Boosted WiFi through LTE Small Cells: The Solution for an All-Wireless Enterprise

David López-Pérez[1], Jonathan Ling[1], Bong Ho Kim[1], Vasudevan Subramanian[1], Satish Kanugovi[1] Ming Ding[2]
[1]Bell Laboratories/Alcatel-Lucent
[2]Data 61, Australia

***Abstract*—Wireless fidelity (Wi-Fi) benefits from large downlink bandwidth but suffers from inefficient uplink contention, while long term evolution (LTE) benefits from efficient scheduling in the uplink but suffers from bandwidth scarcity in the downlink. In this paper, we quantify the performance gain of *Boost*, a new technology that blends the advantages of Wi-Fi and LTE in an efficient manner to overcome their individual disadvantages. By redirecting uplink traffic from Wi-Fi to LTE, *Boost* avoids resource waste due to Wi-Fi contention in the uplink. Since the radio channel is now fully under the control of an scheduler, downlink controlled by Wi-Fi AP scheduler and uplink controlled by LTE scheduler, delays and rates can be guaranteed. Simulation results show that in typical enterprise scenarios, *Boost* significantly decreases packet delay up to 85 % and increases user throughput up to 3.5x, which is beneficial for delay-sensitive and best effort applications, respectively. This positions *Boost* as a strong candidate for realising the all-wireless enterprise.**

I. Introduction

The main challenges that mobile operators face within enterprises is how to meet the demand for voice, video and data with consistent quality no matter whether subscribers connect over Wi-Fi or LTE networks. While both technologies continue to evolve, each has its own challenges, which can significantly impact the user experience. Wi-Fi performance suffers in dense deployments with a large number of devices because of the poor performance of its simple medium access control (MAC) mechanism, i.e., carrier sense multiple access/collision avoidance (CSMA/CA) [1] [2] [3], while LTE capacity is limited due to the scarcity of licensed spectrum [4]. By combining the strengths of Wi-Fi (i.e., large downlink (DL) bandwidth) and LTE (i.e., efficient uplink (UL) MAC), we believe it is possible to overcome the drawbacks of each technology, and offer the enterprise users a wired like experience.

Up to now, operators would switch enterprise user equipments (UEs) between Wi-Fi and LTE to balance the network load through access network discovery and selection function (ANDSF) [5], or users would manually switch between networks to seek optimum performance. These approaches rarely result in the desired quality of service (QoS).

To address this issue, *Boost* has emerged as a new technology that blends for the first time the download and upload capabilities of Wi-Fi and LTE to generate higher network capacity, and give users a more consistent and higher-quality voice, video and data experience [6]. *Boost* allows operators to combine stand-alone Wi-Fi and LTE networks into one unified wireless network for home, office and outdoor environments, and only requires easy-to-implement software updates at the network and the UE [7].

In this paper, we quantify the performance gains of *Boost*, and show its ability to significantly increase the capacity of voice over IP (VoIP) and file transfer protocol (FTP) applications over Wi-Fi in typical enterprise scenarios, where VoWi-Fi is of particular importance due to the recent popularity of enterprise services such as Lync, WebEx, etc. The rest of the paper is organised as follows: In Section II, the low performance of Wi-Fi in dense deployments due to CSMA/CA is discussed, followed by the introduction of *Boost* in Section III. In Section IV and Section V, the system-level simulation tool and the specific VoIP and FTP traffic models used to evaluate the performance of voice and best effort services over *Boost* are presented. The simulation results are discussed in Section VI. Finally, in Section VII, the conclusions are drawn.

II. Poor Performance of CSMA/CA

Wi-Fi takes a decentralised approach to scheduling transmissions, unlike the base station (BS)-centric approach used by LTE. When a Wi-Fi device, an access point (AP) or a UE, wants to make a transmission, it senses the radio channel and performs a clear channel assessment (CCA) check. If no transmissions are detected for a period of time (referred to as DIFS), the transmission begins. Otherwise, the device draws an integer number uniformly at random from 0 to 15 time slots, the so-called contention window, and starts to count down. The counter is paused during periods when the channel is detected busy. When the counter reaches zero, the device proceeds with the transmission. If another device also transmits at the same time, then a collision occurs and the transmission may fail due to poor signal quality. When a transmission fails, which is detected by the absence of a MAC acknowledgment (ACK) from the receiver at the transmitter after a period of time (referred to as SIFS), a new random number is drawn and this process is repeated. The size of the contention window is doubled on each collision, i.e., increasing as 16, 32, 64 *etc* up to a maximum of 1024 time slots. After a successful transmission, the size of the contention window is set to its original value, i.e.,16 time slots.

This simple and polite MAC mechanism, referred to as CSMA/CA [8], is able to prevent collisions in an uncoordinated manner, allowing Wi-Fi to be robust and scalable in unplanned deployments within sparse scenarios, e.g., at homes. However, it is the simplicity of CSMA/CA that causes Wi-Fi's poor performance in scenarios with a large number of devices,

978-1-5090-3254-9/16/$31.00 ©2016 IEEE

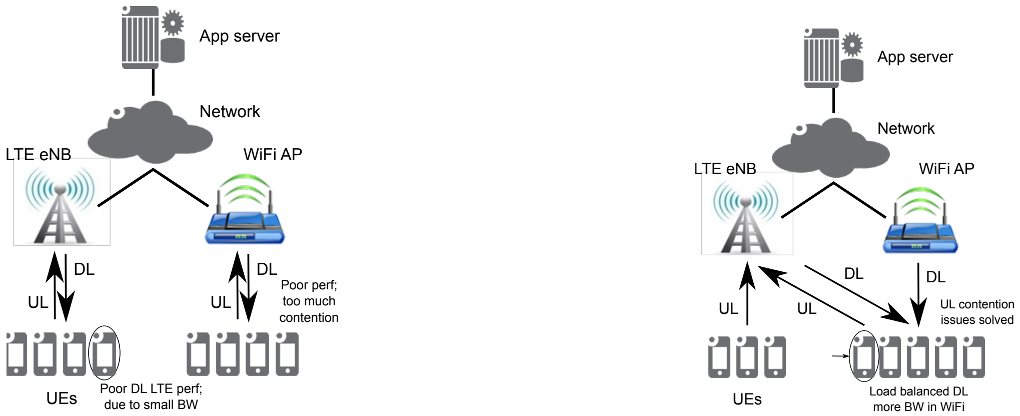

(a) Traditional architecture. (b) *Boost* architecture.

Fig. 1: *Boost* alleviates WiFi network congestion by diverting UL traffic to the LTE network. Moreover, it alleviates LTE DL congestion by moving selected DL traffic to the WiFi network.

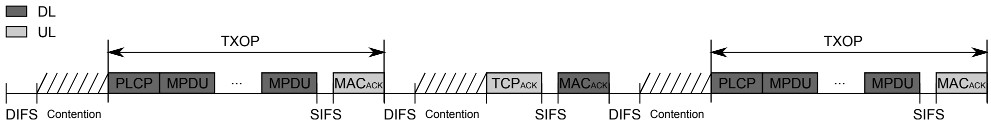

(a) No *Boost*.

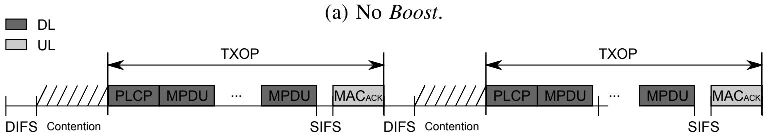

(b) *Boost*.

Fig. 2: Benefit of UL TCP ACK offloads from WiFi to LTE. Shorter time to transmit the same information with *Boost*.

such as shopping malls, airports and enterprises. The probability of multiple devices drawing the same random number and thus the probability of collision grows with the device density, leading to data loss and device back-off, which further damages Wi-Fi performance. Moreover, when devices transmit at the same time, the Wi-Fi AP tends to capture the packet with the highest signal to interference plus noise ratio (SINR), this resulting in a specially poor performance for devices located at the cell-edge, the so-called capture effect. Increasing the contention window mitigates the collision probability, but optimising this parameter does not truly solve the collision problem, only postpone it. This collision problem is further aggravated by the defects in the carrier sensing process, such as the hidden terminal problem [8], and is still an issue in new Wi-Fi version, such as 802.11 n and 802.11 ac. For interested readers, the performance degradation of a Wi-Fi AP with the number of connected UEs is presented in [1] [2] [3].

III. Boost

Boost is oriented towards enhancing the efficiency of Wi-Fi, and as a by-product the performance of LTE, by redirecting UL traffic from the Wi-Fi network (unlicensed band) to the LTE network (licensed band)[1]. Another important advantage of *Boost* is its DL traffic management capabilities. *Boost* can intelligently steer DL traffic between the Wi-Fi and the LTE networks, and benefit from traffic aggregation and/or load balancing to both networks. For example, when the Wi-Fi network is congested, some of its DL traffic can be diverted to LTE and vice versa. Preferably, this traffic diversion should occur when contention begins to impact channel availability in either network. Traffic steering, however, is the subject of future study and it will not be addressed in this paper. Fig. 1 depicts the main advantages of *Boost*.

By routing UL traffic via the LTE network, there is no contention to resolve in the Wi-Fi network using the CSMA/CA protocol, which addresses the issues discussed in Section II. As a result, in a planned network with no co-channel neighbours, Wi-Fi operates only in the DL and works on a fully scheduled basis (DL Wi-Fi traffic is scheduled by the Wi-Fi AP). This allows an efficient use of Wi-Fi's large bandwidth, up to 160 MHz per UE with channel bonding, without the delay introduced by UL contention and in a completely collision-free manner. As a result, DL delays and rates can be guaranteed in the Wi-Fi network. Guaranteed delays can be particularly helpful in providing real-time services, e.g., high quality voice and video. Best effort traffic also benefits from the UL contention avoidance and the scheduled nature of traffic. Moreover, DL transmission control protocol (TCP) applications experience improved throughputs by avoiding UL TCP ACKs over the unpredictable and congested Wi-Fi links. The offload of UL TCP ACKs to LTE allows a back-to-back DL data scheduling in Wi-Fi with the resulting overhead reductions, as shown in Fig. 2, Such offload also avoids the shrinkage of the TCP transmit window due to lost/delayed UL TCP ACKs [7].

In order to allow an implementation of *Boost* that only requires software updates at the network and the UE, the redirection of traffic in both DL and UL is foreseen at the

[1]. Some Wi-Fi traffic with stringent delay constraints will still be sent via the UL of Wi-Fi, e.g., MAC ACKs sent from the receiver and expected at the transmitter within tens of microseconds

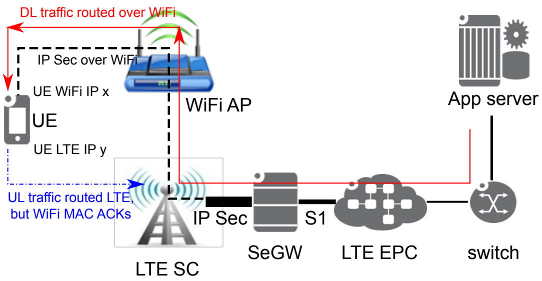

Fig. 3: *Boost* detailed architecture.

internet protocol (IP) layer [9]. With an integration of Wi-Fi and LTE at the IP layer, the *Boost* architecture and data split is easily implementable since network connections can be managed separately without knowledge of each other's presence. Fig. 3 illustrates an example of the *Boost* architecture. In this example, all DL traffic is routed through the evolved packet core (EPC) to a security gateway (SeGW) that acts as endpoint for the IPsec tunnel to the LTE small cell. Then, such DL traffic is routed by the LTE small cell through the Wi-Fi AP, optionally protected by an IPsec tunnel too. All UL traffic is routed over LTE except Wi-Fi MAC ACKs.

IV. System Modeling

In order to evaluate the performance of *Boost*, an in-house Wi-Fi/LTE integrated system-level simulator has been developed. This simulator takes as basis an existing LTE simulator [10], which follows the Third Generation Partnership Project (3GPP) simulation methodology specified in [11] and [12]. On top of it, a Wi-Fi module with 802.11n features has been developed [8]. In the following, the main characteristics of this Wi-Fi module are presented.

A. Carrier Sensing

The core of the Wi-Fi module is the CSMA/CA MAC protocol described in Section II, together with its complementing carrier sensing mechanisms. In this case, when performing a CCA check, two carrier sensing mechanisms are used: energy detection (ED) and virtual carrier sensing.

When using the ED mechanism, if the received energy of the in-band signal at a Wi-Fi device is equal or larger than the ED threshold, its CCA is held busy until the medium energy is smaller than the ED threshold. A typical ED threshold is -62 dBm for a 20 MHz channel.

When using the virtual carrier sensing mechanism, a Wi-Fi device can 'tell' neighbouring Wi-Fi devices for how long it will occupy the channel. This time information, referred to as network allocation vector (NAV), is transmitted in the header of specific frames, e.g., request-to-send (RTS) frames, clear-to-send (CTS) frames, physical layer convergence protocol (PLCP) headers. If the NAV is successfully decoded in a neighbouring device, its CCA is held busy until the end of the specified time. In order to successfully retrieve NAV, the container has to be detected at or above -82 dBm for a 20 MHz channel, and decoded at or above 4 dB SINR.

B. RTS-CTS

The RTS/CTS handshake is a mechanism used by 802.11 to reduce frame collisions caused by the hidden node problem. The transmitter initiates the handshake by sending a RTS frame to the receiver. Upon its reception, the receiver responds to it by sending a CTS frame to the transmitter. The transmitter will not send any data frame before receiving the CTS frame.

Both RTS and CTS may contain NAV information, which alerts other neighbouring devices to hold off from accessing the medium, while the device sending the RTS transmits its data. Since such NAV information is broadcasted from both the transmitter (using RTS) and the receiver (using CTS), the hidden node problem is mitigated. This RTS/CTS handshake is not usually used unless the data frame size exceeds a threshold. In our simulations, the handshake is only enabled for FTP traffic, RTS and CTS packet sizes are 20 bytes and 14 bytes, respectively, and they are sent at the lowest commonly supported data rate to ensure maximum reliability.

C. Data Aggregation

When a Wi-Fi device gains access to the channel, it will hold on to it at most for a transmission opportunity (TXOP) time. Then, it will release the channel and contend for it again, if it has more data to transmit. In this way, selfish behaviours are avoided, allowing for a fair co-existence of Wi-Fi devices. TXOP cannot be larger than the maximum channel occupation time defined by regional regulations, e.g., the maximum channel occupancy time in Europe is 10 ms.

Within its TXOP time, the device sends a burst of information composed of a PLCP header followed by an aggregated MAC protocol data unit (AMPDU). The PLCP header contains information on the rate and length of the rest of the frame, as well as error correction and information on the encoding scheme, and is sent at the lowest commonly supported data rate to ensure maximum reliability. The AMPDU consists of a collection of MAC protocol data units (MPDUs), and each MPDU is comprised of a delimiter (4 bytes), a MAC header (34 bytes) and the actual UE data, e.g., an IP packet (1500 bytes). The number of aggregated MPDUs on an AMPDU depends on the buffered data, the TXOP time and the transmission rate.

D. Rate Control

The Minstrel rate control algorithm is adopted to select the frame rate (modulation and coding scheme (MCS)) [13]. Basically, Minstrel tries all the rates, and selects the one that work best. All rates are tried on a regular basis.

In more detail, each frame has a retry chain, which consists of four rate-count pairs, i.e., r_0/c_0, r_1/c_1, r_2/c_2, r_3/c_3. A frame is first transmitted at rate r_0 for c_0 attempts. If these attempts are not successful, the frame is re-transmitted at rate r_1 for c_1 attempts. This process continues until the frame is successfully transmitted or ultimately discarded after $c_0+c_1+c_2+c_3$ failed re-transmissions.

In order to select the retry chain rates r_0, r_1, r_2 and r_3, statistics per rate on success probability, $P_{\text{success}}(t+1) = a \cdot P_{\text{success}}(t) + (1-a) \cdot \frac{N_s}{N_t}$ and average rate $T(t) = w \cdot P_{\text{success}}(t)$ are computed and updated every 100 ms, where N_t is the number of transmitted frames, N_s is the number of successfully received frames, w is the raw rate of the MCS and $a = 0.25$. Based on these two variables, the retry chain rates are selected according to the scheme presented in Table I of [13]. Note that in such scheme, 1 every 10 packets selects a random rate in order to sample the entire rate space.

E. TCP ACK

TCP is a reliable stream delivery protocol that guarantees that all bits received will be identical to all bits sent, and that they are in the correct order. Since data transfer over the air is not reliable, a technique known as positive ACK with retransmission is used to guarantee the reliability of the transmission at the TCP layer. This technique requires the receiver to respond with an TCP ACK message as it receives the data. The transmitter keeps a record of each frame sent and maintains a timer for each one of them. The timer is needed in case a frame gets lost or corrupted. A frame is retransmitted if its timer expires before the message has been acknowledged.

With the delayed TCP ACK option, one TCP ACKs is generated to acknowledge a number of TCP frames. This number is a tuneable parameter and in many Linux implementations is set to 3. In a typical transmission, a burst of TCP ACKs will be generated at the receiver side in response to a burst of aggregated TCP data at the transmitter side. Such burst of TCP ACKs will be then forwarded down the protocol stack, aggregated and transmitted as an AMPDU [14]. Compressed TCP ACKs occupy 7 bytes per ACK frames.

Also note that TCP effects beyond the transmission of TCP ACKs are not considered here, and thus the contraction of the TCP transmit window due to lost/delayed UL TCP ACKs, and the slow start effect are not accounted for.

V. Traffic Modeling

In this work, we analyse the performance of VoIP and FTP services in enterprise scenarios. It is assumed that a fraction of the UEs carry a VoIP service, while the remaining UEs carry a FTP service. In the following, the modelling of both services is described, as well as how different service types are prioritised at the Wi-Fi AP.

A. Voice over IP

A typical VoIP call is modelled by two states: talking and listening. In the talking state, the UE transmits packets, while in the listening state, the UE receives packets. A two state Markov chain can be used to model this process. The probability of moving from the talking to the listening state equals to 0.5. When the UE is in the talking state, packets are generated at regular intervals with a fix size. The values of these parameters depend on the voice codecs and compression scheme. In adaptive multi rate (AMR) audio codecs, AMR packets are generated at a rate of 20 ms with a payload of 40 bytes. When the user is in the listening state, a comfort noise is generated in order to avoid the confusion of the user between the listening and disconnected states. These packets are generated at a rate of 160 ms with a payload of 15 bytes. user datagram protocol (UDP) packets are used to encapsulate VoIP traffic, which do not require acknowledgement. For more details on this VoIP model, please refer to [15].

B. FTP

FTP is a best effort service. A typical FTP session consists in downloading several files of a given size, one after the other with a reading time in between them. The reading time is the time interval between the end of the download of the previous file and the user request for the next file, and follows an exponential distribution with a given mean. The file size and the reading time define the load per UE. TCP packets are used to encapsulate FTP traffic, which require TCP acknowledgement, as indicated in the previous Section. For more details on this FTP model, please refer to [11].

C. Service Type Prioritisation

Some enterprise Wi-Fi AP support QoS differentiation, prioritising traffic for different application requirements. In this work, we assume that Wi-Fi APs are equipped with 2 QoS queues, one for real-time services (e.g., VoIP) and another for best-effort services (e.g., FTP), the former having a strictly higher priority. Data in the lower priority queue is only scheduled if there is no data to be transmitted in the higher priority queue. This strict priority queue model applies to a larger number of queues, and ensures that real-time service are first in line to mitigate delay. Note that since each UE only carries one service in our simulations, traffic prioritisation only applies to APs in the DL but do not apply to UEs in the UL. Moreover, note that since 802.11e QoS features are not considered here, DIFS times are equal for all service types.

VI. Performance Evaluation

In this section, simulation results are presented to validate the performance of *Boost* in terms of VoIP and FTP capacity. The performance evaluation is conducted over an enterprise scenario of 50 m$\times$120 m, where there is a *Boost* small cell BS located at the centre of it to handle diverted *Boost* UL traffic, and several Wi-Fi APs are deployed within the enterprise. There are no coverage holds in the enterprise. Simulation assumptions in terms of BS and UE deployment as well as antenna gain, path loss, shadowing and multi-path fading modelling follow the 3GPP and LTE-U Forum recommendations [16]. Since the focus is only on the performance of the unlicensed carrier frequency, LTE performance is not presented in this paper. The assumption is that there is enough UL bandwidth to accommodate the diverted traffic. 100 simulation drops are performed and in each drop 10 seconds are simulated.

Wi-Fi AP deployment: 2 Wi-Fi channels of 20 MHz in the 5 GHz band are considered, and 2 or 4 APs are deployed in the enterprise, as indicated in [16], where the inter-AP distance is 60 m and 30m, respectively. Each AP has a transmit power of 24 dBm, and selects upon deployment the channel in which the least load and interference is observed. Two omnidirectional antennas with a 5 dBi gain are considered.

UE deployment: 4, 8, 16, 24 or 32 UE are uniformly deployed within the enterprise, where the minimum AP-to-UE distance is 3 m. Each UE has a transmit power of 18 dBm, and associates to the AP with the strongest pilot, provided that the pilot was detected at or above -82 dBm in the 20 MHz channel. Two omnidirectional antennas with a 0 dBi gain are considered, thus allowing 2$\times$2 MIMO transmissions.

Services: 30 % of the UEs use a bidirectional VoIP service, while the rest use a bidirectional FTP service. The FTP file size is 0.5 Mbytes in the DL and 0.25 Mbytes in the UL, while the mean reading time is 0.1 s (leading to a high demand of 40 Mbps and 20 Mbps per UE in DL and UL respectively).

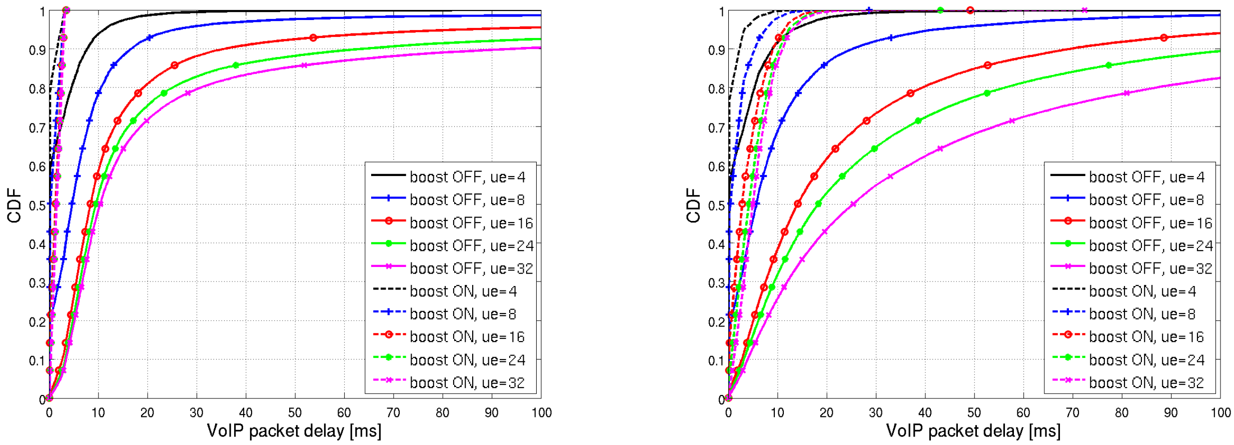

(a) 2 Wi-Fi APs.

(b) 4 Wi-Fi APs.

Fig. 4: CDF of the DL and UL VoIP packet delays.

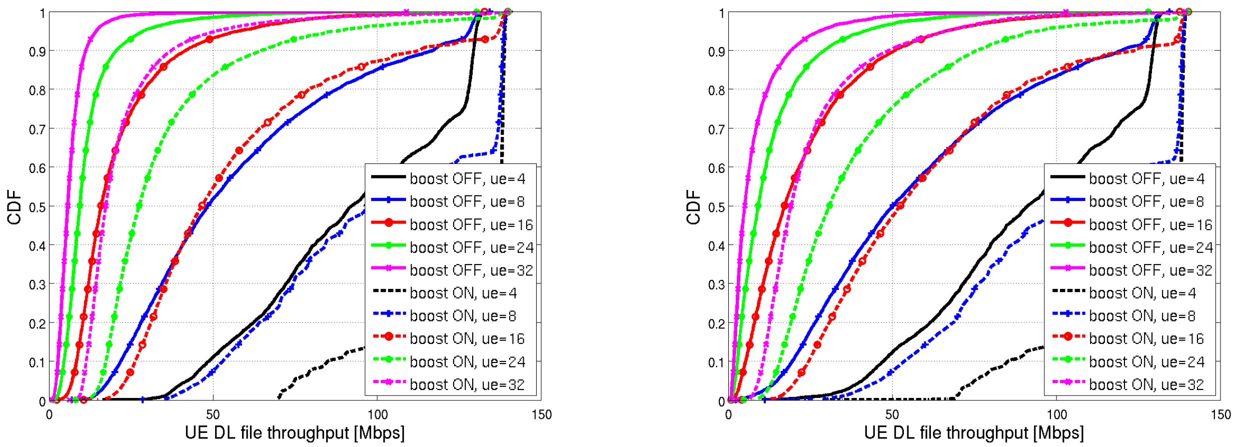

(a) 2 Wi-Fi APs.

(b) 4 Wi-Fi APs.

Fig. 5: CDF of the DL FTP file throughput.

Note that TCP ACKs are generated in response to FTP traffic, where 1 TCP ACKs is sent for every 3 TCP data packets.

Other relevant parameters are set as follows: DIFS= $34\,\mu$s, SIFS= $16\,\mu$s, time slot $= 9\,\mu$s, TXOP=3 ms.

Two system configurations are considered, where traffic prioritisation is enabled in both cases (see Section V-C):

- *Boost OFF*: All traffic is routed through the Wi-Fi APs, this including DL and UL VoIP traffic, DL and UL FTP traffic and the corresponding DL and UL TCP ACKs.
- *Boost ON*: In this case, traffic is split according to the discussion in Section III. DL VoIP and DL FTP traffic are routed over WiFi, while UL VoIP and UL FTP traffic and UL TCP ACKs are routed over LTE. As explained before, WiFi UL MAC ACKs remain in the WiFi network.

A. VoIP Performance

Fig. 4 shows the CDF of the delay that the successfully received DL and UL VoIP packets suffer with and without *Boost* for the scenario with 2 and 4 Wi-Fi APs. The delay of a packet is measured as the time interval between the arrival of the VoIP packet at the buffer and the decoding of the last bit of such packet at the receiver. Regardless the Wi-Fi AP number, Fig. 4 shows that *i)* the VoIP packet delay increases with the number of UEs in the scenario, with and without *Boost*, since the air time per UE decreases; and that *ii)* there is a significant reduction of the VoIP packet delay when *Boost* is activated due to offload of UL traffic from the unlicensed band to the licensed band, and the resulting increased air time per UE. For the 32 UEs case, the median VoIP packet delay reduces by 85 % when activating *Boost* and having 2 Wi-Fi APs in the enterprise; 78 % when having 4 Wi-Fi APs. The reasons for this gain are explained in more detail in the following:

- When *Boost* is not activated, contention and collisions have a major impact in VoIP packet delay. For example, in the 2 Wi-Fi APs case, there are 16 UEs per Wi-Fi AP and channel, and thus the UL traffic of each VoIP UEs only has access to the medium once every 17 opportunities in average. Moreover, FTP UEs introduce a large delay to VoIP UEs since the formers tend to have data to transmit and thus grab the channel for TXOP=3 ms in this case. The case is even worst for the DL traffic since it has to be routed through the AP, and thus the DL traffic of each VoIP UEs only has access to the medium once every $17\times0.3\times16=81.6$ opportunities in average (where 0.3 is the fraction of VoIP UEs, which have priority over FTP UEs). This contention poses a large delay inVoIP

TABLE I: Percentage of outage VoIP UEs [%].

		UEs				
		4 UE	**8 UEs**	**16 UEs**	**24 UEs**	**32 UEs**
2 WiFi APs	**no Boost**	3.0	24.3	66.6	84.5	90.7
	Boost	0	0	0	0	0
4 WiFi APs	**no Boost**	3.0	36.6	81.4	91.4	96.8
	Boost	0	0	0	0	2

packets. Moreover, the inefficiencies of CSMA/CA also lead to a significant number of collisions in these high load scenarios, with the resulting back-off and extra delay on VoIP packets. The capture effect significantly affects cell-edge UEs too.

- When *Boost* is activated, contention is avoided since all UL traffic is diverted to the licensed band, and thus all traffic in the unlicensed band is DL traffic and scheduled by the Wi-Fi AP. As a result, since VoIP packets have priority over FTP packets, they quickly go down the pipe, incurring very little delay.

Also note that going from 2 to 4 Wi-Fi APs in the enterprise increases VoIP packet delay, as shown in Fig. 4. Since all devices within the enterprise are within the energy detection/virtual carrier sense range of each other, deploying one more AP per channel does not increase spatial reuse due to these carrier sensing mechanisms. (one device defers access to the channel if the channel is busy) but results in an increased probability of collision, since two or more UEs may try to access the two Wi-Fi APs simultaneously, and in turn in larger delays. The throughput gain due to the shorter distance among transmitter and receiver cannot be exploited by VoIP traffic because the payload of the packet is too small.

Table I shows the percentage of VoIP UEs in outage in the scenario for the 2 and 4 APs case, where the VoIP service of a UEs is said to be in outage if 2 % of its packets or more have a delay larger than 50 ms [15].

- When *Boost* is not activated, the large delays presented earlier lead to a high number of outages, which significantly degrade VoIP capacity. For the 32 UEs case, 90.7 % of the VoIP UEs are in outage when having 2 Wi-Fi AP in the enterprise; 96.8 % when having 4 Wi-Fi APs.
- When *Boost* is activated, outages mostly disappear due to the very low packet delay.

B. FTP Performance

Fig. 5 shows the CDF of the throughput of the successfully received DL FTP files with and without *Boost* for the scenario with 2 and 4 Wi-Fi APs. As can be seen, *Boost* not only enhances VoIP capacity but also significantly improves FTP performance. For the 32 UEs case, the median throughput increases by 3x when having 2 Wi-Fi APs in the enterprise; 3.5x when having 4 Wi-Fi APs. The reason for such performance increase is again the diversion of the UL traffic to the LTE interface, and the resulting increased air time per UE as well as avoidance of contention and collisions at the Wi-Fi devices, which leads to an all scheduled DL traffic.

Also note that going from 2 to 4 Wi-Fi APs in the enterprise have a different impact in FTP file throughput depending on the traffic load. In this particular scenario, when the number of UEs is low, since collisions are few and UEs are closer to their serving Wi-Fi APs, then the file throughput increases with the more Wi-Fi APs per channel. However, when the number of UEs is high, since contention is also high and collisions are more, then the 5 %-tile and 50 %-tile file throughput decrease with the more WiFi-APs per channel. In contrast, the 95 %-tile file throughput still increases with the more Wi-Fi APs when the number of UEs is high since cell-centre UEs are nearer to the Wi-Fi AP and they tend to capture the cell.

VII. CONCLUSIONS

In this paper, we have discussed Wi-Fi's poor performance in deployments with a large number of devices, and presented how *Boost* can be used to address this issue. The concept, architecture and benefits of *Boost* have been discussed, along with a detailed description of the Wi-Fi system-level simulator used to analyse its performance. Simulations results have shown that by diverting Wi-Fi uplink traffic through the LTE network and avoiding Wi-Fi contention, *Boost* can significantly reduce VoIP delay by up to 85 % in the most challenging scenarios, while also significantly improving FTP throughput up to 3.5x. This positions *Boost* as a strong candidate for providing wireless services in the future enterprises. Our future work will analyse the impact of offloading UL traffic from the unlicensed to the licensed band in other LTE traffic, and the interactions of Boost with TCPs flow control.

REFERENCES

[1] D. Goldman, "Why Wi-Fi sucked at Mobile World Congress," CNN, Feb. 2011.

[2] R. Madan, A. Sampath, and N. Khude, "Enhancing 802.11 Carrier Sense for High Throughput and QoS in Dense User Settings," in *IEEE Personal Indoor and Mobile Radio Conference (PIMRC)*, Sydney, Australia, Sep. 2012.

[3] A. Flores and E. Knightly, "Dual Wi-Fi: Dual Channel Wi-Fi for Congested WLANs with Asymmetric Traffic Loads," IEEE 802.11-13/1409r0, Tech. Rep., Nov. 2013.

[4] Huawei, "White Paper on Spectrum," Feb. 2013.

[5] 3GPP, "Access Network Discovery and Selection Function (ANDSF) Management Object (MO)," 3GPP TS 24.312 v.12.10.0, Sep. 2015.

[6] Alcatel-Lucent. (2015, Mar.) Wireless Unified Networks. [Online]. Available: https://www.alcatel-lucent.com/press/2015/alcatel-lucent-combines-best-wi-fir-and-lte-enhance-mobile-performance-and-offer-consistent-high

[7] J. Ling, S. Kanugovi, S. Vasudevan, and A. K. Pramod, "Enhanced capacity & coverage by Wi-Fi LTE Integration," *IEEE Comm. Mag.*, vol. 53, no. 3, pp. 165–171, Mar. 2014.

[8] E. Perahia and R. Stacey, *Next Generation Wireless LANs: 802.11n and 802.11ac*, 2nd ed. Cambridge University Press, Jun. 2013.

[9] J. Ling, S. Kanugovi, O. Marcé, and S. Vasudevan, "Performance Gains of a Hybrid Wi-Fi/LTE Architecture," in *IEEE Vehicular Technology Conference (VTC Spring)*, Glasgow, United Kingdom, May 2015.

[10] X. Chu, D. López-Pérez, Y. Yang, and F. Gunnarsson, Eds., *Heterogeneous Cellular Networks: Theory, Simulation and Deployment*, 1st ed. Cambridge University Press, May 20013.

[11] 3GPP, "Evolved Universal Terrestial Radio Access; Further Enhancements for E-UTRA physical layer aspects," 3GPP TR 36.872 v.9.0.0, Mar. 2010.

[12] ——, "Small Cell Enhacements for E-UTRA and E-UTRAN," 3GPP TR 36.872 v.12.1.0, Dec. 2013.

[13] D. Xia, J. Hart, and Q. Fu, "Evaluation of the Minstrel rate adaptation algorithm in IEEE 802.11g WLANs," in *IEEE International Conference on Communications (ICC)*, Budapest, Hungary, Jun. 2013.

[14] M. Heusse, S. A. Merritt, T. X. Brown, and A. Duda, "Two Way TCP Connections: Old Problem, New Insight," in *ACM SIGCOM 2011*, Apr. 2011.

[15] R. Irmer, "NGMN Radio Access Performance Evaluation Methadology," NGMN Alliance, Tech. Rep. Jan., 2008.

[16] 3GPP, "Feasibility Study on Licensed-Assisted Access to Unlicensed Spectrum," 3GPP TR 36.889 v.13.0.0, Jul. 2015.