무선랜 환경에서 협동 상향/하향 링크 기회적 스케줄링 기법

(Joint Uplink/Downlink Co-Opportunistic Scheduling Technique in WLANs)

유 준 [†] 김 종 권 [††]
(Joon Yoo) (Chong-kwon Kim)

요 약 무선랜의 보편화와 속도증가로 인해 시스템 성능향상을 위한 기회적 스케줄링 기법의 중요성이 증대되고 있다. Voice-over-IP, peer-to-peer와 같은 새로운 응용 프로그램의 증가와 함께 무선랜 트래픽은 점점 더 상하링크 대칭이 되어가고 있으며, 그에 따라 무선랜 하향링크에 편중된 기회적 스케줄링 기법만으로는 공유된 한 개의 무선 채널의 효율을 높일 수 없게 되었다. 그러나, 무선랜의 채널 특성은 시간에 따른 변화가 크며 또한 상하대칭적이지도 않기 때문에, 무선랜 상향링크에서의 기회적 스케줄링은 결코 쉬운 문제가 아니다. 각 전송 클라이언트는 access point에게 채널상황을 보고해야 하며, 스케줄링 결정은 모든 클라이언트들의 동의를 얻어야만 한다. 본 논문에서는 무선랜 환경에서 협동 상향/하향 링크 기회적 스케줄링 기법 JUDS(joint uplink/downlink opportunistic scheduling for WLANs)을 제안한다. JUDS는 상향과 하향링크 스케줄링을 결합함으로써 채널의 상이점을 최대화하며 오버헤드를 최소화한다. JUDS는 또한 상향과 하향링크 사이에 균등한 채널 공유를 하도록 한다. Qualnet 시뮬레이션을 통해, JUDS는 최대 127%의 성능향상을 보이며, 상향/하향 링크 사이 완벽에 가까운 균등한 채널 공유를 보인다.

키워드 : 무선랜, 기회적 스케줄링 기법

Abstract Recent advances in the speed of multi-rate wireless local area networks (WLANs) and the proliferation of WLAN devices have made rate adaptive, opportunistic scheduling critical for throughput optimization. As WLAN traffic evolves to be more symmetric due to the emerging new applications such as VoWLAN, collaborative download, and peer-to-peer file sharing, opportunistic scheduling at the downlink becomes insufficient for optimized utilization of the single shared wireless channel. However, opportunistic scheduling on the uplink of a WLAN is challenging because wireless channel condition is dynamic and asymmetric. Each transmitting client has to probe the access point to maintain the updated channel conditions at the access point. Moreover, the scheduling decisions must be coordinated at all clients for consistency. This paper presents JUDS, a joint uplink/downlink opportunistic scheduling for WLANs. Through synergistic integration of both the uplink and the downlink scheduling, JUDS maximizes channel diversity at significantly reduced scheduling overhead. It also enforces fair channel sharing between the downlink and uplink traffic. Through extensive QualNet simulations, we show that JUDS improves the overall throughput by up to 127% and achieves close-to-perfect fairness between uplink and downlink traffic.

Key words : Wireless LAN, opportunistic scheduling

1. Introduction

Rapidly increasing speed of wireless local area networks (WLANs)[1] and the dynamics of channel conditions in unlicensed frequency bands have made rate adaptation a critical component for high-performance wireless networking. Based on the channel condition, e.g., the instantaneous signal-to-noise ratio (SNR), an IEEE 802.11a/b/g [2-4] WLAN inter-

---

† 학생회원 : 서울대학교 컴퓨터공학과
joon@popeye.snu.ac.kr
†† 종신회원 : 서울대학교 컴퓨터공학과 교수
ckim@popeye.snu.ac.kr

논문접수 : 2007년 4월 6일
심사완료 : 2007년 8월 7일

정보과학회논문지: 정보통신 제34권 제6호(2007.12)

face can choose among 4 to 8 different rates, ranging from 1Mbps to 54Mbps. The coming new standard 802.11n [5] will offer even more available rates at a larger range, enabling fine grained channel rate control for optimized utilization of the wireless channel. Meanwhile, the increasing number of wireless users has reached the point where a wireless access point always has a choice when scheduling the traffic for multiple wireless clients. As a result, opportunistic scheduling, a technique that was first developed in cellular wireless networks [6], has been recently applied to the downlink of a multi-rate, multi-user WLAN [7]. Opportunistic scheduling leverages multi-user diversity by scheduling the user whose instantaneous channel condition is above the average.

However, the existing work on WLAN opportunistic scheduling [7] is limited on the downlink only. This limit seriously bounds the utilization of the shared wireless channel for the following two reasons. First, while traditional Web traffic is asymmetric with a significant bias on the downlink, modern WLAN traffic is becoming more and more symmetric [8,9] due to the emerging new network applications such as BitTorrent [10], peer-to-peer file sharing, and Voice over WLAN(VoWLAN). Since both downlink and uplink traffic shares a single wireless channel in a WLAN, poor channel unitization at the uplink directly impacts the achievable throughput at the downlink. Second, the dominating IEEE 802.11 MAC, i.e., distributed coordination function (DCF) [1], treats all senders, including both the access point and the clients, equally. Therefore, as the number of clients competing for the shared channel increases, the bandwidth allocation for the downlink (or the access point) decreases proportionally. Without proper scheduling on the uplink, the achievable throughput gain due to downlink opportunistic scheduling diminishes as the number of transmitting clients served by an access point increases. This is in stark contrast to the expectation that opportunistic scheduling achieves higher throughput gain at higher multi-user diversity.

One major challenge of applying opportunistic scheduling at WLANs is the maintenance of the highly dynamic and asymmetric [11] condition of wireless channels defined in the unlicensed frequency bands. Different from the cellular wireless networks where efficient mechanisms for fine grained channel condition feedback are built into the high end hardware, the overhead of channel condition maintenance in a WLAN based on low-end wireless transceivers may be prohibitively high. For example, the evaluation of the opportunistic WLAN downlink scheduling [7] shows that the achievable throughput decreases with more than 3 clients involved, due to the excessive channel probing and feedback. This problem is further aggravated when a naive opportunistic uplink scheduling is applied, since every wireless client has to probe the access point and the access point has to feedback the best rate individually. Furthermore, even with such updated channel conditions established, it is very difficult for all wireless clients as well as the access point to reach the consensus regarding the transmission schedule of every packet originated at different transmitters.

In this paper we present JUDS, the Joint Uplink and Downlink Opportunistic Scheduling for WLANs. JUDS synergistically integrates the opportunistic scheduling at both the uplink and the downlink for maximum channel utilization and significantly reduced overhead in channel condition maintenance and scheduling coordination. In more specific, JUDS doubles the channel diversity while reducing the overall scheduling overhead by half, compared with the scenario where downlink or uplink opportunistic scheduling is conducted alone. Furthermore, JUDS enforces fair bandwidth allocation between the uplink and downlink, regardless of the number of contending wireless clients. This bandwidth allocation policy ensures that the performance gain of opportunistic scheduling increases as the number of clients competing for the shared channel increases, consistent with the expectation.

We make three key contributions in this paper. First, we reveal the fundamental limit of the downlink opportunistic scheduling in a WLAN, and present the first joint uplink/downlink opportunistic scheduling in WLANs. Our design leverages the maximum diversity of the shared wireless channel in a multi-user WLAN. Second, we present the details of the jointly uplink/downlink scheduling protocols. JUDS protocols exploit the unique characteristics of the broadcast wireless channel and the repetitive operating cycles of the scheduling algorithm. As a result, JUDS implementation eliminates a number of unnecessary signaling in popular CSMA/CA wireless MAC, e.g., the per-frame acknowledgements. Although these signaling messages are generally small, the constant per-frame PHY and MAC overhead plus the mandatory interframe spacing leads to an overhead up to 15~70% [1-4]. Finally, the performance of JUDS has been evaluated through simulations in QualNet[12]. Our evaluation shows that the unitization of the shared wireless channel is almost doubled and the fairness between uplink and downlink traffic is close-to-perfect.

The rest of the paper is organized as follows. In Section 2 we will review and compare with the related work. Section 3 gives an overview of JUDS design. Section 4 provides the detailed description of the probing stage of JUDS and Section 5 shows the scheduling and data transmission stage of JUDS. Section 6 shows the performance via simulations. Finally, Section 7 concludes our paper.

2. Related Work

Many MAC protocols have been proposed [13-15] to exploit the multi-rate capability at the physical layer. The Auto Rate Fallback (ARF) [13] is the one typically implemented in commercial 802.11 products. ARF chooses to raise or lower its transmission rate according to consecutive transmission successes or failures, respectively. In the Receiver Based Auto Rate (RBAR) [14], the receiver selects an adequate transmission rate according to the channel quality measured from the received request-to-send (RTS) frame. It then piggybacks the selected data rate in the CTS frame. The Opportunistic Auto Rate (OAR) [15] improves the channel unitization by allowing a client to hold the wireless channel for an extended period when the achievable data rate is high. Those designs enable intelligent rate adaptation between a specific pair of sender and receiver. They do not explicitly leverage the channel diversity due to the increasing number of clients.

Opportunistic scheduling optimizes the utilization of the wireless channel shared among multiple users. The higher the user diversity, represented by a larger number of users and a larger channel variation, the higher the potential performance gain. Opportunistic scheduling was designed and applied in cellular wireless data networks, e.g., HDR [6]. In HDR, the channel conditions are measured by the pilot signal sent by the base station to each individual user. The users then feed back the channel condition simultaneously via the CDMA uplink. In a WLAN based on low-end wireless transceivers, the lack of efficient support for closed-loop channel condition feedback represents the main challenge for opportunistic scheduling.

The Medium Access Diversity (MAD) protocol [7] applies opportunistic scheduling on an IEEE 802.11 WLAN downlink and is backward compatible with the legacy 802.11 DCF. However, since the wireless channel is shared between the uplink and downlink traffic, opportunistic scheduling on the downlink only is not enough.

The opportunistic scheduling proposed in the literature [6,7] and in this paper are all based on signal-to-noise ratio (SNR) measurements to determine the appropriate data rate. Recent measurements [16] using off-the-shelf 802.11 devices showed that the SNR, measured as an average over many packets, may not be a good predictive tool for the successful delivery of a packet. Note that the measurement results do not contradict or invalidate our approach in that our rate adaptation is based on instantaneous SNR feedback. In fact, one contribution of this work is the design of efficient mechanisms for timely channel condition feedback, based on which the opportunistic scheduling runs.

3. Overview

In this section we overview JUDS operation in a WLAN. We first elaborate the challenges and issues, and then present the basic ideas and mechanisms that JUDS employs to address those challenges.

The major challenge of WLAN opportunistic scheduling is the overhead for the senders to collect from the receivers the updated channel condition at the receivers' sides. For downlink scheduling all candidate clients have to feed back their measured downlink channel quality to the access point. The communication overhead is therefore O(k), given k downlink candidate clients chosen by the access point. For uplink scheduling all candidate clients have to probe the access point individually for the access point to assess the uplink channel quality. The complexity is again O(k) given k uplink candidates.

Note that the O(k) probing or feedback overhead is severe in an asynchronous, packet-switched WLAN, where a constant physical and MAC layer overhead is incurred for every frame regardless of the frame size or channel rate. For example, in the context of IEEE 802.11b with 2Mbps basic rate and 11Mbps data rate, the physical and MAC layer overhead accounts for 31% for a frame of 1500 bytes. For higher speed IEEE 802.11a/g with 6Mbps basic rate and 54Mbps data rate, the overhead increases to 68%. To limit the overhead it is therefore more effective to decrease the number of frames, as opposed to reducing the frame size. Although a small number of candidates k lowers the feedback or probing overhead, it also limits the multi-user diversity or the throughput gain of opportunistic scheduling.

JUDS addresses this challenge by the following three mechanisms. First, since JUDS schedules both downlink and uplink traffic simultaneously, it combines the clients' downlink channel condition feedbacks with client's uplink channel condition probing. Therefore, the feedback and probing overhead is cut by half, without compromising multi-user diversity. Second, JUDS exploits the cyclic scheduling between the uplink and downlink traffic, and piggybacks control signals whenever possible. In specific, because the access point and one of the clients alternate in transmitting on the shared wireless channel, the mandatory per-frame acknowledgement can be piggybacked on the next data frame transmission from the other direction. This mechanism removes the 70.1% physical layer overhead and the inter-frame spacing if the acknowledgement were transmitted in a separate frame. Finally, JUDS exploits the broadcast nature of wireless transmission and enables opportunistic probing and feedback. When a client fails to receive a RTS, other clients that detect the idle slot will automatically advance the feedback schedule, therefore further saving the feedback overhead.

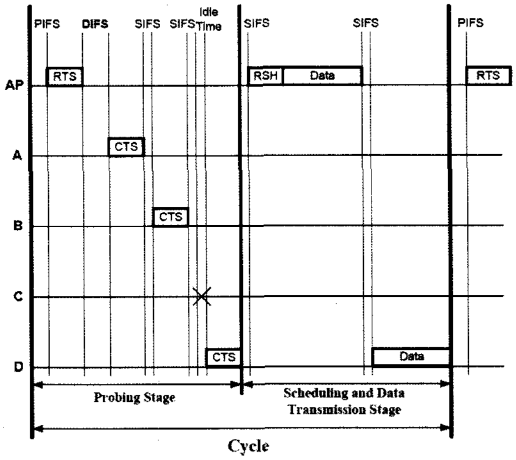

Fig. 1 JUDS operation

Putting it all together, the operation of JUDS can be illustrated by Fig. 1. JUDS divides the time into repetitive cycles, and schedule exactly one downlink transmission and one uplink transmission, if the demand exists, in each cycle. In specific, a cycle consists of two stages: the channel probing stage and the scheduling and data transmission stage. In the probing stage, the access point and the user stations coordinate to probe for channel quality measurement at both downlink and uplink. In the scheduling and data transmission stage, based on the channel quality measurements and the scheduling priority, the access point selects the stations that should participate in uplink and downlink transmission. Once the schedule is determined, the access point sends down a downlink data frame to the selected downlink station and the selected uplink station then transmits afterwards. We present the details of those two stages in the next two sections respectively.

4. Probing Stage

We explain the probing stage in details in this section. The access point and selected candidate clients exchange channel uplink downlink channel conditions in this stage. We first describe the two basic steps: candidate selection and downlink probing in Section 4.1 and the downlink condition feedback and uplink probing in Section 4.2. We finally describe in Section 4.3 how the access point is made aware of the clients that are competing for the uplink.

4.1 A. Candidate selection and downlink probing

The first step that the access point takes in the probing stage, prior to sending the RTS, is selecting the clients that are either to be probed for the downlink or to probe for the uplink. The access point chooses those candidates randomly for long-term fairness. Fig. 1 shows an example of the candidate selection. The access point has full knowledge of those clients who are receiving packets from the downlink, i.e., client A, B, C and E in this example. Through mechanisms presented in Section 4.3 the access point is also made aware of client A, D and F who are competing on the uplink (Clients E and F are not shown in Fig. 1.). The access point then selects up to k candidates from clients A-F. In this example, we assume that k = 4. Suppose the access point selects client A, B, C, and D. A, B, and C will be competing for the downlink, while A, D will be competing for the uplink.

The selected candidate nodes are listed in the destination address (DA) list field of the RTS frame, as shown in Fig. 2. The RTS message is then broadcasted to all clients. A straightforward method of listing the clients in the DA list field in the RTS packet would be to just list the clients' unique 48 bit MAC addresses in the sequence order for CTS transmission. Since this could result in a large overhead we utilize the Association ID (AID)-a 16-bit ID of a client assigned by the access point during the association procedure [1].

AID is effective to identify the list of destination addresses (DAs) in the RTS frame. We only use the 8 least significant bits, or the least significant byte of the AID for the addressing. Therefore to make the one byte AID as diverse as possible, we assume that the access point assigns the smallest available AID to a new client during the association procedure. As shown in Fig. 2, each destination address in the list corresponds to the least significant byte of the AID. Therefore the default value of the Address length subfield in the Destination List Information field will be set to 8, enough for accommodating $2^8 = 256$ clients. Since a single access point is unlikely associated with more than 256 clients in realistic scenarios, this method will be generally feasible. In case there are more than 256 clients, the address length will be changed by setting the Address length subfield appropriately.

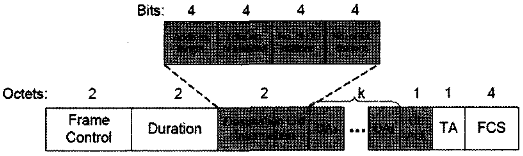

Fig. 2 Modified RTS packet format for downlink probing and uplink candidate selection

The total number of clients, as well as the number of clients competing for uplink and the number of clients competing for downlink will be recorded in the Destination List Information field. For example, if there are a total number of 10 clients and the number of uplink and downlink clients is 3 and 8 respectively, then there is 1 client that has both uplink and downlink frames. The first 2 clients in the list are therefore clients for uplink, the next 1 will the client with both uplink and downlink, and the other 7 will be the clients competing only for the downlink.

4.2 Downlink condition feedback and uplink probing

On receiving the RTS frame, the clients first check if it is included in the DA list field. If yes, the client prepares to send a CTS back to the access point. The format of the CTS packet is similar to that of 802.11 [1], but contains the extra schedule information field (See Fig. 3). The CTS frame works either as a downlink feedback or an uplink probe packet. As for a downlink feedback, it reports the feasible data rate in the 4 bit Rate field of the CTS frame, based on the measured downlink channel quality when receiving the RTS frame. As for an uplink probe, it sets the Rate field 0.

The clients will send CTS in the order indicated in the DA list field in the RTS frame. As the CTS uses the basic data rate and the clients know the

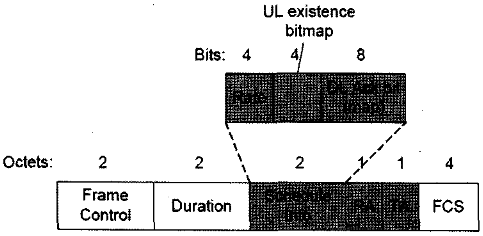

Fig. 3 Modified CTS format for downlink feedback and uplink probing

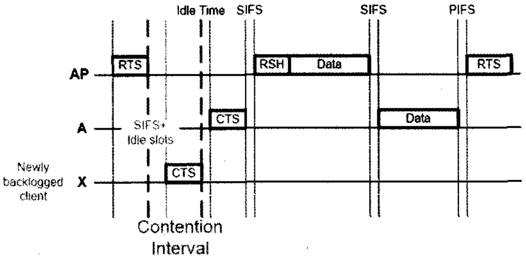

Fig. 4 A newly backlogged client will contend in the Contention Interval

order of the CTS candidates, each client can easily calculate the transmission schedule and their own slot. The first client in the list waits for inter frame space (DIFS), or a period of [busy time + 1 idle slot], and sends its CTS back to the access point (See Fig. 1). There are two reasons why it specifically waits for a DIFS period or a [busy time + 1 idle slot]. One is to give chances to clients that need to contend for the uplink (See Section 4.3). The other reason is to prevent neighboring clients that use IEEE 802.11 DCF from interfering. If the waiting period is larger that DIFS, then neighboring clients that use IEEE 802.11 DCF may dominate the channel.

In case where a candidate client fails to receive the RTS packet, it results as an idle time slot instead of an entire idle CTS duration. When a client that is awaiting to transmit a CTS senses the channel as idle for a time slot, it will realize that a turn has passed and advance its transmission starting time accordingly. For example in Fig. 1, client C fails to receive the RTS. As for client D, by sensing the channel idle for a time slot after B's CTS and the SIFS duration, it will assume that client C's turn is over and transmit right after the idle time slot. This will reduce the idle time by 79% when using IEEE 802.11a [2]. Note that the clients do not need to decode the CTS frame as long as they can carrier sense the transmission as busy.

4.3 Uplink contention

In order for the access point to select uplink probe candidates, a client uses the 4-bit UL existence bitmap in the CTS (See Fig. 3) frame to inform the access point the uplink packet queue status. The first bit represents the existence of a packet in the uplink packet queue, and the next 3 bits represents the length of the queue. For example, if the first bit is set to "1" (true), and the next 3 bits are set to "3", then there are a total number of 4 packets uplink. The access point uses this information to determine the existence of the uplink packets for the corresponding client.

The access point will find out the client's uplink packet existence if the client is probed for the downlink or selected as the uplink candidate. If the client does not have a packet in its uplink buffer and there are no downlink packets destined for the client, then the client has to contend to inform the access point. The contention process is also necessary for a client first arriving. We call this station as a newly backlogged client and this duration for contention as the Contention Interval as shown in Fig. 4.

The newly backlogged clients contending for the uplink will either send a packet after randomly selecting a backoff of contention window size CW, which is similar to the backoff procedure in IEEE 802.11 DCF [1]. In each Contention Interval, only one contenting client is allowed to transmit. After a contending client finishes transmitting, the first client listed in the DA list will send the CTS after an idle time slot. Since the Contention Interval consists of only [SIFS + 2 idle slot], the newly backlogged clients may need to wait multiple cycles to access the channel. In [17], we show through simple analysis that the expected waiting time of a newly backlogged client is around 35.09 ms even when up to 30 clients are simultaneously contending.

5. Scheduling and Data Transmission Stage

In this section we describe the second stage, scheduling and data transmission, of JUDS. Note that after probing stage, the access point collects not only the downlink channel conditions of all candidate clients, but also their uplink channel conditions. The access point then performs opportunistic scheduling for both the downlink and uplink (Section 5.1). The access point first transmits a data frame to the client scheduled on the downlink. The identifier of the client that is scheduled for uplink transmission is piggybacked in the downlink data transmission, together with the feedback of the uplink channel condition and supported data rate. The selected client for uplink overhears the downlink data transmission, and decodes the piggybacked uplink feedback. It then transmits on the uplink afterwards (Section 5.2). We finally present the mechanism that removes the explicit per-frame acknowledgement (Section 5.3).

5.1 Opportunistic scheduling

After the access point receives all the downlink feedback and uplink probes, it schedules two clients for transmissions on the downlink and uplink respectively. In JUDS the access point enforces proportional fairness [18,19] by default. That is, the access point assigns each client a priority according to:

$$\text{Priority}_k = \frac{DR_k(t)}{T_k(t)} \tag{1}$$

The priority of client k is determined by the instantaneous data rate (DRk(t)) over the average throughput ($T_k(t)$). Similar to the HDR [6] system, the access point monitors the throughput of each client in a recent time window. Among the probed clients with backlogged downlink queues, the access point selects the one with the highest priority, defined in Eqn. (1), and schedules the client for downlink transmission. Similarly, the access point selects among the probed clients with backlogged uplink queues the one with the highest priority, and piggybacks the identifier and the supported data rate (according to the measured uplink channel condition at the access point) to the downlink data transmission. Since the candidate clients are chosen randomly, (See Section 5.1), the clients will achieve proportional fairness in the long run. Furthermore, the fairness between uplink and downlink traffic is ensured since exactly one data on the downlink and one data on the uplink are transmitted during each cycle.

5.2 Data Transmission

In order for the uplink scheduling information be piggybacked in the downlink data transmission, two receiver addresses, i.e., the downlink receiver (DL RA) and the uplink transmitter address (UL TA), are defined in the data header, as shown in the data packet format in Fig. 5. The DL RA identifies the destination client of the downlink data transmission, while the UL TA identifies the client that is scheduled for uplink transmission in this cycle. The packet format is used in both downlink/uplink data packets. The UL TA field is always set to 0 for the uplink data packet.

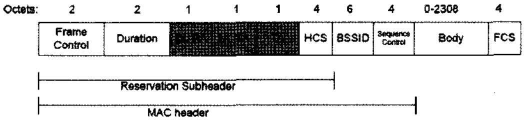

Fig. 5 Data packet format

Note that DL RA and UL TA are transmitted at the basic rate regardless of the data rate at which the data is transmitted. Since the basic rate is supported by all clients, it is guaranteed that the client scheduled for uplink transmission can decode the piggybacked uplink scheduling information. This technique is named the reservation subheader (RSH) [14]. It is used in case the uplink client cannot receive the MAC header on the transmitted downlink data rate, since the transmission range for higher data rates are typically smaller.

If no downlink traffic exists the access point can simply transmit a packet with zero payload. This packet will inform the client who is chosen for uplink transmission. In case an uplink packet does not exist or the client scheduled to transmit the uplink data fails to receive the RSH, an idle duration of [SIFS + time slot] will occur. In turn, the access point will move on to the next cycle immediately after an idle duration is detected.

A final note is that JUDS can be easily extended to schedule multiple small data packets in the same cycle, similar to the opportunistic scheduling developed in [5,15,7,20]. We leave this part as a future extension.

5.3 Implicit Acknowledgement

JUDS exploits the cyclic scheduling between downlink and uplink to eliminate explicit per-frame MAC-layer acknowledgement for throughput optimization. The idea is to piggyback the acknowledgement into the downlink/uplink probing and feedback messages during the probing stage of the next cycle. In specific, we piggyback the acknowledgement at ACK bit(map) in the RTS and CTS (See Fig. 2 and 3).

For the uplink data packet, the access point will piggyback the acknowledgment information in the ACK bit(map) of the RTS in the next cycle. So there will be no difference between the piggybacked acknowledgement and the ACK defined in IEEE 802.11 DCF: every uplink data packet is followed by an immediate acknowledgement piggybacked in the RTS at the beginning of the next cycle.

For timely acknowledgement of the downlink data transmission, the access point always adds the client that receives the downlink data packet in the current cycle into the list of candidate clients for the next cycle, even if the client does not have any other downlink/uplink packet or the client is not randomly chosen as a candidate for the next cycle. This way, the access point is guaranteed to receive an implicit acknowledgement from the client, piggybacked in the CTS message in the probing stage of the next cycle. However, if the client is indeed not randomly chosen as a candidate for the next cycle, the access point will not schedule the client for transmission on either downlink or uplink to avoid skewed proportional fairness.

6. Performance Evaluation

We have implemented JUDS in the Qualnet Simulator [12]. We compare JUDS with ARF [13], RBAR [14] and MAD [7]. We use the IEEE 802.11a physical radio [2] which supports 8 variable data rates, ranging from 6 to 54Mbps. The physical specifications such as the transmission power or the receive sensitivity etc., were adopted from the commercial Cisco Aironet 802.11a/b/g Wireless CardBus Adapter [21]. The channel model we used was free space Rayleigh Fading distribution and Ricean Fading distribution [22]. Most of the simulations were performed on Rayleigh Fading except when the channel stability was measured by varying the Ricean factor (Fig. 7).

We use constant bit rate (CBR) traffic with 1kbyte packet size to control the traffic load intensity. The downlink aggregate traffic load of 54Mbps is given so that packets are always backlogged. For uplink we vary the traffic load so that we can look into both cases where traffic is symmetric (Section 6.1, 6.2 and 6.4) and asymmetric (Section 6.3 and 6.5). The topology used was an infrastructure based WLAN, where all the clients communicate only with the access point. The access point maintains a separate queue and scheduling information for each client to perform opportunistic scheduling. There are 1 access point and 10 stationary clients in the simulations.

We performed 5 set of simulations. First we look into the effect of number of candidates in the probing stage with variable parameters. We then study the performance as the number of active users increase. Then we show the enhanced performance of JUDS compared to the other protocols for variable traffic load and distance. Then we vary the channel coherence time to study the performance. The final set of simulations show that JUDS achieves both close-to-perfect uplink/downlink fairness and also proportional fair among each clients.

6.1 Number of candidate clients

In this set of simulations we study the effect of number of candidate clients (k) selected in the probing stage, when the distance between sender/receiver (d) and Ricean factor (K) [22] are varied. Symmetric overloaded aggregate traffic of 54Mbps is given at both uplink and downlink.

The distance between the access point and all the clients were set to d, where a small d implies good average channel quality. Fig. 6 shows the throughput as a function of k as d is varied. We can see that when d is small (50m), the throughput decreases as the number of candidates increases more than 2. On the other hand, when d is relatively large, the throughput steadily increases with the number of candidates. The reason is that the control packets are transmitted at the basic rate

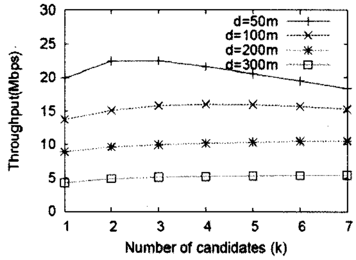

Fig. 6 Network throughput vs. number of candidates when distance(d) is varied

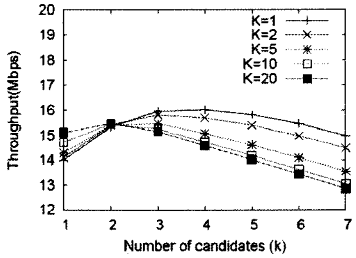

Fig. 7 Network throughput vs. number of candidates when Ricean factor (K) is varied

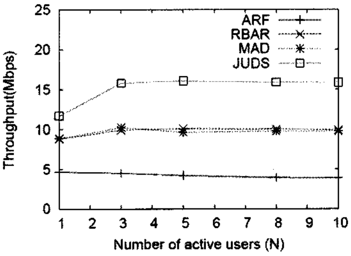

Fig. 8 Average throughput vs. number of active users

while the data packets typically use a higher rate, in turn, the scheduling overhead fraction will increase as the data rate for the data packets gets higher. So, when the average channel is very good, it is better to probe only a few clients.

The Ricean factor, K represents the channel variance; As K increases, the channel variance decreases. Fig. 7 shows that as K increases it is better not to probe more clients, since the channel becomes more stable in terms of variance. This is due to the fact that the multiuser diversity increases with high channel dynamics. [19].

6.2 Throughput vs. number of active users

Fig. 8 plots the network throughput when the number of active users is varied. Each active user has both overloaded symmetric uplink and downlink traffic. When there is only 1 active user, opportunistic scheduling does not take effect. Therefore, MAD and RBAR perform about the same. As the number of active users *N* increases the bandwidth allocation at the downlink decreases, so the effect of downlink opportunistic scheduling for MAD diminishes. In comparison, JUDS benefits from the increased user diversity and achieves higher utilization of the shared wireless channel.

6.3 Throughput vs. traffic load intensity

Fig. 9 depicts the throughput performance of JUDS compared to ARF [13], RBAR [14] and MAD [7]. The x-axis represents the aggregate uplink traffic load originated from 10 clients while the aggregate downlink traffic load is fixed to 54Mbps in each case. As shown in [7], MAD performs best when *k* is set to 3. So for fair comparison, we set *k* to 3 for both MAD and JUDS. The distance between access point and all the clients is set to *d*, which is varied from 50m to 200m.

RBAR and MAD clearly outperform ARF in every scenario due to the better rate adaptation to the fast fading channel. RBAR and MAD perform almost the same when the distance (*d*) is generally small, i.e., the average channel quality is relatively high.

MAD outperforms RBAR when the average channel quality gets lower, since the multiuser diversity on the downlink finally improves. But as the uplink traffic intensity gets higher, the number of clients

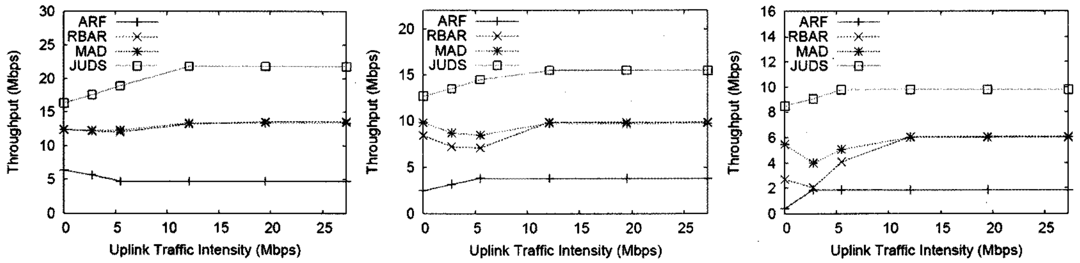

Fig. 9 The aggregate network throughput vs. the aggregate uplink traffic load. The distances between the access point and the all clients (*d*) is varied: (a) *d*=50m, (b) *d*=100m, (c) *d*=200m

competing for the shared channel increases. Therefore, the downlink is dominated by the uplink traffic so that the effect of downlink opportunistic scheduling from MAD diminishes.

JUDS improves the throughput over MAD by up to 127% and average around 62%, when the uplink traffic is high. As the uplink traffic increases, JUDS takes advantage of the higher channel diversity due to uplink scheduling for the increased number of clients competing for the channel. Even when there is only downlink traffic, JUDS shows a throughput gain around 30~56% over MAD due to its effectiveness in reducing overhead such as implicit acknowledgement.

6.4 Throughput vs. channel coherence time

In this section we vary the channel coherence time and study the impact on the throughput. The channel coherence time is represented in terms of the velocity (m/s) of the user. Traffic was overloaded at both the uplink and downlink (54Mbps). Fig. 10 shows that the overall throughput of ARF, RBAR, MAD and JUDS all decrease as the coherence time decreases due to data transmission failures. JUDS clearly outperforms all three due to better channel utilization.

6.5 Fairness

In this section, we examine the fairness issue. Fig. 11 shows the fraction of downlink throughput achieved when the aggregate uplink traffic load is increased. The aggregate downlink traffic load is fixed to 54Mbps in each case. The x-axis represents the aggregate uplink traffic load from 10 clients. *d* was set to 100m. As the uplink traffic load increases, the downlink throughput of ARF, RBAR and MAD all suffer since the increasing uplink flows from 10 clients will dominate the downlink flows from 1 access point. In turn, the throughput gain of downlink opportunistic scheduling will diminish as shown in Fig. 9. RBAR and MAD show better fairness than ARF in low uplink traffic since they have better capacity. JUDS achieves a close-to-perfect uplink/downlink fairness through its cyclic scheduling between uplink and downlink traffic.

Finally Fig. 12 represents the throughput and time fraction shared among the clients for both uplink and downlink. Symmetric overloaded traffic of 54Mbps is given at both uplink and downlink. Nodes 2 to 10 are 20m, 40m, 60m, 80m, 100m, 120m, 140m, 160m and 180m apart from the access point. As we implement the proportional fairness scheme in Eq. 1, user 2 gains the best throughput while user 10 has the worst. On the other hand, observe that the fraction of temporal share of each user is in between 0.097~0.127. This shows that proportional fairness is achieved in JUDS.

7. CONCLUSION

The increasing rate and user diversity in a modern WLAN has made the adoption of opportunistic scheduling highly desirable. However, opportunistic scheduling on the downlink only is fundamentally insufficient, especially when the amount of uplink traffic and the number users are high. In this paper we presented JUDS, the design and evaluation of a joint uplink/downlink opportunistic scheduling algorithm for WLANs. Through synergistic integration of both the uplink and the downlink opportunistic scheduling, JUDS benefits from maximized channel diversity at significantly reduced scheduling over-

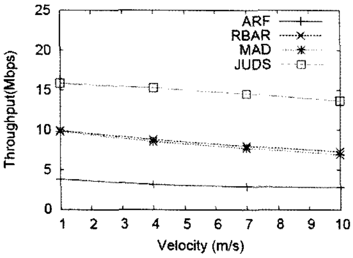

Fig. 10 Average throughput vs. channel coherence time

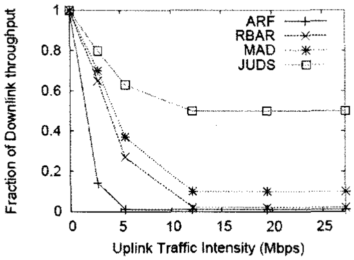

Fig. 11 Uplink/downlink fairness

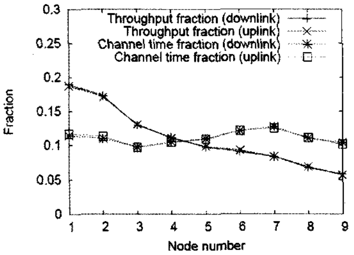

Fig. 12 The throughput share and temporal share of each client

head. Through analysis and extensive QualNet simulations, we show that JUDS improves the utilization of the shared wireless channel by up to 127% and achieves close-to-perfect fairness between uplink and downlink traffic.

References

[1] IEEE, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. IEEE Std 802.11-1999, Aug. 1999.

[2] IEEE Std. 802.11a, 802.11a-1999, 1999.

[3] IEEE Std. 802.11b, 802.11b-1999, 1999.

[4] IEEE Std. 802.11g, 802.11g-2003, 2003.

[5] IEEE P802.11 - TG N - Standard for Enhancements for Higher Throughput, "http://grouper.ieee.org/groups/802/11/ Reports/tgnupdate.htm"

[6] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana, S. Viterbi, "CDMA/HDR: A Bandwidth - Efficient High-Speed Wireless Data Service for Nomadic Users," IEEE Communications Magazine, Vol.38, pp. 70-77, July 2000.

[7] Z. Ji, Y. Yang, J. Zhou, M. Takai and R. Bagorodia, "Exploiting Medium Access Diversity in Rate Adaptive Wireless LANs," Proc. of ACM Mobicom'04.

[8] T. Henderson and D. Kotz and I. Abyzov, "The Changing Usage of a Mature Campus-wide Wireless Network," Proc. of ACM Mobicom'04.

[9] "cachelogic research: The True Picture of peer-to-peer file sharing," http://www.cachelogic.com/research/p2p2004.php

[10] D. Qiu and R. Srikant, "Modeling and Performance Analysis of Bit Torrent-Like Peer-to-Peer Networks," Proc. of ACM Sigcomm'04, 2004.

[11] D. Couto, D. Aguayo, J. Bicket, R. Morris, "A High-Throughput Path Metric for Multi-Hop Wireless Routing," Proc. of ACM Mobicom'03.

[12] QualNet User Manual v3.9. Scalable Network Technologies, Culver City, CA. http://www.qualnet.com

[13] A. Kamerman and L. Monteba, "WaveLAN II: A high-performance wireless LAN for the unlicensed band," Bell Labs Technical Journal, Summer 1997, pp.118-133.

[14] G. Holland, N. Vaidya and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," Proc. of ACM Mobicom'01, 2001.

[15] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic Media Access for Multirate Ad Hoc Networks," Proc. of ACM Mobicom' 02, 2002.

[16] D. Aguayo, J. Bicket, S. Biswas, G. Judd and R. Morris, "Linklevel Measurements from an 802.11b Mesh Network," Proc. of ACM Sigcomm'04, 2004.

[17] J. Yoo, H. Luo, and C. Kim, "Opportunistic Joint Uplink/Downlink Scheduling in WLAN," Technical Report, Report No. UIUCDCS-R-2006-2757, July 2006.

[18] F. Kelly, "Charging and rate control for elastic traffic," European Transactions on Telecommunications '97.

[19] J. Holtzman, "Asymptotic analysis of proportional fair algorithm," Proc. of IEEE PIMRC'01, Vol.2, pp. F33-F37, Sept. 01.

[20] S. Kim, Y. Kim, S. Choi, K. Jang and J. Chang, "A High-Throughput MAC Strategy for Next-Generation WLANs," Proc. of IEEE WoWMoM' 05, 2005.

[21] Cisco Aironet 802.11a/b/g Wireless CardBus Adapter, "http://www.cisco.com/en/US/products/hw/wireless/ps4555/products data sheet 09186a00801ebc29. html"

[22] T. Rappaport, "Wireless Communications Principle and Practice Second Edition," Prentice Hall 2002.

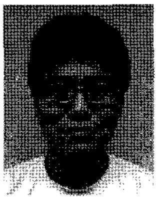

유 준

1997년 KAIST 기계공학과(학사). 2001년~현재 서울대학교 전기·컴퓨터공학부 석박사통합과정. 관심분야는 무선 애드혹망, 무선랜

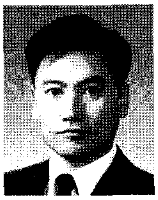

김 종 권

1981년 서울대학교 산업공학과(학사) 1982년 미국 조지아 공과대학교 산업공학과(석사). 1987년 미국 일리노이 대학교 전산학과(박사). 1984년~1987년 IBM 산호세 연구소 연구조원. 1987년~1991년 미국 Belcore 통신연구소 연구원. 1991년~현재 서울대학교 전기·컴퓨터공학부 교수. 관심분야는 차세대 인터넷, 초고속 라우터, 이동통신