

Coexistence of IEEE 802.11n and Licensed-Assisted Access devices using Listen-before-Talk techniques

Claudio Casetti,
Politecnico di Torino, Italy

Abstract—The proliferation of smartphones and related high-bandwidth services such as video streaming or cloud services are leading to increasing traffic demands, hence to the need of more efficient mobile networks. As current 4G networks are nearing their capacity, there is growing interest for the inclusion of new spectrum bands for LTE networks, among which unlicensed spectrum features prominently. Coexistence issues are all but certain, especially in the 5 GHz unlicensed band, where incumbent, ubiquitous WiFi devices are likely to be affected. Many coexistence solutions are currently being investigated, at both physical and MAC layer. In this work, we explore the impact of a MAC-layer solution that is supposed to allow graceful co-channel coexistence, finding that there are indeed some cases where, depending on the choice of parameters and configuration, the performance of WiFi traffic can be seriously degraded.

I. INTRODUCTION

Fourth-generation (4G) cellular networks, embodied by LTE, have now reached full commercial maturity across all regions of the world, providing high-speed, high-capacity mobile access to millions of people. 4G expansion and growth, however, is far from over. According to Cisco's latest forecast [1], traffic from wireless and mobile devices will exceed traffic from wired devices by 2019, scoring an overall 10-fold growth between 2014 and 2019. As more and more users upgrade their hardware to 4G and bandwidth-hungry applications cause a surge in data traffic demands, the identification of new spectrum, earmarked for exclusive or shared use by cellular technologies, becomes crucial. It is thus only natural that regulators and operators alike are eyeing unlicensed spectrum as a potential solution [2]. Such unlicensed bands for mobile devices are located around 2.4 GHz and 5 GHz. Given the plethora of devices and technologies already crowding the 2.4 GHz ISM band, LTE Release 13 has started addressing potential coexistence of LTE and other technologies, namely WiFi, in the unlicensed portion of the 5 GHz spectrum [3]. LTE in unlicensed bands is referred to by Rel. 13 as Licensed-Assisted Access (LAA). Among the regulations foreseen by Rel. 13 [4] are: (i) limitations on transmission power (23 dBm in Europe and 24 dBm in the U.S. for indoor usage); (ii) interference-avoidance mechanisms toward incumbent meteorological radar systems using Dynamic Frequency Selection (DFS, already implemented by newer WiFi Access Points); (iii) Listen Before Talk (LBT) MAC-layer operations for graceful coexistence with the contention-based WiFi DCF protocol. The latter feature is likely to be the most prominent one, given the existing regulatory requirements in Europe and Japan, which specifically call for LBT.

In this paper, we chose to focus on LBT regulations specified by the ETSI EN 301 893 [5] Draft Standard, through what is called the Clear Channel Assessment (CCA) procedure for graceful WiFi-LTE coexistence. Given the lack of LAA devices that could allow experiments on a real testbed, we resorted to simulation to assess the coexistence of the most commonly deployed WiFi interface, IEEE 802.11n, and LAA. We restricted our analysis to a small-cell residential scenario since it is the most likely deployment solution given the transmission power limitations.

Our work falls within the scope of the growing interest surrounding unlicensed LTE and its various proposed solutions, which do not necessarily follow the path laid out by 3GPP. We are the first to provide a comprehensive assessment of the coexistence of LAA-LTE and IEEE 802.11n through a detailed, frame-level simulation using realistic traffic assumptions.

Other authors have investigated coexistence issues in unlicensed spectrum for LTE. For example, in [6], authors examine the use of Almost-Blank Subframes (ABS) for LTE-WiFi coexistence in small-cell networks, while [7] focuses on a time division duplex approach using Reinforcement Learning techniques to equally allocate the spectrum to the two transmission technologies. The authors in [8] have looked at coexistence strictly from the point of view of mutual interference, while a comparison of the performance of LBT in different LTE scenarios (indoors and outdoors) was the focus of [9]. Finally, similar in spirit to our work, [10] chooses an analytical, rather than simulative approach, to assess the joint performance of WiFi and LAA with LBT.

The paper is organized as follows. In Sec. II, we outline the LAA MAC protocol that we are going to investigate, detailing the network model we focus on in Sec. III. Simulation results are presented and analysed in Sec. IV, while Sec. V concludes the paper.

II. MAC PROTOCOL FOR LICENSED-ASSISTED ACCESS

The LAA MAC protocol introduced by [5] aims at preventing an LTE device using shared unlicensed spectrum from transmitting on the channel in the presence of transmissions from other systems, such as WiFi cards. The draft proposes two different MAC protocols aimed, respectively, at Frame-Based Equipment (FBE) and Load-Based Equipment (LBE). The former includes devices where the transmit/receive structure is not demand-driven but, rather, has a fixed timing, while the latter has opposite characteristics. We will address only LBE: the two MAC-layer solutions are quite similar anyway.

The operating guidelines of LBE access are as follows.

- Any transmission by LBE must be preceded by a *Clear Channel Assessment (CCA)*: the channel is observed for at least $20 \mu\text{s}$ (*Channel Observation Time, COT*) and only if no transmission is detected during this time, can the equipment subsequently transmit.
- If the channel is found occupied, an *Extended CCA (eCCA)* check is performed: the channel is monitored for a time computed as the COT multiplied by a random integer N . This procedure, similar to the DCF backoff, results in an idle period that must be observed before transmission can begin. The N integer is chosen at random in the interval $[1, q]$, the selection is performed every time an extended CCA is required and the value is stored in a counter. The q parameter is selected by the manufacturer in the range $[4, 32]$. Again, in a fashion similar to DCF backoff, the counter is decremented every time a COT is deemed to be unoccupied and transmission will start upon the counter reaching zero.
- In another feature, reminiscent of DCF TXOP, transmission by an LBE must not last for more than a *Maximum Channel Occupancy (MCO)* time, which is determined as $13/32 \cdot q$ ms. The transmission shall be followed by an Extended CCA, regardless of whether more transmissions are scheduled.
- Upon reception of a frame intended for it, the LBE can skip CCA and reply immediately with the appropriate control frame (e.g., an ACK).

As we pointed out, there are many similarities to basic DCF as well as to its modifications (EDCA) introduced since IEEE 802.11e onwards. The eCAA is tantamount to DCF backoff, except for the fact that the random values do not include 0. Of course, backoff and eCCA ranges differ also depending on the implementation, the version of 802.11 and the choice of Access Category. It should also be remarked that the exponential nature of the backoff is not replicated in the eCAA: the LBE will not become less aggressive after a collision, and such a behavior is potentially harmful versus WiFi in crowded environments.

The choice of the values for the q parameter can have a relevant impact on the coexistence with WiFi, as will be shown in our results. Indeed, the choice of q is a bit of a tradeoff between increased aggressiveness in capturing the channel and length of the channel occupancy time during a transmission. It is easy to see that, on the one hand, small values of q shorten the wait before the channel is declared idle by the LBE. Thus, on average, the LBE can forestall channel access for lower-priority EDCA Access Categories. A small q , however, would allow the LBE only a limited amount of time to transmit its data once it captures the channel. On the other hand, a high q gives the LBE more transmission time, while diluting its chances to grab the channel ahead of WiFi.

III. NETWORK MODEL

Our study focuses on a single-cell heterogeneous residential network where a LAA Access Point (LAA-AP) is deployed

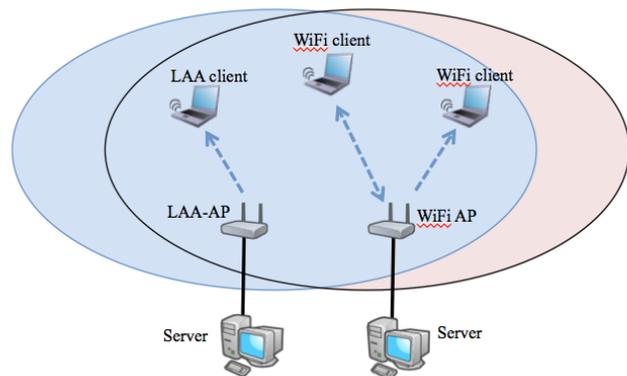


Fig. 1. Residential heterogeneous network scenario.

alongside a WiFi Access Point (WiFi-AP). We are mainly interested in a downlink traffic scenario. Therefore, unless explicitly indicated otherwise, we assume that data is generated by a server connected to the LAA-AP and streamed to one LAA-client. A similar traffic configuration is applied to the WiFi hot spot as well. When simulations involve two types of traffic, an additional WiFi client is associated to the WiFi-AP. This is a reasonable scenario for a residential environment where an LAA-AP and a WiFi-AP may be deployed in neighboring apartments, or where a LAA-AP serving a small-cell is right outside the house. LAA and WiFi devices operate on the same frequency channel and are in each other's transmission range, in order to investigate the worst-case interoperation scenario. However, they are oblivious to each other, i.e., they do not take turns in sharing the channel, but contend for it following the respective MAC protocol specifications. In the case of LAA, as mentioned previously, we implemented the LBE behavior. Figure 1 shows the network topology.

Traffic consists of two different flow types: UDP streaming and VoIP, used individually or together, depending on the scenario, as will be specified in Sec. IV:

- UDP streaming is unidirectional (downlink) and features 1000-byte UDP packets whose inter-generation time is exponentially distributed, with different average values that are used to define the offered traffic. Although such a traffic configuration is hardly realistic, it nevertheless will provide us with the opportunity to investigate the network behavior under near-saturation and saturation conditions.
- VoIP connections are bidirectional and are modelled as On/Off sources whose durations are Weibull-distributed random variables [12]. During a talkspurt, 40-byte packets are generated every 20 ms. At the receiver, a 20-ms playout buffer smooths out the received stream.

PHY and MAC parameters for all wireless cards are set to standard values, and the most meaningful parameters are summarized in Table I. Results are derived for three different 802.11 Access Categories (AC), namely Best Effort (AC_BE), Video (AC_VI) and Voice (AC_VO): their parameters are listed in Table II. They are the most commonly

TABLE I
LIST OF PARAMETERS

PHY/MAC Parameters	Value
PHY data bitrate	135 Mb/s
PHY basic bitrate	13 Mb/s
PHY control bitrate	135 Mb/s
802.11 MAC Slot Time	9 μ s
LAA MAC COT	20 μ s
802.11/LAA MAC retry limit	7 μ s
LAA q parameter	8, 32 μ s
802.11 MAC A-MPDU max size	65535 B
802.11 MAC MPDU spacing	8 μ s

TABLE II
802.11N ACCESS CATEGORY PARAMETERS

AC	CW_{min}	CW_{max}	AIFSN	TXOP
AC_{BE}	15	1023	3	0
AC_{VI}	7	15	2	3.008 ms
AC_{VO}	3	7	2	1.504 ms

found on commercial WiFi cards.

In our work we have derived results for both single and aggregated 802.11n frames. Frame aggregation in 802.11n is considered optional and, when we used it, we just considered MPDU aggregation (whose frames are referred to as A-MPDUs in the 802.11n standard [11]). It consists of several MAC MSDU (typically, IP datagrams), which are given an MPDU delimiter, a MAC header and frame check sequence, and are padded so that each aligns on symbol boundaries in the physical transmission. These subframes are then aggregated into a single frame and provided with a single PLCP header, to form an A-MPDU. All subframes within an A-MPDU must be destined to the same receiver address on the wireless link, which is obviously our case.

IV. SIMULATION RESULTS

We have implemented the LAA MAC protocol and the WiFi 802.11n DCF protocol in the OMNET++ simulator. All simulations were run for an amount of time so that the point estimate of the average throughput was within the 95% interval of confidence. All the metrics we show are plotted against the UDP streaming offered traffic, shown on the x-axis, and computed as the overall UDP packet generation rate from servers sending streaming traffic into the wireless network.

We collect the following metrics:

- *throughput*: computed at the receiver, it accounts for the average of the number of packets correctly received, divided by the simulated interval;
- *end-to-end delay*: average packet delivery delay measured at the application layer;
- *frames per MCO*: average number of frames transmitted by the LAA-AP during the MCO;
- *subframes per A-MPDU*: average number of subframes aggregated into an A-MPDU by the WiFi AP.

A. No frame aggregation

The first set of results refers to our basic scenario, featuring one LAA client and one 802.11n WiFi client receiving

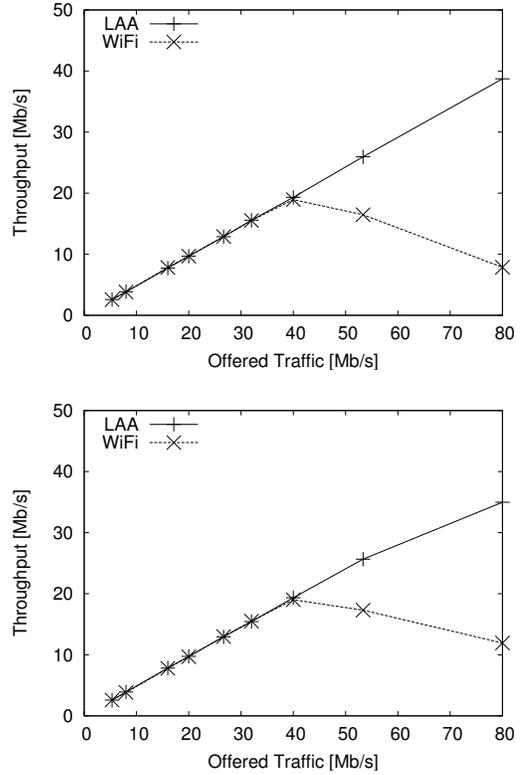


Fig. 2. Throughput as a function of offered traffic, for different values of the LAA ' q ' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as AC_{BE} without MPDU aggregation.

AC_{BE} traffic (without MPDU aggregation). Figure 2 shows the comparison between the received throughputs at the LAA and WiFi client. The network reaches pre-congestion levels around an offered traffic of 40 Mb/s, at which point the WiFi throughput is throttled, while the LAA client is unaffected, regardless of the value of the q parameter. Adverse co-habitation effects of the two technologies actually begin at a lower load, as shown in Figure 3 by the surge of the average packet end-to-end delay for WiFi for loads higher than 25 Mb/s. LAA does not show such a remarkable delay increase because it manages to send more frames per MCO (shown in Figure 4), thus countering its fewer access opportunities and, at the same time, inflicting a longer defer interval to each WiFi access attempt. It is perhaps worth recalling that AC_{BE} WiFi traffic allows only one frame per TXOP.

In the second set of results, all WiFi traffic is sent as AC_{VI} , resulting in a sizable TXOP for it as well. The choice of such an Access Category is motivated by downlink UDP traffic coming from what could well be a video streaming source. As shown by Figure 5, the situation is reversed. No changes occur until the network reaches saturation, since neither the MCO nor frame aggregation are needed. However, when traffic reaches the 50-Mb/s mark, the LAA throughput dips slightly for $q = 8$, and more markedly for $q = 32$. The use of MCO/TXOP, reported in Figure 6, allows the two technologies to equally share the channel even at high

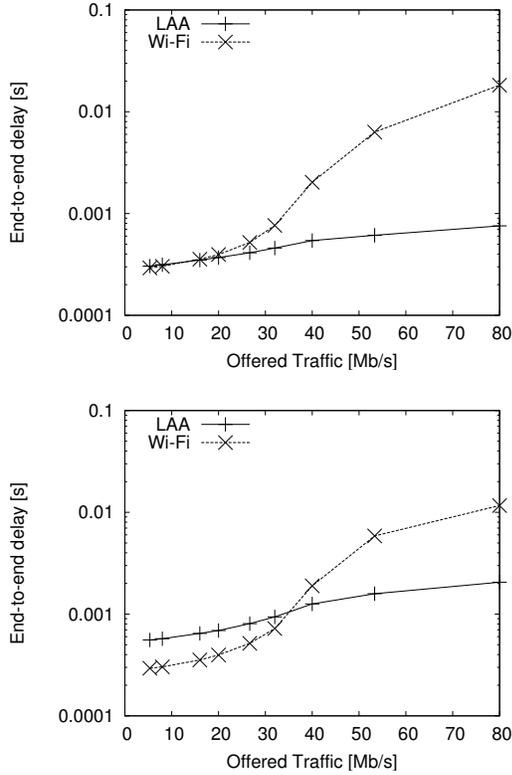


Fig. 3. End-to-end delay as a function of offered traffic, for different values of the LAA 'q' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as AC_{BE} without MPDU aggregation.

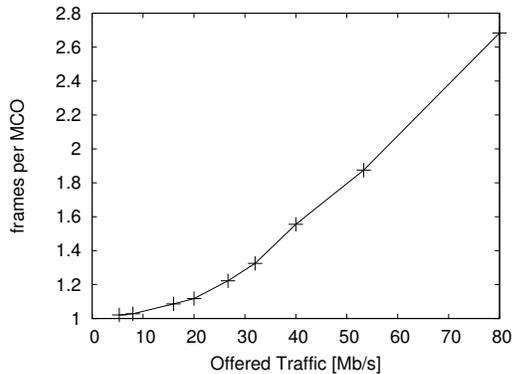


Fig. 4. Average number of frames per MCO transmitted by the LAA-AP as a function of offered traffic; WiFi traffic sent as AC_{BE} without MPDU aggregation.

loads. However, on average, WiFi in AC_{VI} enjoys a smaller backoff interval ($31.5 \mu s$) than the extended CCA performed by LAA, for any of the considered q values ($90 \mu s$ for $q = 8$, $330 \mu s$ for $q = 32$). This translates into WiFi being consistently able to wait a lesser amount of time. In spite of the additional LAA frames forced into each MCO period (Figure 6) by the longer defer period, this does not give LAA the edge versus WiFi at saturation loads.

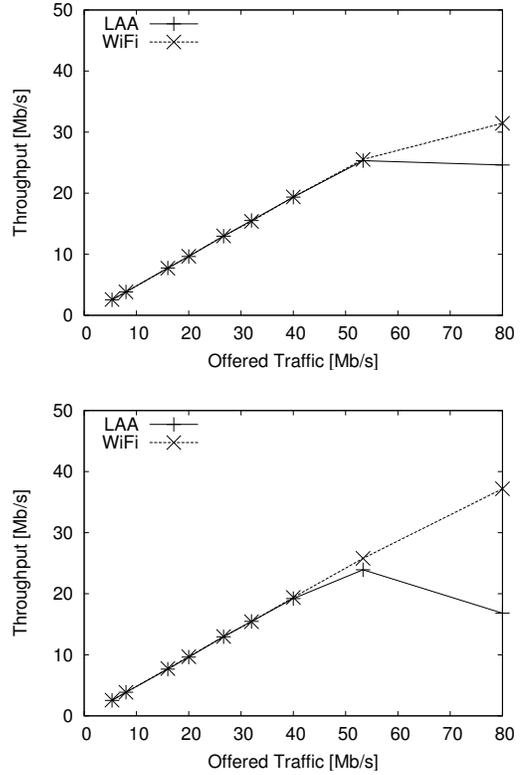


Fig. 5. Throughput as a function of offered traffic, for different values of the LAA 'q' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as AC_{VI} .

B. MPDU aggregation

We now introduce the possibility to aggregate 802.11n MPDUs and repeat some of the simulations already shown in the previous subsection. Specifically, we look again at the basic scenario (one LAA client and one 802.11n WiFi client receiving AC_{BE} traffic). Interestingly, aggregation gives AC_{BE} traffic the same competitive edge versus LAA enjoyed by higher access category traffic, as shown in Figure 7. Such a conclusion is supported by Figure 8, which highlights the utilization of A-MPDUs by WiFi in relation to that of MCOs of LAA and shows a similar pattern as in Figure 6.

C. Coexistence with WiFi VoIP traffic

For our next set of simulations, we break a new factor into our scenario, represented by an additional WiFi client generating *uplink* VoIP traffic toward another client located on the wired network accessible through the same WiFi AP used in previous simulations. Such traffic is carried in AC_{VO} category (whose parameters are also listed in Table II), although in many commercial implementations, VoIP data are transmitted as plain AC_{BE} traffic.

We are of course keen to establish the impact of LAA traffic on VoIP data; specifically, whether it affects the VoIP delivery delay and its standard deviation, which leads to increase in jitter, and consequent degraded quality of experience for the user. Figure 9 compares the above mentioned metrics (mean

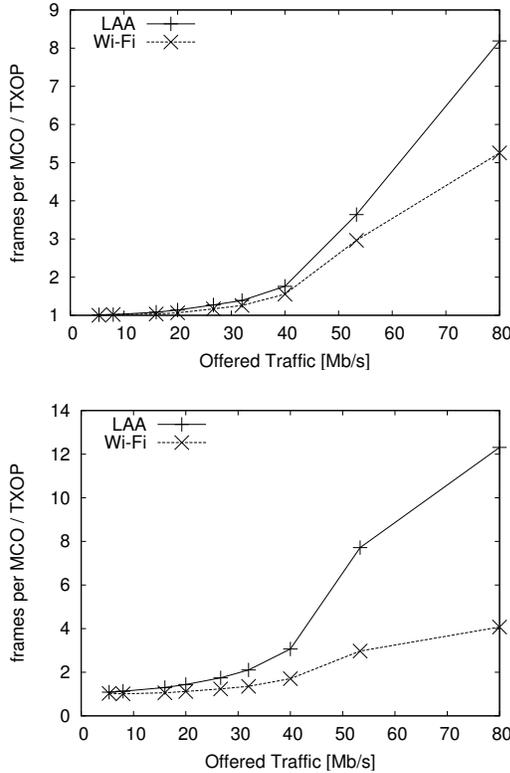


Fig. 6. Average number of frames per MCO transmitted by the LAA-AP and frames per TXOP transmitted by the WiFi AP, as a function of offered traffic, for different values of the LAA 'q' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as *AC_VI*.

delay and its standard deviation) when the VoIP client is set against two competing downlink UDP flows transmitted, respectively, by: i) LAA-AP ($q = 8$) and 802.11n AP using no aggregation [*LAA+11n no aggr* in the figure], or ii) LAA-AP ($q = 8$) and 802.11n AP using A-MPDUs [*LAA+11n aggr*], or iii) two 802.11n AP using A-MPDUs, sending to two different 802.11n clients [*2x11n aggr*]. Note that the latter flow pair features no LAA traffic. VoIP traffic is carried as *AC_VO* category. It should also be noted that, for uniformity of reference, VoIP traffic is not included in the computation of offered traffic on the x-axis, and VoIP load is never increased, but conforms to the parameters specified in the previous section.

The results clearly show that VoIP traffic would be only marginally affected when competing against the two 802.11n AP: not surprisingly, since its access category lends it a decisive priority against *AC_BE* traffic. After a transmission ends and no other transmissions occupy the channel, VoIP flows can “declare” the channel as idle on average after $2 \cdot 9 + 1.5 \cdot 9 = 31.5 \mu\text{s}$, while *AC_BE* traffic can do it after $3 \cdot 9 + 7.5 \cdot 9 = 94.5 \mu\text{s}$ (see Table II). VoIP does not enjoy such an overwhelming advantage when competing against LAA with $q = 8$: LAA can declare the channel as idle after $20 + 4.5 \cdot 20 = 60.5 \mu\text{s}$. Indeed, mean and standard deviation of the end-to-end delay surge around 25 Mb/s of offered traffic when LAA is involved: such a load marks the time when

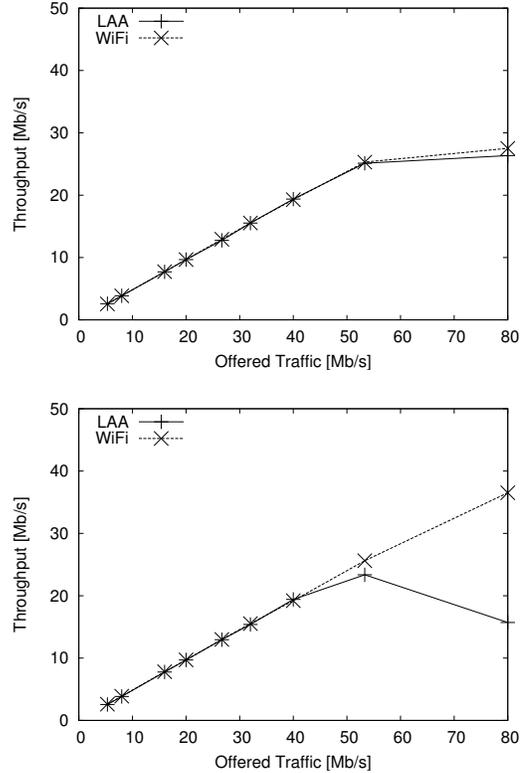


Fig. 7. Throughput as a function of offered traffic, for different values of the LAA 'q' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as *AC_BE* with MPDU aggregation.

LAA starts piling up more and more consecutive frames in its MCO interval, delaying the transmission of scheduled VoIP frames, hence increasing the standard deviation of their delay. This behavior is even more evident for higher values of the q parameter (which results in longer MCO), although we chose not to present the results in detail for reasons of space.

D. Final remarks

Although our simulation study was limited to a selected residential scenario with two types of traffic, we can already glean some interesting behaviors and establish some guidelines for parameter selection for future deployments. WiFi 802.11n has been shown to fairly compete against LAA at high loads only if WiFi traffic is sent in higher Access Categories or if frame aggregation is enabled. These features, however, are not uniformly supported by all 802.11n chipset installed on consumer devices, and even when they are, some chipsets only include a limited-size MAC buffer, which bounds the maximum A-MPDU size. As for VoIP traffic, if protected by the *AC_VO* category, it has the potential to be unaffected at low-medium loads. However, an extended MCO interval by LAA traffic can seriously disrupt the delivery of VoIP packets, leading to jitter increase and lower QoE for users.

V. CONCLUSIONS AND FUTURE WORK

We have examined the coexistence of co-channel WiFi and Licenced-Assisted Access networks in a typical residential

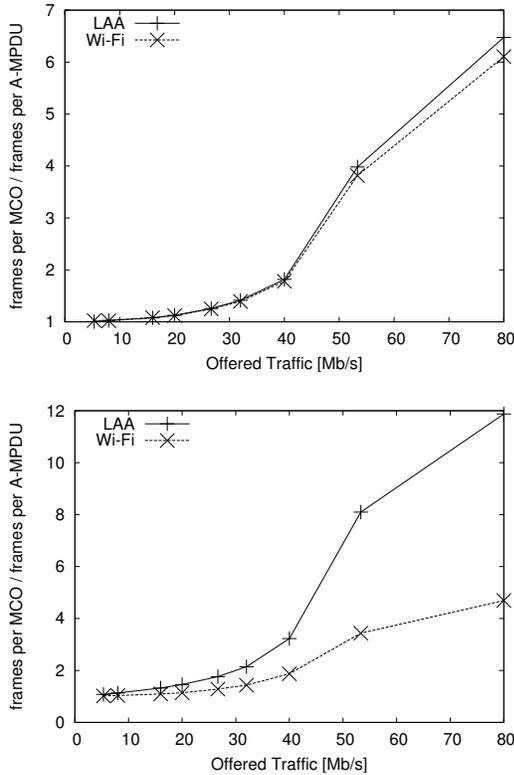


Fig. 8. Average number of frames per MCO transmitted by the LAA-AP and subframes per A-MPDU transmitted by the WiFi AP, as a function of offered traffic, for different values of the LAA 'q' parameter ($q = 8$ top, $q = 32$ bottom); WiFi traffic sent as *AC_BE* with MPDU aggregation.

scenario. Focusing on a residential heterogeneous scenario, we have simulated concurrent LAA and 802.11n WiFi traffic at low, medium and high load. Our results point at good co-existence properties if WiFi and LAA are allowed to compete on similar grounds. This can be achieved either by letting WiFi use higher-access categories or by using frame aggregations, though these features are not uniformly supported by commercial chipsets. We have also investigated the impact of LAA co-channel traffic on VoIP communication, highlighting possible disruptive phenomena at medium-high loads.

Since the timeline of LAA technology adoption is not clear yet, future studies should also include comparisons with 802.11ac to determine coexistence with this standard as well.

ACKNOWLEDGMENTS

This work was supported in part by Telecom Italia - Wireline Devices and Home Network and in part by the European Union through the XHAUL project (H2020-ICT-2014/671598).

REFERENCES

- [1] Cisco Visual Networking Index: Forecast and Methodology, 20142019
- [2] Qualcomm, "Qualcomm research LTE in unlicensed spectrum: Harmonious coexistence with Wi-Fi," available at <https://www.qualcomm.com/media/documents/files/lte-unlicensed-coexistence-whitepaper.pdf>, June 2014.
- [3] 3GPP RP-140808, "Review of regulatory requirements for unlicensed spectrum," June 2014.

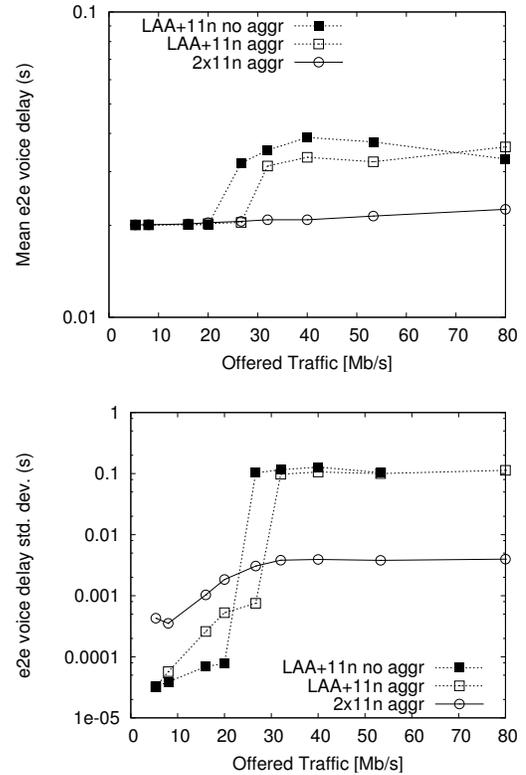


Fig. 9. Comparison among mean end-to-end delay (top) and delay standard deviation (bottom) of VoIP packets in three different network scenarios, as a function of offered traffic; VoIP traffic in *AC_VO*.

- [4] R. Zhang, M. Wang, L. X. Cai, Z. Zheng, X. Shen, L. Xie, "LTE-Unlicensed: The future of spectrum aggregation for cellular networks," *IEEE Wireless Communications Magazine*, IEEE Wireless Communications, Vol. 22, No.3, pp. 150-159, June 2015.
- [5] ETSI EN 301.893 v1.7.1. "Broadband Radio Access Networks (BRAN); 5 GHz high performance RLAN."
- [6] H. Zhang, X. Chu, W. Guo, S. Wang, "Coexistence of Wi-Fi and heterogeneous small cell networks sharing unlicensed spectrum," *IEEE Communications Magazine*, vol. 53, no. 3, pp. 158-164, March 2015.
- [7] N. Rupasinghe, İ. Güvenç, "Reinforcement Learning for Licensed-Assisted Access of LTE in the Unlicensed Spectrum," *IEEE WCNC 2015*, New Orleans, USA, March 2015.
- [8] J. Jeon, Q. Li, H. Niu, A. Papathanassiou, G. Wu, "LTE in the Unlicensed Spectrum: A Novel Coexistence Analysis with WLAN Systems," *IEEE Globecom 2014*, Austin, USA, Dec. 2014.
- [9] R. Ratasuk, N. Mangalvedhe, A. Ghosh, "LTE in Unlicensed Spectrum using Licensed-Assisted Access," *IEEE Globecom 2014*, Austin, USA, Dec. 2014.
- [10] C. Chen, R. Ratasuk, A. Ghosh, "Downlink Performance Analysis of LTE and WiFi Coexistence in Unlicensed Bands with a Simple Listen-before-talk Scheme," *IEEE VTC Spring 2015*, Glasgow, U. K., May 2015.
- [11] "IEEE 802.11n-2009Amendment 5: Enhancements for Higher Throughput," *IEEE-SA*, Oct. 2009.
- [12] J. Seger, "Modelling Approach for VoIP Traffic Aggregations for Transferring Tele-traffic Trunks in a QoS enabled IP-Backbone Environment," *International Workshop on Inter-domain Performance and Simulation*, Salzburg, Austria, Feb. 2003.