DeMiLTE: Detecting and Mitigating LTE Interference for Enterprise Wi-Fi in 5 GHz

Swetank Kumar Saha*, Christina Vlachou+, Dimitrios Koutsonikolas*, Kyu-Han Kim+
*University at Buffalo, SUNY, Buffalo, NY, USA, +Hewlett Packard Labs, Palo Alto, USA
{swetankk,dimitrio}@buffalo.edu,{christina.vlachou,kyu-han.kim}@hpe.com

ABSTRACT

LTE in unlicensed 5 GHz bands is being deployed by mobile operators for increased capacity. In this paper, we conduct an extensive measurement study with commodity LTE and Wi-Fi hardware to identify key coexistence challenges. Our study – the first to include a commercial LAA base station – confirms that LTE interference causes Wi-Fi performance to degrade, harming 802.11ac high-throughput features. We then present DeMiLTE – a system for commodity enterprise WiFi APs that detects, quantifies, and reacts to LTE interference. To our best knowledge, our solution is the first that achieves fair coexistence without modifying the LTE PHY/MAC, while still being fully-compliant to the 802.11ac standard. DeMiLTE's architecture is based on lightweight per-link interference detection and enables Wi-Fi APs to mitigate LTE-induced performance degradation with minimal overhead. Our evaluation results show that DeMiLTE can provide up to 110% throughput gains and alleviate client disruption caused by LTE interference.

CCS CONCEPTS

• **Networks → Mobile networks**.

ACM Reference Format:
Swetank Kumar Saha*, Christina Vlachou+, Dimitrios Koutsonikolas*, Kyu-Han Kim+. 2019. DeMiLTE: Detecting and Mitigating LTE Interference for Enterprise Wi-Fi in 5 GHz. In *The Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing (Mobihoc '19), July 2–5, 2019, Catania, Italy.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3323679.3326501

1 INTRODUCTION

LTE in the 5 GHz unlicensed spectrum, namely LTE-U and LAA technologies, is expected to grow in the future as licensed spectrum becomes increasingly loaded. For cellular operators, this offers a desirable solution enabling them to harness the "free" unlicensed bands along with a primary licensed cell using carrier aggregation. T-Mobile and Verizon have already deployed LTE unlicensed solutions by leveraging commodity end-user compliant devices [8, 9]. LTE-U/LAA small cells are very likely to become ubiquitous in environments such as enterprises, stadiums, auditoriums, campuses.

Mobihoc '19, July 2–5, 2019, Catania, Italy

ACM ISBN 978-1-4503-6764-6/19/07…$15.00
https://doi.org/10.1145/3323679.3326501

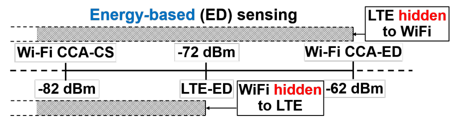

Figure 1: Illustration of LAA or Wi-Fi hidden nodes, due to heterogeneity in CCA.

LTE, however, creates a fundamental challenge for coexistence with existing technologies in these bands, mainly 802.11ac based WLANs. LTE's centralized control architecture and frequency/time-division MAC makes it non-trivial to coexist with Wi-Fi, which utilizes randomized medium access (CSMA/CA). Consequently, both LTE-U and LAA have MAC schemes for achieving fair coexistence. While LTE-U uses a duty cycling approach to allocate some airtime to Wi-Fi, LAA uses Listen-Before-Talk (LBT), similar to CSMA/CA.

Prior coexistence solutions [13, 16] have been shown to often fall short of fair resource allocation on throughput or latency. The underlying reason is the inherent *deafness* (inability to decode) of Wi-Fi and LTE-U/LAA to each other's frames. This leaves energy based sensing as the only available mechanism to enable the two technologies to become aware of each other's presence. Fig. 1 shows the different sensitivity thresholds of LTE and Wi-Fi for clear channel assessment (CCA), and potential hidden-terminal issues due to the interplay among these thresholds. Although LAA's LBT mechanism will lead to better coexistence between the two technologies, LBT still only relies on a static energy detection threshold (LTE ED in Fig. 1) of -72 dBm to detect Wi-Fi transmissions. In a typical WLAN where multiple APs with different transmission-power settings will coexist with LTE-U/LAA, there will exist scenarios where LBT cannot detect Wi-Fi's frames, leading to collisions at clients where LAA transmissions are strong enough to cause interference.

To address LTE hidden terminals, Chai et al. [12] have proposed a modification of the LTE PHY layer, which embeds a Wi-Fi preamble header in LTE frames. Although this solution does comply with the LTE PHY standard, it requires major changes to the LTE PHY implementation to place the required symbols (for a Wi-Fi CTS-to-self) inside LTE transport blocks and, consequently, modifications to hardware and firmware of each LTE small-cell. Moreover, there may be cases (exposed terminal scenarios) where an LTE-generated CTS might unnecessarily suppress data transmission to clients that can decode frames even under simultaneous LTE transmission. Actually, in our experiments, we find that LTE interference to clients displays a whole range from minimal impact to complete disruption.

In this work, we take an entirely different approach from prior work and focus on Wi-Fi based coexistence solutions. Since a Wi-Fi AP has a much more detailed view of the LTE-induced performance degradation at different Wi-Fi clients, it is better equipped to decide the appropriate reaction to LTE's presence, compared to an LTE base station. As a first step, we conduct measurements with a large number of Wi-Fi clients under all possible scenarios depicted in Fig. 1 to understand how diverse clients are affected by LTE interference and the underlying reasons of performance degradation. For the first time, we use commodity LTE-U/LAA hardware to shed light on LTE-induced performance degradation on crucial 802.11ac Wi-Fi features, such as bandwidth and rate control. We find that in scenarios where LAA transmissions are below the Wi-Fi ED threshold, Wi-Fi clients suffer the worst degradation, often experiencing complete disconnection. Based on our observations, we design *DeMiLTE*, a per-client LTE-detection system that does not require any changes to LTE small cells and enables commodity Wi-Fi APs to detect, quantify, and react to LTE interference in an *802.11-standard compliant* way, to ensure their fair share of the bandwidth.

In summary, our work makes the following contributions:
(1) We experimentally evaluate the impact of LTE on Wi-Fi performance (§3) across different clients and device hardware. To our knowledge, this is the first investigation of LTE-induced degradation using commodity LTE/Wi-Fi hardware, including commercial LAA. We conduct our measurements in a dense enterprise Wi-Fi environment with active WLAN users and human mobility.
(2) We design *DeMiLTE*, the first **Wi-Fi based coexistence solution** that boosts performance for clients experiencing degradation due to LTE without affecting clients that do not suffer (§4,5.1).
(3) We present a lightweight detection strategy that can run on any commodity enterprise Wi-Fi AP (relying on readily available loss metrics) and diagnose LTE-induced interference for each client (§4).
(4) We implement and evaluate *DeMiLTE* in commodity APs, and show its performance gains over the default Wi-Fi rate control (§5.2).

2 BACKGROUND

LTE activates a 5 GHz channel only when LTE downlink traffic demand is high, with a vendor-specific channel selection algorithm. In this work, we assume that LTE and Wi-Fi share the same channel, as typically dense WLANs have few available channels.
LTE-U is the first protocol for LTE in 5 GHz, formed by a group of enterprises [4]. It does not respect Wi-Fi transmissions, as it (de)activates the unlicensed cell at predefined intervals with a transmission duration of up to 20 ms. Its duty cycle is adapted to sensed Wi-Fi channel utilization by a carrier sense adaptive transmission (CSAT) algorithm. We study Wi-Fi performance degradation as a function of the CSAT duty cycle.
LAA is the LTE in unlicensed standard by 3GPP [1]. Contrary to LTE-U, LAA implements listen-before-talk (LBT) in 5 GHz, which is similar to Wi-Fi CSMA/CA. LAA's parameters, such as energy thresholds and contention-window values, are similar to those of Wi-Fi for fair coexistence. However, *LAA does not enforce Wi-Fi preamble detections or control packet transmissions* (e.g., RTS/CTS), causing severe hidden terminal issues. LAA deploys different parameters based on the country of operation (e.g., maximum transmission duration TXOP is 4 ms for Japan vs. 10 ms for USA).
802.11ac Wi-Fi: To enable Gbps rates, IEEE 802.11ac uses various very-high-throughput (VHT) features. APs dynamically adapt VHT features based on packet error rate (PER) which can be affected by aggressive collisions. We explore how VHT features deteriorate due to LTE interference and how to reverse LTE effects on VHT.
Rate Control: Commodity APs use rate adaptation (RA) algorithms similar to Minstrel [5]. RA selects bandwidth, number of spatial streams (SS), and modulation and coding scheme (MCS) index, which all determine the transmission rate. RA assumes monotonicity of PER vs. PHY rates and can be affected by collisions or hidden terminals. Our AP enables RTS/CTS upon a packet failure for the first retransmission. RTS/CTS can *help under hidden Wi-Fi terminals but not LTE ones.* RA also determines the frame aggregation limit, i.e., the number of sub-frames to be included into a single frame transmission (nFrames). We measure the sub-frame error rate (SFER) – the fraction of corrupted sub-frames in a single frame transmission. PER computations in the firmware are based on SFER.
Dynamic Bandwidth: 802.11ac enables dynamic bandwidth adaptation by bonding multiple 20-MHz channels for 40/80/160 MHz operations. Bandwidth control can fall back to smaller bandwidth, if one of the secondary 20 MHz channels is sensed occupied by hardware or if PER increases in the firmware RA.

3 DIAGNOSING COEXISTENCE ISSUES

We conduct an extensive experimental study to quantify the impact of LTE unlicensed on Wi-Fi. To our knowledge, this is the first measurement over a diverse environment of more than 10 clients. Previous studies [25] have not investigated dense WLAN issues, i.e. diversity in client performance degradation, and have not included LAA protocols. They also have not delved into the underlying issues (e.g., rate control's reaction to collisions) that cause Wi-Fi performance degradation.

3.1 Experimental Setup

LTE: We use two commodity LTE femtocells from different vendors, which implement the full L1–L3 LTE stacks (i.e. PHY to RRC layers) and are compliant with 3GPP Releases 12-13. They support 2x2 MIMO with up to 20 MHz channels yielding 150 Mbps peak rates on downlink. Their transmission power is set to 13dBm. Our femtocells support LTE-U and LAA protocols. Our LTE client, Samsung Galaxy Note8, supports secondary cells in 5 GHz. The 4G eNodeB (eNb) itself needs an evolved packet core (EPC) to operate. Our testbed uses OpenEPC [6] with the full stack running on virtual machines.
Wi-Fi: We use commodity Wi-Fi APs, equipped with 4x4 MIMO-capable 802.11ac radios. Our AP supports up to 80 MHz channel bandwidth and up to 256-QAM modulation level, with 1733.3 Mbps peak PHY rate. For our experiments, we have modified the firmware of our AP to collect per-client PHY rates (MCS, number of SS), channel bandwidth, and SFER. Our Wi-Fi clients are commercial 802.11ac mobile phones and laptops with Qualcomm, Broadcom, and Intel chipsets. Although we conducted experiments with several client devices, in the interest of space, only the measurements with Samsung Galaxy S6 are presented unless the results differ across devices. Unless stated otherwise, all throughput values reported in the paper refer to downlink (AP to client) TCP throughput.
Topology: We select multiple locations for AP and clients in an office area of 84 ft x 91 ft. Fig. 2 shows the floor plan used for our measurement study with the positions of LTE femtocells (eNB), LTE client (L1), Wi-Fi APs (AP1–AP4) and 9 Wi-Fi client locations (C1–C9) marked. The grey rectangles shown in the floor plan are office cubicles. AP positions are chosen such that they create two scenarios based on LTE signal's received power at the Wi-Fi AP: (*i*) Above ED threshold (AP1) and (*ii*) Below ED threshold (AP2–4).

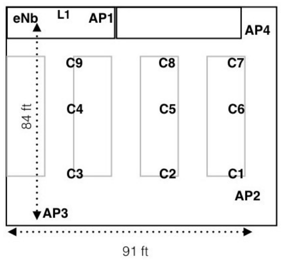

Figure 2: Testbed plan: eNb, APs and clients.

3.2 LTE Above Wi-Fi ED

When LTE has a sufficient transmission power such that its received signal strength is above the Wi-Fi energy-detection level, one would expect 802.11 based APs, which rely on CSMA/CA for channel access, to be able to detect the LTE transmissions and defer their transmissions during LTE operation. However, 802.11n/ac APs' use of channel bonding can create scenarios where, depending on whether LTE operates on the primary or secondary Wi-Fi channel, Wi-Fi may be unable to fully sense LTE transmissions even though the AP and the femtocell are located within each others sensing range. For instance, 802.11ac performs the CCA operation for a much shorter time on the secondary channels (when using a bonded 40/80 MHz channel) as compared to the primary channel and the secondary channels do *not* support decoding preambles of the detected packets and setting of the NAV (Network Allocation Vector) accordingly [22, 27]. Further, CCA thresholds are different for primary and secondary channels, resulting in diverse behaviors to the same LTE received power. Optionally, 802.11ac defines an enhanced RTS/CTS mechanism per 20 MHz in addition to CCA on secondary channels. However, if the client is far away and not able to listen to LTE transmissions, client can falsely send CTS and LTE can still create *hidden terminal conditions*. Hence, to fully characterize the coexistence issues when using 802.11's bonded channels, we consider two sub-cases based on LTE's operating channel: (*i*) LTE (20 MHz) on the Wi-Fi AP's (80 MHz) primary channel (*ii*) LTE (20 MHz) on one of the Wi-Fi AP's (80 MHz) secondary channels.

3.2.1 LTE on Primary Wi-Fi Channel. To fully capture LTE's effect on different clients in a WLAN, we consider AP1 and clients C1–C9, with C9 having the best channel quality to AP1 while C1 the worst. We repeat the experiments with both LAA and LTE-U, expecting a fair share of the channel. LAA has saturated downlink traffic and we run tests with two different TXOP, 4 and 8 ms. LTE-U duty cycle is set to 50% to emulate a fair protocol.

In our tests (omitted for brevity), we observe that clients C1–C9 suffer 40-60% throughput reduction when LTE-U is ON, as expected. The Wi-Fi AP is able to utilize its fair share of the channel capacity, given that LTE-U is occupying the other 50%. Throughput reduction of higher than 50% (60% max.) can be attributed to collisions and retransmission overheads. In fact, we find the PHY rate and SFER are not adversely affected by LTE interference. Similarly, LAA does not harm VHT features thanks to the LBT protocol. However, we find that long LAA TXOP durations, can lead to unfairness to Wi-Fi: with 8 ms LAA TXOP, Wi-Fi has a throughput reduction of 74%, because Wi-Fi typically uses TXOP smaller than 4 ms and it has to wait for long LAA transmissions when deferring with CSMA/CA. In contrast, we find that Wi-Fi throughput reduction is 34% with 4 ms LAA TXOP, and Wi-Fi gets more than 50% of the airtime.

Finding: LTE on primary Wi-Fi channel does not harm VHT features if transmission power triggers ED. Low SINR clients might have a slight performance degradation due to collisions. We recommend no Wi-Fi reaction when LTE is detected on primary channel. LAA could lead to higher throughput reduction than LTE-U and LAA should deploy 4 ms TXOP for fair coexistence with Wi-Fi.

3.2.2 LTE on Secondary Wi-Fi Channel. Here, we use a similar setup as described in §3.2.1 but with the LTE operating on a secondary channel (20 MHz, channel 44) of the Wi-Fi AP's 80 MHz bonded channel (primary channel 36). Wi-Fi AP is expected to use 40 MHz when LTE is active on channel 44. We specifically focus on the performance of two clients, C9 and C7 (high and low SNR).

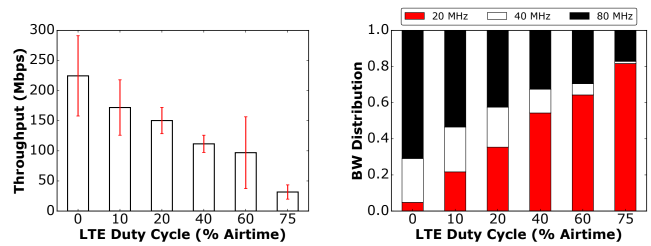

(a) Throughput (C7) **(b) Bandwidth (C7)**

Figure 3: Above ED: LTE-U on secondary channel.

Client C9 achieves a throughput of ~525 Mbps, when LTE-U is OFF[1]. When LTE-U operates with duty cycle 10%, 20%, or 40%, the Wi-Fi throughput reduction (from the LTE-U OFF value) is indeed in proportion indicating that Wi-Fi is able to gain access to the channel when LTE-U does not transmit. We observe that as the LTE-U duty cycle increases, the proportion of packets sent at 40 MHz bandwidth also increases and the number of packets transmitted over the full 80 MHz bandwidth decreases. For example, with 75% LTE-U duty cycle, 40 MHz is used for more than 80% of the time. Hence, high SNR clients are not affected by LTE.

Next, we look at the performance of client C7 which has a bad link quality (compared to C9) under the same setup (Fig. 3a). In contrast to C9, we observe higher LTE-U duty cycle operation causing severe Wi-Fi throughput degradation, significantly higher than expected after accounting for LTE-U's airtime usage. For example, under 75% LTE-U operation, the throughput reduction is much higher than the expected value of 25% of the baseline 80 MHz value. Fig. 3b illustrates the reason behind the lower than expected performance. In contrast to C9, where bandwidth was reduced to 40 MHz, here the AP instead resorts to 20 MHz bandwidth most of the time, especially under higher LTE-U duty cycle. We also observe that SFER increases with increasing LTE-U duty cycle, reaching 40% at 75% duty cycle, as opposed to the case for C9, where SFER always stays below 10% (plot omitted). Higher SFER results in rate-control not only selecting lower MCS values (e.g., MCS 0 under 75% LTE-U) most of the time, but also dropping the number of SS to 1. An associated effect is a drop in nFrames, which results in a less efficient MAC. All of these factors combined contribute to a low data rate for C7.

Prior works [22, 27] have shown similar results where 80 MHz links sharing one of their secondary 20 MHz channels with other Wi-Fi APs can be starved.[2] The explanation given was that CCA

[1] We omit LAA results as we find similar performance.

[2] These works have placed the APs at most 16 ft or 5m apart [27], similar to our LTE/Wi-Fi setup.

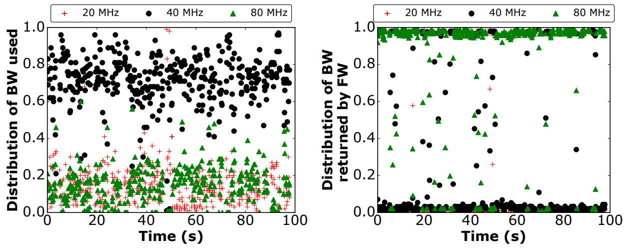

Figure 4: Chosen bandwidth by Wi-Fi AP in hardware (left) and firmware (right), when LTE is on secondary channel.

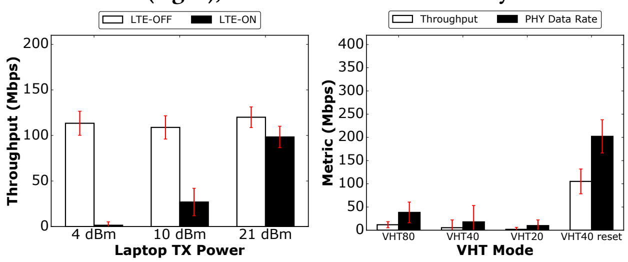

(a) Near-far effect (b) SG Note8

Figure 5: LTE-U on secondary channel, duty cycle 75%: Tests with different clients and VHT modes at location C7.

malfunctions and fails to sense the other AP resulting in collisions. However, if CCA was indeed failing on 80 MHz, the 802.11ac rate control would simply converge to 40 MHz bandwidth (see §2). We believe that, irrespective of LTE or Wi-Fi interferers on secondary channels, the *Wi-Fi AP can indeed sense the interferer contrary to observations by previous works*. To show this, we consider a sample link in our topology with the AP at location C9 (Fig. 2) and client at C7 that suffers from LTE interference. We modify the AP firmware to log the fraction of packets whose maximum bandwidth was chosen to be 20/40/80 MHz by the rate-control (in the firmware) and compare it with the fraction of the bandwidth that was actually used by the hardware for transmission. Fig. 4 shows that the AP is able to sense the medium: the rate-control/firmware almost always returns 80 MHz maximum bandwidth, whereas approximately 80% of packets are transmitted at 40 MHz, due to sensing LTE operating on secondary channel at 75% duty cycle.

Client Diversity: To understand our previous findings, we repeat the experiment with a diverse set of APs and client devices (laptops and smartphones from different vendors) for the AP1-C7 link. All APs show similar issues hence, results are omitted. However, we find that performance varies greatly across client devices. The laptop (HP), Xiaomi phone, and the iPhone8 Plus are not severely affected by LTE on secondary channel, whereas Samsung Galaxy Note8 and 6 are severely affected. After investigating further, we discover that a significant difference in received signal strength (RSS) of transmissions from the clients to the AP is the reason behind the difference in client behavior. Diverse antenna designs and transmission powers yield different RSS values. Note that, although all measurements have downlink traffic, uplink data still exists due to 802.11 Block ACKs and TCP ACKs. Loss of uplink packets for clients with low RSS at the AP causes performance degradation. Essentially, a *near-far effect* is created where LTE on secondary channel causes the weak uplink (e.g., AP1-C7) to be distorted at the analogue front-end of the AP. To prove the near-far effect and the performance impact of uplink SINR, we reduce the TX power of the laptop at C7. In Fig. 5a, we clearly observe that downlink TCP performance strongly depends on the uplink SINR. When LTE is above ED on secondary channel 48, and AP1 operates at channel 36 VHT80 with a laptop at location C7, we find that low uplink SINR starves the downlink.

Reducing Wi-Fi Bandwidth to Combat LTE on Secondary Channel: To explore this potential strategy, we look again at the AP1-C7 link with Note8. We set different bandwidth modes by configuring the maximum bandwidth for the rate-control (in the firmware): VHT40 and VHT20 for 40 MHz and 20 MHz, respectively. As we observe in Fig. 5b, *decreasing bandwidth in firmware does not help* in this case. Note that ideally by reducing our bandwidth to 40/20 MHz, we move into a separate collision domain from LTE and should not experience any interference. Interestingly, we find that resetting the radio interface after setting the VHT mode[3] indeed helps to alleviate performance degradation (marked as VHT40 reset in the figure). We suspect that *dynamically setting VHT modes, although it sets up the downlink bandwidth appropriately, it does not actually alter the receiver spectral-mask and the AP continues to receive over the entire 80 MHz channel.* Hence, strong LTE, operating on any one of the secondary channels, creates strong interference and severely harms the uplink causing almost complete starvation.

Finding: LTE on secondary Wi-Fi channel and at most 16 ft away from AP (above ED) can completely starve low SINR clients, due to near-far effects that completely starve the uplink, hence preventing block ACKs and other control messages from being transmitted. LAA and LTE-U have similar effects. We recommend that LTE eNbs be placed at large distance ($> 20ft$) from APs or that eNbs avoid secondary Wi-Fi channels when the Wi-Fi beacon RSSI is higher than -50dBm. In addition, as a reaction solution we recommend Wi-Fi AP to change either channel or VHT mode. As this can cause short performance disruption due to client re-association, we propose this reaction to be triggered only if the majority of clients suffer.

3.3 LTE Below Wi-Fi ED

We now consider cases where Wi-Fi AP cannot sense LTE transmissions. This is an even more challenging environment for Wi-Fi as it is unaware of LTE's presence and does not have any mechanism to achieve co-existence. We set the LTE-U duty cycle to 50% for all following experiments and use AP locations AP2–4, where LTE signal strength is below AP's ED threshold. In case of LAA, we set the frame TXOP to 4 ms and then experiment with both backlogged and a limited source traffic rate (that corresponds to LAA using 50% airtime for comparison with LTE-U).

3.3.1 LTE on Primary Wi-Fi Channel. First, we run tests on 20 MHz Wi-Fi links in order to isolate the MCS and SS control from bandwidth adaptation. We use the topology shown in Fig. 2 and consider three links: AP2-C9, AP3-C6, and AP4-C3.

Fig. 6a shows the TCP throughput for each of three links under three cases: LTE-U OFF, LTE-U ON, and for comparison a case where the LTE-U eNB is replaced by another Wi-Fi AP. For all links, Wi-Fi throughput degradation is extremely severe considering LTE-U operating at 50%. For example, the throughput of the AP4-C3 link under LTE-U is ~5 Mbps – a 93% decrease from link's baseline

[3]A majority of the APs need to re-associate all clients after modifying channel and VHT mode and resetting radio.

throughput of ~78 Mbps. Further, all links perform significantly better when they share the channel with another Wi-Fi AP (at the same location and with same transmission power as the eNB), indicating that LTE-U's presence *unfairly* (compared to a another fair Wi-Fi AP) does not allow Wi-Fi to fully utilize its share of the channel capacity. The performance difference between LTE-U and Wi-Fi interference is due to Wi-Fi's CCA thresholds for non-Wi-Fi and Wi-Fi interferers: the AP defers at LTE signal strength larger than -62dBm vs. -82dBm for a Wi-Fi interferer (see Fig. 1). Hence, *LTE-U acts as a hidden terminal whereas a second Wi-Fi AP does not.*

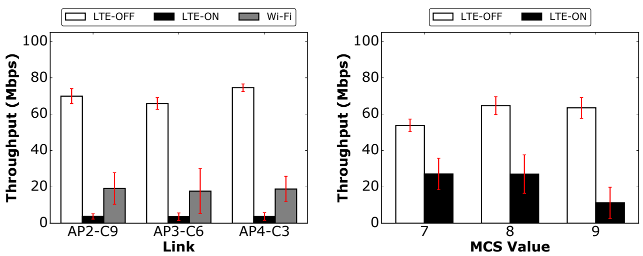

(a) Default rate control (C9, C6, C3) **(b) Static MCS (C9)**

Figure 6: Below ED: AP2-C9, AP3-C6 and AP4-C3.

To further understand the cause behind the lower than expected Wi-Fi throughput under LTE-U, we look at the link's PHY/MAC metrics. In the interest of space, we focus on the AP2-C9 link, however the observations are similar for the other two links. We find that the Wi-Fi rate-control under LTE-U drops the MCS to 0/1 more than 60% of the time and resorts to using MCS 2/3/4 in the rest of the time; whereas it only uses MCS 8/9 when Wi-Fi is operating alone. Throughput is further hurt by the reduction in size of the aggregated frames to only 2 subframes from 12 (in the median case). Both MCS and size of aggregated frames are also affected when competing with another Wi-Fi AP, however the magnitude in reduction is much less resulting in better performance.

Nonetheless, using a lower MCS does not help to deal with LTE-U's presence as the losses are not channel-induced but a result of collisions due to the inability of the two technologies to sense each other's presence. To further confirm our observations, we modify the rate-control to statically fix the MCS to a given value, irrespective of losses, for the AP2-C9 link and see if Wi-Fi performance improves. Fig. 6b shows the throughput for three static MCS cases (7, 8, and 9) with and without LTE. We observe that with both MCS 7 and 8, Wi-Fi indeed can achieve better throughput (~30 Mbps) compared to default rate-control (~5 Mbps), which amounts to *6x* gain. In fact, even with LTE-ON, Wi-Fi indeed achieves its fair share of 50% throughput (given LTE-U's usage of 50% airtime) compared to the baseline throughput with LTE-OFF.

Next, we extend our study with AP2, bandwidth 80 MHz, and clients C1-C9 with diverse link SNRs that would exist in a typical enterprise WLAN scenario, to better understand how links of different strength are affected by LTE-U and LAA. Such a large scale study is necessary before we can formulate a mitigation strategy.

LTE-U: Fig. 7a shows that there is large *diversity* in performance degradation caused by LTE to different clients. High SNR clients are barely affected whereas low SNR clients cannot achieve their fair share, given 50% LTE-U. Moreover, we observe the entire range of performance degradation, from clients that are not affected (C1, C6 & C7) to clients that are mildly affected (C2 & C5) to those that are severely affected (C3, C4, C8 & C9). In Fig. 7b, which plots SFER with LTE-ON for each client, we observe that clients that display severe performance degradation suffer from high SFER, causing the link adaptation to drop the MCS, nFrames, etc. to low values resulting in lower throughput. Different clients, owing to differing SINR, suffer diverse levels of packet error rates; thus, their link adaptation mechanisms react in accordance.

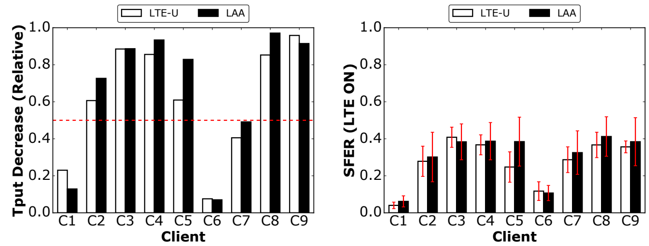

(a) Throughput **(b) SFER**

Figure 7: Below ED: AP2, LTE on primary channel.

LAA: Under saturated/backlogged traffic conditions, we observe that LAA can be *much worse* than LTE-U and in fact completely starves TCP connections (Wi-Fi TCP connection cannot be established). Since LAA cannot sense Wi-Fi transmissions, it ends up occupying 100% of the airtime. LTE-U, on the other hand is always bound by the duty cycle configurations. We then repeat the measurements with a limited source rate (that emulates LTE-U operating at 50% duty cycle) and observe (Fig. 7a) similar diversity, in performance degradation, as with LTE-U. Moreover, under this setting, clients show throughput degradation levels that almost mirror LTE-U. *Our observations are in stark contrast to the popular belief that LAA would be more fair to Wi-Fi. We show that it degrades performance of Wi-Fi clients at least as much as LTE-U and can end up using the entire channel capacity under backlogged traffic.*

Finding: LTE can cause hidden terminals even within a single office space, due to high Wi-Fi ED thresholds. LTE hidden terminals on primary channel cause the most severe performance degradation. VHT features such as MCS, SS, and bandwidth adaptation suffer.

3.3.2 LTE on Secondary Wi-Fi Channel. We set LTE-U on secondary Wi-Fi channel 48 and repeat the above experiments. Wi-Fi is expected to use 40 MHz bandwidth when LTE is active. Fig. 8 presents our results. All clients are able to get at least 90 Mbps TCP throughput. Although some clients are affected by dropping PHY rates at most 50% (e.g. C8), no client is starved. Analyzing the distributions of used bandwidth hints that AP used 40 MHz more than in LTE-OFF case for low SNR clients (C3–5, C8-9).

Finding: LTE hidden terminals on secondary Wi-Fi channels impose less severe Wi-Fi throughput degradation than on primary channel, thanks to the rate control that converges to optimal maximum bandwidth based on PER statistics per bandwidth.

***Conclusions*:** An effective reaction to LTE should encourage rate adaptation to not drop MCS under LTE interference. Moreover, any reaction needs to detect the case when LTE is below ED threshold. Even if we detect LTE below ED threshold at the AP, we cannot trivially react as there are clients (with high SNR) for which we should not modify the regular rate control mechanism. Hence, in addition to detecting LTE presence at the AP, we need to know on a per-client basis whether it is being affected by LTE or not.

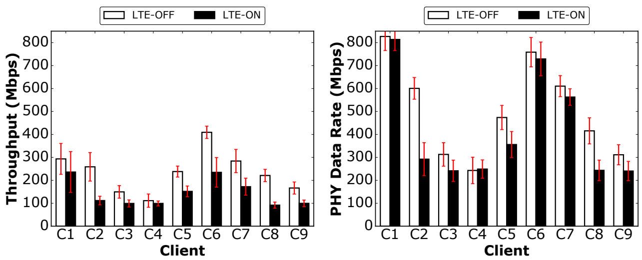

(a) Throughput (b) PHY Data Rate

Figure 8: Below ED: AP2, LTE on secondary channel.

4 *DeMiLTE*: INTERFERENCE DETECTION

Our study indicates that APs can mitigate LTE performance degradation by intelligent link adaptation. To this end, LTE interference and its impact on different Wi-Fi clients need to be detected.

AP-Level Detection: Commodity enterprise-grade APs use spectrum analysis for non-Wi-Fi interference detection [7]. With tens of classifiers, APs can distinguish interferers such as microwave ovens, cordless bases, and frequency hoppers. Having a spectrum monitoring AP is not enough for interference detection and for quantifying the impact on every Wi-Fi link. To overcome this limitation, WiFiNet [20] employs two-radio APs that are used for communication and spectral analysis simultaneously. The dual-radio approach offers high accuracy for per-link interference detection. However, it is far from being practical, as typically network administrators would not double the cost of their deployments just for spectral analysis. LTE interference detection by commodity APs has recently been addressed by LTERadar [25], which uses spectrum analysis. Enterprise WLANs may exploit dedicated spectrum tools to detect LTE, and then inform all APs in the WLAN about LTE presence, channel, and airtime utilization.

While AP-level detection may act as a first step in detecting LTE, our measurements in §3 clearly indicate that for a given AP there is high variation in the impact of LTE on different clients. These observations direct us to the need for a *per-link detection strategy* to avoid harming unaffected clients. AP-level detection tools can always be used in conjunction with *DeMiLTE*, whereby hints from spectrum monitors trigger our system, which then starts a per-link detection for each served client to quantify the impact of LTE.

Client-Level Detection: Building a per-link detection system is challenging as it needs to be lightweight and run in real-time. Further, it should work without requiring additional hardware (e.g., a dedicated monitoring radio per AP), using only metrics available on a typical enterprise Wi-Fi AP. Towards this end, our detection strategy relies on using fine-grained loss metrics directly accessible in the Wi-Fi AP's rate-control module. In essence, we observed that the values of these metrics increase under LTE's presence for clients that are indeed affected severely and need to use *DeMiLTE*'s mitigation system. More importantly, these metrics show little to no increase for scenarios when default rate-control should be allowed to control link parameters for best performance. We first introduce our metrics followed by an evaluation in a realistic WLAN.

4.1 LTE Detection Metrics

We look at several fine-grained loss metrics listed in Table 1, measured over a time interval (100 frames) to detect LTE. These statistics are collected by the rate-control firmware at the Wi-Fi AP for each served client, hence are available without any overheads.

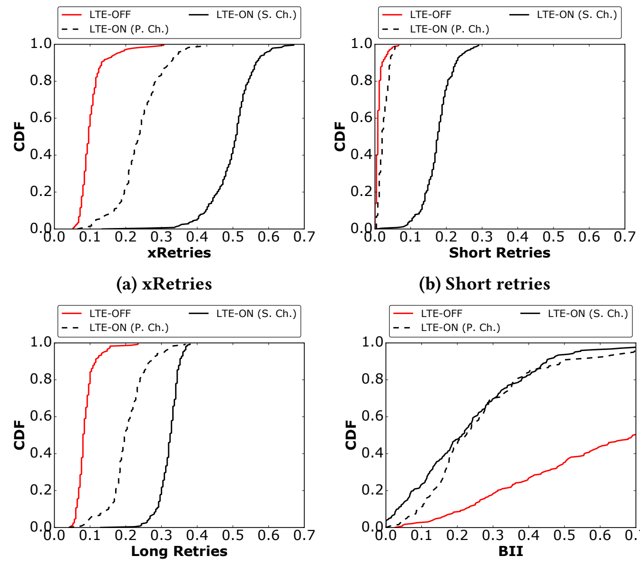

(a) xRetries (b) Short retries

(c) Long Retries (d) Burstiness Index Inverse

Figure 9: LTE-U per-link detection (Above ED): AP1-C7.

Table 1: Detection Metrics

Metric	Description
xRetries (xR)	*Number of times rate-control failed to send a frame even after multiple re-transmissions*
Short Retries (sR)	*Number of retries experienced by short/small packets (e.g., RTS-CTS frames)*
Long Retries (lR)	*Number of retries experienced by regular/large packets (e.g., data frames)*
Burstiness Index Inverse (BII)	*Indicates number of consecutive sub-frame losses in a Wi-Fi frame (lower value indicates more burstiness)[10]*

Detection Above ED: Fig. 9 plots the CDFs of our metrics for link AP1-C7 under three scenarios: LTE-OFF, LTE on primary channel (LTE-ON (P. Ch.)) and LTE on secondary channel (LTE-ON (S. Ch.)). For each of the xRetries, short retries, and long retries, we can find clear thresholds for separating LTE-ON-S from the LTE-OFF and LTE-ON-P scenarios. Based on our results in §3.2, LTE operating on the primary channel does not require *DeMiLTE*'s reaction measures. Thus, we do not seek to differentiate between LTE-OFF and LTE-ON-P scenarios. Our detection mechanism should not report LTE to be active in this case. For comparison, Fig. 10 plots the same CDFs for the AP1-C9 link which represents a case where *DeMiLTE* should not be engaged. Indeed, all three metrics show low values under LTE-U on both primary and secondary channel scenarios. We observe that xRetries, short retries, and long retries are upper bounded by 0.3, 0.06 and 0.24, respectively. In contrast, for the AP1-C7 link, the metrics in the LTE-ON (S. Ch.) case almost always had values higher than these specific thresholds.

Detection Below ED: Fig. 11 shows the CDFs of the three metrics under LTE-OFF, LTE-ON (P. Ch.) scenarios for a client that suffers performance degradation under LTE-U presence. We observe that thresholds can be easily selected for each of the *xR*, *lR* and *sR* metrics to differentiate between LTE-OFF and LTE-ON (P. Ch.) cases. However, note that we might need to use slightly different

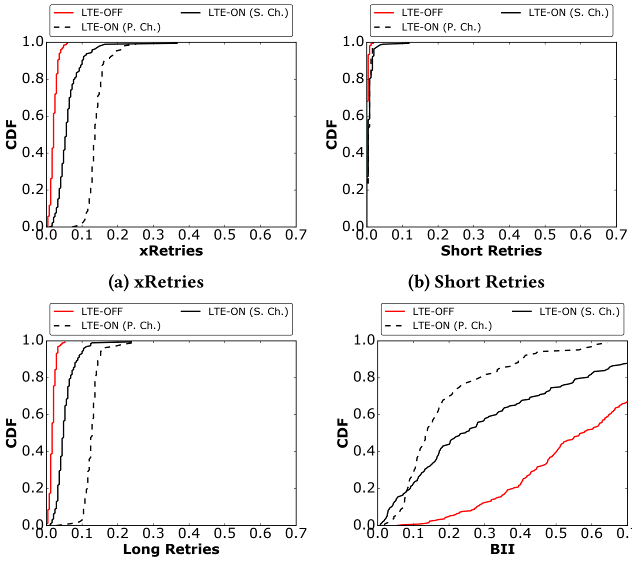

(a) xRetries (b) Short Retries (c) Long Retries (d) Burstiness Index Inverse

Figure 10: LTE-U per-link detection (Above ED): AP1-C9.

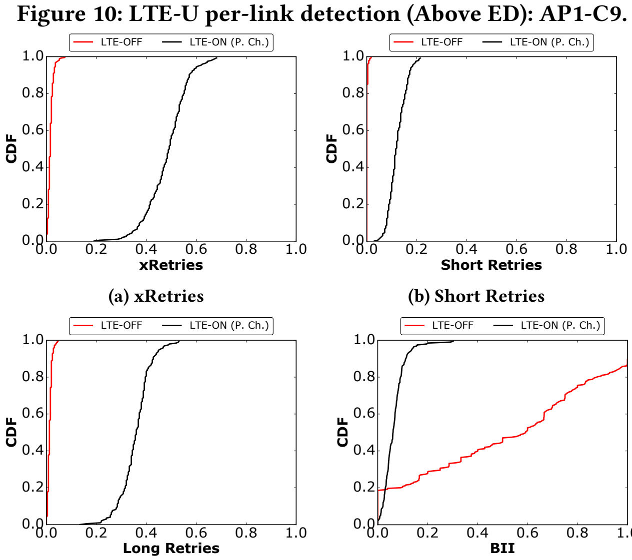

(a) xRetries (b) Short Retries (c) Long Retries (d) Burstiness Index Inverse

Figure 11: LTE-U per-link detection (Below ED): AP2-C8.

thresholds than the Above ED scenario. Nonetheless, a clear distinction exists between the OFF and ON (P. Ch.) scenarios. Further, we notice that for non-affected clients, the xR, lR and sR do not show elevated values (plots omitted). Lastly, when LTE-U operates on the secondary channel, we observe that for most clients the PHY data rates are barely affected (see §3.3.2), hence the metrics look similar to the Above ED primary channel case discussed before.

Detection Condition: To arrive at a detection mechanism/condition, utilizing the above discussed metrics, which can provide both high accuracy and low false positives, we first look at learning-based solutions. Specifically, we try decision trees which are particularly suited to the type of our classification task. Training the decision trees on dataset, we observe the following: (i) *BII* has extremely low feature importance, indicating that it is not a useful metric for LTE detection (ii) The trained trees assign extremely high feature importance to one of xR, lR and sR, indicating the classification almost entirely depends on one of the metrics. Using such a classifier can be prone to high rates of false positives as it does not take into account the changes in the values of metrics that might occur under regular link operation, when LTE is not transmitting. Specifically, with another Wi-Fi AP as a hidden interferer, xRetries and short retries can display similar elevated levels as in the presence of LTE. However, long retries should remain low: as soon as the Wi-Fi AP under test switches on RTS/CTS protection, data-packet loss rates decrease. Further, the probing mechanism built into the rate control, which occasionally sends packets at higher MCS before it converges to optimal rate, might temporarily cause a large number of long retries, but in such cases short retries will remain low. Lastly, under less aggressive LTE operation, it is possible that the number of xRetries might remain low, as one of the several multiple re-transmissions might indeed succeed, but the number of both long and short retries will still be high.

Based on the hints from decision trees and above considerations, we use a combined metric to differentiate between affected and non-affected clients. Specifically, we engage reaction mechanisms for a client if Condition 1 below is true. The three thresholds xR_{thresh}, lR_{thresh} and sR_{thresh} are determined experimentally (details are discussed in §4.3) by using the dataset described in §3 and multiple additional measurements.

$$(xR \geq xR_{thresh}) \parallel ((lR \geq lR_{thresh}) \;\&\&\; (sR \geq sR_{thresh})) \quad (1)$$

Our detection condition does not make use of BII as we found it does not help with detection under any of the scenarios, a fact also confirmed by decision trees. In the above ED scenario (Fig. 9d and 10d), BII in the LTE-OFF case shows great variation; hence, no clear threshold can be selected to distinguish it from the LTE-ON cases. In the below ED scenario (Fig. 11d), any selected threshold will lead to high false positives, making the metric unnecessary.

4.2 LTE Detection: Ground Truth

Before we train our thresholds and evaluate the detection strategy, we need to determine the ground truth for each client in our measurement study. This ground truth can be used for future LTE detection studies or online learning in a WLAN. Our goal here is to identify cases that experience more than expected performance deterioration and where engaging the link adaptation reactions would significantly improve the performance. To accurately annotate the ground truth of our experiments for training our metrics, we pose certain conditions under which a client is considered as affected. Let $R_{LTE-OFF}$ be the PHY data rate when LTE is not operating and R_{LTE-ON} be the PHY rate achieved by the client under LTE interference of airtime α. Similarly, we define $T_{LTE-OFF}$ and T_{LTE-ON} as the application-layer throughput without and with LTE presence, respectively. Given that we employ different link adaptation mechanisms for the cases where LTE operates on Wi-Fi AP's primary channel vs. secondary channel, we define the ground truth differently for the two scenarios.

LTE on Wi-Fi's Primary Channel: We define the ground truth as positive/react if Condition 2 is true. When there are no collisions between LTE and Wi-Fi transmissions, Wi-Fi links should ideally be able to utilize the channel for the percentage of airtime when LTE is not transmitting. However, we found certain cases where

TCP performance is severely affected (due to consecutive packet losses), although the underlying PHY data rate is not affected. In such cases, we cannot/should not react as the link adaptation is working optimally. To account for such cases, we impose a further condition that the AP reacts only if the PHY data rate under LTE operation is much lower than link's natural PHY rate. We expect that some collisions reduce the PHY rate of the client. However when interference is severe (such as hidden terminals), we observe that affected clients' rates are much lower. In such cases, the application throughput would be affected as well and our reaction mechanisms come into effect.

$$\left(\frac{T_{LTE-ON}}{T_{LTE-OFF}} \ll 1-\alpha\right) \&\& \left(R_{LTE-ON} \ll R_{LTE-OFF}\right) \quad (2)$$

Since LAA uses LBT and does not have a fixed duty-cycle, we consider the ground truth to be positive if (2) is true with $\alpha = 50\%$. In other words, we consider it fair for LAA to occupy half of the airtime but no more, as any other fair Wi-Fi AP would.

LTE on Wi-Fi's Secondary Channel: Here, the reaction involves reducing the bandwidth used by the WiFi link so as to not utilize the channel (as part of the 802.11ac's bonded channel) used by LTE, so we base our decision whether to react or not on comparing the expected PHY data rate after reducing bandwidth and the PHY data rate under LTE operation. Let BW_{max} be the maximum supported bandwidth for a given link and $BW_{optimal}$ be the optimal bandwidth under LTE interference.[4] Further, we assume that when we reduce bandwidth from 80 MHz, to say 20 or 40 MHz, then PHY rate also reduces by the same factor of 4 or 2, respectively. We realize that this is an approximation, as when reducing bandwidth from BW_{max} to $BW_{optimal}$, a higher (by 1-2 indexes) MCS might be supported. However, an LTE affected client typically drops the MCS much more than just 2 indexes due to strong LTE interference, hence moves down beyond these higher supported MCSs. Thus, to identify affected clients, we require that Condition 3 is met.

$$\frac{R_{LTE-ON}}{R_{LTE-OFF}} \ll \frac{BW_{LTE}}{BW_{max}} \quad (3)$$

4.3 LTE Detection: Evaluation

To determine the thresholds and evaluate our detection methodology, we augment the dataset used in the previous sections with additional measurements and include measurements from multiple devices to account for client diversity. Our dataset consists of 27975 LTE-U samples (16304 above ED and 11671 below ED) and 15980 LAA samples (6022 above ED and 9958 below ED), collected under diverse links operating in a typical enterprise WLAN.

To find the optimal thresholds, we do an exhaustive search over the range of possible values: [0, 1] with a step size of 0.01 for each of the three thresholds. We use five-fold stratified cross-validation and for each fold pick the thresholds that maximize the sum of *f1-score* (weighted average of the precision and recall) and *accuracy* over the training data-set in each fold. Table 2 lists the optimal thresholds obtained. Note that the thresholds selected under different folds are similar. The thresholds generalize well across links, as our metrics are highly correlated with PER and as a result the rate control, which degrades PHY rates similarly for any link, under heavy interference. To further confirm that the chosen thresholds are generic, we obtain the thresholds by training only on the LTE-U dataset and see how well they perform on the unseen LAA dataset (Table 3).

[4]Based on which secondary channel LTE is operating on, the AP can reduce bandwidth from 80 to 20 or 40 MHz.

Table 2: Trained thresholds for our Detection Metrics

Metric	Above ED	Below ED
xRetries (xR)	0.45	0.45
Short Retries (sR)	0.17	0.09
Long Retries (lR)	0.10	0.16

Next, we look at the trade-off between averaging consecutive samples to improve the detection system's performance. Table 3 presents the f1-score and accuracy of the *DeMiLTE*'s detection over different averaging windows ranging from $n = 100$ to $n = 500$ frames for the above ED and below ED scenarios. In the above ED scenario, we observe that increasing the window size, which in turn increases the detection latency, does not offer significant gains. In general, we find that with window size 200, we can achieve ~ 0.9 accuracy (and ~ 0.9 f1-score) for LTE-U. Interesting, under LAA, we can achieve even higher accuracy and f1-scores although the thresholds were never explicitly trained for LAA, indicating that the detection condition and thresholds perform well even across LTE technologies. Lastly, increasing the window size further offers minor to no improvements.

Table 3: Detection Performance (f1-score & Accuracy)

	LTE-U/Above ED		LTE-U/Below ED		LAA/Above ED		LAA/Below ED	
	f1	*Accuracy*	*f1*	*Accuracy*	*f1*	*Accuracy*	*f1*	*Accuracy*
$n = 100$	0.88	0.90	0.89	0.88	0.91	0.93	0.90	0.92
n = 200	**0.89**	**0.90**	**0.92**	**0.91**	**0.93**	**0.94**	**0.92**	**0.93**
$n = 300$	0.88	0.90	0.91	0.90	0.94	0.95	0.92	0.93
$n = 400$	0.89	0.91	0.92	0.91	0.94	0.95	0.92	0.93
$n = 500$	0.89	0.91	0.90	0.89	0.94	0.95	0.92	0.93

Discussion: We have introduced a per-client LTE interference detection system. Although we have trained our metrics in realistic environments with thousands of samples, we highlight that commodity APs should not be tied to the thresholds in Table 2 and should be updated with online learning algorithms. Enterprise WLANs already maintain statistics similar to ours [2]. While implementing *DeMiLTE* in an online fashion, we recommend using the following data: (i) LTE detection (channel and airtime from a spectrum monitor); (ii) PHY rates with LTE ON/OFF; (iii) AP-side device mobility detection [23], to disable *DeMiLTE* sample collection during periods of device mobility; (iv) *xR*, *lR*, *sR*. Metrics (ii)–(iv) should be maintained on a per-client basis; (i)–(iii) are needed for computing the ground-truth of the learning process by applying Conditions 2 and 3, whereas (iv) should be used for training.

5 *DeMiLTE*: INTERFERENCE MITIGATION

Our LTE/Wi-Fi coexistence study indicates that LTE can harm crucial 802.11ac VHT features. We introduce an LTE interference mitigation scheme that reverts LTE degradation on affected clients. Our design satisfies the following goals: 1) compliance with the 802.11ac standard 2) lightweight implementation (required given our limited firmware space, tens of connected clients per AP, and high 802.11ac rates) 3) seamless operations for non-affected stations 4) fairness (at least 50% airtime) to LTE by either channel adaptation or airtime reduction (as a result of PHY rate improvements).

5.1 System Design

We implement our solution on enterprise-grade APs. Our interference mitigation code lies in the rate control class and is about 300 lines. For each Wi-Fi client, in addition to its rate control state variables, we introduce six variables to track the long/short/excessive retries, number of transmissions, LTE detection decision, and bandwidth tried under LTE detection. We verified that our code does not add any overhead by running exhaustive, saturated experiments at Gbps rates. In the rate control class, we add functions and variables that read channel state information (CSI) from AP hardware registers and estimate the SNR of any downlink.

To reduce false positives under no LTE presence, a spectral engine (e.g., LTERadar [25]) passes LTE channel and duty cycle (DC) to VHT modules in firmware. As spectral classifications typically run in AP's user space [19, 25], we implement a system that handles LTE spectral detection to firmware though `ioctl` calls and Atheros wireless module interface (WMI) messages [3], as shown in Fig. 12. The same figure shows the detection and reaction flow running on firmware and rate control. We distinguish two cases: LTE above and below ED. The AP receives such information from spectral analysis (running either on same AP or a spectrum monitor).

Reaction Above ED: We exploit *DeMiLTE*'s interference detection presented in §4, and spectral analysis that yields the center frequency of the interferer. When LTE is on the primary Wi-Fi channel, it does not harm 802.11ac features. When LTE operates on a secondary Wi-Fi channel, the AP should keep track of the number of the affected stations. Our reaction solution above ED involves *only efficient spectrum allocation* when LTE is on a secondary channel. If the proportion of affected clients is significant (e.g., 50%), the AP should blacklist the LTE channel and either reduce bandwidth or change channel. As we observe in Fig. 5, we can achieve close to 10x performance gains (throughput of VHT40 is 105 Mbps after mode reset vs. 11 Mbps for VHT80) for affected clients with LTE on secondary channel by switching channel or VHT mode.

Reaction Below ED: We implement *interference-aware rate control* in the AP's firmware. Similarly to prior work [18], we need to detect collisions that increase SFER and force the firmware to drop the PHY rates. Severe collisions with LTE are hinted by high excessive retries (> 41%), as we discuss in §3. However, excessive retries in rate control can also result from infeasible selection of MCS/SS/bandwidth and we need to identify such cases. Our interference-aware rate control uses CSI to determine the maximum supported MCS, similarly to previous protocols [24] and works as follows: 1) If monitored metrics hint LTE interference on primary channel, the rate control starts requesting CSI samples. To avoid significant overhead of CSI requests, we request 10 samples per detection interval (e.g. 300 packets) and determine the optimal MCS from IEEE 802.11ac MCS/SNR tables. SNR is estimated as the median SNR amplitude over all subcarriers. 2) After excessive retries, commodity firmware adds a severe negative bias for the specific tried rate (e.g., 30% additional PER). *DeMiLTE* modifies this bias and does not increase PER if excessive retries result from a rate with lower than the supported MCS. In addition, *DeMiLTE* does not penalize the bandwidth used in the same transmission, contrary to default rate control. *DeMiLTE* highlights that next-generation enterprise WLAN can achieve high gains by using CSI information, especially in dense environments with dynamically varying interference.

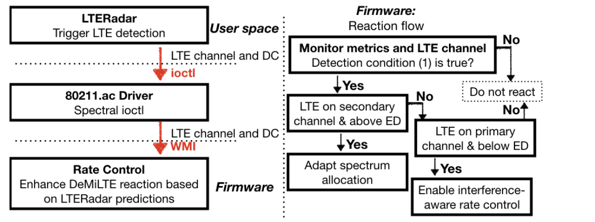

Figure 12: System architecture with LTERadar (left) and reaction flow of *DeMiLTE* (right).

5.2 Performance Evaluation

We evaluate the performance of *DeMiLTE* under diverse scenarios. First, Fig. 13a presents the relative gains of *DeMiLTE* vs. Autorate (default 802.11ac rate control) when LTE's transmission power is below the Wi-Fi ED threshold and its channel coincides with the Wi-Fi primary one. All results from our affected clients (C3–9) report gains for *DeMiLTE*. In particular, the TCP and UDP throughput gains are up to 49% and 81%, respectively. UDP throughput gains are higher compared to TCP ones due to multiple reasons: 1) TCP has multiple timeouts during LTE-ON periods, hence dropping the congestion window and WLAN traffic demands. 2) TCP acknowledgements from the Wi-Fi client can also severely suffer.[5] 3) PHY-rate gain under TCP experiments is lower than UDP one. During TCP experiments, the AP is less aggressive with transmissions than with UDP due to low TCP congestion window, hence the rate adaptation does not converge to optimal MCS/SS/bandwidth as fast as with UDP. The SFER during all experiments was very similar (30–45%) between Autorate and *DeMiLTE* due to collisions with LTE, hence it is omitted. It is important to note that *DeMiLTE* not only increases throughput, but also decreases airtime thanks to higher PHY rate. Average airtime reduction is 43% for TCP tests and 57% for UDP ones[6] and it is beneficial for both WLAN and LTE networks that can use the remaining airtime.

DeMiLTE has high gains for static scenarios where maximum supported MCS does not vary. However, typically clients move and supported MCS varies. If a client moves from a location where it is affected to one where it does not suffer from LTE, then our detection metrics turn off the interference-aware reaction. In the reverse case though, our CSI approach ensures that the AP has information on maximum supported MCS when the client starts being affected by LTE. To illustrate the efficiency of *DeMiLTE*, we conduct a mobility experiment from location C1 (closest to AP2) to C9 (farthest from AP2 and closest to LTE eNb). We stay at C1 for 70 s until we start walking towards C9 (following the path C1->C6->C7->C8->C9), where we stay until the end of the experiment. Fig. 13c presents UDP throughput of two identical tests run back-to-back, with Autorate and *DeMiLTE*. In addition, we show the outcome of per-link LTE detection, the median MCS used, and client's SNR over time. We observe the following: 1) *DeMiLTE* leaves performance at location C1 unaffected. 2) Although *DeMiLTE* does not target improving performance under mobility, *DeMiLTE* throughput results are very close to Autorate ones during mobility. 3) *DeMiLTE* is able to dynamically adapt the maximum supported MCS and

[5] We cannot control the rate control at the client side given the hardware diversity.
[6] We approximate airtime reduction as the PHY rate gain, as airtime is inversely proportional to transmission rates.

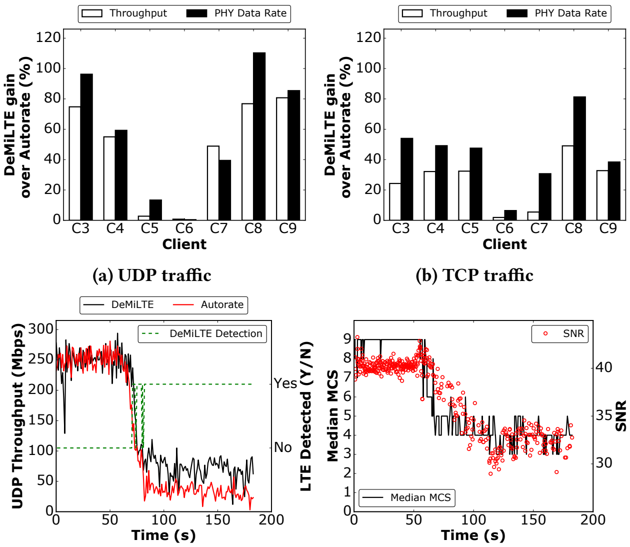

(a) UDP traffic

(b) TCP traffic

(c) Mobility

Figure 13: *DeMiLTE* with LTE below ED (AP2), on primary channel and 50% duty cycle.

interference-aware rate control based on SNR, and to provide 110% throughput gains for link AP2-C9 in a realistic scenario.

6 RELATED WORK

There are several *LTE/Wi-Fi coexistence* studies debating whether LTE harms Wi-Fi [11, 14]. Our study with commodity hardware shows that LTE can cause severe performance degradation. Several efforts have been made towards fair protocols/signalization on LTE eNodeBs [12, 15, 26]. However, these protocols are vendor specific and do not guarantee ubiquitous implementation. WiPLUS [17] uses ED registers and modifies the client scheduling at Wi-Fi AP to detect LTE-U. However, ED registers cannot identify the channel of LTE operations and modifying client scheduling increases delay and unfairness. All previous works either use simulations or software defined radios, making the findings oblivious to 802.11ac VHT Wi-Fi features or LTE L1–L3 stacks. LTERadar [25] uses commodity hardware to introduce per-AP LTE detection, in any 802.11ac channel. However, LTERadar does not introduce any per-client mechanism. Contrary to aforementioned studies, we conduct measurements with commodity hardware and we propose *DeMiLTE*, which is a lightweight, Wi-Fi based, and 802.11ac-compliant per-client solution for LTE detection and mitigation.

Non-Wi-Fi interference detection and mitigation has been extensively studied in the literature, especially for the crowded 2.4 GHz bands. Airshark [19] uses FFT spectrum analysis for interferer detection. WiFiNet [20] uses Airshark and two radios per AP to simultaneously detect and mitigate interference. The main drawback of WiFiNet is the two-radio requirement, which increases cost and complexity, as tight synchronization is needed.

There is prior work on *HT/VHT optimization*, such as dense WLAN channelization [21], 802.11n/ac rate control improvements [18], and packet loss detection [10]. Prior loss-differentiation approaches [10, 18] neither detect nor react to LTE interference. We find that proposed metrics [10, 18] do not perform well under LTE interference.

7 CONCLUSION

LTE unlicensed in 5 GHz is being deployed by several mobile operators and is a key building block of 5G networks. Nevertheless, it can harm Wi-Fi operations in dense enterprise and residential networks. Compared to Wi-Fi interferers, LTE neither respects RTC/CTS nor can it decode Wi-Fi frames, hence posing new, unique challenges in WLAN rate and bandwidth control. We have conducted an extensive measurement study to shed light on which scenarios LTE harms 802.11ac key features. Building on this study, we proposed *DeMiLTE*: the first per-link detection mechanism for aggressive LTE interference and the first Wi-Fi-based reaction algorithm that helps combat interference. *DeMiLTE* exhibits up to 110% gains in realistic environments and we recommend its adoption in dense enterprise environments where LTE base stations may be installed.

REFERENCES

[1] [Online]. 3GPP Release 13, TS 36.213.
[2] [Online]. Aruba Wi-Fi Statistics. https://goo.gl/dWoXmc.
[3] [Online]. ATH WMI. http://bit.ly/2XBdHrt.
[4] [Online]. LTE-U Forum. http://www.lteuforum.org/.
[5] [Online]. Minstrel Rate Adaptation. https://wireless.wiki.kernel.org.
[6] [Online]. OpenEPC. http://www.openepc.com/
[7] [Online]. Spectrum Analysis - Aruba Networks. https://goo.gl/awcDFY.
[8] [Online]. T-Mobile launches LTE-U. https://t-mo.co/2IDZKFA.
[9] [Online]. Verizon launches LAA. https://goo.gl/Rx8WqX.
[10] Ruwaifa Anwar, Kamran Nishat, Mohsin Ali, Zahaib Akhtar, Haseeb Niaz, and Ihsan Ayyub Qazi. 2014. Loss differentiation: Moving onto high-speed wireless LANs. In *IEEE INFOCOM*.
[11] Dean R. Brenner and John W. Kuzin. 2015. Technical Response to Parties Opposed to LTE Unlicensed. In *Qualcomm TR*.
[12] Eugene Chai, Karthik Sundaresan, Mohammad A. Khojastepour, and Sampath Rangarajan. 2016. LTE in Unlicensed Spectrum: Are We There Yet?. In *ACM MobiCom*.
[13] Y. Gao, Q. Huang, S. Xu, H. Li, Z. Li, and W. Tang. 2016. Experimental Performance Evaluation and Analysis of LAA and Wi-Fi Coexistence in the Unlicensed Spectrum. In *IEEE Globecom Workshops*.
[14] Nihar Jindal and Donald Breslin. 2015. LTE and Wi-Fi in Unlicensed Spectrum: A Coexistence Study. In *Google TR*.
[15] Ramanujan K Sheshadri, Karthikeyan Sundaresan, Eugene Chai, Amir Khojastepour, Sampath Rangarajan, and Dimitrios Koutsonikolas. 2017. BLU: Blueprinting Interference for Robust LTE Access in Unlicensed Spectrum. In *ACM CoNEXT*.
[16] N. Makris, A. D. Samaras, V. Passas, T. Korakis, and L. Tassiulas. 2017. Measuring LTE and WiFi coexistence in Unlicensed spectrum. In *European Conference on Networks and Communications*.
[17] Michael Olbrich, Anatolij Zubow, Sven Zehl, and Adam Wolisz. 2017. WiPLUS: Towards LTE-U Interference Detection, Assessment and Mitigation in 802.11 Networks. In *European Wireless*.
[18] Ioannis Pefkianakis, Yun Hu, Starsky H.Y. Wong, Hao Yang, and Songwu Lu. 2010. MIMO Rate Adaptation in 802.11N Wireless Networks. In *ACM MobiCom*.
[19] Shravan Rayanchu, Ashish Patro, and Suman Banerjee. 2011. Airshark: detecting non-WiFi RF devices using commodity WiFi hardware. In *ACM IMC*.
[20] Shravans Rayanchu, Ashish Patro, and Suman Banerjee. 2012. Catching whales and minnows using WiFiNet: deconstructing non-WiFi interference using WiFi hardware. In *USENIX NSDI*.
[21] Shravan Rayanchu, Vivek Shrivastava, Suman Banerjee, and Ranveer Chandra. 2011. FLUID: Improving Throughputs in Enterprise Wireless LANs through Flexible Channelization. In *ACM MobiCom*.
[22] Ljiljana Simić, Janne Riihijärvi, and Petri Mähönen. 2017. Measurement study of IEEE 802.11 ac Wi-Fi performance in high density indoor deployments: Are wider channels always better?. In *IEEE WoWMoM*.
[23] Li Sun, Souvik Sen, and Dimitrios Koutsonikolas. 2014. Bringing Mobility-Awareness to WLANs using PHY layer information. In *ACM CoNEXT*.
[24] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2016. Practical MU-MIMO User Selection on 802.11Ac Commodity Networks. In *ACM MobiCom*.
[25] Christina Vlachou, Ioannis Pefkianakis, and Kyu-Han Kim. 2018. LTERadar: Towards LTE-Aware Wi-Fi Access Points. In *ACM SIGMETRICS*.
[26] Sangki Yun and Lili Qiu. 2015. Supporting WiFi and LTE Co-existence. In *IEEE INFOCOM*.
[27] Y. Zeng, P. H. Pathak, and P. Mohapatra. 2014. A first look at 802.11ac in action: Energy efficiency and interference characterization. In *IFIP Networking*.