A measurement based channel aware scheduler to lessen VoIP capacity degradation in 802.11 networks

Rosario G. Garroppo, Stefano Giordano, Stefano Lucetti, Luca Tavanti
Department of Information Engineering
University of Pisa
Via Caruso, 16 – Pisa, 56122 – Italy
{r.garroppo, s.giordano, s.lucetti, luca.tavanti}@iet.unipi.it

***Abstract*—In the paper we introduce and evaluate the performance of the Deficit Transmission Time (DTT) scheduler, a centralized channel aware scheduling algorithm for IEEE 802.11 networks. The scheduler has been designed to face two inconveniences of 802.11: the performance anomaly and the "inter-flow blocking" problem. DTT splits outgoing frames into multiple queues and uses gross transmission times as a measure of the link quality. DTT guarantees each downlink flow an equal time share of the channel, which translates into the desirable properties of proportional fairness and flow isolation. The performance of the scheduler has been measured in terms of the number of VoIP calls that can coexist in the system with an acceptable perceived speech quality. This assessment has been done on the basis of the E-model, an ITU-T standardized computational tool. The results definitely point out the improvement over the plain FIFO policy.**

***Keywords*—IEEE 802.11; channel aware scheduler; performance anomaly; VoIP.**

I. Introduction

The IEEE 802.11 standard [1] has rapidly become the most popular technology for indoor and outdoor broadband wireless local area networking. This fact has promptly encouraged the research in increasing system capacity. Nevertheless, one of the most critical factors driving the efficiency of 802.11 networks still remains the ability to overcome the obstacles imposed by the wireless channel. The fairly unpredictable capacity of the links leads to unreliable frame delivery time, hence it dramatically hampers network performance, especially with regard to real-time applications such as voice. The same 802.11 DCF mode of operation, joined with the plain *First In First Out* (*FIFO*) strategy adopted at the access point (AP), may even give a boost to this problem.

DCF provides a fair sharing of the medium in terms of channel access probability. In ideal conditions, this translates into an equal portion of the overall effective bandwidth for each user [2]. However, as soon as some frame is not correctly received, 802.11 provides for the retransmission of the frame, with the possibility of adopting more robust but less efficient modulation schemes. Such retransmission attempts occupy the channel at the expenses of all other stations, which perceive a sensible reduction of their available bandwidth regardless of the condition of their own links. In other words, the performance of the network is driven by the station with the worst link. This phenomenon is known as the *performance anomaly* of 802.11 [3].

As for the queuing discipline, most commercially available APs adopt a single FIFO queue. All the frames stored after the currently served one must wait for it to be dequeued. This occurs either after a successful transmission (which may take a long time, due to retransmissions and rate adaptation policies) or when the retransmission limit has been reached. As a consequence, the transmission delay experienced by each frame is also a function of the physical position of the other stations and the number of frames queued in front of it. We refer to this problem as the *inter-flow blocking* problem.

A way to overcome these two hurdles has been reckoned by many to reside in a smart scheduling algorithm to be deployed at the AP (i.e. a *centralized* scheduler). In fact, several centralized schedulers have already been proposed [4][5][6]. Most of them rely on a model of the wireless channel. The links between the base station and the user devices are independent of each other and are subject to bursty errors according to a two-state Markov model: link quality is either good (i.e. error-free) or bad (faulty). Unfortunately, this model is sometimes far from the actual channel behavior, and a system based on it can easily become inefficient. For that reason, a more reliable solution would be centering scheduling decisions on a real measure of the channel.

Starting form these observations, we have proposed a centralized scheduling algorithm that works on a per station basis and takes into account the real channel state. The scheduler is briefly outlined in Section II. A thorough description can be found in [7], together with some experimental tests using a prototype implementation with unidirectional downlink traffic. In the present paper we further analyze its behavior in comparison with the plain FIFO discipline employed in commercial APs focusing on bidirectional VoIP traffic. We carried out a series of simulative tests within the framework defined by the E-model. This approach has already been proved to be practical by Coupechoux et al. [8], who studied the VoIP capacity of an 802.11b network operating in the DCF mode. In their work they showed that the capacity of the network is highly dependent on the position of the terminals, thus proving the

This work has been carried out in the framework of the TWELVE Project, funded by the Italian MIUR Ministry by means of the COFIN research program.

effectiveness of the E-model in revealing the performance anomaly. In our study we adopt a similar approach to demonstrate that the scheduler we propose is successful in mitigating the already mentioned issues.

II. THE DTT SCHEDULER

The architecture of the framework is illustrated in Figure 1. A classifier, several queues and a scheduler engine are inserted between MAC and network layers. The classifier splits outgoing traffic into the queues according to some predefined rule. In our framework, this is performed on the basis of the destination MAC address, hence one queue is needed for each associated station. Note however, that nothing prevents the classifier to support other rules, e.g. based on different user priorities, as in 802.1Q and 802.11e. Each queue is then coupled with a bucket that accounts for the air-time usage of the previous frames. Air time is converted into "tokens" that are used to fill/drain the buckets according to the rules described hereafter.

Once a frame transmission has been completed (either successfully or not), the scheduler computes the *Cumulative Frame Transmission Time* (*CFTT*) that includes all retransmission attempts, backoff and idle periods occurred since the frame has reached the head of the transmission queue at MAC level. The bucket associated to the destination of the dequeued frame is then drained by a number of tokens equal to the CFTT. Next, this same value is divided among all non-empty queues and used to fill their buckets. All the buckets associated to an empty queue are then cleared (set to zero). We have chosen to clear the buckets of empty queues and not to fill them when other queues are served to avoid that some queues, idle for a long period of time, receive too many credits. If that happens, once they have some new frames to transmit, they would grab and monopolize the access to medium in the short-term.

After these tasks have been completed, the next frame to be transmitted is picked from the queue whose associated bucket has the largest number of tokens. If more buckets are at the same level, the scheduler chooses randomly among them. This frame is then passed to the MAC layer, which provides for the physical delivery. Note that all frames are stored at scheduler level, while MAC own queue holds at most the frame under transmission.

The CFTT, in opposition to most channel models, is a deterministic measure (not an estimate, nor a prediction) of the link state: more retransmissions, possibly at lower bit rates, are carried out in the attempt to deliver the frame to stations whose channel quality is poor. Such destinations get their buckets emptied by more tokens, whereas easily reachable destinations get their buckets only lightly drained. In other terms, stations in unfavorable positions have to wait longer before being chosen again, while the bucket with more tokens is also the one whose associated queue has transmitted less. In this way our scheduler can achieve the long term fairness in air time usage: all stations are granted the same amount of downlink transmission time. By the way, from here it also comes the name of the scheduler: Deficit Transmission Time (DTT).

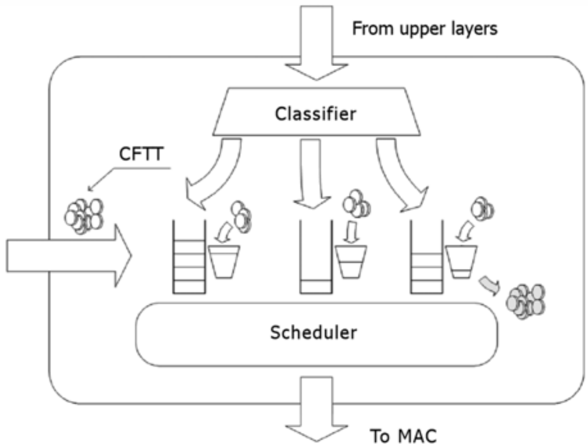

Figure 1. Architecture of the DTT scheduler.

III. THE E-MODEL

The E-model [9] is an ITU-T standardized computational method for the assessment of the quality of voice connections, as perceived by an average user. It allows the designer to calculate the expected speech quality given the transmission characteristics of the network and terminal equipment. The E-model takes into account many parameters, such as the effects of room noise, quantizing distortion, delay, and impairments due to codec and packet loss. The primary output of the model is the scalar rating factor R, that is calculated as:

$$R = R_0 - I_S - I_D - I_{Eeff} + A. \tag{1}$$

In this equation R_0 represents the basic signal-to-noise ratio, including e.g. circuit noise and room noise, I_S is a combination of all impairments occurring simultaneously with the voice signal, e.g. quantizing distortion, I_D accounts for the deterioration caused by delay of voice signals, I_{Eeff} represents impairments caused by the equipment and packet losses, and the advantage factor A allows for compensation of the other factors when the user is likely to accept some degradation of the speech quality due to the adopted technology.

The terms R_0, I_S, I_D and I_{Eeff} are further subdivided into more specific factors, for a total of about twenty atomic parameters. Describing them all is outside the scope of this paper (the interested reader may refer to [9]); we will just mention the ones that are directly related to our study.

One of the most important element is the total delay T_a undergone by each voice packet since its creation. This time can be split into the packetization time T_{pack}, the voice encoding/decoding time T_{DSP}, the network delay T_{nw} and the dejittering time T_{jit}. Network delay T_{nw} can be further divided in two parts, one depending on the wireless LAN (T_{WLAN}), and the other representing the time to traverse the wired portion of the connection (T_{fixed}). Therefore:

$$T_a = T_{pack} + T_{DSP} + T_{fixed} + T_{WLAN} + T_{jit}. \tag{2}$$

A second critical factor is the packet loss ratio P_{pl}. In our system, losses mainly occur in three situations: packets dropped at the AP due to buffer overflow (L_{of}), frames dropped after the retransmission limit has been reached (L_{ch}), and packets arriving at destination with a delay that cannot be compensated by the dejittering buffer (L_{jit}), thus resulting useless for the smooth reconstruction of the speech. This last term is inversely related to the temporal size of the dejittering buffer: the greater T_{jit} the smaller L_{jit} (but also the greater the total delay T_a). Packet losses have impact on the I_{Eeff} factor, and can be mitigated by the robustness of the codec (B_{pl} factor). For a more in depth discussion about the dejittering issues, see again [8].

The rating factor R ranges between 0 and 100, being 100 the better possible connection quality. In most cases, a value of R higher than 70 represents an acceptable level of user satisfaction. Therefore this value will be our threshold for the analysis of network capacity.

IV. Performance Analysis

The proposed scheduler has been evaluated via simulation. In the following Subsections we describe the tool we made use of and the simulation environment. The results are presented in Subsection IV.C and discussed in Section V.

A. Simulation Tool

The tool we have chosen to employ is the OMNeT++ simulator, version 3.0b1 [10]. Since the core library comes with no modules beyond the bare minimum, we have integrated it with the Mobility Framework (version 1.0a3) developed at the Technical University of Berlin [11], and with an 802.11b MAC layer that we have built from scratch.

Given that both the simulator and its parts are not yet deeply established in literature, we carried out some validation tests, comparing the results with known models. Specifically, our baselines were an analytic formula for the maximum achievable throughput and Bianchi's performance study [12].

As for the first model (the details can be found in [13]), the maximum observed difference between the theoretical value and the simulation is in the order of 0.1%, which can be considered negligible. A very similar behavior has been observed also with regard to Bianchi's model. Figure 2 presents the results for a network of 50 stations that works in saturation conditions. As it can be seen, our simulator matches pretty closely the theoretical values.

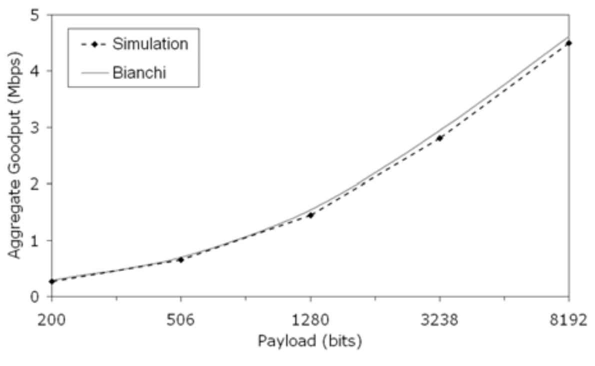

Figure 2. Aggregate goodput (throughput as seen by network layer) versus frame payload for a network of 50 nodes.

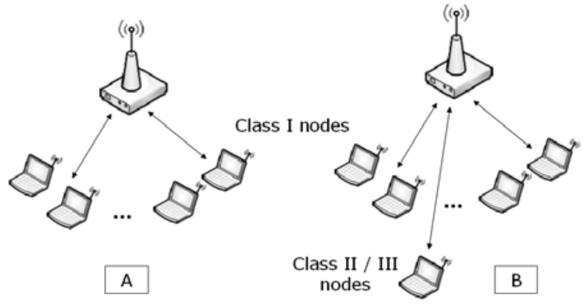

Figure 3. Simulation scenarios.

B. Simulation Environment

We simulated two kinds of scenario (see Figure 3). In the first (let us call it "A"), all stations (or *nodes*) are at the same distance from the AP and experience good channel quality, i.e. all frames are received correctly (apart from collisions). All stations transmit at the highest rate, that is 11 Mbps. In the second scenario (say "B") one or more nodes are far from the AP and therefore their links are very poor, causing frequent failures and the need of some retransmissions in order to deliver each frame. The far station(s) have been placed at two different distances, so that two levels of bad quality are simulated. It is therefore convenient to distinguish the terminals into three classes: class I, comprising the nodes with good link quality; class II, consisting of the nodes at half the way; and class III, that refers to the farthest nodes.

In our simulator, following to the many existing solutions, we also implemented a simple automatic rate adaptation mechanism, that scales down the transmission rate after each failed frame transmission. The bit rate is decreased to 5.5 Mbps after the first attempt, then to 2 and, finally, 1 Mbps. This last value is kept until an ACK is received or the retransmission limit is reached. A summary of the features of the simulated scenarios is reported in Table I. Note that the number of retries only counts those due to the poor channel quality, neglecting collisions (which depend on the number of transmitting stations).

TABLE I. Features of the scenarios

Frame delivery failure rate and mean number of retries		
Class I nodes	0 %	0
Class II nodes	0.15%	1.9
Class III nodes	29.1%	3.2
Class of nodes per scenario		
Scenario A		I only
Scenario B-1		I and II
Scenario B-2		I and III

Note that scenario A is the ideal case, with all nodes working at the best of their capacities; on the contrary, scenario B-2 illustrates a very critical situation: while class III nodes have no possibility to sustain any voice service (around 29% of the frames is never received), they nevertheless spoil the chances of the other nodes. We therefore expect a significant reduction in the system capacity, event that in fact occurs. Scenario B-1 represents a compromise, as class II nodes do encumber the network, but they can still carry voice traffic.

In every scenario all the stations are involved in a bidirectional voice call in which the other end is represented by a remote terminal connected to the AP along a wired network. Voice frames are produced by a GSM-EFR encoder and encapsulated into an IP packet, which is in turn transported by the RTP protocol. Each voice source is modeled according to the ITU-T recommendation P.59 for artificial conversational speech: the source alternates on and off periods, whose length is described by an exponential distribution with mean 1 and 1.35 seconds respectively. During the on periods the source transmits at the GSM-EFR nominal rate (12.2 kbps), during the off periods it is silent. Table II resumes all the values.

As for the parameters of E-model, the choice of the codec and the network topology determined the values of some of the involved factors, while we assigned the remaining (in fact, the greatest part) the default values. A summary of the most meaningful factors, with their assigned value, is reported in Table III. These are the typical figures for modern equipment (see e.g. [8][14]). Note that, since our attention is focused on the wireless access, we have given T_{fixed} a constant value, in order to let the features of the WLAN emerge more clearly. The results of the simulations gave the values of T_{WLAN} and P_{pl} for each run.

Finally, the proposed scheduler has been located at the AP, while all the stations always work with the plain 802.11 cards. The maximum number of queued frames is the same for both the FIFO and the DTT schedulers. For the latter, it sums up the frames in all queues.

TABLE II. SIMULATION PARAMETERS

Codec GSM-EFR	
Size of voice frames	244 bit
Frame interval	20 ms
RTP / UDP / IP header	320 bit
MAC IEEE 802.11b	
RTS / CTS	disabled
Retransmission limit	4
Max. no. of packets in all queues	150

TABLE III. E-MODEL PARAMETERS

Packetization time T_{pack}	20 ms
Voice encoding/decoding time T_{DSP}	10 ms
Wired network delay T_{fixed}	50 ms
Dejittering time T_{jit}	40 ms
Robustness to packet losses (B_{pl})	10
Equipment impairment factor (I_E)	5
Advantage factor A	0
Rating factor threshold	70

C. Simulation results

For each simulation scenario we tried to evaluate the maximum number of stations allowed in the network with all the users experiencing a satisfactory speech quality, i.e. having $R \geq 70$. All voice calls starts at the beginning of the simulation and last until the end. The simulation time has been set to 210 seconds. For each scenario we averaged the results obtained in five simulation runs with different seeds for the random number generator and collected the statistics for the stations with the worst quality in each class.

Figure 4 depicts the rating factor for the first scenario with a different number of user stations (N). As it can be noted, the insertion of our scheduler in this context does not change much the behavior of the network, being the values very close to the ones for the plain FIFO strategy. Having set the threshold for the rating factor R at 70, up to 24 stations can be supported under this configuration with either the standard FIFO discipline or the DTT scheduler.

In accordance with [8] and [14], we also noted that the R factor is much more sensitive to packet losses than to increased delays. Having a look at Table IV, where T_{WLAN} and L_{jit} are reported (L_{of} and L_{ch} are both zero), we can find the confirmation of the previous statement. In the case of 24 stations, L_{jit} registered for the FIFO policy is greater than that for the DTT, while T_{WLAN} is smaller. The resulting R factor promotes the DTT, which gains 1.4 points over the plain FIFO. These position holds throughout all the scenarios.

TABLE IV. RESULTS FOR THE SCENARIO "A"

N	Scheduler	T_{WLAN} [ms]	L_{jit} [%]	*R*
23	FIFO	2.38	0.44	81.4
	DTT	2.13	0.21	83.4
24	FIFO	3.80	1.76	71.7
	DTT	3.93	1.55	73.1
25	FIFO	7.10	3.87	60.1
	DTT	8.81	3.99	59.4

As soon as a station in a bad position joins the network, the two analyzed schedulers start showing some significant divergences. Figure 5 reports the results for the B-1 scenario with one and two class II nodes. It should be remembered that in this case the number of stations reported along the abscissa also includes the terminal(s) in bad position.

The plain FIFO policy clearly does not distinguish among the stations, therefore all users experience almost the same quality of the worst station. Consequently we can see that no more than 18 users can be supported when there is just one class II node (plot on the left). The slight difference in the R factor comes from the longer frame transfer delays and the small amount of lost frames.

On the contrary, the adoption of DTT results in a sharp separation of the two classes. The scheduler is able to preserve a very good speech quality to near stations by penalizing the far

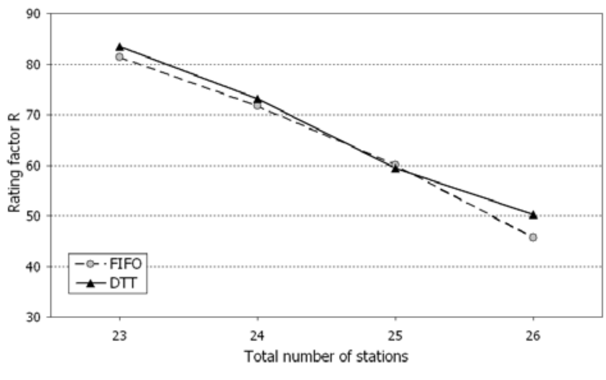

Figure 4. The rating factor R versus the number of stations in the network for the scenario A.

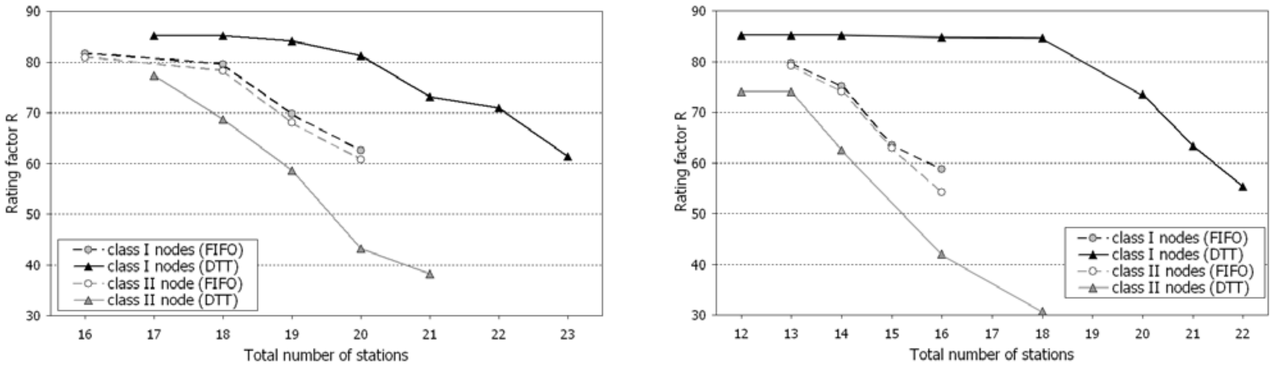

Figure 5. The rating factor R versus the number of stations in the network for the scenario “B-1” with one class II node [left] and two class II nodes [right]. Black lines refer to nodes close to AP (class I), gray lines to nodes far from AP (class II).

nodes. Therefore we can see how the system can support up to 22 users before they all experience a critical connection decline. If we do not count the single class II node that is impaired by the AP, the DTT provides a net gain of three users over the plain 802.11 discipline.

The gap between the two strategies gets larger and larger as the number of far nodes increases. As an example, the plot on the right of Figure 5 illustrates the results for the B-1 scenario with two class II stations. Under the FIFO governance, only 14 users are allowed in the network, with a drop of 10 units with respect to the ideal case. On the other hand, the DTT reduces the loss to 6 nodes, 4 due to network saturation and 2 due to the serving policy.

Finally, we studied the performance under the scenario B-2, when all nodes but one are in a good position (class I), and the last node is subject to a very faulty link (a class III node). Note that the class III terminal, due to the high frame loss percentage, can never keep up a voice connection with an acceptable level of quality. Therefore we can limit the capacity analysis to class I nodes only.

Figure 6 reports these results; the rating factor is plotted for class I stations only. In this case the adoption of the DTT scheduler has a major impact on system performance. The capacity is increased from 12 (under the FIFO policy) to 20 class I stations, i.e. the capacity is almost doubled. Moreover, when the FIFO-based AP is already in a critical condition, DTT can still offer a very high speech quality. This happens e.g. with 14 class I nodes: for the DTT scheduler we have registered the highest *R* (85.2), while the FIFO policy is already well below the threshold ($R = 47.0$).

V. Discussion

From the simulation trials, it has clearly emerged that the insertion of a smart scheduler on top of the plain 802.11 card has brought a significant improvement in network capacity. As already mentioned in Section I, the presence of just one station with a poor link to the AP damages all the associated stations. This is because the simple FIFO-driven AP, in the effort of supplying all stations with the same bandwidth, reduces the overall network performance. Unfortunately this kind of try has no benefit neither to the farthest users, whose connection is intrinsically hampered by propagation conditions, nor to the other users, whose speech quality is dragged down to the values of the worst one. This is visible in all the presented figures, where *R* has the same trend for all nodes, and further emphasized in Table V: the average delay and the number of packets discarded due to excessive latency are comparable for both classes of stations.

On the contrary the DTT scheduler sharply separates the behavior of the two classes of stations. While the far user sees its bandwidth severely reduced, the closer ones can still perceive the channel in a good state, thus being able to sustain a voice connection until the network saturates (see again Table V and Figure 5). Remarkably, in all “B” scenarios this event occurs almost at the same stage as for the ideal “A” case. For example, in the B-2 scenario (Figure 6) the DTT allows the network to host up to 20 calls before the speech quality of any station (apart from the far one, of course) decays below the satisfaction threshold. Employing the FIFO policy would have granted only 12 calls to coexist in the network, with a 50% drop from the ideal scenario. The DTT reduces this gap to just 4 users (i.e. 17%).

TABLE V. Some results for the scenario “B-2”

Scheduler	N	Class (# nodes)	T_{WLAN} [ms]	L_{jit} [%]	L_{of} [%]
FIFO	13	I (12)	4.07	1.19	0
		III (1)	10.9	1.37	
FIFO	14	I (13)	6.08	3.12	0
		III (1)	12.9	2.93	
FIFO	15	I (14)	13.5	7.26	0.01
		III (1)	21.6	8.14	
DTT	21	I (20)	7.61	1.22	0
		III (1)	510	33.9	
DTT	22	I (21)	8.85	2.12	0
		III (1)	576	33.3	

It should again be remarked that the main reason of capacity degradation is not network saturation, but the increasing value of the jitter, that makes more and more packets useless for speech reconstruction. Table V also shows L_{of}. This parameter, that gives a measure of congestion through the number of frames dropped at the AP due to buffer overflow, is zero, or close to zero, in all cases. At the same time, L_{jit} more than doubles for every added user under the FIFO scheme, whereas it increases more smoothly when the DTT is working .

Such an outstanding gain can be explained with the principle lying behind the design of our scheduler. By granting each flow the same amount of channel occupancy time, the DTT achieves the valuable property of *proportional fairness* in bandwidth usage. As already stated in [15], this quality represents a good compromise between two extreme goals: maximum network efficiency and maximum fairness (the one of the plain 802.11 strategy). In practice, as shown in [2], proportional fairness in a multi rate environment, such as the one we studied, translates into equal usage of air time for all stations. But [2] also proved that a key effect of applying this principle is the isolation of the flows: each station gets the maximum of its bandwidth share independently of the others.

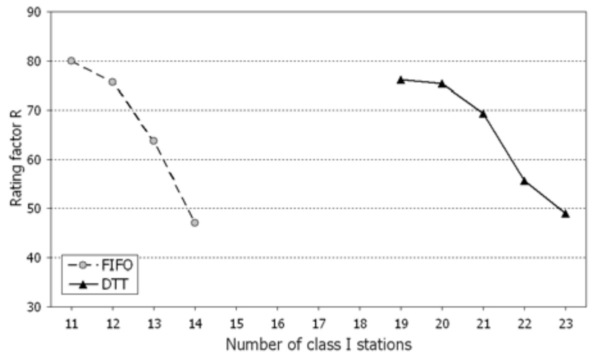

Figure 6. The rating factor R versus the number of users in the network for the scenario "B-2". Both refer to class I stations only.

VI. CONCLUSIONS

In this paper we have presented a simple centralized channel-aware scheduler, whose main feature is the achievement of flow isolation. This property has been attained by splitting outgoing traffic into multiple queues and using an indirect but reliable measure of the link quality: the CFTT. This variable accounts for all the events that occurs during each frame transmission cycle. With this approach, our scheduler manages to reduce the effects of both the performance anomaly and the inter-flow blocking problem. We have verified the effectiveness of the proposed solution via simulative trials, on the basis of the standardized E-model. The outcome of these tests has highlighted that our scheduler is capable of exploiting the available resources to a greater extent than the simple FIFO discipline of standard 802.11 networks. The proof of that is in the maximum number of tolerable users, which keeps pretty close to the ideal case in all the examined scenarios.

REFERENCES

[1] ANSI/IEEE Std 802.11, 1999 Edition. Part 11: wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, 1999.

[2] L. B. Jiang, S. C. Liew, "Proportional Fairness in Wireless LANs and Ad Hoc Networks", IEEE Wireless Communications and Network Conference, March 2005.

[3] M. Heusse, F. Rousseau, G. Berger-Sabbatel, A. Duda, "Performance Anomaly of 802.11b", Proc. of IEEE Infocom 2003, San Francisco, April 2003.

[4] Y. Cao, V.O.K. Li, "Scheduling algorithms in broadband wireless networks", Proceedings of the IEEE, Vol. 89, Issue 1, Jan. 2001, pp. 76-87.

[5] M. Bottigliengo, C. Casetti, C. F. Chiasserini, M. Meo, "Short-term Fairness for TCP Flows in 802.11b WLANs", Proc. of IEEE Infocom, Hong Kong, March 2004.

[6] X. Liu, E. K. P. Chong, N. B. Shroff, "Opportunistic Transmission Scheduling With Resource-Sharing Constraints in Wireless Networks", IEEE Journal on Selected Areas in Communications, Vol. 19, Issue 10, Oct. 2001, pp. 2053–2064.

[7] R. G. Garroppo, S. Giordano, S. Lucetti, and L. Tavanti, "Providing Air-time Usage Fairness in 802.11 Networks with the Deficit Transmission Time (DTT) Scheduler", Wireless Networks, in press.

[8] M. Coupechoux, V. Kumar, L. Brignol, "Voice over IEEE 802.11b Capacity", ITC Specialist Seminar on Performance Evaluation of Wireless and Mobile Systems, Antwerp, Belgium, August 2004.

[9] ITU-T Recommendation G.107, "The E-model, a computational model for use in transmission planning", March 2003.

[10] The OMNeT++ discrete event simulation system. Available at: http://www.omnetpp.org.

[11] Mobility Framework for OMNeT++. Available at: http://mobility-fw.sourceforge.net/hp/index.html.

[12] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function", IEEE Journal on Selected Areas in Communications, Volume 18, Issue 3, March 2000, pp. 535–547.

[13] R. G. Garroppo, S. Giordano, S. Lucetti, F. Russo, "IEEE 802.11b Performance Evaluation: Convergence of Theoretical, Simulation and Experimental Results", Networks 2004, Wien 2004, volume 1, pp. 405-410.

[14] D. De Vleeschauwer and J. Janssen, "Voice Performance over packet-based networks", Alcatel White Paper, Oct. 2002.

[15] B. Radunovic, J. Le Boudec, "Rate Performance Objectives of Multihop Wireless Networks", IEEE/ACM Trans. on Mobile Computing, Vol. 3, Issue 4, pp. 334–349, Oct.-Dec. 2004.