This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/LCOMM.2015.2456911

Performance Increase for Highly-Loaded RoF Access Networks

A.Panagiotakis, P. Nicopolitidis, G.I. Papadimitriou and P.G Sarigiannidis

***Abstract*—This letter presents a Medium Access Control (MAC) mechanism that allows for the effective use of the bandwidth of highly-loaded 60 GHz Radio-over-Fiber (RoF) networks. We extend a previous approach proposed in the literature which is shown to have problems in highly-loaded cases. We propose a protocol that makes use of a memory buffer at the Central Office (CO) at the providers side that allows for the reduction of the number of polling packets required to identify all the active wireless nodes that need to transmit data. This greatly increases the overall throughput of the network at a maximum percentage of 42% for the examined configurations. The high throughput provided by our proposal is achieved without the need to update current infrastructure and only needs a trivial amount of memory at the providers side.**

***Index Terms*—Radio over Fiber, MAC protocols, Wireless networking**

I. Introduction

Radio-over-Fiber (RoF) networks [1] provide increased bandwidth in comparison to conventional wireless networks by making use of the 60 GHz range that can allow for bitrates beyond 100 Mbps per wireless node. They allow for the bridging of wireless and optical communications by means of modulating the optical signal by a radio one, allowing for the transmission and reception of all radio signals by a single antenna. This has the effect of minimizing the cost of deploying and maintaining Remote Antenna Units (RAU) because only the minimum required logic is implemented at the RAU for the successful communication between wireless nodes in the range of the RAUs that are controlled by the same Central Office (CO) at the providers side. Thus the MAC logic must be naturally implemented in the CO.

Early works in the field of MAC protocols in RoF implementations ([2]- [4]) used two separate protocols to allocate bandwidth in the optical and wireless links, so as to grant it to the wireless nodes, with IEEE 802.11 MAC being commonly used in the wireless part. A recent proposal, the MT-MAC Protocol [5], seamlessly allocates bandwidth in both optical and wireless media via a two-stage contention scheme and can achieve satisfactory performance. However we will show that its main drawback is waste of bandwidth in high-load environments due to the design nature of the second stage of its contention scheme.

In this letter we propose a method to achieve increased performance in such environments by altering the way wireless nodes are identified by the RAU. Simulation results show that the performance of our approach significantly surpasses that of the MT-MAC protocol. This is achieved without the need to update current infrastructure and only needs a trivial amount of memory at the providers side in the CO.

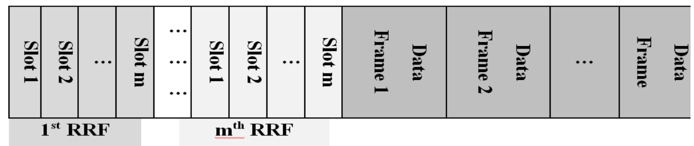

Fig. 1: The MT-MAC superframe

II. The MT-MAC Protocol

In the MT-MAC protocol the communication of the CO and the wireless nodes that are under the coverage of a RAU is done over a pair of wavelengths, carrying downlink and uplink traffic respectively. We will refer to these wavelengths as channels in the rest of the letter. The network can employ a fixed number of channels for carrying data traffic and an additional channel used for the allocation of the available channels, called the control channel, that is shared by all RAUs.

MT-MAC uses two parallel contention periods to allocate the capacity of the network. During the first contention period, which takes part in the optical part of the RoF network, a short optical beacon pulse is emitted by the CO via the control channel followed by a number of slots equal to the number of RAUs in the network. The pulse is propagated by the RAUs to the wireless nodes in their respective ranges. If any of these wireless nodes have outstanding packets waiting for transmission they respond by emitting a short pulse of the same duration as the pulse send by the CO, or stay silent otherwise. The RAU that receives any number of pulses propagates them in its respective time slot to the CO. After reception of the pulses from all the RAUs, the CO allocates the available channels to the RAUs. If the number of channels is less than the number of active RAUs the allocation is done in a Round Robin Fashion, ensuring the fair distribution of the available bandwidth between RAUs. The allocation of the channels is done by first emitting a pulse of length $\log_2 n + 1$ bits where n is the number of channels, followed by a series of time slots equal to the number of RAUs. In each time slot, the CO transmits the code of the channel assigned to the respective RAU, or 0 otherwise.

After wavelengths are assigned to the RAUs, the second contention period takes place in the wireless part of the RoF network and comprises superframes (SFs), as shown in Figure 1. Each SF contains a constant number of frames that are a mix of Resource Requesting Frames (RRF) and Data Frames (DF). First, a series of RRFs are transmitted by the CO. RRFs are used to identify the active wireless nodes in the range of a RAU. In each RRF there is a constant number of slots, m. In the beginning of a RRF all active wireless nodes choose a random number in $[0, m)$ corresponding to the slot in which they will try to identify themselves. In each slot of the RRF the CO sends a POLL packet. Any wireless node that has chosen the current slot to identify itself, responds with an ID packet when it receives the POLL packet. If only one wireless node has chosen the current slot, then the ID packet is correctly received by the CO and an ACK packet is send as a response. If more than one wireless nodes transmit during the current slot, then a collision occurs. This results in the reception of no packet to the CO and the CO does not transmit an ACK packet, forcing the wireless nodes that participated in the collision to try to identify themselves again at the next RRF. RRFs continue to be transmitted until no collisions occur. After the last RRF is transmitted, the second contention period ends.

After the two stages, the CO has knowledge of all the wireless nodes that have pending packets and transmits DFs for the rest of the frames inside the SF, polling each wireless node for its data in a Round Robin fashion. If there are no more wireless nodes to poll during the SF period or the SF has reach its maximum length while other RAUs await for transmission, the channel serving the RAU is assigned to another RAU and the current RAU must wait for another channel to be assigned to it. On the other hand, if there are no other RAUs having wireless nodes awaiting for transmission of data and there are still wireless nodes to be polled, the SF is extended by continuing transmission of DFs.

A. Throughput Analysis of MT-MAC

The maximum throughput can only be obtained when a channel is assigned for every RAU. By having a channel for each RAU we eliminate the need for the first contention period. This, in effect, means that the throughput of the protocols relies solely on the number of DFs available for each SF in contrast to the RRFs used per SF.

Since we assume that every RAU is assigned its own channel, the SFs will always be extended so that data from each of the active nodes is transmitted. This means that the number of DFs in a maximum load situation will be constant and equal to the number of wireless nodes in the range of a RAU. As a consequence to this, the maximum attainable throughput is the number of DFs per SF to the number of total frames in a SF, both RRFs and DFs.

To find the number of RRFs required for correctly identifying all active wireless nodes under a RAU for the MT-MAC protocol, for a typical RRF of 10 slots [5], we ran a set of simulations of a million SFs and computed average values for various numbers of active nodes under the RAU. Based on the obtained average number of RRFs required to identify all wireless nodes and the actual number of active wireless nodes, we calculated the maximum attainable normalized throughput for the MT-MAC protocol. This is shown in Figure 2. We see that normalized throughput reaches a maximum and then saturates, since the number of RRFs needed to identify every wireless node increases exponentially.

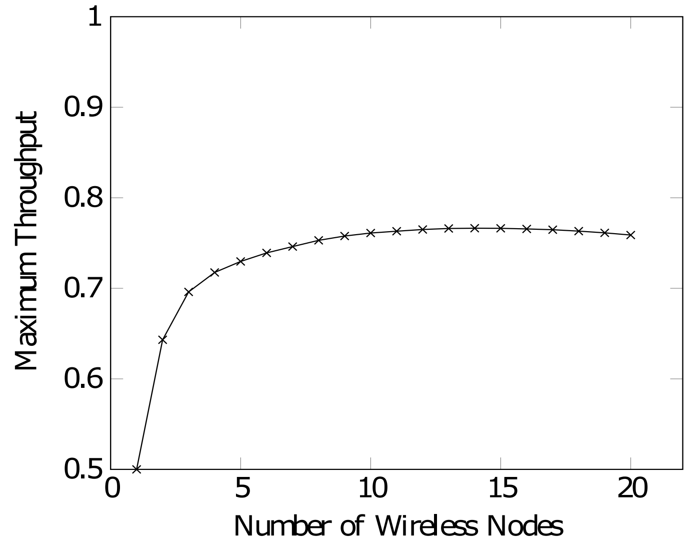

Fig. 2: Maximum normalized throughput versus the number of wireless nodes per RAU for MT-MAC.

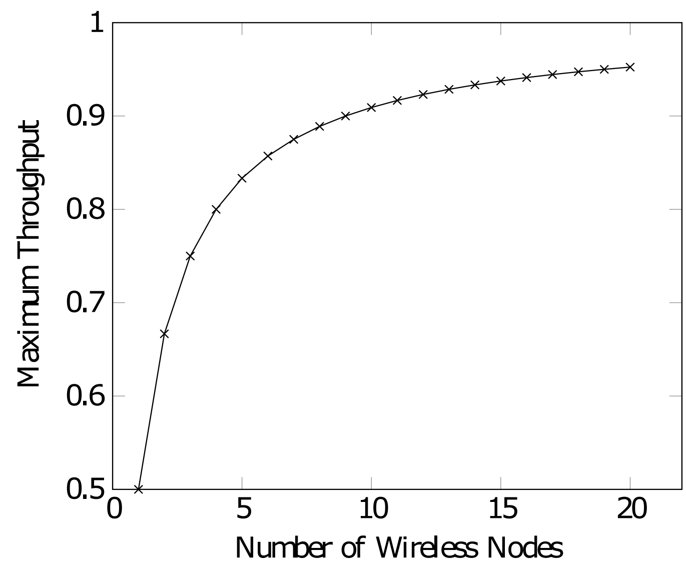

Fig. 3: Maximum normalized throughput versus the number of wireless nodes per RAU for the proposed approach.

III. The Proposed Approach

Our proposal uses a memory buffer in the CO to keep track of the wireless nodes that were active recently. The memory buffer values hold the number of SFs for which each wireless node will be considered active (including the current SF) for polling. When a wireless node successfully identifies itself during an RRF, its corresponding value in the memory buffer of the CO is set to a constant number L. The wireless node also stores its value in its own memory buffer so that it knows when it is no longer considered active for polling and needs to re-identify itself.

For the next L SFs, the wireless node does not have to identify itself during a RRF and will be polled by the CO regardless of having data to transmit or not. If the wireless node does not have a packet for transmission during a SF its

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/LCOMM.2015.2456911

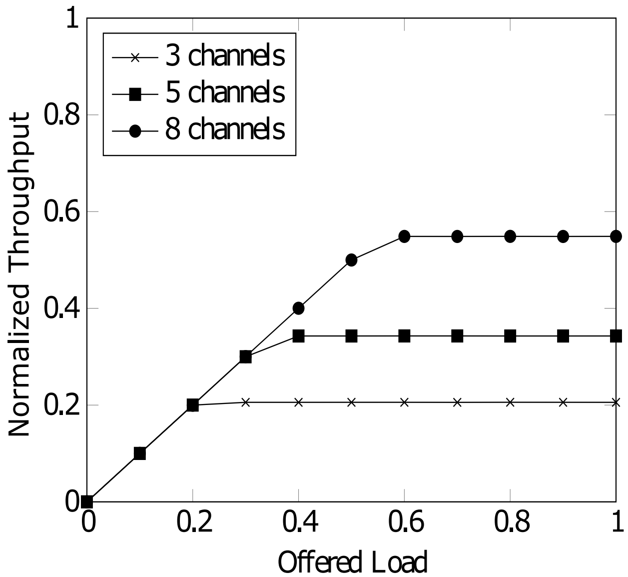

Fig. 4: Throughput versus Offered Load characteristic for MT-MAC.

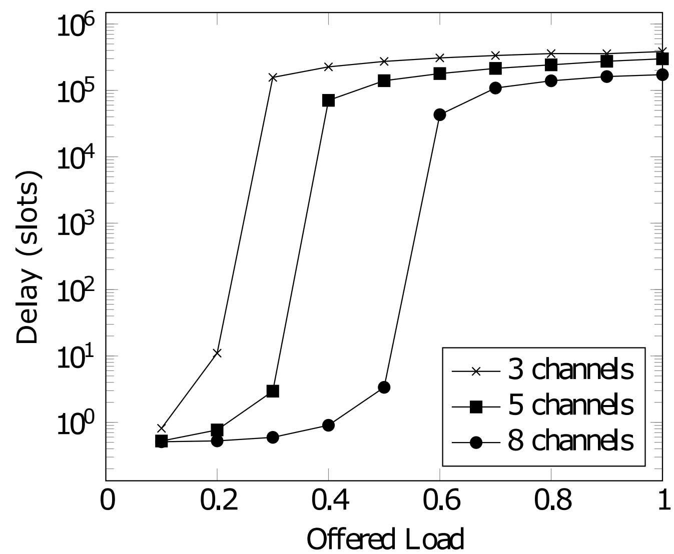

Fig. 6: Delay versus Offered Load characteristic for MT-MAC.

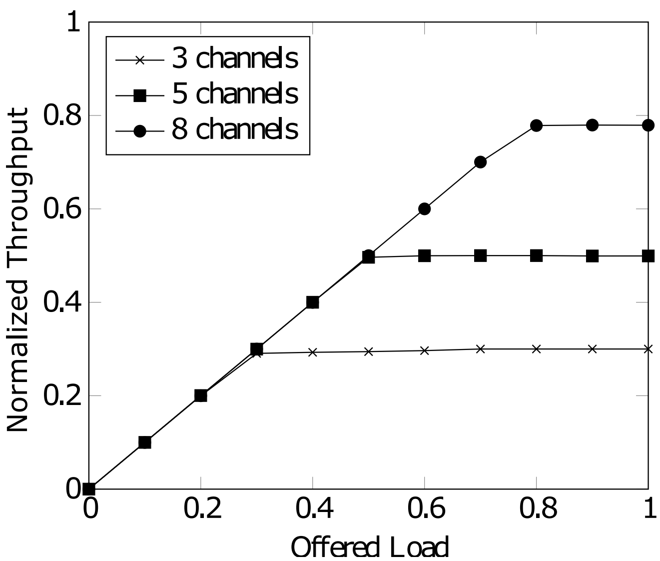

Fig. 5: Throughput versus Offered Load characteristic for the proposed approach.

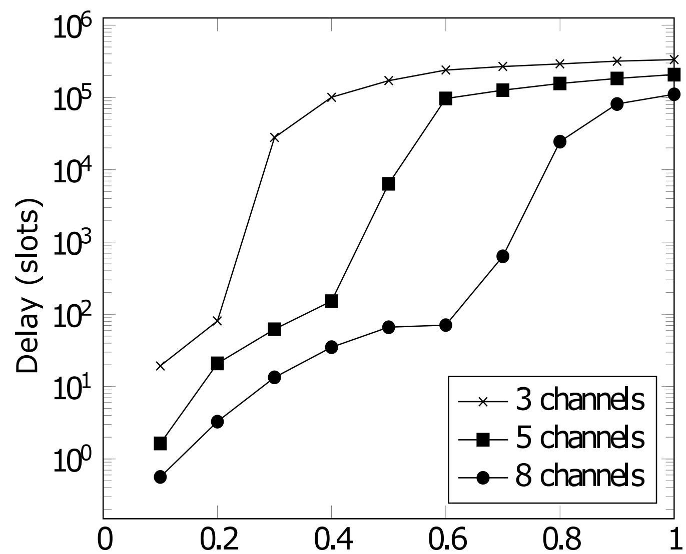

Fig. 7: Delay versus Offered Load characteristic for the proposed approach.

corresponding value in the memory buffer of the CO, as well as the value in its own memory buffer, is reduced by one. If this value reaches zero, the wireless node stops being considered active for polling and has to re-identify itself in a subsequent RRF to be considered for polling again. On the other hand, if a wireless node has a packet for transmission during a SF, the value in the memory buffer of the CO, as well as in its own memory buffer, is set again to L.

Our scheme allows wireless nodes that have many packets queued for transmission to not need to identify them-selves in each SF, reducing the number of RRFs required and increasing the available bandwidth for data transmission. This provides a far better utilization of the available bandwidth than that of the MT-MAC protocol, since the minimum needed RRFs are used in high load conditions.

A. Throughput Analysis of the proposed approach

For our proposed approach, in a highly-loaded network where all wireless nodes always have buffered packets, we do not need RRFs to identify active nodes after their first identification. That means that our approach can fully use the available bandwidth for data transmission and thus, achieve nearly 100% utilization of the available channels at high loads, as every wireless node will be registered in the memory buffer of the CO and we will only need one RRF per SF to check for any newly-active wireless nodes. That means that throughput grows constantly with the number of active wireless nodes as shown in Figure 3, something that shows the scalability potential of our proposed approach, contrary to the case of MT-MAC, which as can be seen from Figure 2, saturates for

This is the author's version of an article that has been published in this journal. Changes were made to this version by the publisher prior to publication.
The final version of record is available at http://dx.doi.org/10.1109/LCOMM.2015.2456911

more than 10 nodes per RAU. One can also see from Figures 2 and 3 that our approach significantly outperforms that of MT-MAC.

IV. Simulation Results

The simulation setup we used was that of [5]. Thus, we used a network of 10 RAUs in a bus topology. Each RAU covers 10 wireless nodes. We used 10 frames per SF, 10 slots per RRF and L=3. The fiber propagation delay for 200 m was set to 1 sec, the wireless propagation delay to 0.16 sec and the data bit rate to 155 Mbps. We did not simulate wireless channel errors and the reason for doing so is that we wanted to assess the relative performance increase of the proposed approach compared to [5] due to the protocol difference. The simulation also accounts for the contention in the optical part of the network. We run the experiments for 3, 5 and 8 available channels at the optical part. Our results are obtained for load values ranging from 10% up to 100% of the maximum theoretical network capacity. Throughput is measured as the packets delivered per time slot, where a time slot is the time required for a data packet to be transmitted. Delay is measured in time slots, as the average delay between the arrival of a packet in a wireless nodes buffer and its successful transmission. The simulations were run for a million time slots.

Figures 4 and 5 show the normalized throughput versus normalized offered load characteristics of the MT-MAC and our approach respectively, for a different number of channels. It clearly shows the superiority of our proposed protocol versus the MT-MAC protocol, with performance gains up to 42%. Specifically, for 3 channels MT-MAC throughout converges to 0.2 while that of the proposed approach converges to 0.3. For 5 channels the throughput of MT-MAC converges to 0.35 while that of the proposed approach converges to 0.5. Finally, for 8 channels MT-MAC throughout converges to 0.55 while that of the proposed approach converges to 0.8. Throughput decreases for a decreasing number of channels for both protocols, as when fewer channels need to be round-robin distributed to the 10 RAUs, access delay for the wireless nodes is increased.

Figures 6 and 7 show that, for small values of the offered load, our proposed protocol produces a larger average value of delay, something that is explained by the assumption made by the CO that there exists in a RAU an active wireless node that has packets to transmit because of its previous successful attempt to transmit data in a previous SF. This does not hold for low offered loads, as there is only a small probability that a wireless node will have many packets queued for transmission and, thus, will also have a packet for transmission at a subsequent SF after its successful transmission. On the contrary, when the offered load increases to mid and high levels and the probability of falsely identifying a wireless node as active decreases, our proposed protocols delay values increase at a slower rate than the MT-MAC protocol, producing smaller values of delay through the rest of the offered load range. This is of course due to the fact that our mechanism is not burdened with reidentifying the active wireless nodes in each SF, as it carried his knowledge between successive SFs.

Finally, worthy of note is that the use of small value of L=3, for which a wireless node will be considered active for polling, suffices for a satisfactory performance. Thus, small L values, allow for insignificant needs of memory for the implementation at the CO, since a CO with ten thousand nodes would only need about 2.5 kilobytes of storage space. On the contrary, increasing L to large values will increase storage space requirements and not bring performance improvements as the CO will continue polling for more time wireless nodes that are obviously not active any more.

V. Conclusion

This letter proposes a MAC protocol for RoF networks that is able to provide increased performance when the network load is high. Its throughput exceeds that of a previously proposed protocol of the same kind up to 42% and also yields lower delays for high network loads. The implementation overhead requires insignificant amounts of memory at the CO.

References

[1] Sarigiannidis, A., Ilioridou, M., Nicopolitidis, P., Papadimitriou, G., Pavlidou, F-N. Sarigiannidis, P., Louta, M., Vitsas, V.: "Architectures and Bandwidth Allocation Schemes for Hybrid Wireless-Optical Networks", IEEE Communications Surveys and Tutorials, in press, 2015.

[2] Crisp, M. J., Li, S. , Watts, A., Penty, R. V., White, I. H.: "Uplink and downlink coverage improvements of 802.11g signals using a distributed antenna network", IEEE Journal of Lightwave Technology, 2007, 25, (11), pp. 33883395.

[3] Das, A., Mjeku, M., Nkansah, A,. Gomes, N. J.: "Effects on IEEE 802.11 MAC throughput in wireless LAN over fiber systems", IEEE Journal of Lightwave Technology, 2007, 25, (11), pp. 33213328.

[4] Mjeku, M., Gomes, N.J.: "Analysis of the Request to Send/Clear to Send Exchange in WLAN Over Fiber Networks", IEEE Journal of Lightwave Technology, 2008, 26, (15), 2531 2539.

[5] Kalfas G., Pleros, N.: "An Agile and Medium-Transparent MAC Protocol for 60 GHz Radio-Over-Fiber Local Access Networks", IEEE Journal of Lightwave Technology, 2010, 28, (16), pp. 2315-2326.