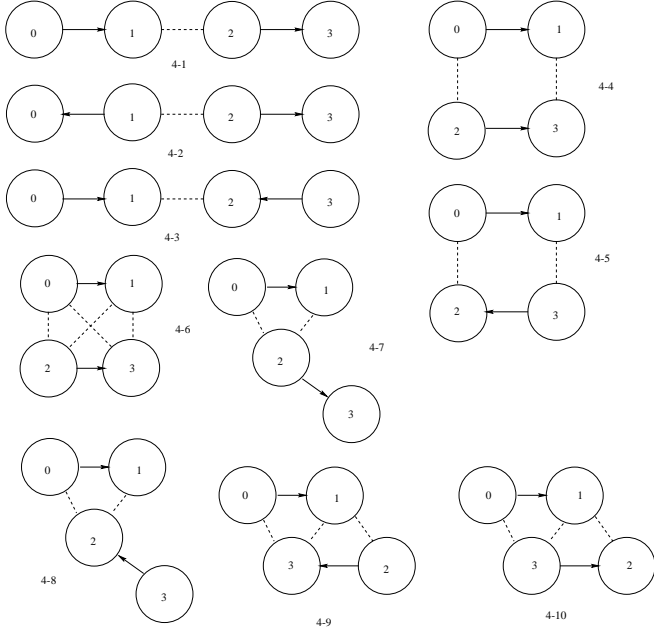


Fig. 4. Networks with 4 nodes

We use GloMoSim 2.0 [14] as the network simulator and implement the new hybrid scheme based on its implementation of the IEEE 802.11 MAC protocol for fair comparison. Direct sequence spread spectrum (DSSS) parameters are used throughout the simulations, which are shown in Table I. The raw channel bit rate is 2Mbps.

We investigate the performance of the original IEEE 802.11 MAC protocol and the new hybrid scheme under both UDP- and TCP-based traffic. In the first set of simulation experiments, there are two competing UDP-based flows. For each flow, one node keeps sending data packets to the other at a constant bit rate, such that the sending queue is always non-empty. UDP is the underlying transport layer, thus no acknowledgment packet is sent back to the initiating node. Simulation results are shown in Tables II and III.

In Table II, the performance of the original IEEE 802.11 MAC protocol is shown. It is clear that, for configurations 4-1 and 4-8, some nodes are almost denied access to the shared

TABLE I  
IEEE 802.11 PROTOCOL CONFIGURATION PARAMETERS

RTS	CTS	data	ACK	DIFS	SIFS
20-byte	14-byte	1460-byte	14-byte	50 $\mu$ sec	10 $\mu$ sec
contention window		slot time	sync. time	prop. delay	
31-1023		20 $\mu$ sec	192 $\mu$ sec	1 $\mu$ sec	

channel and suffer severe degradation in throughput. For other configurations, it is unnecessary to use the new hybrid scheme. In Table III, the performance of the original IEEE 802.11 MAC protocol and the hybrid scheme is shown.<sup>2</sup> It is clear that the fairness problems in configurations 4-1 and 4-8 are alleviated significantly without a sacrifice in throughput. The RI mode is triggered unnecessarily only in three other configurations and has almost no negative effect on throughput. It is also worth noting that the aggregate throughput in all these network configurations remains almost the same despite the fact that the fairness problem exists in some configurations. This shows the importance of considering individual node's throughput in the performance evaluation of these MAC protocols.

In the second set of simulation experiments, there are two competing TCP-based flows. We use the FTP/Generic application provided in GloMoSim, in which an FTP client simply sends data packets to an FTP server without the server sending any control information back to the client other than the acknowledgment packets required by TCP. Only after a packet is indicated success of delivery by the transport layer (TCP), will the client send the next data packet. Here it should be noted that the acknowledgment packet from TCP is still regarded as a normal data packet from the view of MAC layer. Hence, due to the peculiarities of the application, it is disadvantageous for the MAC layer to transmit more than one packet at a time. When this is applied to the hybrid scheme, it means that it is more desirable for a node and its peer to leave RI associated mode just after a CTS-Data-ACK handshake is done so that they can switch the roles of sender and receiver timely. So in the implementation of the hybrid scheme, we make the necessary changes to take this into account. Simulation results are shown in Tables IV and V.

It is clear from Table IV that the fairness problem is much more severe for two competing TCP-based flows than for the case of UDP-based flows. In some cases, such as configurations 4-1, 4-7 and 4-8, one FTP flow is denied access to the shared channel for most of the time. Here throughput of zero does not mean that TCP connection is not set up. Instead, because the throughput is extremely low (on the order of a few kilobytes per second) for these flows, it is meaningless to show the statistics here. When the hybrid channel access scheme is used, in some cases it is triggered and performs almost the same as the original IEEE 802.11 MAC protocol while in some

<sup>2</sup>When the RI mode is not triggered in some network configurations, the hybrid scheme is the same as the original IEEE 802.11. For simplicity, performance of both schemes in these configurations is not shown here.

TABLE II  
FAIRNESS PROBLEMS IN THE ORIGINAL IEEE 802.11 – TWO CBR FLOWS

Config #	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
2-1	0 → 1	8.06e+05	1 → 0	7.99e+05	1.60e+06
3-1	0 → 1	8.06e+05	2 → 1	7.97e+05	1.60e+06
3-2	0 → 1	7.97e+05	1 → 2	8.07e+05	1.60e+06
3-3	0 → 1	7.61e+05	2 → 1	7.83e+05	1.54e+06
3-4	0 → 1	7.69e+05	1 → 2	8.39e+05	1.61e+06
<b>4-1</b>	0 → 1	8.34e+04	2 → 3	1.50e+06	1.58e+06
4-2	1 → 0	8.20e+05	2 → 3	8.14e+05	1.63e+06
4-3	0 → 1	6.88e+05	3 → 2	7.09e+05	1.40e+06
4-4	0 → 1	8.24e+05	2 → 3	8.08e+05	1.63e+06
4-5	0 → 1	8.08e+05	3 → 2	7.95e+05	1.60e+06
4-6	0 → 1	8.07e+05	2 → 3	7.95e+05	1.60e+06
4-7	0 → 1	7.83e+05	2 → 3	8.24e+05	1.61e+06
<b>4-8</b>	0 → 1	1.55e+06	3 → 2	2.81e+04	1.58e+06
4-9	0 → 1	7.34e+05	2 → 3	8.09e+05	1.54e+06
4-10	0 → 1	7.81e+05	3 → 2	8.26e+05	1.61e+06

TABLE III  
THROUGHPUT COMPARISON FOR THE IEEE 802.11 AND THE HYBRID SCHEME (WITH RI MODE) – TWO CBR FLOWS

Config #	Scheme	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
3-3	802.11	0 → 1	7.61e+05	2 → 1	7.83e+05	1.54e+06
	+RImode	0 → 1	7.94e+05	2 → 1	7.74e+05	1.61e+06
<b>4-1</b>	802.11	0 → 1	8.34e+04	2 → 3	1.50e+06	1.58e+06
	+RImode	0 → 1	3.69e+05	2 → 3	1.23e+06	1.60e+06
4-3	802.11	0 → 1	6.88e+05	3 → 2	7.09e+05	1.40e+06
	+RImode	0 → 1	6.65e+05	3 → 2	6.43e+05	1.31e+06
<b>4-8</b>	802.11	0 → 1	1.55e+06	3 → 2	2.81e+04	1.58e+06
	+RImode	0 → 1	1.28e+06	3 → 2	3.19e+05	1.60e+06
4-9	802.11	0 → 1	7.34e+05	2 → 3	8.09e+05	1.54e+06
	+RImode	0 → 1	8.15e+05	2 → 3	7.42e+05	1.56e+06

TABLE IV  
FAIRNESS PROBLEMS IN THE ORIGINAL IEEE 802.11 – TWO FTP FLOWS

Config #	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
2-1	0 → 1	4.66e+05	1 → 0	4.68e+05	9.34e+05
3-1	0 → 1	4.72e+05	2 → 1	4.73e+05	9.45e+05
3-2	0 → 1	4.56e+05	1 → 2	4.79e+05	9.35e+05
3-3	0 → 1	4.92e+05	2 → 1	3.84e+05	8.75e+05
<b>3-4</b>	0 → 1	3.52e+05	1 → 2	5.48e+05	9.00e+05
<b>4-1</b>	0 → 1	0	2 → 3	9.26e+05	9.29e+05
<b>4-2</b>	1 → 0	(4.88±1.03)e+05	2 → 3	(4.53±1.02)e+05	9.42e+05
<b>4-3</b>	0 → 1	(5.30±4.32)e+05	3 → 2	(3.92±4.38)e+05	9.22e+05
4-4	0 → 1	4.49e+05	2 → 3	4.36e+05	8.84e+05
4-5	0 → 1	4.75e+05	3 → 2	4.74e+05	9.49e+05
4-6	0 → 1	4.75e+05	2 → 3	4.74e+05	9.49e+05
<b>4-7</b>	0 → 1	9.28e+05	2 → 3	0	9.30e+05
<b>4-8</b>	0 → 1	9.29e+05	3 → 2	0	9.30e+05
4-9	0 → 1	4.27e+05	2 → 3	4.49e+05	8.76e+05
<b>4-10</b>	0 → 1	3.76e+05	3 → 2	5.26e+05	9.02e+05

TABLE V

THROUGHPUT COMPARISON FOR THE IEEE 802.11 AND THE HYBRID SCHEME (WITH RI MODE) – TWO FTP FLOWS

Config #	Scheme	Flow #	Throughput (bps)	Flow #	Throughput (bps)	Aggregate (bps)
3-3	802.11	0 → 1	4.92e+05	2 → 1	3.84e+05	8.75e+05
	+RImode	0 → 1	4.08e+05	2 → 1	4.66e+05	8.73e+05
4-2	802.11	1 → 0	(4.88±1.03)e+05	2 → 3	(4.53±1.02)e+05	9.42e+05
	+RImode	1 → 0	(4.39±0.99)e+05	3 → 2	(5.02±0.98)e+05	9.40e+05
4-3	802.11	0 → 1	(5.30±4.32)e+05	3 → 2	(3.92±4.38)e+05	9.22e+05
	+RImode	0 → 1	(3.97±0.71)e+05	3 → 2	(4.55±0.78)e+05	8.52e+05

other cases it is not triggered at all. Table V shows only the results when there exist differences between these two schemes. It is clear that the hybrid scheme performs slightly better than the original 802.11 MAC scheme for configurations 3-3 and 4-2 while it performs much better in terms of fairness in configuration 4-3 though there is about 8% degradation in throughput. It is more difficult to improve fairness of TCP-based flows than UDP-based flows due to the flow control and congestion avoidance functions in TCP. A node that suffers excessive packet loss or delay decreases its sending rate according to TCP, which can aggravate the fairness problem already existing at the MAC layer. In such cases, even the hybrid scheme can lose its effectiveness. It can be reasoned that, without more explicit information exchange among nodes, the fairness problem cannot be solved conclusively.

#### IV. CONCLUSION

We have proposed a new hybrid channel-access scheme that includes both sender-initiated and receiver-initiated collision avoidance. This is based on the observation that a receiver-initiated scheme is more appropriate when receivers are more knowledgeable of the contention around themselves and can compete for the channel more effectively. By adaptively sharing the burden of initiating the collision-avoidance handshake between the nodes that experience different levels of contention, better fairness may be achieved with almost no degradation in throughput. An attractive feature of the new scheme is that it is a simple extension to the existing IEEE 802.11 MAC protocol and maintains compatibility with the standard. Simulation experiments of the IEEE 802.11 MAC protocol and the new scheme show that, although the proposed hybrid scheme does not solve the fairness problem conclusively, it does alleviate the fairness problem in some cases without sacrificing much throughput and simplicity. The difficulty of improving fairness for TCP-based flows is also shown. A promising topic for future work consists of combining the new hybrid scheme

with some proposed mechanisms [9, 12, 13] that try to approximate fair queueing for ad hoc networks to achieve some QoS assurances.

#### REFERENCES

- [1] IEEE Computer Society LAN MAN Standards Committee, ed., *IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. IEEE Std 802.11-1997, The Institute of Electrical and Electronics Engineers, New York, 1997.
- [2] C. L. Fullmer and J. J. Garcia-Luna-Aceves, "Solutions to Hidden Terminal Problems in Wireless Networks," in *Proc. of ACM SIGCOMM '97*, 1997.
- [3] F. Talucci and M. Gerla, "MACA-BI (MACA by Invitation): A Receiver Oriented Access Protocol for Wireless Multihop Networks," in *Proc. of PIMRC '97*, 1997.
- [4] J. J. Garcia-Luna-Aceves and A. Tzamaloukas, "Receiver-initiated Collision Avoidance in Wireless Networks," *ACM Wireless Networks*, vol. 8, pp. 249–263, 2002.
- [5] F. A. Tobagi and L. Kleinrock, "Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple-access Modes and the Busy-tone Solution," *IEEE Trans. on Communications*, vol. 23, no. 12, pp. 1417–1433, 1975.
- [6] T. Ozugur, M. Naghshineh, P. Kermani, C. M. Olsen, B. Rezvani, and J. A. Copeland, "Balanced Media Access Methods for Wireless Networks," in *Proc. of ACM/IEEE MOBICOM '98*, pp. 21–32, Oct. 1998.
- [7] T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC Layer Fairness in Wireless Packet Networks," in *ACM Mobicom 2000*, (Boston, MA, U.S.), Aug. 2000.
- [8] N. H. Vaidya, P. Bahl, and S. Gupta, "Distributed Fair Scheduling in a Wireless LAN," in *ACM Mobicom 2000*, (Boston, MA, U.S.), Aug. 2000.
- [9] H. Luo, S. Lu, and V. Bharghavan, "A New Model for Packet Scheduling in Multihop Wireless Networks," in *ACM Mobicom 2000*, (Boston, MA, U.S.), Aug. 2000.
- [10] B. Bensaou, Y. Wang, and C. C. Ko, "Fair Medium Access in 802.11 Based Wireless Ad-Hoc Networks," in *IEEE/ACM MobiHoc Workshop*, Aug. 2000.
- [11] X. Huang and B. Bensaou, "On Max-min Fairness and Scheduling in Wireless Ad-Hoc Networks: Analytical Framework and Implementation," in *ACM MobiHoc '01*, 2001.
- [12] H. Luo and S. Lu, "A Topology-Independent Fair Queueing Model in Ad Hoc Wireless Networks," in *IEEE ICNP 2000*, (Osaka, Japan), Nov. 2000.
- [13] H. Luo, P. Medvedev, J. Cheng, and S. Lu, "A Self-Coordinating Approach to Distributed Fair Queueing in Ad Hoc Wireless Networks," in *IEEE INFOCOM 2001*, Apr. 2001.
- [14] X. Zeng, R. Bagrodia, and M. Gerla, "GloMoSim: a Library for Parallel Simulation of Large-scale Wireless Networks," in *Proc. of the 12th Workshop on Parallel and Distributed Simulations*, May 1998.