Fair Time Sharing Protocol: a Solution for IEEE 802.11b Hot Spots

Anelise Munaretto[1,2], Mauro Fonseca[1,4], Khaldoun Al Agha[2,3], and Guy Pujolle[1]

[1] LIP6 Laboratory, University of Paris VI, 75015 Paris, France
[2] LRI Laboratory, University of Paris XI, 91405 Orsay, France
[3] INRIA Laboratory, Rocquencourt, 78153 Le Chesnay, France
[4] PPGIA, PUC-PR, 80215-901 Curitiba, Parana, Brazil

Abstract. To adapt the data rate in accordance with the quality of the link, the IEEE 802.11b standard proposes the variable rate shifting functionality. This intrinsic functionality of the 802.11b products progressively degrades the bit rate when a host detects unsuccessful frame transmissions. Furthermore, the basic CSMA/CA channel access method guarantees that the long-term channel access probability is equal for all hosts. When one host captures the channel for a long time because its bit rate is low, it penalizes other hosts that use the higher rate, inciting a performance anomaly. This paper aims at avoiding this performance anomaly and the consequent waste of bandwidth. We propose the Fair Time Sharing (FTS) approach to perform real fair sharing among the active hosts in the hot spot, thus avoiding the performance degradation caused by one or more slow hosts. This paper presents the FTS architecture and its performance evaluation, showing the improvement achieved.

1 Introduction

The economic feasibility and an easy deployment of IEEE 802.11b for wireless local area networks, leading to the wireless mobile Internet, incite its widely adoption as future access networks. The raise of the bit rate in the wireless networks technologies offers to these networks a comparable performance with respect to wired local area networks. However, the wireless nature of the radio networks introduces many problems that were unknown in the wired networks. These are mainly due to the transmission of a signal over the radio medium. Furthermore, there are various effects due to fading, interference from other users, and shadowing from objects, all them degrading the channel performance [JA94]. Measurements of a particular radio LAN appear in [DU92], showing that the packet-error rates critically depend on the distance between the transmitter and receiver, and surprisingly, not monotonically increasing with the distance. Because of such constraints, transmission in radio networks is not reliable enough to allow the construction of a control based access mechanism for the network management. Likewise, several performance evaluation studies of the IEEE 802.11 DCF [CH04] [CR97] [WE97] [BI00] show that performance is very sensitive to the number of competing stations on the channel, especially when the Basic Access mode is employed. To adapt the data rate in accordance with the quality of the link, the IEEE 802.11 standard [ST99] proposes the variable rate shifting functionality.

This intrinsic functionality of the 802.11b products progressively degrades the bit rate from the nominal 11 Mbps to 5.5, 2, or 1 Mbps when a host detects unsuccessful frame transmissions. The multiple data transfer rate capability performs dynamic rate switching with the objective of improving performance. Nevertheless, to ensure coexistence and interoperability among multirate-capable stations, the standard defines a set of rules that shall be followed by all stations. A performance anomaly causing a performance degradation in a hot spot due to the influence of a host with lower bit rate, was first analyzed in [HE03]. The authors examine the performance of 802.11b showing that the throughput is much smaller than the normal bit rate. They also analyze how a host with a lower bit rate influences the throughput of the other hosts that share the same radio channel. This results in a performance anomaly perceived by all hosts. The authors solely analyze the performance anomaly without introducing any solution. Our paper introduces the new *Fair Time Sharing* approach to avoid the performance anomaly in a dynamic and adaptive manner. Such an approach is based on the attribution of time slot for each host, to perform a real fair sharing among the active heterogeneous hosts in the hot spot.

2 Fair Time Sharing (FTS)

The fair access to the channel provided by CSMA/CA causes a slow host transmitting at 1 Mbps to capture the channel for a period eleven times longer than hosts transmitting at 11 Mbps [HE03]. Such a procedure induces, at the long-term, a non equitable time sharing, because the active hosts degrade their bit rate to the smaller bit rate in the hot spot, e.g. around 1 Mbps. This characterizes a performance anomaly. We introduce the *Fair Time Sharing* approach to avoid the performance anomaly. In our approach we fairly share the access time destined for each host according to the configured bit rate. We attribute a specified time slot to each slow host, shaping their maximum output traffic to prevent the unfair share due to the data rate degradation. The assessment of the time slot is based on the number of competing hosts within the hot spot and the useful throughput. The Fair Time Sharing (FTS) architecture proposes a simple architecture composed by three components over the UDP protocol: FTS-Server, FTS-Client, and Traffic shaper. To achieve a better bandwidth optimization in the hot spot, two classes are implemented: Slow class and Fast class. In the slow class, all client stations that can compromise the hot spot performance are included, i.e. client stations having a slow data rate. In such a class, the resources are shared based on the number of clients and a resource reservation is performed. In the fast class, all client stations work in fast data rate. The FTS-Server is located on the access point. The access point logs the number of connected users in the hot spot and their corresponding bit rate. The FTS-Server broadcasts this number of connected users and sets the fast class threshold. It works over UDP to avoid connection state management. The FTS-Client is located on the client station, it receives the FTS-Message and implements the traffic shaper. It works over UDP remaining simple and connectionless. The Traffic shaper is located on the client station and responsible for adapting the outgoing traffic to network bandwidth availability. This traffic shaping policy is implemented using the Traffic Control (`TC`) program from the “iproute2” package, detailed in [HOTC]. It works at network level.

Each FTS-Client enforces the calculated allowed time slot using the traffic shaper element. The FTS-Client does not limit the time slot when it belongs to the fast class, using as much bandwidth as possible. This means that a FTS-Client in fast class uses the default fair sharing method, allowing this class of clients to also consume the idle bandwidth left by slow class clients. In order to optimize the bandwidth utilization, the FTS-Message sets a threshold for the fast class. For instance, when there is no host working at 11 Mbps, the threshold for the fast class is set to the highest data rate in the hot spot. Therefore, this field allows changing the maximum data rate used in the fast class in a dynamic manner. Then, each node *i* maintains a local control of the bandwidth in setting its traffic shaper.

3 Performance Measurements

In this section, we present the performance measurements of the proposed FTS protocol. The measurements are performed using the **iperf** tool [IPER]. The **iperf** represents a TCP/UDP bandwidth measurement tool. We set up a testbed platform to measure the throughput that hosts can obtain when sharing a 11 Mbps 802.11b wireless hot spot. We have used three notebooks (host1, host2, and host3) with 802.11b cards (Lucent Orinoco). The wired part of the network is connected by the access point. We implement the FTS protocol comparing the results with a classical hot spot. The first step is to measure the real useful throughput in data rate degradation. These measurements depicted in Table 1, are stored in the host to gauge its allotted time slot and are then used to set the traffic shaper. Table 1 presents the measurements of the useful throughput through our performance evaluation results.

Table 1. The useful throughput obtained for one single host in a hot spot with UDP traffic.

Bit Rate	*Useful throughput*	*Proportion of useful throughput*
11 Mbps	6.05 Mbps	55%
5.5 Mbps	3.81 Mbps	69.27%
2 Mbps	1.68 Mbps	84%
1 Mbps	0.891 Mbps	89.1%

The FTS protocol uses a bandwidth value to set the traffic shaper to avoid the degradation of the average throughput. This bandwidth is calculated according to the number of connected clients in the hot spot and the maximum useful throughput, presented in Table 1. In the measurement environment, the choice of the destination nodes for the flows is somewhat arbitrary and any destination could have been chosen for each flow without affecting the results. Unless otherwise specified, the following assumptions are made: Each flow is active throughout the duration of the experimental. All packets on all flows contain 1534 bytes. The time duration of our implementation is of 400 s. The results of the first experimentation are depicted in Fig. 1(a). Observe that this experimentation shows the effect perceived by improving the mobility of one host in the hot spot.

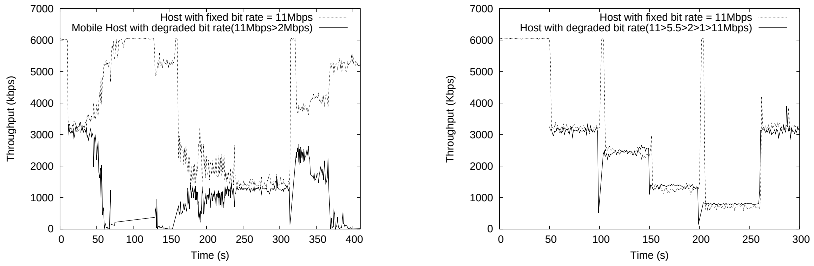

(a) Two hosts in mobility behavior and bit rate degradation.

(b) The throughput degradation according to the bit rate in two hosts.

Fig. 1. Measurements results

In time 0 s, only one station (host1) is active in the hot spot with 11 Mbps as bit rate and approximately 6 Mbps as useful throughput. After approximately 10 s, another station (host2) also with 11 Mbps as bit rate arrives in the hot spot. The useful throughput decreases to 3 Mbps for the two hosts. Afterwards, the host2 starts to move away from the Access Point, decreasing its useful throughput, although it keeps the same bit rate. At the other hand, we can see the throughput increase of host1. From 50 s until 150 s, the useful throughput perceived by host2 is strongly degraded. At about 160 s, we set the bit rate to 2 Mbps only for the host2. In this instant, the slower host increases its useful throughput to about 500 kbps. Thus, the fast host decreases drastically its useful throughput to about 2.5 Mbps. In such a case, we can verify that the degradation of the bit rate in one mobile host imposes the degradation in the other host. However, as the slower host gets closer the Access Point, and consequently near the fast host, the useful throughput becomes the same. For the two hosts this stabilization is perceived between time 240 s and time 310 s, when the two stations achieve approximately the same throughput. Afterwards, about the time 320 s, we set the bit rate to 11 Mbps for the host2. In such an instant, the faster host increases its useful throughput to about 4 Mbps and the slower host to about 2.5 Mbps. Then, the useful throughput of the host2 starts again to decrease because of the mobility. Next experimentation measures the throughput of the competing hosts if there are 2 heterogeneous hosts in the hot spot. The results presented in Fig. 1(b), show the final throughput according to the bit rate degradation. At the beginning, only one host is active in the hot spot, thus achieving 6 Mbps as useful throughput. At 50 s, a second host arrives in the hot spot, fairly sharing the useful throughput. At 100 s, the bit rate of the first host is degraded to 5.5 Mbps. Note that the degradation of the 2 hosts reaches nearly the same throughput, about 2.5 Mbps, despite that only one host has the bit rate degraded. The same anomaly is depicted each time that the bit rate of a single host is degraded. At 260 s, the bit rate of the slower host is once again set to 11 Mbps. Then, we can verify that the useful throughput perceived by the two hosts is again of 3 Mbps. The effectiveness of the proposed FTS approach is demonstrated in Fig. 2. This scenario demonstrates that our proposed technique is able to track abrupt variations in the network state, while keeping a very high level of accuracy.

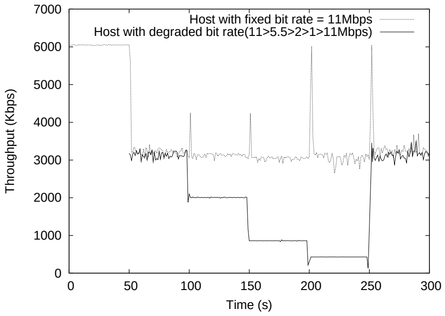

Fig. 2. The Fair Time Sharing improvement avoiding the throughput degradation.

In these experimentations, our FTS-Server broadcasts each 5 s the FTS-Message that represents the number of connected clients (2 clients) and the fast class threshold (11 Mbps). At 100 s, the bit rate of the first host, is degraded to 5.5 Mbps. In this point, the FTS-Client checks its bit rate and the corresponding useful throughput according to Table 1. In such a case, as we have 5.5 Mbps as bit rate configuration, we use the 3.81 Mbps value. Then, the FTS-Client divides this value with the last received number of connected clients cad 1.905 Mbps. This value is applied in the traffic shaper limiting the output bandwidth as shown in Fig. 2. We verify that the FTS approach prevents the degradation in faster host's throughput. Only the hosts belonging to a slow class have their throughput degraded. Likewise, the same improvement occurs when the bit rate of the slower host is degraded to 2 Mbps. The FTS-Client checks its bit rate and the corresponding useful throughput according to Table 1. As we have 2 Mbps as bit rate configuration, we use the 1.68 Mbps value. Then, the FTS-Client divides this value with the last received number of connected clients, that is equal in this case to 0.840 Mbps. Therefore, this value will be applied in the traffic shaper limiting the bandwidth output as stated in Fig. 2. The FTS-Client checks its bit rate and the corresponding useful throughput according to Table 1. In such a case, as we have 1 Mbps as bit rate configuration, we use the 0.891 Mbps value. Then, the FTS-Client divides this value with the last received number of connected clients, that is equal in this case to 0.445 Mbps. Therefore, this value will be applied in the traffic shaper limiting the output bandwidth as depicted in Fig. 2. After 250 s, the bit rate of the slower host is once again set to 11 Mbps. Then, we can verify that the useful throughput perceived by the two hosts returns to 3 Mbps. A drawback of the FTS architecture is the need to deploy a FTS-Server application in the Access Point in order to broadcast the number of clients and a FTS-Client application in every host to set the shaper. Therefore, we assume that all stations in the hot spot support the FTS architecture. Our findings show the improvement achieved by our architecture: the faster host keeps its throughput unchanged regardless of the presence of a slower competing host.

4 Conclusion

In this work, we have analyzed the performance of a 802.11b network. This performance is very sensitive to the number of competing stations on the channel. Through our analysis, we verify that a lower bit rate host influences the throughput of other hosts that share the same radio channel. This paper has discussed the impact of bit rate degradation in an 802.11b wireless LAN that characterizes a performance anomaly. We have then introduced a new approach, *Fair Time Sharing* to avoid the performance anomaly. The proposed FTS approach implements a real fair sharing scheme among the active heterogeneous hosts in the hot spot, avoiding the performance degradation caused by one or more slow hosts. We attribute a specified time slot to each slower host, shaping their maximum traffic during the degradation period. The FTS approach presents the cure for the performance anomaly of 802.11b.

References

[JA94] William C. Jakes, ”Microwave Mobile Communications, Wiley-IEEE Press., May, 1994.

[DU92] Dan Duchamp and Neil F. Reynolds, ”Measured Performance of a Wireless LAN”, In Proc. of 17th Conf. on Local Computer Networks, Minneapolis, 1992.

[CH04] H. Chaouchi and A. Munaretto, ”Adaptive QoS Management for IEEE 802.11 future Wireless ISPs”, ACM Wireless Networks Journal, Kluwer, Vol.10, No. 4, July, 2004.

[CR97] B. P. Crow, I. Widjaja, J. G. Kim, and P. T. Sakai, ”IEEE 802.11 Wireless Local Area Networks”, IEEE Communications Magazine, September, 1997.

[WE97] J. Weinmiller, M. Schlager, A. Festag, and A. Wolisz, ”Performance Study of Access Control in Wireless LANs IEEE 802.11 DFWMAC and ETSI RES 10 HIPERLAN”, Mobile Networks and Applications, Vol.2, 1997.

[BI00] G. Bianchi, ”Performance Analysis of the IEEE 802.11 Distributed Coordination Function”, IEEE Journal of Selected Areas in Telecommunications, Wireless Series, Vol.18, March, 2000.

[ST99] IEEE 802.11: ”Wireless LAN medium access control (MAC) and physical layer (PHY) specifications”, 1999.

[HE03] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda., ”Performance anomaly of 802.11b”, In Proceedings of IEEE INFOCOM 2003, San Francisco, USA, 2003.

[IPER] http://dast.nlanr.net/Projects/Iperf/

[HOTC] http://tldp.org/HOWTO/Adv-Routing-HOWTO/