Unicast-Friendly Multicast in IEEE 802.11 Wireless LANs

Nakjung Choi, Jiho Ryu, Yongho Seok, Yanghee Choi and Taekyoung Kwon
School of Computer Science and Engineering
Seoul National University, Seoul, Korea
Email: {fomula, jhryu, yhseok, yhchoi, tk}@mmlab.snu.ac.kr

***Abstract*— The IEEE 802.11 protocol has become the de facto standard in wireless LANs. However, it reveals the unfairness problem between unicast flows and multicast flows since multicast packets are not subject to binary exponential backoff. To prevent a multicast flow from overwhelming the wireless link bandwidth is a crucial issue. This paper seeks to achieve fairness between unicast and multicast flows by introducing Unicast-Friendly Multicast (UFM). The central idea behind the UFM algorithm is to dynamically change the contention window size for multicast packets to limit the bandwidth share of a multicast flow equal to that of a unicast flow. The proposed UFM algorithm adjusts the multicast contention window size depending on the number of competing stations. We present two versions of UFM: the first one calculates the multicast window size by inferring the average contention window size of unicast flows, while the second one maintains the mapping table between the number of competing stations and the corresponding multicast window size given by related analysis. Simulation reveals that both versions of UFM achieve the fairness by providing almost the fair share of bandwidth to each flow regardless of unicast or multicast under the saturated network conditions.**

I. Introduction

Multicast is the capability to support point-to-multipoint communications between a single sender and multiple receivers. This capability can be employed at the different layers such as application layer, IP layer and MAC layer in a network protocol stack. Numerous multicasting routing algorithms are proposed at the IP layer (e.g. DVMRP and PIM). Also, to support multicast routing protocols, Internet routers use a group membership protocol (e.g., IGMP for IPv4). However, multicasting at the IP layer requires all intermediate Internet routers to upgrade. Therefore, more realistic proposals focus on multicasting overlay at the application layer, called *end-to-end multicast*. On the other hand, due to the broadcasting nature of link layer communications, there have been little research works on multicasting at the MAC layer in the literature.

In the IEEE 802.11 wireless LANs, multicasting is specified as a simple broadcasting mechanism at a fixed rate with no ACK. Unlike unicasting, there is no binary exponential backoff process in multicast packets, which may allow the multicast traffic to dominate the wireless link. Accordingly, when reliable unicast flows and unreliable multicast flows coexist, multicast flows will get more channel access chances than unicast flows with binary exponential backoff. This unbalance or unfairness causes the degradation of the aggregate throughput of unicast flows. Therefore, we propose a novel MAC protocol, called Unicast-Friendly Multicast (UFM), to provide a fair bandwidth share to every flow regardless of unicast or multicast.

The rest of this paper is organized as follows. Section II explains the unfairness problem between unicast flows and multicast flows in IEEE 802.11 wireless LANs. Then, we define the problem and propose two types of Unicast-Friendly Multicast (UFM) in Section III and Section IV, respectively. Our proposal makes multicast flows to dynamically adjust their backoff windows to prevent the throughput degradation of unicast flows. The performance evaluation using ns-2 simulation is presented in Section V, and then we conclude this work in Section VI.

II. Background

In this section, we first explain the fairness problem between unicast and multicast flows and review several proposals in the literature.

A. Unfairness Problem between Unicast and Multicast

In [2], the authors investigated three bandwidth allocation policies: Receiver Independent (RI), Linear Receiver Dependent (LinRD), and Logarithmic Receiver Dependent (LogRD). These policies differ in how to allocate the available bandwidth between unicast flows and multicast flows to prevent unicast flows from starving. RI allocates a fair share of bandwidth to each flow regardless of unicast or multicast. LinRD and LogRD allocate a multicast flow a bandwidth share, which is linearly or logarithmically proportional to the number of receivers downstream.

TCP-Friendly Multicast Congestion Control (TFMCC) was presented in [3]. TFMCC is an equation-based multicast congestion control mechanism that extends the TCP-friendly TFRC protocol [1] from the unicast to the multicast domain. A major contribution is the feedback mechanism to inform the multicast source of network traffic conditions, so that the multicast traffic cannot overwhelm other unicast flows. The feedback mechanism consists of scalable round-trip time measurements, appropriate feedback suppression, both of which ensure that feedback delays in the control loop do not adversely affect fairness with contending flows.

B. Reliable Multicast in IEEE 802.11 Wireless LANs

Most researches on multicasting in IEEE 802.11 wireless LANs have focused on how to improve the reliability of

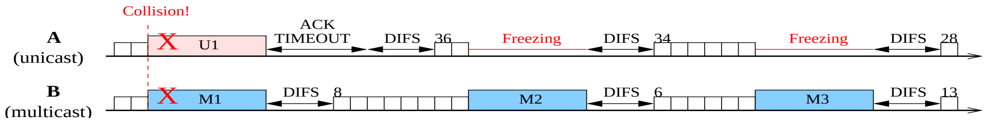

Fig. 1. Contention for a unicast flow and a multicast flow

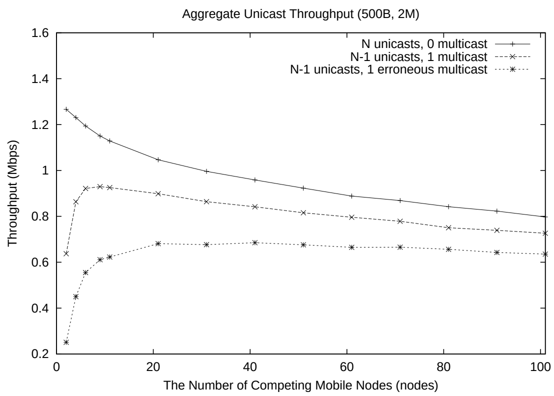

Fig. 2. Aggregate Throughput of Competing Stations

multicast flows, which are based on *RTS/CTS* and/or *ACK*. The authors in [4] introduced a naive multicast acknowledgement protocol in which every recipient should reply with an ACK packets, which causes the ACK implosion problem. In [5], a leader-based protocol utilizing RTS/CTS/ACK/NAK exchange was proposed to provide a partial reliability since only one of the recipients will reply with ACK/NAK. This approach solves the ACK implosion problem but the reliability is not perfect at all. On detection of collisions, the above proposals perform the binary exponential backoff mechanism for multicast flows like unicast flows.

A novel approach called Early Multicast Collision Detection was proposed in [8] to increase the reliability of a multicast packet in the case of frequent collisions. This algorithm operates in three phases. The first phase is the probe transmission, in which a sender with a multicast packet transmits the early multicast collision detection packet of a small size. Right after the transmission comes the second phase, the listening period, during which the sender detects whether the channel is idle or not. If the channel is idle, the sender will transmit the multicast packet. Otherwise, the sender detects collisions and transmits more jamming packets to make other contending nodes detect collision, which is the last phase.

III. Problem Definition

Currently, the basic operation of IEEE 802.11 DCF specifies only a simple multicasting functionality at the MAC layer. With an ACK mechanism at the MAC layer, a station is able to infer the successful receptions of unicast packets if ACKs are returned. Otherwise, it detects transmission failures and performs binary exponential backoff. Until its contention window reaches the maximum size, its contention window size is doubled at each round. When the backoff timer expires, it retransmits the unsuccessfully transmitted unicast packet. If the number of retransmissions exceeds a pre-defined *retry counter*, the station drops the unicast packet and attempts to transmit the next unicast packet. However, a station transmitting multicast packets performs neither the ACK mechanism nor the binary exponential backoff mechanism. Therefore, it attempts to transmit the next multicast packet immediately after the backoff counter expires. In this backoff process, the contention window size of a unicast flow is doubled on every collision; however, the contention window size of a multicast flow is always fixed at the minimum value due to the absence of the ACK mechanism. Because of these different operations, in the congested networks, there will be the unfairness problem between unicast and multicast flows. As the number of contending unicast and multicast flows is increasing, the collisions between unicast and multicast packets will occur more frequently. On each collision, a station with the unicast packet performs binary exponential backoff while another station with the multicast packet will keep the same size of backoff window(so-called post-backoff). Because the contention window size of a unicast flow after binary exponential backoff is larger than that of a multicast flow, it is highly likely that stations with multicast packets will get the channel access much more than ones with unicast packets.

Fig. 1 illustrates the result of a collision between a unicast flow and a multicast flow. We assume that a unicast packet (U1) and a multicast packet (M1) collide at the first time. A station that transmitted U1 waits for an ACK from the receiver during ACK timeout but there is no ACK, resulting in binary exponential backoff. In the case of IEEE 802.11b, it therefore picks up a random number from the doubled range [0, 63], which is 36 in this example. However, the station that transmitted the M1 is not aware of the collision and performs post-backoff. A random number in the post-backoff stage is chosen from the range [0, cw_{min}-1], which is 8 in this example. Therefore, the station with a multicast flow obtains the channel access and transmits the next multicast packet. On the other hand, the station with the unicast flow is waiting for the channel access until the large backoff timer expires.

Recently, there are several proposals like UDP-Lite [9] to exploit the effect of erroneous packets for multimedia services. However, erroneous multicast packets can worsen this unfairness. According to the IEEE 802.11 specification, stations receiving an erroneous packet should defer EIFS (EIFS is usually much longer than DIFS.) instead of DIFS. If stations with unicast flows receive erroneous multicast packets, the probability that they can acquire the channel is even less. Fig. 2 shows the effect of a multicast flow to the aggregate throughput under the saturated network condition.

IV. UNICAST-FRIENDLY MULTICAST

To improve the fairness between unicast and multicast flows, we propose two versions of *Unicast-Friendly Multicast*, aiming to limit the bandwidth share of a multicast flow equal to that of a unicast flow. The main reason for this unfairness is the high channel access probability of a multicast flow due to the absence of the collision detection mechanism and binary exponential backoff in the transmissions of multicast packets. Therefore, UFM makes a multicast flow to perform an one-stage backoff to make the rate of the channel access attempts of a multicast flow equal to that of a unicast flow. To balance the attempt rate of the channel access among unicast and multicast flows, UFM adapts the multicast backoff window (cw_m) to time-varying traffic-dependent network conditions. Because most multicast traffic is not throughput-sensitive but delay-sensitive, the collision detection and retransmissions are not needed. The outline of UFM is as follows. UFM periodically calculates cw_m for multicast packets depending on the number of competing stations. Then, it performs random backoff in the range of [0, cw_m-1]. cw_m will increase as more stations with unicast flows compete to guarantee the fair share of unicast flows. We assume that each station can estimate the number of other competing stations by some mechanisms such as extended Kalman Filter coupled with a change detection mechanism [10].

A. MFI: Multicast Fairness Index

We first introduce a new index, the multicast fairness index (MFI). Because it is difficult to clearly define the aggregate throughput of multicast flows, we instead measure the aggregate throughput of unicast flows affected by multicast flows. With all stations saturated, $MFI_{(m,n)}$ is defined by the ratio of the aggregate throughput of m unicast flows when there are m unicast flows and n multicast flows to the aggregate throughput of m unicast flows when there are m+n unicast flows. In other words, $MFI_{(m,n)}$ indicates the throughput loss of m unicast flows due to n multicast flows. If there are m unicast flows and n multicast flows under the saturation, $MFI_{(m,n)}$ can be given as follows.

$$
\begin{aligned}
MFI_{(M,N)} &= \frac{[A_2 - M_2]_{(M,N)}}{[A_1 - U_1]_{(M+N,0)}} \\
&= \frac{\left[(\sum_{j=1}^{M} u_j + \sum_{k=1}^{N} m_k) - \sum_{k=1}^{N} m_k \right]_{(M,N)}}{\left[\sum_{i=1}^{M+N} u_i - \frac{N}{M+N} \sum_{i=1}^{M+N} u_i \right]_{(M+N,0)}} \\
&= \frac{\left[\sum_{j=1}^{M} u_j \right]_{(M,N)}}{\left[\frac{M}{M+N} \sum_{i=1}^{M+N} u_i \right]_{(M+N,0)}}
\end{aligned}
$$

In an above equation, A_1 and U_1 are the aggregate throughputs of m+n unicast flows and n unicast flows in the saturated network, respectively. A_2 is the aggregate throughput of M unicast flows and N multicast flows while M_2 is the aggregate throughput of N multicast flows in the saturated network. u_i, m_k, and a tuple (a, b) represent the bandwidth occupied by the i^{th} unicast flow, the bandwidth occupied by the k^{th} multicast flow, and a pair of the number of unicast flows, the number of multicast flows. $MFI_{(m,n)}$ becomes 1 if n multicast flows use the same bandwidth as n unicast flows. That is, the closer to 1 the MFI is, the better fairness is achieved. In a wireless LAN, only an access point is able to transmit multicast packets, so n is set to one.

B. UFMv1: Random Collision Inference

In the first version of UFM (UFMv1), each station transmitting multicast packets infers the packet collision probability, based on the estimate of the number of other competing stations. With the assumption that the average number of packet collisions until successful transmission is known a priori, the station with a multicast flow virtually performs binary exponential backoff like unicast flows until it reaches the inferred average backoff stage (or average number of packet collisions) of other stations, called *virtual backoff*. At each backoff stage, it picks up a random number from its contention window corresponding to that of a unicast flow. Finally, it has cw_m equal to the recursive sum of all the selected backoff times, each of which is multiplied by its collision probability. Then it performs a random backoff in the range of [0, cw_m-1]. cw_m is expressed by:

$$
\begin{gathered}
cw_m = cw_{m,k} \\
cw_{m,0} = rand[0, cw_{u,0} - 1] \\
cw_{m,i+1} = (1 - p_{i,n}) \cdot cw_{m,i} + p_{i,n} \cdot rand[0, cw_{u,i+1} - 1] \\
cw_{u,i} = 2 \cdot cw_{u,i-1} = 2^i \cdot cw_{u,0} \\
p_{i,n} = 1 - (1 - \frac{2}{cw_{u,i}+1})^{n-1}
\end{gathered}
$$

where i, k and n represent the backoff stage, the given average backoff stage and the estimate of competing stations. The current multicast contention window cw_m is an approximation to the multicast contention window $cw_{m,k}$ of the average (k^{th}) backoff stage. Given the average backoff stage, UFMv1 can derive the collision probability $p_{i,n}$ with the estimate n of the number of competing stations and the contention window $cw_{u,i}$ of the station with unicast flows at the i^{th} backoff stage. With the collision probability $p_{i,n}$, the multicast contention window $cw_{m,k}$ of the average backoff stage is calculated iteratively. After calculating cw_m, the station with multicast packets performs a random backoff in the range of [0, cw_m-1].

After empirical study by simulations we observed that the second or third backoff stage is appropriate for the average backoff stage of a multicast flow as the number of nodes increases.

C. UFMv2: Analytical Collision Inference

The second version of UFM (UFMv2) takes a different approach. To guarantee the same transmission attempt rates of a unicast flow and a multicast flow, UFMv2 chooses the multicast contention window cw_m for an one-stage random backoff from the related analysis. According to the performance analysis of the IEEE 802.11 DCF in [11], a station has an attempt rate τ as follows:

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

where p, m and w represent the collision probability, the maximum backoff stage and the minimum contention window size. The collision probability p can be derived from the number of competing stations. In IEEE 802.11b, the transmission attempt rates of a unicast flow τ_u and a multicast flow τ_m are given by:

$$\tau_u = \frac{2\cdot(1-2p)}{(1-2p)\cdot(32+1)+p\cdot 32\cdot(1-(2p)^5)}$$
$$\tau_m = \frac{2}{cw_m+1}$$

Because a multicast flow has no binary exponential backoff, m is zero, resulting in an attempt rate independent of p. To make $\tau_u = \tau_m$, some samples of the solution table of the multicast contention window cw_m is given in Table I.

TABLE I
MULTICAST CONTENTION WINDOW cw_m

# of STAs	cw_m	# of STAs	cw_m
10	55	50	131
20	77	60	146
30	96	70	161
40	114	80	175

All stations maintain this Multicast Contention Window Table (MCWT). Therefore, a station to transmit a multicast packet first determines the multicast contention window cw_m from its MCWT. The one-stage random backoff in the range of [0, cw_m-1] prevents multicast flows from overwhelming unicast flows.

D. Deployment Issues

In general, *multicast* is used for multimedia services, so it may require adequate QoS parameters such as throughput and delay. If all incoming unicast/multicast calls are accepted, each flow in the service suffers from the severe performance degradation. To resolve this problem, admission control techniques should be introduced in wireless LANs. When the number of unicast/multicast flows in the service goes across the predefined threshold of admission control techniques, new flows should be rejected or existing flows are dropped depending on a priority. To alleviate the performance degradation of multicast flows due to additional unicast flows, the maximum multicast contention window can be given. This prevents multicast flows from experiencing too little share of network resources due to many unicast flows.

Another possible optimization of UFM is to utilize multirate multicast capability. Currently, only static single-rate multicast is supported in the wireless LANs, and the rate is normally set to one of the basic rates. In other words, the lowest rate is used for multicast to reach as many stations as possible. However, this configuration causes the problem similar to performance anomaly in the [12]. The aggregate throughput of unicast flows with high transmission rates is degraded to the low aggregate throughput of multicast flows tuned to the lowest transmission rate. If we can know which stations subscribes to the multicast group through IGMP snooping techniques, multicast packets can be transmitted by a higher rate, so improving network performance.

V. PERFORMANCE EVALUATION

We evaluate the proposed schemes, UFMv1 and UFMv2 by using the NS-2 simulator. The network topology in this simulation consists of n stations and one access point in wireless part and n fixed nodes in wired part with varying n. Each station sends unicast packets to the corresponding fixed node via the access point and an access point sends multicast packets to all stations in the wireless network. n stations and one access point compete to transmit data packets, making the wireless network saturated. We use the CBR traffic model with the packet size of 500(byte) for unicast flows and multicast flows. To eliminate other effects but backoff, unicast and multicast packets are transmitted by the same transmission rate equal to 2(Mbps). Total simulation time is 300(s) and we compare two metrics between 802.11 DCF, UFMv1 and UFMv2: aggregate throughput and multicast fairness index.

As shown in Fig. 2 in IEEE 802.11 DCF, compared to N unicast flows, a multicast flow grabs the large share of network resources in the existence of N-1 unicast flows and 1 multicast flow, hence degrading aggregate unicast throughput substantially. This situation becomes worse due to a long EIFS value in the cast of an erroneous multicast flow. Fig. 3(a) shows aggregate unicast throughput achieved by adopting UFMv1 with the expected number of packet collisions, k, being 2. By assuming the adequate number of packet collisions in advance, UFMv1 guarantees that a multicast flow gets the share of network resources equal to a unicast flow. Fig. 3(b) shows performance improvement achieved by UFMv2 similar to UFMv1. By adapting the multicast contention window depending on the number of competing stations, both UFMv1 and UFMv2 are able to provide the fair share of network resources to a multicast flow regardless of a normal multicast flow or an erroneous multicast flow.

Fig. 4 shows the Multicast Fairness Index (MFI) depending on the number of competing stations including a multicast flow. In the case of a normal multicast flow, MFI is given in Fig. 4(a). As more stations contend for wireless channel access, MFI in the 802.11 DCF decreases to 0.9 while MFI in both UFMv1 and UFMv2 is sustained around 1. Fig. 4(b) shows the case of an erroneous multicast flow, in which we can acquire more fairness improvement. When sufficient stations compete, MFI in the 802.11 DCF is sustained around 0.7, meaning the unfairness between unicast flows and multicast flows. However, both UFMv1 and UFMv2 improve MFI to almost 1, that is, preventing a multicast flow from overwhelming the wireless network. UFM can effectively guarantee the fair share of network resources to a multicast flow.

VI. CONCLUSION

In this paper, we first pointed out the unfairness problem between unicast flows and multicast flows in the IEEE 802.11 DCF. To quantify the fairness, we introduced Multicast

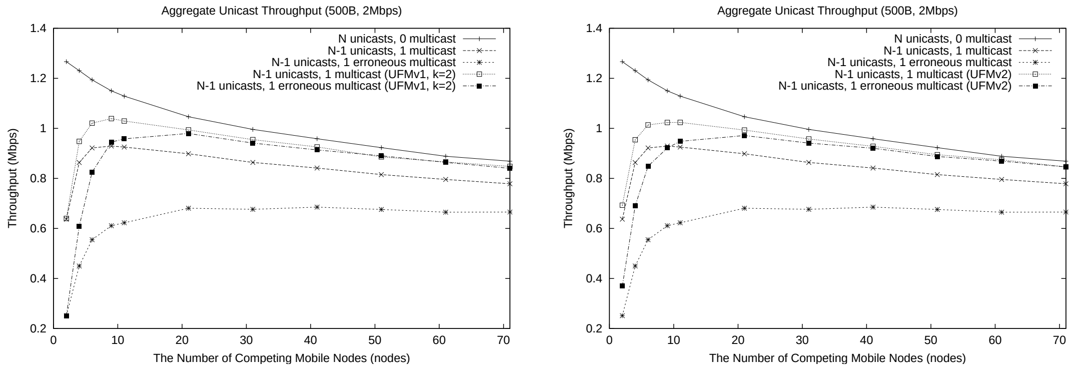

(a) UFMv1 (b) UFMv2

Fig. 3. Aggregate Unicast Throughput

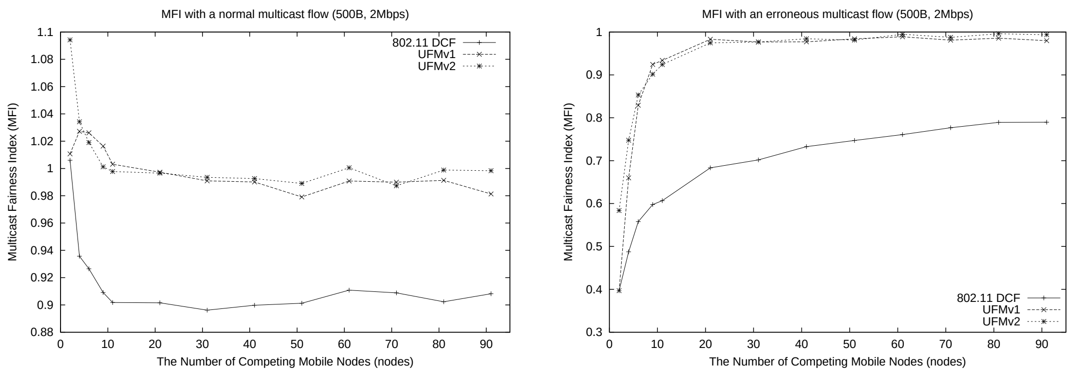

(a) $\mathrm{MFI}_{(m,1)}$ on a normal multicast flow (b) $\mathrm{MFI}_{(m,1)}$ on an erroneous multicast flow

Fig. 4. Multicast Fairness Index (MFI)

Fairness Index (MFI) to evaluate how fairly multicast and unicast flows are competing. Then, we proposed Unicast-Friendly Multicast (UFM), which provides the fair share of bandwidth to each flow regardless of unicast or multicast. The central idea of UFM is to make a multicast packet perform one-stage random backoff, whose contention window size is dynamically adjusted depending on the number of competing stations. We presented two versions of UFM: the first one calculates the multicast window size by inferring the average contention window size of unicast flows, while the second one maintains the mapping table between the number of competing stations and the optimal multicast window size given by related analysis. Simulation reveals that MFIs of the both versions of UFM are almost one, which means unicast and multicast flows obtain the equal share of bandwidth under the saturated network conditions. Furthermore, simulation verifies that the bad effect of erroneous multicast packets is also solved by UFM.

VII. Acknowledgment

This work was supported in part by the Brain Korea 21 project of the Ministry of Education, 2005, Korea.

References

[1] Handley, M.; Floyd, S.; Padhye, J.; Widmer, J.“TCP Friendly Rate Control (TFRC): Protocol Specification,” Request for Comments (RFC) 3448, The Internet Society January 2003.
[2] Arnaud Legout, Jorg Nonnenmarcher and Ernst W. Biersack, “Bandwidth Allocation Policies for Unicast and Multicast Flows,” IEEE INFOCOM March 1999.
[3] Jorg Widmer and Mark Handley, “TCP-Friendly Multicast Congestion Control (TFMCC): Protocol Specification,” Internet Draft draft-ietf-rmt-bb-tfmcc-04.txt, October 2004.
[4] Richard Akester, “Reducing Multicast Collision Loss for Digital TV Over 802.11 Wireless Networks,” ICOMIV September 2004.
[5] Joy Kuri and Sneha Kumar Kasera, “Reliable Multicast in Multi-access Wireless LANs,” CEDT January 1999.
[6] S. K. S. Gupta, V. Shankar and S. Lalwani, “Reliable Multicast MAC Protocol for Wireless LANs,” IEEE ICC May 2003.
[7] Igor Kozintsev and Jeff MC Veigh, “Improving last-hop multicast streaming video over 802.11,” BroadWIM October 2004.
[8] T. Nilsson, G. Wikstrand and J. Eriksson, “Early Multicast Collision Detection in CSMA/CA Networks,” IEEE MWCN September 2002.
[9] Larzon, LA., Degermark, M., Pink, S., Jonsson, LE., and G. Fairhurst,“The Lightweight User Datagram Protocol” Request for Comments (RFC) 3828 July 2004.
[10] G. Bianchi and Ilenia Tinnirello, “Kalman Filter Estimation of the Number of Competing Terminals in an IEEE 802.11 network,” IEEE INFOCOM April 2003.
[11] G. Bianchi, “Performance analysis of the IEEE 802.11 distributed coordination function,” JSAC vol. 18, pp. 535-547, March 2000.
[12] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance anomaly of 802.11b,” IEEE INFOCOM April 2003.