

# Multimedia multicast performance analysis for 802.11n network cards

Nicolas Letor, Wim Torfs, Chris Blondia

University of Antwerp - IBBT,

Middelheimlaan 1, B-2020 Antwerp, Belgium

Email: {nicolas.letor, wim.torfs, chris.blondia}@ua.ac.be

**Abstract**—The technical evolution of Wi-Fi networks has made it possible to attain very high speeds for delivering content to users. However multicast transmissions never received the same level of attention as unicast transmissions in the IEEE802.11n standard. In this paper, we experimentally assess the viability of employing commercially available IEEE802.11n hardware to multicast real-time multimedia content, requiring high bandwidth and low delay. We verify the performance of the leader based multicast scheme and compare it to the standard multicast scheme.

## I. INTRODUCTION

Wi-Fi networks have become increasingly popular in multimedia applications, which can be attributed to their low cost, constantly increasing bandwidth and relative ease of deployment. The IEEE802.11n standard enables multimedia streaming thanks to its higher bandwidth and is now being deployed in applications where previously proprietary radio systems were required. To distribute multimedia streams in real-time to multiple users over a wireless link, either the streams can be duplicated and unicast to each individual participant, or the stream can be multicast to all the participants. The former lacks efficiency due to the duplication of the streams and poses a scalability issue.

IEEE802.11 provides a MAC multicast mechanism, however, it is a well known fact that it has some issues [1]. First, the standard does not provide means to acknowledge received multicast packets. Therefore, the sender is unable to adapt to changing channel conditions by modifying the contention window. Due to the lack of acknowledgements, there are no retransmissions. A lot of research focuses on this topic, for example, [2] consider the lack of RTS/CTS message exchange, for IEEE802.11 multicast, as an obstacle to provide reliable transmissions. To resolve this issue, both proposed an extension for the IEEE802.11 MAC, where RTS frames were modified to contain extra information, while CTS frames are sent in a controlled manner. These protocols provide a significant contribution, however they cannot be verified with commercially available hardware as the RTS and CTS frames and transmission mechanism are implemented in hardware.

The lack of acknowledgements forces the transmission rate of multicast packets to be fixed to the lowest possible physical rate, therefore, limiting multicast transmissions to a lower bandwidth. The IEEE802.11n standard does provide enhancements that allow higher bandwidths, however, these optimizations are specifically designed for unicast transmissions. Some

research, such as [3], avoid the limitations of the available hardware and search for a functional multicast transmission scheme. The scheme implements a leader based method by sending unicast packets to a single station and having all other stations listen in promiscuous mode. We based our experimentation setup on the leader based implementation from this paper.

The goal of this research is to verify whether it is possible to use a commercially available IEEE802.11n wireless network card to multicast multimedia traffic, while considering requirements such as throughput and latency. Note that, the paper neither provides any performance analysis of multicast protocols, nor does it propose a new multicast protocol. The discussion in Section II provides details regarding our methodology, results can be found in Section III and we conclude in Section IV.

## II. METHODOLOGY AND TEST SETUP

The use case on which we base our test setup is a single AP, which sends the multimedia content to multiple users who are located up to 30 meters from the AP. According to [4], the maximum modulation and coding scheme (MCS) that can be used in such an environment is MCS4 (single stream), or MCS12 (dual stream MIMO).

The key requirement for this use case is a guaranteed and stable latency. Due to several factors, the 802.11n standard does not provide reproducible results. First, the automatic rate selection algorithm selects the most optimal rate for each packet depending on previous transmission reports. Second, if a packet is not delivered successfully, the sender tries to retransmit the packet up to seven times. This sequence can be repeated for a number of rates, each lower than the previous in case of recurrent failure. Such behavior does not allow a precise quantization of the resulting throughput, latency and number of missed packets based on the specified parameter settings. Therefore, an open source kernel, such as the Linux kernel, which allows the modification of the network drivers at the AP side, is essential for multimedia scenarios and the collection of reproducible measurement results.

The part of the Linux kernel that is relevant to our research is the IEEE802.11 MAC layer. To exclude some exotic behavior of a certain network card, we investigate two different IEEE802.11n chipsets, Atheros AR9220 (ath9k driver) and Ralink RT2860 (rt2800pci driver). These drivers use the

provided rate selection information of the MAC layer to configure the hardware on a per packet base. Our modification consists of replacing the values of the automatic rate selection algorithm with a configurable fixed MCS rate value, setting the transmission count to one and removing any fallback rates from the tx descriptor. The MAC layer module is modified such that traffic to non-associated stations is forwarded. The resulting transmission behavior is such that packets are sent only once at a fixed transmission rate, while both retries and the automatic rate fallback scheme are disabled.

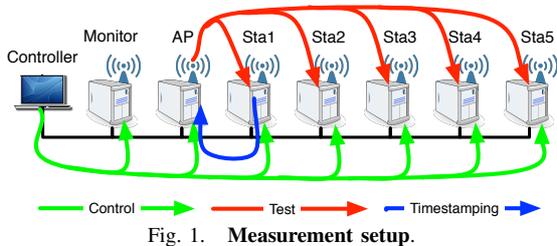


Fig. 1. Measurement setup.

The resulting measurement setup is depicted in Figure 1. A single Device Under Test (DUT) is operating in AP mode, while the others are operating in station mode. The AP sends packets over the wireless link to the stations, which are relayed back to the AP over a wired link by a single station in order to determine the Round Trip Time (RTT). The latency over the wired link is at least hundreds of orders smaller than the latency of the wireless link and is therefore negligible in the calculation of the latency. Packet arrival statistics are collected by all stations, of which afterwards the average throughput is calculated. We defined three test cases. In the first test case, the AP sends data in broadcast packets. In the second case, a single station is associated with the AP and unicast packets are sent from the AP towards this associated station, while all other stations listen in promiscuous mode. This is called a leader based multicast scheme [3]. In the last test case, no station is associated with the AP and the AP sends unicast packets to a fictional station, while all stations are in promiscuous mode. This scheme is a variation of the previous leader based multicast scheme, preventing the use of the IEEE802.11n aggregation improvements.

TABLE I  
PERFORMANCE TEST PARAMETERS

Parameter	Values
MCS rate	4, 12
Channel	6, 44
Channel width	20Mhz, 40Mhz
Iperf tx rate	8, 12, 16, 20, 24, 28, 32Mbps
UDP packet size	500, 1000, 1500 bytes

The performance measurements were collected using iperf by a sending bandwidth limited UDP stream during 120 seconds. The parameter settings that can be found in Table I. The minimal throughput requirement equals 32Mbps and the maximum latency requirement equals 20ms.

The tests were performed in a low interference environment, the 5GHz band, as well as an environment with heavy background traffic, the 2.4GHz band. All tests are performed with several packet sizes, which can be regarded as packet aggregation at application level. In order to mitigate any circumstantial effects, the tests were executed five times.

### III. MEASUREMENT RESULTS

The IEEE 802.11n standard mandates the use of a basic rate for transmission of broadcast and multicast frames. Typically, a Wi-Fi access point transmits these frames at the lowest possible modulation scheme in order to guarantee reception on all stations. The driver modifications enable us to lock the transmission rate for broadcast data to MCS4 (single stream) and MCS12 (double stream MIMO).

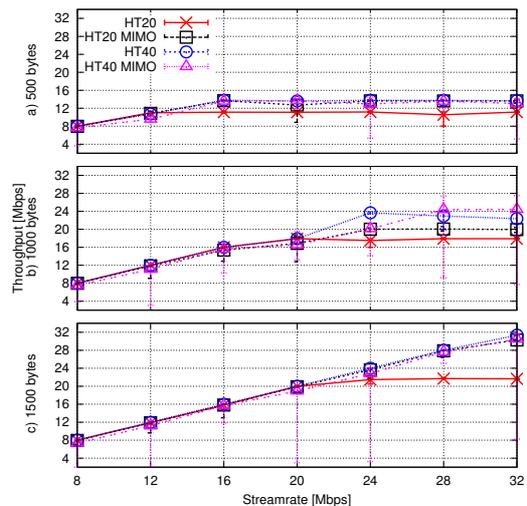


Fig. 2. Atheros Linux 3.2.0 5Ghz Broadcast throughput performance

Figure 2 shows the throughput performance for the broadcasting scheme of the Atheros chipset. The x-axis depicts the stream rate, that is, the rate at which the packets are sent, while the y-axis depicts the measured throughput. We observe that for small packets (i.e. 500 bytes) the throughput reaches a maximum at 11Mbps for the single stream HT20 case. With MIMO and/or HT40 enabled, a maximum of around 14Mbps can be obtained. Increasing the packet size has a beneficial impact on the throughput performance, lifting the capacity to 17.9Mbps using 1000 byte packets and 21.6Mbps using 1500 byte packets for HT20, single stream. By employing a HT20 dual stream (MIMO) provides us with a small improvement with packets of 1000 byte packets and enables almost the full bandwidth up to 32Mbps with 1500 byte packets. Using HT40 either single stream or dual stream, is even more interesting, where the measurements show a maximum throughput of 24Mbps for 1000 byte packets and the full bandwidth by using 1500 byte packets. Using larger packet sizes diminishes the overhead induced by accessing the wireless medium, which translates in a higher throughput efficiency.

The average latency, depicted in Figure 3, climbs to a maximum of 90ms at a streamrate of 12Mbps for HT20 and at a streamrate of 16Mbps for all other cases when sending

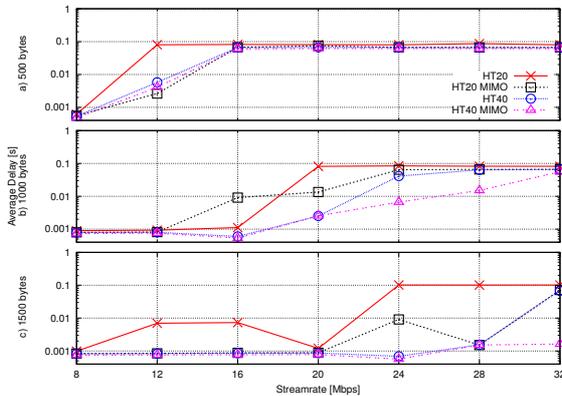


Fig. 3. Atheros Linux 3.2.0 5Ghz Broadcast average latency

500 byte packets. This also indicates that the network card is having issues accessing the medium for so many packets in such a short time. Larger packet sizes, MIMO or the use of HT40 have a beneficial impact on the latency, however, the latency crosses the 20ms boundary at streamrates higher than 28Mbps and 32Mbps for 1000 bytes and 1500 bytes packets respectively. We observe that the maximum latency is not guaranteed to be below 20ms, even for the combination of HT40 and MIMO, which was the most promising case.

The throughput results of the same tests performed with a Ralink card, indicate that it is able to sustain a higher stream rate for broadcasting than the Atheros card at the expense of packet loss, i.e. the setup is becoming saturated. By observing the latency, we notice a similar behavior as with the Atheros card, the higher the packet size, the higher the streamrate can be before the latency starts to climb to its maximum, which is close to 100ms.

When performing these broadcasting tests in the 2.4Ghz band, we noticed that the throughput is somewhat lower as in the 5Ghz band, but with a similar behavior. The maximum latency, however, drifts around 100ms continuously due to interference from neighboring networks, as the 2.4Ghz band is rather overcrowded.

The second test scenario evaluates the unicast performance of the IEEE802.11n compliant hardware for the throughput and latency. The test setup consists of a single AP that sends unicast data towards one of its stations, while all other stations listen in promiscuous mode. Since the packet transmission is a regular unicast link, it benefits from all optimizations in the IEEE802.11n standard and acts as a leader based multicast system. Automatic retransmissions and rate control are disabled for reproducibility.

As expected, the results indicate that the throughput is equal or better in comparison with the throughput of the broadcast. Due to the A-MPDU aggregation, the hardware has no problems handling the transmissions of small packets. The average latency of the packets is around 2ms when transmitting packets of only 500 bytes, except for the case of HT20. Note that there also is no noticeable stagnation of the throughput.

Although the Ralink card performs better than Atheros concerning broadcast transmissions, it shows a similar perfor-

mance for the unicast transmissions. The overall performance of the test setup is quite satisfactory, however there still is an issue with the maximum latency which can go up to 100ms. The results of both devices in the 2.4Ghz band can be regarded as similar to the results in the 5Ghz band. The throughput reaches a little sooner its maximum, which is lower than in an interference free environment and the latency also is a little higher. However, the differences are not so substantial as in the broadcast scenario.

In a final test scenario we investigated the influence of aggregation. This scenario consisted of sending the multicast stream to a non existent leader station, preventing the use of the 802.11n aggregation mechanisms. The collected data exhibited similar behavior to the measurements in the broadcast scenario (Figure 2), indicating that packet aggregation is of great importance for performance.

#### IV. CONCLUSION

In this paper we verified the performance of three methods for multicasting real-time multimedia content to Wi-Fi clients when using commercially available hardware. We tested each solution on two types of Wi-Fi cards. By simple modifications in the drivers, higher rates can easily be achieved for broadcast. However, the experienced throughput was significantly lower than the expected throughput, which manifested itself in the results of the latency. The latency was too high when using a simple broadcast scheme.

The unicast leader based multicast scheme, without retransmissions nor adaptive rate selection, performed better, since it benefited from the unicast improvements of the IEEE802.11n standard, such as aggregation, which is not supported while sending broadcast packets or unicast packets towards non-associated stations. This explains the lower experienced throughput while using the broadcast or non-associated unicast scheme. However, the leader based unicast is insufficient to support multicast multimedia in real-time, since there is no guaranteed maximum latency.

#### ACKNOWLEDGMENT

This research is partly funded by the Flemish Interdisciplinary Institute for BroadBand Technology (IBBT) through the OMUS project.

#### REFERENCES

- [1] R. Chandra, S. Karanth, T. Moscibroda, V. Navda, J. Padhye, R. Ramjee, and L. Ravindranath, "DirCast: A Practical and Efficient Wi-Fi Multicast System," in *Proceedings of 17th IEEE International Conference on Network Protocols, Princeton, NJ, USA*, October 2009.
- [2] W. Chaohui, L. Yuhong, H. Li, and M. Jian, "A new reliable multicast scheme for multimedia applications in wireless environment," in *Proceedings of 2nd International Conference on Broadband Network and Multimedia Technology, Beijing, China*, October 2009.
- [3] D. Duvojne and T. Turletti, "Multicast in 802.11 WLANs: An Experimental Study," in *Proceedings of the 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems, Torremolinos, Malaga, Spain*, October 2006.
- [4] W. Joseph, E. Tanghe, D. Plets, and L. Martens, "Prediction of range, power consumption and throughput for IEEE 802.11n in large conference rooms," in *Proceedings of the 5th European Conference on Antennas and Propagation (EUCAP), Rome, Italy*, April 2011.