Achieving Fairness in Wireless LANs by Enhanced IEEE 802.11 DCF *

Nagesh S. P. Nandiraju†, Hrishikesh Gossain‡, Dave Cavalcanti†, Kaushik R. Chowdhury†, Dharma P. Agrawal†

† Center for Distributed and Mobile Computing, Dept. of ECECS, University of Cincinnati - Cincinnati,OH 45221
(nandirns, cavalcdt, kaushir, dpa)@ececs.uc.edu

‡Mesh Research and Development Group
Motorola, Inc., Maitland, Florida
Hrishikesh.Gossain@motorola.com

Abstract—**Over the past few years, Wireless Local Area Networks (WLANs) have gained an increased attention and a large number of WLANs are being deployed in universities, companies, airports etc. Majority of the IEEE 802.11 based WLANs employ Distributed Coordination Function (DCF) in Wireless Access Points (AP) to arbitrate the wireless channel among Wireless Stations (STAs). However, DCF poses serious unfairness problem between uplink and downlink flows. To overcome this unfairness problem, we propose a simple enhancement to the IEEE 802.11 DCF which provides priority to the AP and thus enables it to acquire a larger share of the channel when required. We have demonstrated the unfairness problem through systematic measurements in an experimental test bed of WLAN using the legacy 802.11 DCF. We also developed analytical models to calculate the throughput of AP and the STAs and verify these results through thorough simulations in ns-2. We observe that our simulation results find in good agreement with our analytical models. Results show that our proposed enhancement achieves a fair distribution of bandwidth and improves the throughput (by nearly 300%) for the downlink flows as compared to the DCF, without severely affecting the performance of uplink flows.**

Keywords: *Fairness, MAC protocols, Performance Evaluation, Test beds, WLANs.*

1. INTRODUCTION

In the last decade, Wireless LANs experienced a proliferating growth due to their flexibility and ubiquitous nature. In particular, WLAN hotspots are typically found in universities, companies, airports, shopping malls, etc. With the increasing interest in the integration of various wireless networks (4G networks), WLANs are gaining more popularity than ever. The capacity of WLANs has rapidly increased from 2Mbps to 54Mbps and proposals (IEEE 802.11n) to achieve nearly 100Mbps are also underway. Many efforts are also being made to make QoS provisions for real time traffic (IEEE 802.11e).

Most of the current WLAN implementations are based on the IEEE 802.11 [1] standard, which supports two basic mechanisms for channel arbitration: Distributed Co-ordination Function (DCF) and Point Coordination Function (PCF). The implementation of DCF in IEEE 802.11 compliant devices is mandatory while provision of PCF is optional. DCF is based on the traditional CSMA/CA paradigm and provides equal channel access privileges to all participating Wireless Stations (STAs). In contrast, PCF is a centralized scheduling algorithm. It requires a point coordinator (PC) at the AP to control the channel access. The default scheduling algorithm of IEEE 802.11 PCF is a round robin scheme and may not always be ideal. Due to the inherent complexity involved with the deployment of PCF [2], most of the current implementations of IEEE 802.11, even in hot spot scenarios, use DCF access mechanism.

However, DCF poses serious unfairness problem between uplink and downlink flows. With DCF, the channel share of the AP would be a fraction of total number of transmitting STAs in its service area. All STAs, including the AP, have the same channel access privileges. As a result, the share of the channel obtained by the AP is nearly equal to the share of any other STA under its coverage. This results in unfair sharing of the bandwidth among uplink and downlink flows. All the downlink flows (flows that are destined for wireless stations) have to utilize the AP's channel share while the uplink flows originating from different STAs enjoy a larger share. With the increase in the number of STAs under the AP's coverage, the downlink flows would suffer from relatively low share of the available bandwidth.

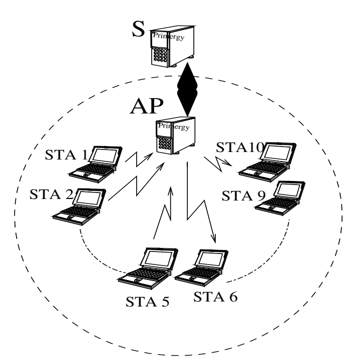

Figure 1: IEEE 802.11 DCF based WLAN Access Scenario

In order to overcome this unfairness problem, we propose a simple MAC layer enhancement to IEEE 802.11 DCF, called Bidirectional-DCF (BDCF). We specifically address the uplink/downlink unfairness by providing the AP with more contention free transmission opportunities when high

*This work was supported by the Ohio Board of Regents' Doctoral Enhancement Funds

load is experienced. In particular, if an AP's MAC receives a DATA packet, instead of transmitting a regular MAC layer ACK, it checks the buffer for an outstanding packet to any of the STAs in its Basic Service Set (BSS). If a packet is found, then it will send the DATA with a piggybacked ACK after SIFS time, thus eliminating the need for a fresh channel contention to transmit this packet. In this way, the AP gets a preferential treatment resulting in a relatively higher bandwidth share as compared to its STAs. Clearly, this kind of preferential treatment for the AP is desirable in hot spots scenarios. This is because most of the users accessing internet in these hotspots use applications (e.g. email, web browsing, Internet radio, etc.) that typically generate large volume of downlink traffic for a single uplink request. It is worthwhile to note that if the AP does not have traffic to send in downlink direction, BDCF works exactly as DCF.

Our specific contributions through this paper are:

- Experimental demonstration of the unfairness problem in an IEEE 802.11b based test-bed.
- A simple enhancement to DCF for overcoming the unfairness problem in AP based networks.
- An analytical model to evaluate the throughputs of AP and STAs complemented by extensive simulation study.

Our analytical and simulation results show that BDCF has a better throughput and delay performance for downlink flows and also has a fair channel sharing in both directions as compared to DCF. In addition, we compare BDCF with DCF+ [5], which uses a similar idea to reduce MAC layer overhead and increase the throughput, while it does not give any priority to the AP.

The outline for the rest of this paper is as follows. In section 2, we experimentally demonstrate the unfairness problem caused by DCF in WLAN hot spot scenarios. We describe our proposed BDCF mechanism in section 3. Next, we develop an analytical model for the uplink and downlink throughputs with BDCF in section 4. Section 5 provides comprehensive simulation results, comparing BDCF with DCF and DCF+. In section 6, we discuss the related work, and finally, we conclude the paper in Section 7, highlighting some open problems and future research directions.

2. THE DCF UNFARINESS PROBLEM IN WLAN HOT SPOTS

In this section we illustrate the unfairness among downlink/uplink flows when DCF is employed in a WLAN. We have set up an experimental test bed to model the typical WLAN hot spot scenario as shown in Figure 1.

We configured an infrastructure based IEEE 802.11b network with a data rate of 11Mbps and connected the AP to a desktop PC (with an Intel Pentium 4 processor). We start the server S in the same machine, so that the connection between the AP and the server is not a bottleneck. Furthermore, the AP uses the *hostap* driver package [12], and the STAs use IEEE 802.11b USB adapters, Netgear WG311 or Atheros wireless cards. We have systematically examined the TCP and UDP throughput performance with a symmetric and an asymmetric traffic configuration. TCP and UDP traffic is generated by Iperf v1.7 [13] which runs in both the client and server modes.

We first consider an asymmetric traffic scenario where 3 UDP flows originating at the different wireless nodes destined to the server (called uplink flows) and 7 UDP flows from the server towards the wireless nodes (called downlink flows). Each flow generates traffic at a rate of 3Mbps, which is enough to saturate the wireless link. We have conducted 10 different runs for each traffic scenario and measured the throughputs of individual flows. The throughputs of flows in the same direction do not have much variation. Thus, for the sake of brevity, we only show the aggregate throughputs of uplink and downlink flows for all the four scenarios in Figure 2. As can be seen from the plot (in Figure 2 (a)), in the case of asymmetric UDP traffic, the three uplink flows obtain a throughput of 2249.3 Kbps (749.6 Kbps per flow on an average) while the 7 downlink flows obtain a throughput of only 1798.14 Kbps (256.85 Kbps per flow on an average). It can be observed that the throughput of an individual downlink flow is nearly 1/3 of the throughput achieved by any uplink flow. Clearly the MAC level fairness achieved by DCF leads to an undesirable situation. For example, consider a typical WLAN hotspot scenario where a couple of students (who are a part of peer-to-peer file sharing network) are uploading songs/movies through the wireless network. Their applications (are uplink) consume fairly large share of the wireless channel and thus limit the bandwidth for the downlink traffic. Thus other users who are checking their e-mail or using other predominantly downlink traffic based applications experience larger download delays and increasing frustration.

These results motivate the need for a preferential treatment to the AP in order to allot a fair share of bandwidth for the downlink flows. We propose a simple enhancement to DCF in the next section, which prioritizes the AP without the requirement of any additional information and achieves the required fairness among the uplink and downlink flows.

3. BIDIRECTIONAL DISTRIBUTED CO-ORDINATION FUNCTION (BDCF)

In order to handle the unfair bandwidth availability to the downlink traffic in WLAN hot spot scenarios, we propose a Bidirectional DCF (BDCF), which provides preferential treatment to the downlink flows at the AP. For implementing BDCF, we modified IEEE 802.11 DCF to support piggybacking of ACK packets in the DATA transmission from the AP. Similar to IEEE 802.11 DCF, BDCF supports both basic (DATA-ACK) and 4-way handshake (RTS-CTS-DATA-ACK) channel access mechanisms. In the remainder of this section we describe the details of BDCF.

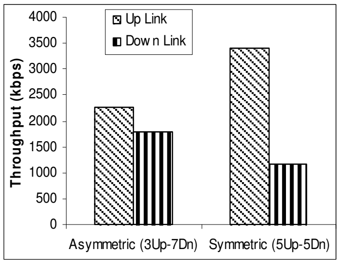

Figure 2 (a) UDP Traffic

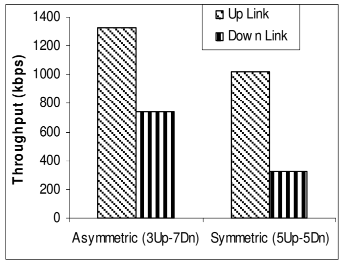

Figure 2 (b) TCP Traffic

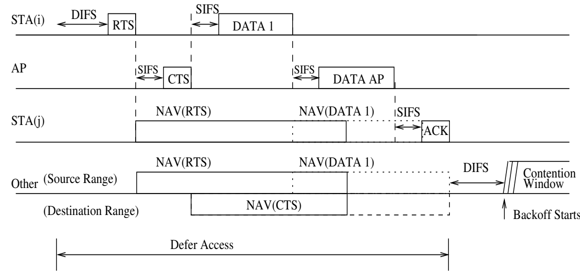

Figure 3. Timing diagram of BDCF with RTS-CTS handshake

3.1 BDCF with RTS-CTS handshake

We first start by describing the operation of BDCF when RTS-CTS handshake is enabled. Figure 3 shows the timing diagram of the BDCF operation. Initially all STAs and the AP contend for the channel with equal privilege. During the contention resolution period, if the AP gets the access to the channel, BDCF works exactly like DCF. The difference comes when a STA wins the channel contention. Figure 3 illustrates this case in detail. STA(*i*) initially sends an RTS to the AP in order to reserve the channel for its DATA transmission. The AP responds this request of STA(*i*) by sending a CTS. Upon reception of the CTS packet, STA(*i*) sends its DATA packet ($DATA_i$) to the AP. After receiving $DATA_i$, the AP checks its MAC buffer for any packet to transmit. If no packet is found, the AP simply sends the corresponding ACK to STA(*i*). However, if the AP has an outstanding packet for any of its serving STA(*j*) (where *j* may be same as *i*), it transmits that DATA packet (indicated by $DATA_{AP}$ in Figure 3) with a piggybacked ACK, after a SIFS time period. If STA(*i*) senses transmission from the AP after a SIFS period, it will implicitly recognize the ACK sent by the AP. As long as the STAs are within the coverage region of the AP, they are able detect the piggybacked ACK (STA(*i*) was waiting for an ACK for DATA_i from AP). Since all the STAs in the serving BSS should be within the coverage of AP our assumption is valid.

After a successful reception of the DATA_{AP}, the destination STA(*j*) will send back an ACK to the AP in the usual way. However, if the AP does not receive the initial data frame ($DATA_i$) correctly, it does not send any $DATA_{AP}$ or ACK, thus STA(*i*) becomes aware of its unsuccessful transmission and schedules a retransmission. It is important to note that BDCF allows piggybacking of DATA packets only at the AP, thus avoiding the chance of forming any cycles (phenomenon where AP and any one of STAs repeatedly access and capture the channel).

Furthermore, when the AP transmits the DATA packet with piggybacked ACK ($DATA_{AP}$), it freezes its backoff timer, such that the transmission of $DATA_{AP}$ is totally transparent to the regular channel contention at the AP.

The changes made at the MAC layer for implementing BDCF does not require any changes to the upper layers and thus totally transparent to the upper layers. As we will show in the next section, by using this simple enhancement, BDCF can provide fair access to the wireless channel in both the directions, irrespective of the type of transport layer (UDP or TCP) under consideration. It should also be noted that with BDCF, we avoid sending ACK, RTS, and CTS packets while sneaking a $DATA_{AP}$ packet from the AP, thus reducing MAC layer control overhead as compared to DCF. Moreover we also reduce the time wasted in any channel contention and backoff mechanism.

In order to avoid any negative effect of BDCF on the uplink flows when there are fewer downlink flows than uplink flows, we adopt a dynamic piggybacking strategy. With this strategy the AP records the number of STAs transmitting the uplink flows and the number of STAs receiving downlink flows over a time window. The AP piggybacks a DATA packet only with a probability equal to the ratio of downlink and uplink flows. This way, we ensure that the downlink flows do not get any undue advantage and influence fairness of the system.

The working principle of BDCF in the basic access mechanism (without RTS-CTS handshake) is similar to its operation when RTS-CTS handshake is enabled.

4. ANALYTICAL EVALUATION OF BDCF

As pointed out in section 3, BDCF ensures that the downlink flows gets a fair share of the system bandwidth. We validate this claim by deriving analytical expressions for the throughput of the AP (downlink) and the remaining STAs (uplink) when BDCF is used with the RTS-CTS handshake mechanism. We assume *n* fixed STAs and one AP, each of them having a packet to transmit at all times (saturation traffic conditions with equal number of uplink and downlink flows). We also assume perfect channel conditions and no hidden terminals. We define the throughput by the following equation

$$S = \frac{E[\textit{payload information transmitted in a slot time}]}{E[\textit{length of slot time}]} \quad (1)$$

We first study various events that can occur in any arbitrary slot, the time durations for the individual events, and finally the total transmission time. It is assumed that each station including the AP transmits in a randomly chosen slot time with probability τ and collisions occur with a constant probability *p,* irrespective of the number of previous collisions before a successful transmission. We consider the slot time to be of length σ. Let E_p, H and δ denote the average packet payload size, the packet header size (calculated as $PHY_{hdr} + MAC_{hdr}$) and the propagation delay, respectively. E_p and H are measured in time units. For simplicity we define, W = CW_{min} and m as the maximum backoff slots so that, $CW_{max} = 2^m * CW_{min}$ .The probability of transmission (τ) in a given time slot as derived in [14] is

$$\tau = \frac{2(1-2p)(1-p)}{(1-2p)(W+1)+p(W(1-(2p)^m)}$$

In what follows, we identify the various events that can occur in an arbitrary time slot under BDFC operation and calculate their probabilities of occurrence. Recall from section 3 that in BDCF, the AP transmits a data packet under two conditions: by contention or by piggybacking.

I. When AP contends for the channel

In this case, the AP contends for the channel afresh with all the other nodes (this event is denoted by *tr1*). In this scenario, BDCF works similar to DCF and collisions can occur when two or more RTS packets are transmitted simultaneously. The probability of a transmission by the AP in any arbitrary slot is given by:

$$P_{tr1(AP)} = \tau \qquad (2)$$

This transmission is successful only if none of the other STAs transmit in the same slot. Thus, the probability of success given that the transmission has occurred is:

$$P_{s1(AP)} = \frac{\tau(1-\tau)^n}{P_{tr1(AP)}} = (1-\tau)^n \qquad (3)$$

Once the AP wins the channel, the time spent for a successful transmission of the DATA packet by the AP is calculated as:

$$T_{s1(AP)} = O_{rts} + \{H + E_p + \delta + SIFS + ACK + \delta + DIFS\}, \quad (4)$$

Where,

O_{rts} = {RTS +δ+ SIFS + CTS+ δ + SIFS} is the time required for RTS/CTS exchange.

In case of a collision, when RTS/CTS scheme is used, the collision time is given by

$$T_{c1(AP)} = RTS + \delta + SIFS + CTS + \delta + DIFS \qquad (5)$$

II. When the AP Piggybacks

Whenever the AP receives a DATA packet from a STA, it will have the option of sending a DATA packet with the piggybacked ACK after SIFS time period. We denote this transmission as *tr2*. With the assumption of saturation traffic conditions, the AP always has a DATA packet to send to any one of the STAs. As the channel is reserved for the ACK transmission from the AP, this transmission (DATA with piggybacked ACK) is guaranteed (assuming no hidden terminals) and there is no need for contention. Thus the collision probability in this case is 0.

The transmission probability for the AP ($P_{tr2(AP)}$) in this case is the same as the probability of successful transmission by any STA in the prior data transfer stage, which is given by

$$P_{tr2(AP)} = (1-\tau)\times n\tau(1-\tau)^{n-1} = n\tau(1-\tau)^n. \qquad (6)$$

$$P_{s2(AP)} = 1. \qquad (7)$$

The time for which the channel is busy (we have to consider the time spent by the STAs as well) is given by:

$$T_{s2\,(AP)} = O_{rts} + \{H + E_p + \delta + SIFS + H + E_p + \delta + SIFS + ACK + \delta + DIFS\} \qquad (8)$$

We now proceed with the calculations for the STAs. STAs can only transmit by contending and winning the channel; and the probability that at least one STA transmits is given as

$$P_{tr\,(STA)} = [1-(1-\tau)^n] \qquad (9)$$

The probability of successful transmission given that there is transmission from a STA is the probability that the AP did not transmit and only one STA transmitted,

$$P_{s\,(STA)} = \frac{(1-\tau)\times n\tau(1-\tau)^{n-1}}{P_{tr(STA)}} = \frac{n\tau(1-\tau)^n}{1-(1-\tau)^n} \qquad (10)$$

Each time a STA undergoes a successful transmission, the AP piggybacks another packet immediately and hence the total time for successful transmission accounts for both the packet durations and is the same as $T_{s2(AP)}$

$$T_{s(STA)} = T_{s2(AP)}. \qquad (11)$$

Similarly, the time spent in a collision in this case is same as the case 1 when the AP contends for the channel.

$$T_{c(STA)} = T_{c1(AP)}. \qquad (12)$$

The time duration in which the system is active is given by:

$$T_{total} = (1-\tau)^{n+1}\sigma + \tau(1-\tau)^n \times T_{s1(AP)} + n\tau(1-\tau)^n \times T_{s(STA)} + [1-(1-\tau)^{n+1} - (n+1)\tau(1-\tau)^n] \times T_{c1(AP)} \qquad (13)$$

For the AP, the time spent in useful transmission considering both the cases in which transmissions occur is given by:

$$T_{AP}(E_p) = P_{tr1(AP)} \times P_{s1(AP)} \times E_p + P_{tr2(AP)} P_{s2(AP)} \times E_p$$
$$= [\tau(1-\tau)^n + n\tau(1-\tau)^n] \times E_p$$
$$= [\tau(1-\tau)^n(n+1)] \times E_p. \qquad (14)$$

Similarly, evaluating the useful transmission time for the remaining STAs we have,

$$T_{STAs}(E_p) = P_{tr(STA)} \times P_{s(STA)} \times E_p$$
$$= n\tau(1-\tau)^{n-1} \times E_p. \qquad (15)$$

We finally calculate the AP's throughput using (1), (13) and (14) as:

$$S_{AP} = \frac{[\tau(1-\tau)^n(n+1)] \times E_p}{T_{Total}}. \qquad (16)$$

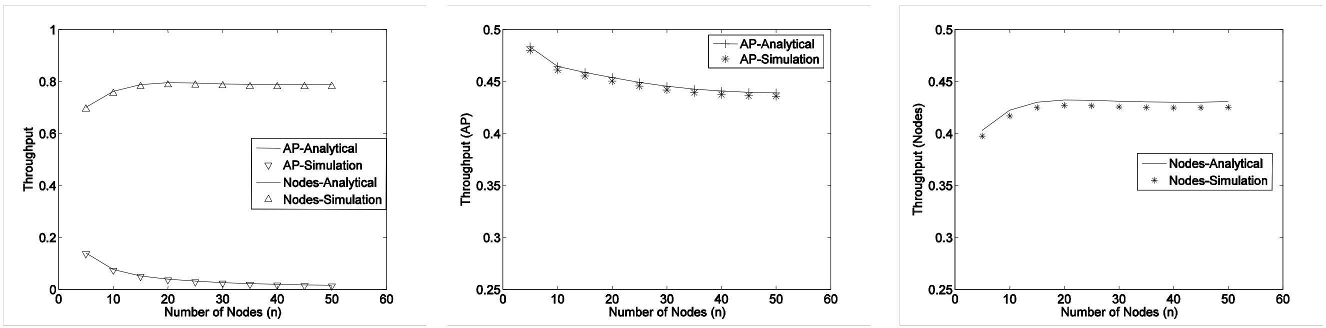

(a) Throughput comparison of AP and STAs with IEEE 802.11 DCF

(b) Throughput comparison of AP with BDCF

(c) Throughput comparison of STAs with BDCF

Figure 5. Analytical and Simulation throughput comparisons of BDCF

From (1), (13) and (15) the throughput of the STAs is given by:

$$S_{node} = \frac{n\tau(1-\tau)^{n-1} \times E_p}{T_{Total}} \qquad (17)$$

We next evaluate the throughput for the STA and AP for 802.11 DCF based on the model presented in [21]. The time, T_s, for which channel will be busy due to a successful transmission is same as $T_{s1(AP)}$. Also the collision time remains constant in both BDCF and DCF. Thus,

$$T_s = T_{s1(AP)} \text{ and } T_c = T_{c1(AP)} \qquad (18)$$

Let P_{tr} denote the probability of at least one transmission in a considered slot time and P_s denote a successful transmission. As the AP and the *n* STAs contend for the channel with equal privileges, the share of the AP in the system throughput is $S/(n+1)$ while that of the remaining *n* STAs is $S \times [n/(n+1)]$ where *S* is the system throughput as derived by Bianchi [14]:

$$S = \frac{P_{tr} \times P_s \times E_p}{(1-P_{tr})\sigma + P_{tr}P_sT_s + (1-P_s)P_{tr}T_c} \qquad (19)$$

Figure 5(a-c) shows the numerical results of the throughput as a function of number of stations and validate our analysis through simulations. We consider 1 Mbps channel and other parameters are same as described in next section. Figure 5(a) shows the obtained throughput by AP and STAs with DCF. Clearly, we can observe a substantial unfairness between the throughputs of AP (downlink flows) and STAs (uplink flows). This is because of the CSMA/CA mechanism, which provides equal access privileges to the AP and STAs. Thus the AP only gets 1/(n+1) of the total available bandwidth (with n STAs and an AP), while the STAs obtain a higher share (n/(n+1)). On the other hand, with the preferential treatment for AP in BDCF the AP achieves fair share of the bandwidth (almost equal to that obtained by the STAs).

5. SIMULATION RESULTS

In this section, we compare the performance of BDCF with DCF and DCF+ [5]. We have implemented BDCF and DCF+ in the ns-2 simulator (version 2.26) [6]. For our simulations, we have used the scenario described in Figure 1, where an AP is serving 10 stationary STAs. We placed the STAs such that all were in the transmission range of each other, thus avoiding any hidden terminal problem. The bandwidth of the wireless channel is set to 2 Mbps and the AP is connected to a server through a wired link with bandwidth of 100 Mbps and 2ms of propagation delay.

We have considered both symmetric and asymmetric traffic patterns: For symmetric traffic we considered 5 uplink and 5 downlink flows, denoted by 5Up-5Dn and for asymmetric traffic we again consider: 3 uplink flows and 7 downlink flows, denoted by 3Up-7Dn. We assume that no two flows start/end at the same node. We have done a comparative analysis of the aggregate throughput, and per stream fairness of IEEE 802.11 DCF, BDCF and DCF+ for the aforementioned traffic patterns. All results presented here are averaged over 10 simulation runs with different seed values.

5.1 Analysis of UDP traffic

5.1.1 Aggregate Throughput

We first compare the performance of the aggregate uplink/downlink throughputs achieved by the three schemes (BDCF, DCF and DCF+). We vary the DATA rate of the CBR traffic running over UDP from 100Kbps to 800Kbps and measure the aggregate throughput achieved by the flows in either directions (uplink and downlink). Considering the two different traffic patterns mentioned above, Figures 6(a-b) and 6(c-d) show the aggregate uplink and downlink throughputs respectively.

We observe a huge difference between the uplink and downlink throughputs when DCF/DCF+ is employed. As shown in Figure 6(a-b), at low loads all the three protocols have similar performance. This is because the channel is relatively free and all nodes can transmit their packets without much contention. With DCF/DCF+, as we increase the traffic load, the aggregate uplink throughput rapidly increases and downlink flows starve. Even for marginal increase in the traffic load, with IEEE 802.11 DCF or DCF+, the uplink flows obtain a very high throughput and completely dominate the access to the channel.

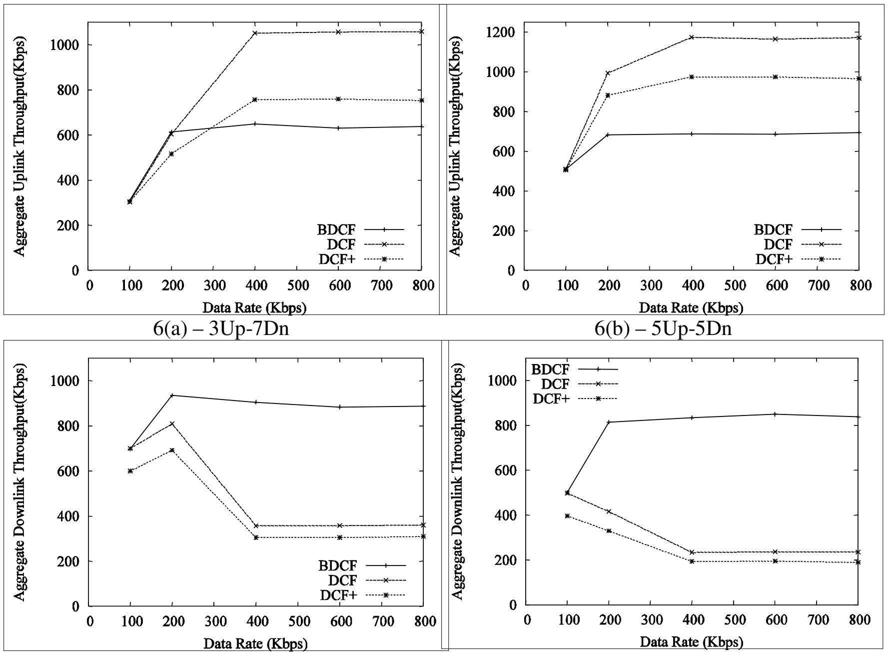

6(a) – 3Up-7Dn

6(b) – 5Up-5Dn

6(c) – 3Up-7Dn

6(d) – 5Up-5Dn

Figure 6. Aggregate Throughput for UDP Traffic

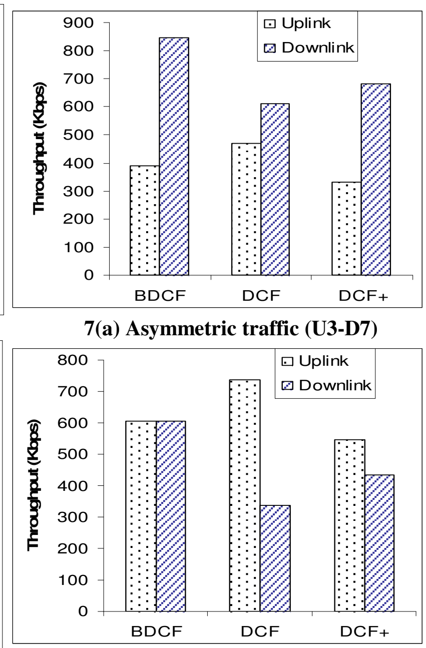

7(a) Asymmetric traffic (U3-D7)

7(b) Symmetric traffic (U5-D5)

Figure 7. TCP Aggregate Throughput

Downlink flows experience severe congestion resulting in drastic reduction of the aggregate downlink throughput. One important thing to note from Figure 6(d) is that, till 200Kbps load, the throughput of downlink traffic increases as there is sufficient bandwidth to accommodate all uplink and downlink flows.

On the other hand, with BDCF the available bandwidth is fairly distributed among uplink and downlink traffic. The downlink flows achieve a fair share of the total bandwidth, without affecting the performance of uplink flows. This can be attributed to the preferential treatment for the AP. AP obtains substantial channel share to accommodate all the downlink flows and thus limiting the uplink traffic from the nodes.

In DCF, the throughput of individual uplink flows is nearly 4 times the throughput of individual downlink flows. However when BDCF is employed, this sharp difference does not occur. The throughput of individual uplink flows is almost equal to any of the individual downlink flows. The aggregate downlink throughput is improved by nearly 300% when compared with DCF/DCF+. Hence, by providing preferential access to the AP, BDCF ensures fair sharing of bandwidth among uplink and downlink flows.

We also analyzed the performance of TCP traffic both the traffic patterns. Once again with DCF/DCF+ we observe similar unfairness among the uplink/downlink flows. Figure 7 shows the throughput performance of TCP traffic. From figures 7(a-b), we can observe the disgraceful performance of the downlink TCP traffic flows when legacy 802.11 DCF is employed.

In contrast, when BDCF is employed, we observe fair sharing of the available bandwidth among uplink and downlink flows for all traffic patterns. For instance, with BDCF in the asymmetric traffic (3Up-7Dn) scenario, the downlink flows achieve 846.18 Kbps (120 Kbps per flow); while the uplink flows obtain around 389 Kbps (129 Kbps per flow). However, with DCF and asymmetric traffic (3Up-7Dn), the downlink flows get only a meager 610 Kbps (around 87 Kbps per flow), while the aggregate uplink throughput (Figure 7(b)) is as high as 470 Kbps (around 157 Kbps per flow).

The poor performance of TCP with DCF can be explained as follows. TCP generates bi-directional traffic. TCP ACKs of the uplink flows and TCP DATA packets of the downlink flows compete for the limited channel share of the AP.

This increased downlink traffic load at the AP leads to overflowing of the link layer queues at the AP and resulting in excessive packet drops. TCP congestion control algorithm further worsens the situation by reducing the congestion window of a flow when it detects a packet drop. If a TCP DATA packet of a downlink flow is dropped at the AP, timeout occurs at the source and hence the congestion window is decreased, leading to a lower throughput. However, if a TCP ACK of an uplink flow is dropped, due to the cumulative nature of TCP ACKs, eventually another ACK with a higher sequence number will reach the uplink flow source and the TCP congestion window will not be reduced. Thus the downlink traffic experiences severe

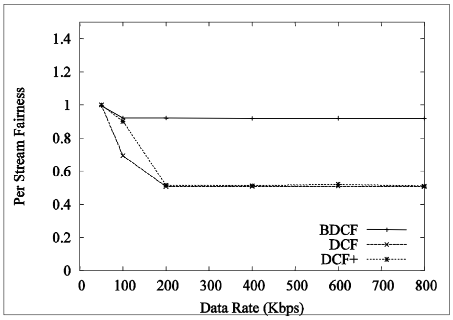

8(a) – UDP Traffic (3Up-7Dn)

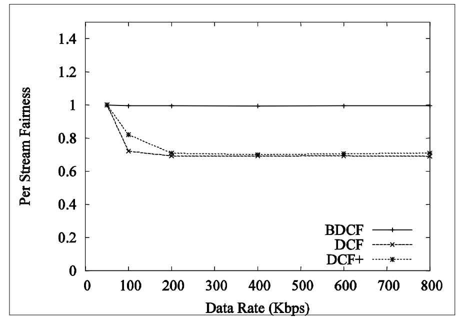

8(b) – UDP Traffic(5Up-5Dn)

Figure 8. Fairness Index

congestion control while the uplink traffic is not affected at the same rate. Consequently the uplink flows enjoy higher net throughputs when compared to downlink flows.

In summary, downlink flows are not only affected by the limited channel availability for the AP, but also due to TCP-DATA packet drops at the AP. We carefully observed the congestion window growth of all flows and notice that the congestion window of downlink flows experience frequent cut down due to packet drops at the AP (and thereby leading to timeouts). In contrast, the congestion window for all the uplink flows grows continuously. Similar observation is also reported by the authors in [4].

Once again BDCF overcomes this problem by giving a preferential treatment to the AP. Whenever a packet from the STAs is received by the AP; it can immediately transmit a packet (recall from section 3) from its buffer without any need for channel contention. In this way, at high traffic loads, the AP can avoid large queues and overflows by immediately transmitting either the TCP-ACKs/TCP-DATA packets. As more and more downstream DATA packets are transmitted the downlink flows achieve acceptable throughputs without affecting the uplink flows.

5.1.2 Fairness Index

In this section we compare the fairness among all the flows. We have used the Jain's fairness index (*f)* [8] to measure the fairness among the flows. It is given by:

$$f = \left(\frac{\left(\sum_{i=1}^{n} x_i \right)^2}{\left(n * \sum_{i=1}^{n} x_i^2 \right)} \right) \qquad 1 \le i \le n,$$

Where, there are n flows in the network and x_i is the throughput achieved by flow i. The fairness index is always positive and when it approaches one, it implies that all the flows are getting equal share of the available bandwidth. When the fairness index drops or has negative slope, then it indicates that the available bandwidth is not fairly shared among the flows. The fairness indexes for asymmetric (3Up-7Dn) and symmetric traffic (5Up-5Dn) configurations are shown in Figure 8(a) and 8(b), respectively. The fairness index is near constant and is close to one with BDCF in both traffic configurations. This is because all the flows obtain a fair share of the available bandwidth. However, this is not the case when legacy DCF and DCF+ are employed, as the downstream flows achieve very low throughputs when compared to the throughputs of uplink flows.

6. RELATED WORK

Fairness provisioning in wireless networks has been an attractive area of research and has been explored at various layers. Most of the research addresses the problem of unfairness observed in the upper layers. The unfairness problem between uplink and downlink TCP flows in an AP based WLAN was initially reported by Ramjee et al. [2].They studied the interaction between TCP and IEEE 802.11 MAC protocol and identified the buffer size at the AP as the cause for unfairness, then they proposed receiver window size manipulation in the TCP ACK to govern the access of wireless link. Bottigliengo et al. [3] studied the unfairness due to channel unavailability. They analyzed the channel conditions between the STA and the AP, choosing the STA with best channel condition for transmission and compensating the other STAs later with a burst transmission.

The concept of reverse channel reservation in WLANs has been first introduced in the DCF+ scheme [5]. The goal of DCF+ is to improve the performance of TCP in WLANs (using DCF) by an implicit reservation of the channel. However, as our simulation results have indicated, DCF+ also has the unfairness problem because the implicit reservation is used by any STA. Recently, attempts to improve performance of TCP in multi-hop ad hoc networks by similar techniques were proposed in [10] and [11]. In [10], the authors try to reduce intra and inter flow contention. Intra flow contention is encountered between TCP-DATA and TCP-ACK packets of a single TCP flow while Inter-flow contention is experienced between TCP-DATA packets of different TCP flows. They propose two schemes Quick exchange, in which the sender reserves the reverse channel for the TCP-ACK packets and Fast-forward, in which the sender allows the receiver to forward the currently received packet towards the destination. Kuang et al. [11] proposed a multi-channel MAC and reverse channel reservation for mitigating the DATA-DATA and DATA-ACK contentions. However they have not addressed the fairness problem for the infrastructure based networks. Moreover, it should also be noted that the above schemes are designed for MANETs. Furthermore, this ability to sneak a packet is enabled for any receiving STA that has a packet to transmit. Hence, when applied to AP based WLANs, these schemes can not distinguish between an AP and other STAs and will not be able provide a preferential treatment to AP packets only.

Kim et al [7] propose a utilization based uplink/downlink fairness mechanism. The AP counts the number of STAs based on the unique MAC addresses and calculates the utilization of uplink and downlink traffic. If the downlink traffic is less in a specific time window then the AP starts transmitting the DATA frames after deferring for PIFS time duration following the ACK from a STA. Thus the regular channel access policy of deferring for DIFS duration and exchanging the RTS-CTS is avoided. However in their scheme some downlink DATA frames may still suffer higher delays as it takes some time to detect the unfairness. In contrast with our scheme the AP sneaks DATA frames whenever an uplink transmission is sensed thus achieving higher downlink throughputs.

7. CONCLUSION

We experimentally demonstrate the existing unfairness problem in typical WLAN hotpots and proposed a simple enhancement to DCF for overcoming the unfairness problem. We developed analytical models to evaluate the throughputs of AP and STAs and verify these models through extensive simulation study. Our proposed BDCF protocol enables the AP to access the channel more frequently by granting a preferential treatment. In addition to this, our protocol also reduces the time wasted in channel contention and backoff mechanism at the MAC layer. We notice that the analytical results find in good agreement with the simulation results. Comprehensive simulations are also conducted on various traffic patterns for both TCP and UDP traffic. It has been demonstrated that BDCF successfully solves the unfairness problem along with substantial improvements in the throughputs and reduction in the end-to-end delays of the downward traffic (around 300%).

REFERENCES

[1] IEEE Std. 802-11. "IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification," ISO/IEC 8802-11:1999 (E), Aug., 1999.

[2] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, and P. Sinha, "Understanding TCP Fairness over Wireless LAN," IEEE INFOCOM, 2003.

[3] M. Bottigliengo, C. Casetti, C. F. Chiasserini and M. Meo, "Short-term Fairness for TCP Flows in 802.11b WLANs," IEEE INFOCOM 2004.

[4] I. Aad and C. Casette, "Differentiating mechanisms in IEEE 802.11," IEEE INFOCOM 2001.

[5] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma, "Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement," IEEE INFOCOM 2002.

[6] UCB/LBNL/VINT Network Simulator (ns-2), Available at http://www.isi.edu/nsnam/ns/index.html.

[7] S. W. Kim, B.S.Kim, Y. Fang, "Downlink and uplink resource allocation in IEEE 802.11 wireless LANs," IEEE Trans. on Vehicular Technology, Jan. 2005.

[8] R. Jain, *The Art of Computer Systems Performance Analysis*, John Wiley and Sons, 1991.

[9] H. Balakrishnan and V. Padmanabhan, "How Network Asymmetry affects TCP performance," IEEE communication magazine, 60-67, April 2001.

[10] D. Berger et al., "TCP-Friendly Medium Access Control for Ad-Hoc Wireless Networks: Alleviating Self-Contention," MASS '04, Oct 24-27, 2004.

[11] T. Kuang and C. Williamson, "A bidirectional multi-channel MAC protocol for improving TCP performance on multihop wireless ad hoc networks," Proc. of ACM/IEEE MSWiM, October 2004.

[12] J. Malinen, Host AP Driver package for Linux, Available at http://hostap.epitest.fi/.

[13] NLANR/DAST Iperf, Available at http://dast.nlanr.net/Projects/Iperf/

[14] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Select. Areas Communication*, Mar. 2000.