LTE in Unlicensed Spectrum: Are We There Yet?

Eugene Chai, Karthik Sundaresan, Mohammad A. Khojastepour, Sampath Rangarajan
Mobile Communications and Networking
NEC Laboratories America
{eugene, karthik, amir, sampath}@nec-labs.com

ABSTRACT

In this work, we explore the potential and impact of unlicensed LTE on WiFi in unlicensed spectrum. Our experiments demonstrate that the large *asymmetry* in the channel access methodologies employed by WiFi and LTE (carrier sensing/notification in WiFi, energy sensing alone in LTE-U), can result LTE-U completely blocking WiFi transmissions, and causing significant degradation to either technologies from collisions.

We address this critical sensing asymmetry with ULTRON, a LTE-WiFi co-existence solution that integrates WiFi's carrier sensing and notification mechanisms into LTE, without any modifications to the LTE PHY standard. ULTRON operates at the LTE base station and consists of two key components: (a) *WiFi embedding* that embeds appropriate data into the LTE-subframes through an intelligent reverse-engineering of the LTE PHY, so as to realize a WiFi PLCP preamble-header transmission over the air directly using the LTE PHY; and (b) *scalable WiFi sensing* that employs a single WiFi interface and maximizes its carrier sensing benefits to all the unlicensed channels operating at the LTE node. Our evaluations demonstrate that ULTRON can increase the WiFi and LTE throughput by 5× and 6× respectively, resulting from a sharp reduction in LTE-WiFi interference.

1. INTRODUCTION

As cellular networks evolve to 5G to support next generation interactive services like AR/VR, spectrum has un-deniably been the achilles heel in this evolution. While a move to higher frequencies like mmWave (≈ 30 GHz and higher) offers higher bandwidth, it is also accompanied by higher attenuation (less coverage), sensitivity and deployment cost. Hence, there has been considerable interest recently in exploring all available options under 6 GHz, of which the un-licensed bands (5 GHz) form a promising candidate.

LTE in unlicensed spectrum: Operating in unlicensed bands requires LTE to coexist fairly with the incumbent WiFi, the dominant technology in 5 GHz spectrum. The industry has devised two modes for LTE operation in un-licensed spectrum, namely LTE Unlicensed (LTE-U) [1] and License Assisted Access LTE (LTE-LAA) [2]. In both modes, LTE is expected to perform energy sensing before transmission and would leverage its carrier aggregation feature to aggregate its licensed carriers (channels) with channels from the unlicensed spectrum. The key difference between the two modes is that while LTE-U (de)activates LTE in the unlicensed channels at coarse time scales (≈ 20 ms duration), LTE-LAA is expected to operate at much finer time-scales (1-10 ms) that are comparable to WiFi transmissions. With LTE-U more closely resembling existing LTE (i.e. continuous operation), it can be realized with today's LTE networks and hardware (e.g. Verizon's trial of LTE-U in Austin [3]). However, its coarse time-scale operation results in short-term unfairness and increased latencies for WiFi and hence has been the subject of intense debate in the US recently [4]. On the other hand, LTE-LAA is poised to provide better fairness and coexistence to WiFi (mandatory mode in Europe and Japan). However, it requires changes to the LTE air interface and is hence undergoing 3GPP standardization.

Challenge: With WiFi being the ubiquitous technology that users rely on, we believe it is critical (as a community) to understand the implications of this co-existence problem and help shape its course. Our studies reveal that *we are not there yet in terms of an efficient, fair co-existence solution between unlicensed LTE and WiFi.* To this end, we focus on LTE-U, and evaluate its interactions with WiFi using real-world experiments.

There exists a subtle but critical problem that arises from the asymmetric channel access mechanisms employed by the two technologies, with the potential to significantly degrade either technology's performance. Specifically, while both WiFi and LTE employ energy sensing to detect strong signals over a threshold (e.g. -62 dBm), WiFi alone employs carrier sensing (preamble+header detection) and notification to detect and inform other WiFi nodes of signals that would otherwise cause interference, with a higher sensitivity (e.g. over -82 dBm). Given that both WiFi and LTE are well-equipped to operate on links over -82 dBm SNR, the lack of carrier sensing/notification feature in LTE creates two issues. First, it leads to collisions (e.g. signals in [-82,-62] dBm range) in numerous situations, where LTE and WiFi cannot detect each other. Indeed, in our experiments, such interference can degrade performance for both LTE and WiFi by as much as 40% even if the SINR is above 10dB. Further, the impact of this problem is not just restricted to LTE-WiFi interactions but also between LTE nodes of different operators. Second, the lack of channel reservation notification (e.g. Physical Layer Convergence Protocol (PLCP) header in WiFi) in LTE leads to un-necessary back-offs and hence reduced throughput for WiFi.

Naive solutions do not work: Two straightforward solutions to this challenge are (a) lowering the energy-detection (CCA) threshold (e.g. to -80dBm), and (b) integrating WiFi PHY into LTE. However, reducing the CCA threshold does not address the problem as

MobiCom'16, October 03-07, 2016, New York City, NY, USA

DOI: http://dx.doi.org/10.1145/2973750.2973781

it would only lead to under-utilization of the channel as the LTE-U node will back off for weak, spurious signals in the un-licensed channel that may not pose an interference threat to either LTE or WiFi. Also, integrating a WiFi PHY means that LTE (a 3GPP standard) will now be dependent on the WiFi specification (an IEEE standard), and this is not necessarily welcomed by cellular operators. Furthermore, other protocols besides WiFi operate in the unlicensed spectrum (e.g. Bluetooth, Zigbee etc) and co-existence with these non-WiFi protocols should also be achieved within a single cohesive framework.

Hence, in order for LTE to coexist fairly and efficiently with WiFi, we argue that *it is critical for LTE to homogenize its channel access mechanism with that of the incumbent WiFi by incorporating the latter's preamble detection and notification capabilities*. However, this faces two issues. First, a single LTE interface can transmit on (up-to) five channels (i.e. component carriers) concurrently. Having a dedicated WiFi sensing device per channel results in a significant hardware overhead. Current *multi-mode LTE small-cells* [5, 6] are typically equipped with a single WiFi interface. Second, it is not feasible to expect the LTE standard to incorporate another technology's control signal (WiFi's PLCP) notification and detection. *Hence, the key question we seek to answer is whether such capability can be realized without requiring any changes to the WiFi and LTE specifications*. Such an approach would homogenize the channel access procedure and hence provide a fair coexistence mechanism between both LTE-WiFi and LTE-LTE without degrading channel utilization efficiency.

ULTRON: We design a novel, standards-compliant **U**nlicensed **LT**E **R**adi**O** **N**ode (ULTRON) to address this co-existence challenge. ULTRON operates at the LTE base station and seamlessly works with legacy LTE clients. It employs two key design mechanisms: (i) *WiFi embedding* that allows an LTE BS to transmit a WiFi control signal (specifically, PLCP preamble+header) prior to its transmission by masking itself as LTE data, and (ii) *WiFi sensing* that allows the LTE BS to detect WiFi control signals in packets transmitted by other WiFi nodes and LTE BSs.

However, realizing ULTRON's mechanisms in practice faces several challenges. First, LTE and WiFi operate at different bandwidths. For example, a 10MHz LTE channel actually consists of a 15.36MHz LTE signal along with the rest being unused. ULTRON has to embed the WiFi's PLCP in its Orthogonal Frequency Division Multiple Access (OFDMA) frame in a manner that will make it appear as though it originated from a WiFi node. This requires a careful reverse-engineering of the LTE PHY. Second, ULTRON has to equip LTE with WiFi carrier sensing for every unlicensed channel it operates on without incurring hardware overhead or any changes to the LTE PHY specification.

ULTRON's mechanisms address these challenges in a cost-effective manner through a novel design. For every unlicensed channel that the LTE BS operates on, ULTRON determines the frequency domain equivalent of WiFi's CTS-to-self packet and embeds it as data in the OFDMA frame. This is achieved through a careful reverse-engineering of the LTE PHY, such that the time domain version of the LTE signal accurately delivers WiFi's CTS packet to any node that employs WiFi's carrier sensing. This allows ULTRON to seamlessly notify channel reservation on every un-licensed channel in a manner similar to WiFi.

Enabling WiFi carrier sensing without LTE PHY modifications on every un-licensed channel is more complicated. Here, ULTRON leverages the growing trend of LTE small cells also being equipped with a WiFi interface (for dual connectivity [5, 6]) to re-purpose the latter for WiFi sensing alone. However, a single WiFi interface (spanning 20-40 MHz) is not sufficient to cover the numerous un-licensed channels (up-to 5) that LTE can operate on simultaneously (i.e. up-to 100 MHz in carrier aggregation [7, 8]). Hence, ULTRON intelligently leverages the single WiFi interface to maximize the benefits of carrier sensing to all the un-licensed channels. It accomplishes this by periodically sensing on each un-licensed channel to estimate the fraction of WiFi traffic that is subject to impact (i.e. falls within the carrier sense threshold, say [-82,-62] dBm). ULTRON then uses this information to solve a joint optimization problem that helps determine (a) the particular un-licensed channel where the WiFi interface's carrier sensing feature would be most beneficial and hence must be deployed, and (b) the appropriate amount of LTE traffic load that must be placed (through flexible LTE scheduling) on each of the un-licensed channels. The joint optimization of LTE traffic placement and WiFi interface assignment to the un-licensed channels works synergistically to minimize the impact from lack of WiFi carrier sensing on all-but-one un-licensed channels.

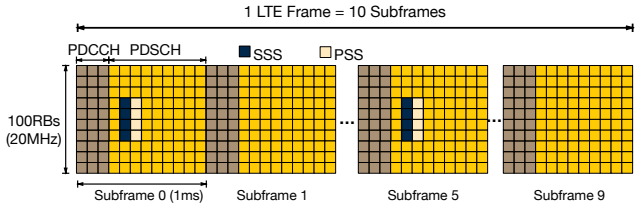

Figure 1: LTE frame structure showing the PSS and SSS.

We implement ULTRON and its mechanisms using a mixture of USRPs and WARPs that enables LTE-U operation in the 5 GHz band. Our co-existence experiments demonstrate the substantial benefits and hence the need for bridging the sensing asymmetry between LTE and WiFi – ULTRON increase WiFi and LTE throughput by up to 5× and 6× respectively, resulting from a sharp reduction in WiFi-LTE interference.

Our contributions in this work are multi-fold.

- We quantify the co-existence issue between LTE-WiFi and LTE-LTE in unlicensed spectrum that stems from the lack of carrier sensing and notification features in LTE, and degrades both WiFi and LTE performance even when operating in the coexistence-friendly LTE-U mode.
- We propose the design of a novel LTE base station node called ULTRON that realizes the benefits of WiFi PLCP notification and detection in the un-licensed channels in a scalable, cost-effective manner, without requiring any changes to the LTE specification.
- We build a prototype of ULTRON and demonstrate its critical role in fair and efficient coexistence between LTE-WiFi in un-licensed spectrum through real world evaluations.

2. BACKGROUND

2.1 LTE Background

LTE Frame Structure. Fig. 1 shows an example of a 20MHz LTE frame. A LTE frame consists of 10 *subframes*, each with a duration of $1ms$. The subframes are made up of *resource blocks (RBs)*. These RBs are grouped into two main physical channels: the control and the data channel. On the downlink, the control channel is known as the Physical Downlink Control CHannel (PDCCH) while the data channel is called the Physical Downlink Shared CHannel (PDSCH). PUCCH and PUSCH are the corresponding uplink control and data channels respectively. The PDCCH carries control

information such as RB-to-UE (user) assignments. The synchronization signals, Primary and Secondary Synchronization Signals (PSS/SSS), are carried in the PDSCH of subframes 0 and 5 (as shown in Fig.1) and contain cell information (group and unique cell IDs). The UE uses the PSS/SSS signals to achieve and maintain precise time and frequency synchronization with the eNodeB, so that the other downlink subframes can be decoded. Each LTE frame also carries a broadcast channel (BCH) which carries the Master Information Block (MIB). The MIB contains key parameters that are needed for a UE to decode the downlink transmissions.

LTE Component Carriers. Each LTE-advanced eNodeB and UE can transmit on up to five (5) distinct channels, known as *component carriers (CCs)* simultaneously through *carrier aggregation*. Each CC can have a bandwidth of 1.4, 5, 10 or 20MHz. One of these CCs is designated as the *Primary Component Carrier*, which is always active and carries the PSS/SSS signals for UEs to attach to the eNodeB.The other four CCs are known as *Secondary Component Carriers* and can be activated/deactivated as needed.

2.2 LTE in Unlicensed Spectrum

Unlicensed LTE leverages the LTE carrier aggregation feature to aggregate its licensed carriers with channels from the unlicensed spectrum, called *unlicensed component carriers*. LTE-U (de-)activates LTE in the unlicensed channels at coarse time scales ($\sim$ 20 ms duration) through a duty-cycling approach, and rely on subframe blanking to co-exist with WiFi by introducing short time gaps into the LTE-U transmission burst. LTE-LAA is expected to operate which much shorter on durations (1-10 ms) that are comparable to WiFi transmissions, and rely on WiFi-like Listen-Before-Talk (LBT) to achieve fine-timescale channel access.

3. CHALLENGES TO LTE-WIFI COEXISTENCE

We operate LTE and WiFi on the same unlicensed channel in real-world experiments to demonstrate the challenges to LTE-WiFi coexistence. Our study of LTE and WiFi shows that (a) energy sensing is not sufficient for LTE-WiFi coexistence. LTE can severely degrade the performance of WiFi, even if the LTE interference power at a WiFi device is as low as -80 dBm (i.e. below the -62 dBm CCA[1] threshold of WiFi and LTE); and (b) simple duty cycling of the LTE transmissions (i.e. LTE-U) is not sufficient to ease the degradation of WiFi, even if the LTE interference is within the CCA threshold of WiFi.

In order to more accurately characterize the LTE-WiFi co-existence behavior, we adopt an unlicensed LTE design from the recent LTE Release 13 specification [9]: each LTE subframe contains a PSS/SSS, as shown in Fig. 2. The UE uses these PSS/SSS symbols to maintain synchronization with every subframe separately, thus increasing its resilience to WiFi interference. Note that without such a design, the UE can lose synchronization with the eNodeB if subframes 0 and 5 are lost due to WiFi interference. This design is explained in greater detail in §5.1. We use the MATLAB LTE System Toolbox to generate this re-designed LTE PHY for all our LTE transmissions. Note that the MIB is not needed for our experiments as all PHY parameters are already known beforehand.

We make use of one LTE eNodeB/UE and one WiFi AP/STA pair for this motivational study. USRP B210s are used for the eNodeB and UE, while WARPv3 devices are used for the WiFi devices. The eNodeB transmits the MATLAB-generated LTE signal, while the UE logs the received I/Q signals and decodes it offline using MATLAB. All LTE and WiFi devices operate at 10MHz bandwidth

[1]Clear channel assessment refers to energy sensing in this work.

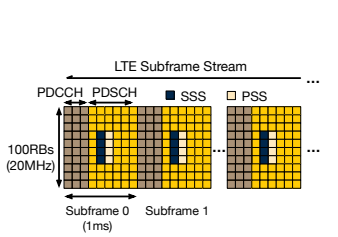

Figure 2: Unlicensed LTE PHY design.

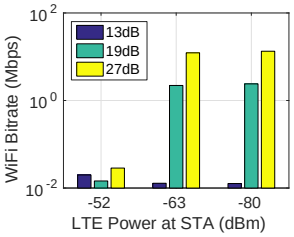

Figure 3: WiFi bitrate with LTE interference. (WiFi throughput without LTE interference is 20Mbps).

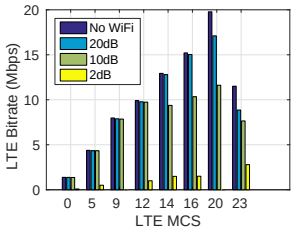

Figure 4: LTE bitrate with WiFi interference

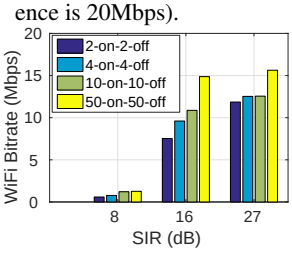

Figure 5: WiFi bitrate under 50% duty-cycled LTE

and are placed within an office environment with $1m$ separation between adjacent nodes.

LTE-WiFi Fairness. We consider LTE to be "fair" to WiFi if the impact of LTE on the WiFi network is similar to the impact seen if the LTE devices are replaced with WiFi nodes.

3.1 Energy-Sensing has Limited Play

We first study the impact of interference energy between downlink LTE and WiFi transmissions. In this experiment, LTE uses a continuous downlink transmission (i.e. without duty-cycling), allowing UEs to decode each sub-frame independently. The signal and interference power at the UE is measured using the I/Q data logged from the B210, while the RF power at the WiFi nodes is directly reported by the WARP 802.11 reference design.

WiFi Performance Under LTE Interference. The transmit power of the eNodeB is varied to ensure that the received LTE interference power at the AP is below the WiFi CCA threshold of -62dBm. The WiFi transmit power is adjusted to achieve a Signal-to-Interference Ratio (SIR) at the WiFi STA of 13, 19 and 27dB.

The maximum achievable WiFi throughput without any interfering LTE transmission is 20Mbps. Fig. 3 shows the achievable WiFi performance when the measured LTE interference power at the AP is -62dBm (the CCA threshold). When the SIR at the WiFi STA is 27dB, the achievable downlink throughput at the STA is 12.6Mbps, a 37% reduction from its maximum throughput. As the SIR decreases further to 19 and 13dB, the WiFi throughput reduces to 6 Mbps and 35 Kbps respectively. Since WiFi cannot detect a LTE carrier, its energy sensing may not detect the LTE signal, even when the latter is close to the CCA threshold. Although such a scenario may not be frequent, when this happens, the strong interference from LTE can degrade WiFi throughput significantly.

However, the impact on WiFi is severe even at lower LTE interference levels (-80 dBm) as observed in Fig. 3. In contrast to the previous scenario, even when the LTE interference is low, the inability of the WiFi device to detect the LTE signal (being below the CCA threshold) increases the frequency of collisions at WiFi significantly. Unfortunately, reducing the energy-sensing threshold to WiFi's preamble detection threshold (-82dBm) will not address the problem, but will only increase the false detection rate, leading

to under-utilization of the channel [10]. Our results validate similar conclusions made by other industry sources [4].

LTE Performance Under WiFi Interference. Our experiments with LTE indicate that LTE's PHY (robust coding and retransmissions) is more resilient to interference from WiFi. However, it still suffers a throughput reduction in the presence of WiFi interference. In this experiment, we fix the WiFi interference power at the eNodeB to be -70dBm, and adjust the eNodeB transmit power to achieve different SIRs at the UE. We observe that when WiFi interference is not present, LTE can achieve a maximum of 20Mbps at a Modulation and Coding Scheme (MCS) of 20. When the SIR is 20dB, this throughput is reduced to 17Mbps. As the SIR reaches 2dB, almost no data can be successfully transmitted over the LTE channel. Hence, interference power below the energy-sensing CCA threshold will have a severe impact on LTE performance as well.

These results highlight that *energy-based CCA is not sufficient to ensure an efficient coexistence between LTE and WiFi and can degrade their performance even at low interference powers*.

This is particularly challenging given the prevalence of WiFi frames that are received below the CCA threshold. Fig. 7 shows the CDF of the WiFi received power measured using our testbed (described in §6), when all APs transmit at 10dBm. Up to 55% of all WiFi frames are received below the -62dBm threshold.

3.2 LTE Duty-Cycling is Insufficient

Simple on-off duty cycling has been proposed as one possible solution to the LTE-WiFi coexistence problem [11]. However, our experiments show that such an approach does not address the coexistence challenges completely. Furthermore, even for the same duty cycle, the impact of LTE on WiFi performance depends on the duration of each on-cycle. This highlights the difficulty in selecting an appropriate duty cycle.

Fig. 5 shows the performance of WiFi under 50% LTE duty cycling. We compare the performance of WiFi under four different duty-cycling configurations, all of which have the same 50% duty cycle: 2 subframes on, 2 subframes off; 4 subframes on, 4 subframes off; 10 on, 10 off; and 50 on 50 off. While the larger on-cycles (tens of subframes) are representative of typical LTE-U, we also study the impact of shorter on-cycles (2-4 sub-frames) on LTE-WiFi co-existence. We note that even though our study focuses on LTE-U, these shorter subframe durations are on the same granularity of LTE-LAA transmissions.

Observe that over a wide range of SIR values, the WiFi throughput is dependent on the duration of the on and off periods — the longer the on period (and correspondingly the off period), the higher the throughput.

To understand why this is so, recall that WiFi uses an exponential backoff procedure whenever it encounters collisions — the AP will double its maximum contention window upon each frame loss event. During the periods of WiFi and LTE interference, we observe that the contention window can be increased to as much as 300 time slots, or $1.8ms$. During the LTE off periods, the contention window sizes can be reduced to 7 time slots. When the LTE transitions from its on to off period, the WiFi AP must still complete its long backoff (incurred during the LTE on period) before it can commence interference-free transmissions. Hence, WiFi cannot adequately take advantage of short LTE off periods as it will spend a large proportion of that time in a backoff state. This indicates that LTE-U, which employs long LTE off periods, is more beneficial to WiFi at the outset. We note that similar results have also been presented in [4].

However, we caution that throughput does not reveal the complete picture, and present the impact on latency. For any WiFi

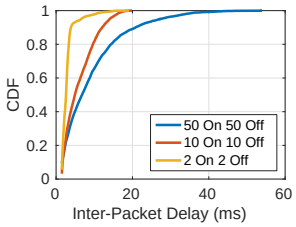

Figure 6: CDF of WiFi inter-packet delay

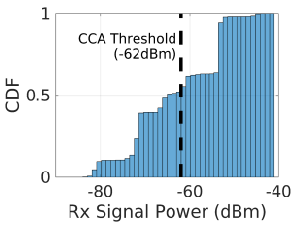

Figure 7: Receive signal power distribution.

device that is within the energy-sensing CCA threshold of LTE, its transmission will be blocked by an LTE on period. Hence, for the same duty cycle, a longer LTE off period (and consequently, a longer LTE on period), the larger the latency in the WiFi transmission. Fig. 6 shows the inter-packet delay of only the WiFi packets received by the STA during the LTE on periods. The 80^{th}-percentile delay of a WiFi during 50-on-50-off LTE is $14.2ms$ and can be as long as $55ms$. This is significantly longer then the inter-packet delays of 2-on-2-off and 10-on-10-off LTE. In all of these cases, the WiFi inter-packet delay during LTE off periods does not exceed $1ms$. The large WiFi inter-packet delay jitter will have a even bigger impact on the corresponding mean overall packet delays.

Thus, even in the absence of hidden/exposed terminals, LTE-U cannot provide both good WiFi efficiency, while maintaining short-term fairness and latency. Unfortunately, these are factors that are critical for co-existence. This deficiency of LTE-U is due to its inability to detect WiFi carriers and notify channel reservations, and plagues the efficiency of the coexistence process.

These observations, when taken together with those in §3.1, clearly emphasize the *need for homogeneous carrier sensing/notification feature LTE and WiFi to address the limitations of energy sensing.*

4. ULTRON DESIGN

In this section, we present the design of a novel, standards-compliant **U**nlicensed **LTE** **R**adi**O** **N**ode, or ULTRON. ULTRON operates at the LTE base station and works seamlessly with legacy LTE clients. ULTRON employs two key components: (i) *WiFi embedding* that allows an LTE BS to transmit a WiFi signal (channel reservation notification) prior to its transmission on each unlicensed channel by masking itself as LTE data, and (ii) *WiFi sensing* that allows a single WiFi interface to work synergistically with the LTE traffic manager/scheduler to maximize the benefits of WiFi carrier sensing to all the unlicensed channels.

Fig. 8 shows a high-level design of ULTRON. ULTRON is designed for downlink-only transmission from multi-mode unlicensed LTE small cells that are equipped with a LTE and a single WiFi interface. WiFi sensing switches the single WiFi interface between each of the LTE CCs. The WiFi interface dwells on each CC for some T amount of time, during which it collects WiFi traffic statistics (average number of contending nodes, and channel occupancy). These statistics are used by the LTE eNodeB for coarse-timescale scheduling of LTE UEs across the multiple CCs. Before each transmission on a CC, the eNodeB uses the WiFi embedding component to embed a CTS-to-Self into an LTE subframe to reserve the channel for the duration of the LTE transmission.

In the rest of this section, we present the various challenges facing ULTRON's design and the mechanisms to address them.

4.1 Embedding WiFi Signals in LTE

WiFi embedding enables a LTE eNodeB to notify other WiFi/LTE nodes in each unlicensed channel of its impending transmission as

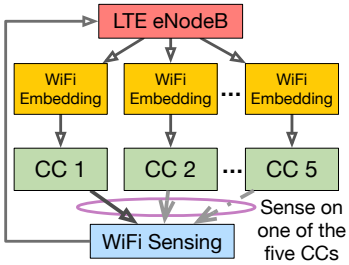

Figure 8: ULTRON architecture with 5 unlicensed component carriers (CCs).

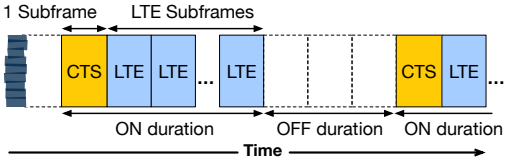

Figure 9: Embedding a WiFi CTS into LTE.

well as its duration. WiFi embedding is protocol-independent — it can be implemented by updating only the scheduler, and does not require changes to the LTE standard. This will ensure that the LTE standard is independent of the WiFi specification, thus simplifying the adoption of ULTRON in existing LTE devices.

Specifically, when ULTRON transmits on the unlicensed channel, it embeds the WiFi CTS-to-Self (hereafter called CTS) packet into the first sub-frame (as shown in Fig. 9). The CTS packet will carry a NAV value that reserves the channel for the remaining duration of the LTE transmission (including the current sub-frame). Before this can be realized in practice, two critical challenges must be overcome: (i) *what data bits need to be sent to the PHY from the scheduler*, and (ii) *where and how to place this data in the OFDMA frame*, such that when the bits are passed through the LTE encoding process, it will generate the desired time-domain, analog CTS waveform that is recognized by any WiFi node? In the rest of this section, we will show how ULTRON uses the LTE scheduler to carefully construct (reverse-engineer) the time-frequency bits, and induce a standard LTE PHY to transmit a WiFi CTS packet.

4.1.1 LTE PHY Encoding Process: A Primer

We describe the encoding process of an LTE PHY for a 20MHz channel. We emphasize that a 20MHz LTE channel is actually

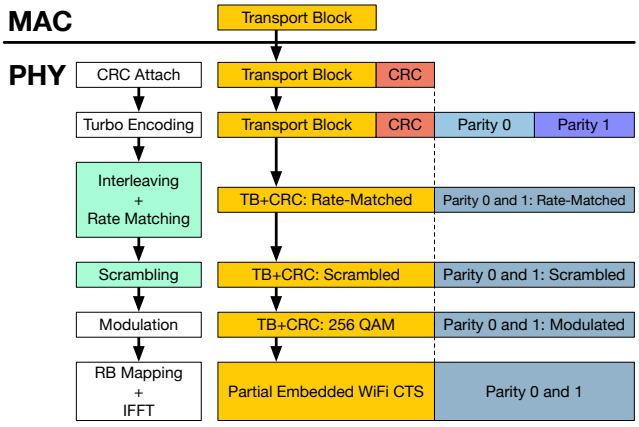

Figure 10: LTE PHY encoding for each UE.

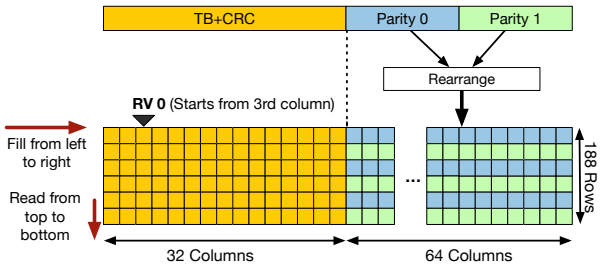

Figure 11: Interleaving and rate matching.

transmitted at 30.72MHz using 2048 OFDM subcarriers. Of these 2048, only 1200 subcarriers (spanning 20MHz) are used for data transmission.

Fig. 10 shows the steps taken by the LTE PHY to encode a single transport block (TB) into the resource blocks (RBs) allocated to a single UE. **(1)** The scheduler sends a TB of a specific size (e.g. 5992 bits), to be transmitted using a specific MCS (e.g. MCS 26 corresponds to 256-QAM), to the PHY[2]. The PHY first generates and appends a 24-bit CRC to this TB, and sends the TB+CRC to the turbo encoder. **(2)** The turbo encoder encodes the TB+CRC with a fixed rate of 1/3 — two parity bits are generated for every TB+CRC bit. The PHY uses a systematic turbo code [12], hence the output includes the original TB+CRC along with the two groups of parity bits (i.e. Parity 0 and Parity 1 in Fig. 10). **(3)** The TB+CRC, along with the parity bits, are passed to a interleaver and rate matcher. The interleaver permutes (i.e. "mixes") the TB+CRC and parity bits to guard against burst errors over the wireless channel. The rate matcher then discards selected bits from both the TB+CRC and parity bits, using a process known as *puncturing*, so that the number of transmitted bits is suitable for the intended bitrate of the LTE PHY. **(4)** The interleaved/rate-matched bits are then scrambled using a subframe-specific scrambling sequence. This is a bitwise multiplication between the interleaved/rate-matched bits and the scrambling sequence. The output of this scrambling step is then modulated into I/Q symbols using the MCS selected data, and mapped (placed) into the subcarriers of the OFDM symbols. After this mapping step, the PHY performs an IFFT, cyclic prefix construction and windowing to obtain the time-domain LTE signal.

Interleaving and Rate Matching: Interleaving and rate matching are critical to the construction of the CTS embedding process. We now describe them in detail with an example (see Fig. 11).

The 6016 bits of TB+CRC data (5992 TB and 24 CRC bits) are first rearranged into a matrix of size 188×32. The TB+CRC is written to the matrix row-wise, starting from the top left-hand corner and moving left to right, one row at a time. The LTE PHY interleaves the TB+CRC bits by exchanging pairs of columns in the 188×32 TB+CRC matrix. In the parlance of the LTE PHY, this step is known as *sub-block interleaving*, and the inter-column permutation pattern can be found in [13]. The Parity 0 and 1 are separately permuted in a similar manner, then rearranged into a combined 188×64 matrix. The matrices corresponding to the TB+CRC and Parity bits are then concatenated as shown in Fig. 11.

Rate matching is then achieved by discarding selected columns from this permuted matrix. The output of the rate-matching step is obtained by reading this concatenated matrix column-wise. The starting column depends on the *redundancy value (RV)* used by the PHY. This RV is a parameter that controls the LTE Hybrid ARQ (HARQ) processing, and is determined by the LTE scheduler. Since we are only interested in embedding a CTS frame, we

[2]The TB size depends on the MCS and the number of RBs. The table of TB sizes can be found in 3GPP specification TS36.213.

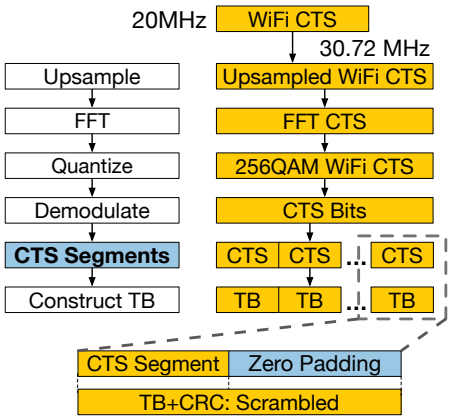

Figure 12: Reverse-engineering LTE PHY encoding.

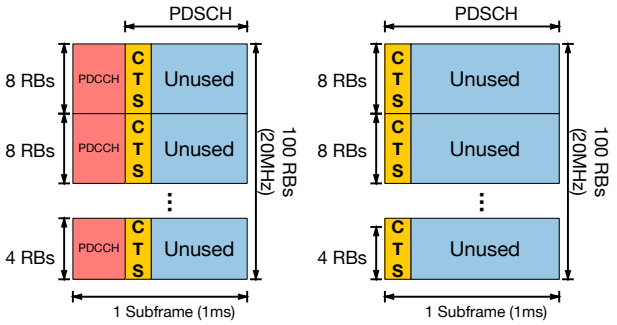

(a) No cross-carrier scheduling. (b) With cross-carrier scheduling.

Figure 13: WiFi CTS embedded in the first LTE symbol of PDSCH.

can select any convenient RV value. Here, we select RV $= 0$. With this RV value, the LTE PHY will start reading the bits in the matrix column-wise, starting from the top of the third column and moving top to bottom, left to right. This rate-matched output is then sent to the scrambler.

4.1.2 *Reverse-Engineering the Time-Domain CTS*

The time-domain signal that is output by the LTE PHY encoder contains the WiFi CTS frame. Essentially, this time-domain signal appears to a WiFi device as a regular WiFi CTS frame. It should thus be decodable by this WiFi device using a standards-compliant WiFi protocol stack without any additional effort. In order to achieve this, ULTRON must *reverse-engineer* the LTE PHY stack. Given this desired time-domain signal (i.e. the CTS frame), we must determine the exact bit sequence from the scheduler that will result in this time-domain signal after it has been encoded by the LTE PHY as described in §4.1.

CTS in Frequency-domain: Fig. 12 shows the steps adopted by ULTRON to embed the WiFi CTS. We begin with the I/Q data of a 20 MHz time-domain WiFi CTS waveform, that is constructed using 781 I/Q samples. Note that since we are using this to notify our channel reservation to neighboring WiFi devices, only the NAV parameter in the CTS is important. These time-domain I/Q samples are then up-sampled to 30.72MHz, equal to the actual sampling rate of a 20MHz LTE channel, to obtain 1200 I/Q samples. Recall that the LTE PHY uses 1200 out of 2048 subcarriers/frequency-domain samples for data transmission. Hence, we zero-pad this upsampled I/Q stream to obtain 2048 time-domain samples and perform a 2048-bin FFT to obtain the frequency domain representation, of the up-sampled CTS.

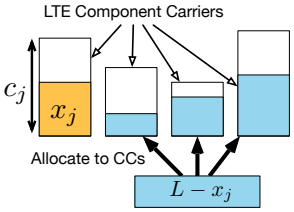

Figure 14: ULTRON traffic allocation.

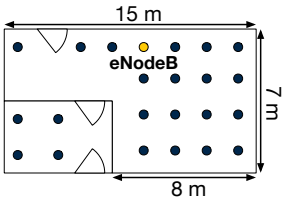

Figure 15: $15 \times 7m$ test environment.

Constructing the LTE Time-Domain Signal: The upsampled CTS is not yet a valid LTE time-domain signal as its I/Q values do not conform to any modulation constellation. We thus need to quantize each I/Q value of this upsampled CTS to the nearest constellation point. In order to minimize quantization noise, we quantize the CTS signal to a 256-QAM constellation point. We then demodulate this 256-QAM data to obtain the corresponding bits.

Constructing the Transport Blocks: In order for the CTS to be correctly embedded into the first symbol of the PDSCH, a total of 13 *transport blocks*[3] are needed. This requirement is due to the data block limits imposed by LTE, and will be explained in §4.1.3. We thus partition these demodulated bits into 13 *CTS segments*, one segment for each transport block.

A transport block is constructed from each CTS segment by reverse-engineering the encoding process. Fig. 12 shows how a CTS segment from Fig. 12, corresponds to the scrambled TB+CRC output (yellow box) from the scrambling step in Fig. 10. The original TB can be constructed by inverting three of the PHY stages, namely scrambling, rate matching and permutation, and interleaving, in that order.

De-scrambling: Recall that the scrambling sequence is subframe specific. ULTRON regenerates this subframe-specific scrambling code, and recovers the pre-scrambled signal by performing a bit-wise addition in GF2 (i.e. a Galois field of 2 elements) of the this scrambling code and the zero-padded CTS segment.

De-rate-matching: ULTRON fixes the redundancy value at RV $=$ 0. We write the de-scrambled output to the first 32-columns of the rate-matching matrix (in Fig. 11), starting from the top of the third column, and moving from top to bottom, left to right. We fill the first two columns with zeros. Note that we begin filling the matrix from the third column (not the first) since the CTS will be embedded with the parameter RV $= 0$.

De-interleaving: The interleaver is a one-to-one map using a fixed permutation pattern. Hence, we can easily construct a reverse map from the interleaved TB+CRC back to the original TB. The de-interleaved TB+CRC can be recovered by reading the matrix row-wise, starting from the top left-hand corner of the matrix, and moving from left to right, top to bottom. At this point, the ULTRON scheduler has constructed the TB that needs to be sent to the PHY for generating a CTS packet.

Impact of an Unknown CRC: Observe that ULTRON is unable to recover all the correct bits in the TB. This is due to the fact that the interleaver and rate matcher mixes the TB and CRC bits, and the entire 24-bit CRC is unknown to the scheduler. However, the actual number of bits affected is minuscule. Note that each CTS segment is only 768-bits long and occupies four full columns and the top 16 bits of the fifth column. Using the permutation pattern from [13], we can determine that only 3 bits will be affected — a mere 0.39% of the CTS segment. Furthermore, the original WiFi CTS frame is BPSK-modulated, and thus is highly resilient

[3]A *transport block* is synonymous with a WiFi MAC PDU, and is the smallest unit of data bits sent by the LTE scheduler.

to noise. Our experiments indicate that over 98% of the embedded CTS frames can be received at SNRs as low as -83dB. Hence, the impact of the unknown CRC is negligible.

Encoding Overhead. Each TB is subject to one of ten subframe-specific scrambling codes. Since ULTRON is restricted in the number of subframes per transmission (e.g. $\leq$4), all possible TBs can be pre-computed.

4.1.3 Placement of the CTS

Fig. 13a shows how the 13 constructed transport blocks (corresponding to the CTS) are placed in a 20MHz LTE subframe. Only one WiFi CTS is embedded in each subframe. This CTS is placed only in the first OFDMA symbol of the PDSCH. The rest of the subframe is left unused, since any data placed there will be corrupted by the embedded CTS. The PDCCH of the LTE subframe carrying the embedded CTS is configured to use a single OFDMA symbol, which is the for the PDCCH. This ensures that the CTS is transmitted as soon as possible in the second OFDMA symbol, which does not carry any LTE reference signals. Hence, all subcarriers in this second symbol will be used to embed the CTS frame.

Why use 13 TBs? We use 13 TBs, one TB per virtual UE, to embed the CTS. Note that these UEs can be virtual UEs as the LTE PHY does not validate the UE IDs. One might consider generating a single large TB (covering the entire CTS) destined to a single UE. The modulated output of this large TB will be mapped to all 1200 subcarriers of the LTE channel. However, when the size of the TB+CRC exceeds the LTE *codeblock* limit of 6144 bits, the PHY will partition the TB+CRC (not just the TB) into smaller code-blocks. Each code-block undergoes Turbo encoding, interleaving, rate-matching and scrambling by the PHY (i.e. similar to the process in Fig. 12). The key difference is that multiple scrambled codeblocks, along with their parity bits, are concatenated before being modulated and mapped to the subcarriers. As a result, *a single OFDMA symbol will contain both CTS and parity bits*. This will severely distort the final time-domain CTS waveform.

12 of these virtual UEs have 8 RBs and use TBs of 5992 bits, which is the largest allowable TB size that is below the 6144-bit limit. This ensures that the ratio of number of TB bits (which ULTRON has control over) to the 24 CRC bits (which ULTRON has no control over) is the highest, thus minimizing the impact of noise from interleaved CRC bits. The last UE has only 4 RBs, and thus has a TB size of only 2984 bits. However, the additional impact of CRC noise from it on the entire CTS (13 TBs) is limited.

4.1.4 Reducing CTS Latency

Transmit Latency. In a standard LTE subframe, as shown in Fig. 13, the first 1 to 3 OFDM symbols is occupied by the PDCCH. This means that the CTS can be transmitted with a latency equivalent to 3 WiFi CTS frame airtimes after the start of the LTE transmission.

This latency can be reduced using the *cross-carrier scheduling* feature of LTE. Cross-carrier scheduling enables the allocation of unlicensed RBs to be carried in the PDCCH of the licensed channel. Note that cross-carrier scheduling is supported from LTE Release 10 onwards, and can be used with without any modifications to the design of ULTRON.

With cross-carrier scheduling, the unlicensed CC will no longer contain a PDCCH and the PDSCH will begin from the very first OFDM symbol of the unlicensed LTE subframe. This is illustrated in Fig. 13b. Hence, the CTS will be transmitted with no latency after the start of the unlicensed LTE transmission.

Encoding Latency. ULTRON uses an unmodified LTE PHY. The PHY encoding process is typically implemented in the FPGA and thus has a negligible overhead.

4.2 Scalable WiFi Sensing in LTE

ULTRON has to equip LTE with WiFi carrier sensing for every unlicensed channel it operates on without incurring any changes to the LTE PHY specification. While providing a WiFi sensing module (through a WiFi interface) for every un-licensed channel is an option, this is not a scalable solution. Note that LTE-advanced can currently aggregate up-to five carriers/channels (licensed and/or unlicensed) through carrier aggregation, which in future 5G networks will only increase in number. On the other hand, LTE small cells are increasingly being equipped with a WiFi interface (although just one) for features such as dual connectivity [5] (communicating with both macro and small cell concurrently), which in turn could be re-purposed for WiFi carrier sensing. However, a single WiFi interface is not sufficient to cover the numerous un-licensed channels that LTE can operate on simultaneously (i.e. up-to 100 MHz in carrier aggregation). Even newer 802.11ac interfaces with larger bandwidths can cover at most two 20MHz LTE CCs at a time. ULTRON assumes that each WiFi interface can cover only one LTE CC, but the results here can be generalized to 802.11n/ac interfaces that can support up to two CCs.

In addressing this challenge, ULTRON intelligently leverages the single WiFi interface to maximize the benefits of carrier sensing to all the un-licensed channels. It accomplishes this through a two-step process. First, it deploys the WiFi interface on each of the un-licensed channels periodically (for t_s seconds) to sense and estimate the potential impact of not having WiFi carrier sensing on that channel. Second, it uses this information to solve a joint optimization problem that helps determine the particular un-licensed channel, where the WiFi interface's channel sensing feature would be most beneficial, along with the appropriate amount of LTE traffic load that must be placed on each of the un-licensed channels. ULTRON operates using this optimized configuration for an epoch of T seconds ($T >> t_s$) before sensing each of the channels again so as to track the traffic variations in the network and determine the appropriate configuration for the next epoch.

4.2.1 Estimating the Impact of Carrier Sensing

When the WiFi interface is deployed on an un-licensed channel i, its carrier sensing feature is used to listen to the WiFi packets on the channel and collect the following coarse-grained statistics over the measurement period of t_s seconds: (i) average number of contending nodes, n_i; (ii) channel occupancy fraction, y_i; and (ii) channel occupancy fraction due to traffic that cannot be detected with just energy sensing but can cause interference, $\hat{y}_i$ (i.e. those packets falling below the energy sensing threshold but over the carrier sense threshold, say [-80,-60] dBm).

4.2.2 Maximizing the Benefits of Carrier Sensing

Problem Formulation: While the WiFi interface can provide carrier sensing capability to only one of the un-licensed channels in real-time, the information collected by it during the measurement period can be used to minimize the impact from lack of carrier sensing in other channels. This is achieved through an intelligent allocation of LTE traffic to the unlicensed channels, which uses LTE's flexible scheduling mechanism that allows traffic to be jointly scheduled across multiple channels. The WiFi interface assignment as well as the LTE traffic allocation to all the unlicensed channels is made jointly through the following optimization prob-

lem, Interface Assignment and Traffic Allocation (IATA).

$$\text{IATA:} \quad \max \sum_i z_i x_i + (1 - z_i) \max\{0, x_i - \hat{y}_i T R_i\}$$

$$\text{s.t.} \quad \sum_i z_i \leq 1, \sum_i x_i \leq L,$$

$$x_i \leq R_i \max\{T(1 - y_i), \frac{T}{n_i}\}, \forall i$$

where R_i denotes the average LTE transmission rate (bits/s) on channel i as measured during LTE transmissions from the previous epoch, and L is the total LTE load (in bits per epoch T s) to be sent across the unlicensed channels. y_i and $\hat{y}_i$ are as defined in Section 4.2.1. z_i and x_i are output variables indicating the allocation of WiFi interface to channel i ($z_i \in \{0, 1\}$) and the amount of LTE traffic allocated to channel i respectively. The objective is to maximize the successful transmission of the LTE traffic load on the un-licensed channels. While the LTE traffic assigned to the channel that receives the WiFi interface will be protected ($z_i x_i$), the lack of the WiFi interface on a channel will result in collision and hence loss of LTE traffic up-to an amount ($\hat{y}_i T R_i$) determined by the aggregate un-detected, interference-causing traffic, $\hat{y}_i T$) in the worst case. The first two constraints indicate the binary nature of the interface assignment variable, and the bound on the LTE traffic that needs to be served on the unlicensed channels respectively. The third constraint captures the maximum amount of LTE traffic that can be allocated to a channel based on its fair share of channel occupancy (i.e. $\frac{T}{n_i}$ when channel is saturated) as well as un-used channel occupancy (i.e. $T(1 - y_i)$ when not saturated).

Algorithm: The above non-linear optimization problem can be easily shown to be NP-hard. Even when the WiFi interface assignment is given, the remaining problem of LTE traffic allocation is NP-hard and corresponds to a modified version of a knapsack problem. Hence, ULTRON adopts the following greedy heuristic that is computationally light-weight but is highly effective in practice.

Step I: Solve IATA M times, where M is the number of un-licensed channels. In each iteration j, solve IATA assuming the WiFi interface is assigned to channel j (i.e. $z_{i=j} = 1; z_{i \neq j} = 0$) and determine the LTE traffic allocation to all the unlicensed channels. Let IATA(j) be the resulting objective value with the corresponding output being $(\mathbf{z}_j, \mathbf{x}_j)$, where $\mathbf{z}_j = [z_1, \ldots, z_M]$ and $\mathbf{x}_j = [x_1, \ldots, x_M]$.

Step II: Pick the solution $(\mathbf{z}^*, \mathbf{x}^*)$ that yields the highest objective value, i.e.

$$(\mathbf{z}^*, \mathbf{x}^*) = \arg \max_{(\mathbf{z}_j, \mathbf{x}_j)} \text{IATA}(j) \tag{1}$$

LTE Traffic Allocation: We now elaborate on Step I, where given the assignment of the WiFi interface to one of the un-licensed channels (say j), we aim to determine the LTE traffic allocation across all the unlicensed channels. Note that with the protection of WiFi carrier sensing on channel j, it is straight-forward to see that maximum traffic that is allowed by that channel (or limited by the input traffic) would be allocated to it, i.e. $x_j = \min\{L, R_j \cdot \max\{T(1 - y_j), \frac{T}{n_j}\}\}$. The problem now reduces to determining the allocation for the remaining un-licensed channels that do not have the carrier sensing feature.

$$\begin{aligned} \text{IATA}(j): \quad & \text{Maximize} \ \textstyle\sum_{i \neq j} \max\{0, x_i - \hat{y}_i T R_i\} \\ & \text{subject to} \ \textstyle\sum_{i \neq j} x_i \leq L - x_j \\ & x_i \leq R_i \cdot \max\{T(1 - y_i), \tfrac{T}{n_i}\}, \ \forall i \neq j \end{aligned} \tag{2}$$

The problem is visually presented in Fig. 14, where we aim to pack (allocate) the remaining LTE traffic ($L - x_j$) into bins (channels) with capacity $c_i = R_i \max\{T(1 - y_i), \frac{T}{n_i}\}$. When a bin i is chosen, it incurs a cost (loss due to lack of carrier sensing) of $\ell_i = \hat{y}_i T R_i$. Hence, the goal is to pick a subset of the bins such that maximum amount of LTE traffic can be successfully served, while accounting for the loss. Being a variant of the minimum cost knapsack problem, it is hard to solve this problem optimally. Hence, we design the following heuristic, where we sort the channels based on their normalized traffic loss, and use it to determine an appropriate subset of channels for allocation.

Step 1: Sort the channels in non-decreasing order based on the metric $\frac{\ell_i}{c_i}$. Channels that provide higher capacity with low loss will appear earlier in the order. Let the ordering be, $\ell_1/c_1 \leq \ell_2/c_2 \leq \ldots \leq \ell_M/c_M$.

Step 2(a): Let m_1 be the index, where the aggregate capacity of channels in the ordering just exceeds the input traffic, i.e.

$$\sum_{i=1}^{m_1} c_i < L - x_j \leq \sum_{i=1}^{m_1+1} c_i$$

Define $A_1 = (1, 2, \ldots, m_1)$. Hence, channels corresponding to indices $(1, 2, \ldots, m_1, m_1 + 1)$ form a potential solution with a net successful LTE traffic allocation of $L - x_j - \sum_{i=1}^{m_1+1} \ell_i$.

Step 2(b): Similar to $m_1 + 1$, if the subsequent indices ($k \in \{m_1 + 2, m_1 + 3, \ldots\}$) provide net capacity that exceeds the input traffic, i.e. $\sum_{i=1}^{m_1} c_i + c_k \geq L - x_j$, then each of the channels corresponding to k are also potential solutions, whose value is recorded.

Step 2(c): Let m_2 be the first next index, when net capacity falls below the input traffic, i.e. $\sum_{i=1}^{m_1} c_i + c_{m_2} < L - x_j$. Let m_3 be the subsequent index, where the net capacity again first exceeds the input traffic (similar to 2(a)),

$$\sum_{i=1}^{m_1} c_i + \sum_{i=m_2}^{m_3} c_i < L - x_j \leq \sum_{i=1}^{m_1} c_i + \sum_{i=m_2}^{m_3+1} c_i$$

Let $A_2 = (m_2, m_2 + 1, \ldots, m_3)$. Now, $A_1 \cup A_2 \cup \{m_3 + 1\}$ forms a potential solution as well.

This process continues by repeating Step 2 over the remaining channels in the ordered list. Once the entire list is traversed, the best of all the potential solutions is picked as the solution delivering the maximum successful LTE traffic allocation. The algorithm is light-weight and incurs a time complexity of $O(M \log M)$.

The above joint optimization demonstrates the synergistic role of LTE traffic placement (scheduling) and WiFi interface assignment to minimize the impact from lack of WiFi carrier sensing on all-but-one un-licensed channels.

5. IMPLEMENTATION

We evaluate ULTRON using both real-world experiments with USRPs and WARP SDRs, and trace-based simulations. Fig. 15 shows the $15 \times 7m$ office environment along with the UE/WiFi node locations used to conduct the measurements and trace collection. All LTE and WiFi nodes are operated using 20MHz WiFi channels in the 5GHz ISM band. Note that a 20MHz LTE channel has a sampling rate of 30.72MHz, with the middle 20MHz used for carrying data, while the remaining 12.36MHz are unused.

5.1 Our Unlicensed LTE PHY Design

Our LTE PHY differs from that of existing LTE in two key aspects: (a) the LTE PHY supports a per-subframe activation and de-

activation, and (b) the UE can commence decoding from any subframe, rather than only subframes 0 and 5 as described in §2.

Unlicensed LTE PHY Encoder. Fig. 2 illustrates our design of a LTE PHY. Each subframe in our LTE PHY carries both a PSS and SSS signal. We assign a different cell ID to each subframe so that a UE can uniquely synchronize with and identify each subframe in a transmission. This remove the dependence on sub-frames 0 and 5 for synchronization signals. The encoding/decoding of the PDCCH (control channel) and PDSCH (data channel) of each subframe follows the LTE Release 10 standard. The PHY is generated using the MATLAB LTE System Toolbox [14] to ensure strict compliance with this LTE specification.

Unlicensed LTE PHY Decoder. Each UE is a USRP B210. The UE logs the received I/Q data from the unlicensed band. This I/Q data is then processed offline by the MATLAB LTE System Toolbox to recover the transmitted LTE subframes. Our testbed does not support real-time UE decoding.

Why Is This Design Necessary? These features are required because our LTE testbed only consists of unlicensed bands, and does not support licensed LTE channels. In actual unlicensed LTE networks, the licensed band provides always-on control signaling for the unlicensed band. In our testbed, we must ensure that the unlicensed channels alone provide all necessary signaling while being resilient to WiFi interference.

WiFi frames can be shorter than a LTE subframe. For example, a 1.4KB WiFi data frame transmitted at 54Mbps has an airtime (including MAC overhead) of under $0.5ms$. Hence, it is likely that WiFi interference across a single LTE transmission burst of several subframes ($1ms$ per subframe) is time-varying. If subframes 0 or 5 are undecodable due to WiFi interference, the UE will lose time and frequency synchronization with the LTE transmission. Several subsequent subframes will thus be undecodable by the UE, even if they are free of WiFi interference. Our LTE design ensures that each subframe can be individually decoded at the UE, regardless of interference in other subframes. The LTE throughput (over the unlicensed channel) measured on our testbed will thus be as close as possible to that obtained on a real LTE network with both licensed and unlicensed bands.

5.2 Testbed Design

Our testbed uses USRP B210s as the LTE eNodeB and UEs, and WARPv3 nodes, running the 802.11 Reference Design v1.4.5 [15], as WiFi APs and STAs. A total of 3 B210s (1 eNodeB and 2 UEs), and 7 WARPv3 devices (different combinations of APs and STAs) are used. The eNodeB is fixed at a single location, shown in Fig. 15, while the other USRP UEs are placed at the other indicated locations. We only use two USRP UEs at any one experimental run, and we repeat the experiments, each time with the UEs placed at a different pair of locations.

Unlicensed LTE. The eNodeB transmits the LTE PHY that is generated as described above. ULTRON prepends each LTE transmission with a specially constructed LTE frame that embeds a WiFi CTS frame, as described in §4. Due to the limitations of our testbed, we do not support a real-time energy-sensing CCA mechanism. Instead, we emulate the CCA behavior using a random backoff technique. Before every transmission, ULTRON selects a random backoff duration of $b \in [0, 1000)\mu s$. The LTE transmission commences at the end of this backoff period. ULTRON maintains time-synchronous nature of LTE subframe time boundaries.. This means that this transmit opportunity can occur between subframe time boundaries. Hence, ULTRON transmits a reservation signal from the time of a transmit opportunity until the next subframe time boundary. The unlicensed subframe(s) will be transmitted thereafter. The number of LTE subframes in each transmission is randomly selected between 1 and 4. This static backoff is a is sufficient for us to understand the behavior of ULTRON that arises due to the lack of symmetric channel sensing. The exact backoff duration does not affect the nature of the results.

WiFi. Our testbed uses two WiFi APs, one with three STAs and the other with two (for a total of seven WARP devices). We place the APs and STAs randomly on the different marked locations of Fig. 15. During experiments, these two WiFi networks share the same wireless contention domain. We use constant bitrate, fully backlogged downlink UDP traffic for all our testbed experiments.

6. EVALUATION

We evaluate ULTRON using two approaches. Real-world testbed measurements are used to valuate the performance of CTS-to-Self embedding in an LTE subframe, as described in §4. Our testbed only allows small-scale evaluations using one eNodeB, two UEs and two WiFi networks. We thus augment our evaluations by using larger scale trace-based simulations to study the full ULTRON design including both the WiFi embedding and sensing features.

SINR Measurement. The SINR at a UE is obtained by separately measuring the signal, interference and noise powers at the UE from the logged I/Q data. These same quantities are directly reported by the 802.11 Reference Design.

Rate Adaptation. Our LTE and WiFi networks do not support real-time rate-adaptation. We measure the throughput of WiFi and LTE by repeating the experiments over multiple MCS values, and selecting the highest throughput obtained.

6.1 Real-World Evaluation

6.1.1 Accuracy of CTS Embedding

We evaluate the decodability of the CTS embedding within an LTE subframe. The eNodeB is placed at a fixed location as shown in Fig. 15 and continuously broadcasts LTE subframes, without backoff, with embedded CTS packets. We decode the CTS packets using a WARPv3 WiFi STA placed at twenty different locations throughout the testbed. No interfering WiFi is present.

Fig. 16 shows the probability at which the CTS frames are successfully received by the WiFi STA. When the receive power at the WiFi STA is greater than -52dBm, all CTS frames are correctly received. Observe that even when the average received signal energy at the STA falls to -85 dBm, which is over 20dBm below the energy-sensing threshold, over 95% of the embedded CTS frames are still correctly received and decoded. Hence, *the embedded WiFi CTS frames in an LTE subframe is highly resilient and can suppress transmissions of WiFi devices outside the energy-sensing threshold.* The CTS frame embedding process is thus an effective mechanism for channel reservation for unlicensed LTE transmissions.

6.1.2 Benefits of CTS Embedding on WiFi

We evaluate the benefits of CTS embedding using both the implementation of ULTRON together with the two WiFi networks, as described in §6. Fig. 17 shows the average WiFi throughput, measured at the STAs, when LTE is operating with and without ULTRON. When LTE is operating without ULTRON, LTE subframes are transmitted after a random backoff, but do not contain the embedded CTS-to-Self. In this configuration, LTE subframes interfere with on-going WiFi transmissions. In the face of such overwhelming LTE-WiFi collisions, the WiFi backoff window sizes increase dramatically, resulting in a sharp drop in WiFi throughput. On the other hand, ULTRON forces the WiFi transmitters to freeze for the duration of the NAV value, without any impact on its backoff win-

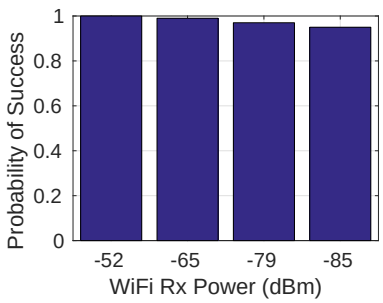

Figure 16: CTS reception success probability.

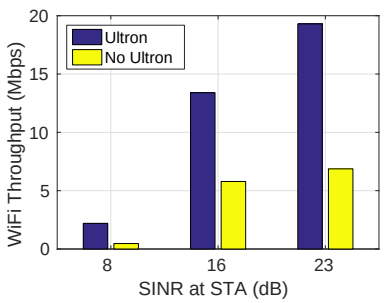

Figure 17: WiFi throughput (WiFi MCS 4).

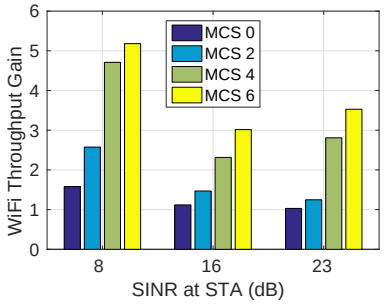

Figure 18: WiFi throughput gain.

dow. The WiFi backoff window sizes can thus be up to 20× smaller with ULTRON.

Fig. 18 shows that ULTRON can increase the WiFi throughput by up to 2.5× (from 7 to over 18Mbps) at 23dB SINR. Further, the gains have a dependence on the WiFi MCS, and can reach over 5× at 8dB SINR with MCS 6. For a given SINR value, greater gains are seen at higher WiFi MCS values since WiFi performance becomes more sensitive to interference at these higher bitrates. Hence, ULTRON *achieves effective WiFi-LTE co-existence and significantly increases WiFi throughput even in the presence of LTE.*

6.1.3 *Benefit of CTS Embedding on LTE*

Fig. 19 shows the mean LTE throughput achieved with and without ULTRON (i.e. without CTS embedding). The mean LTE throughput increases with increasing SINR at the UE. However, observe that the LTE throughput with ULTRON is significantly greater than without. When no CTS is embedded, transmissions from the eNodeB will collide with WiFi transmissions that are out of energy-sensing range. With CTS embedding, ULTRON can improve the achievable LTE throughput by up to 6× — at 3dB SINR, ULTRON increases the LTE throughput from 1.3 to 7Mbps. Generally, the throughput gains achieved by ULTRON is greater at lower SINR values. This is low SINR at a LTE UE is due to high interference power from a neighboring WiFi device. Suppressing these strong interfering WiFi transmissions will give the UE a large improvement in channel quality. This behavior is observed even across different LTE MCS values, as seen in Fig. 20.

6.2 ULTRON On a Larger Scale

We evaluate the performance of ULTRON scalable carrier sensing, when used together with WiFi CTS embedding, in a larger system using trace-based simulations. The full ULTRON system is compared to four other protocols:

(A) One-WiFi-Per-CC utilizes the CTS embedding and LTE traffic allocation policies of ULTRON, but has one WiFi sensing interface for each LTE CC. This system offers more hardware resources for LTE-WiFi co-existence and provides an upper bound on the achievable performance of ULTRON.

(B) ULTRON-EqAlloc utilizes CTS embedding but allocates an equal share of the LTE throughput to each of the five LTE CCs. Comparing ULTRON with ULTRON-EqAlloc performance highlights the benefits of traffic allocation on a larger scale.

(C) Max-WiFi uses the CTS embedding feature of ULTRON. However, it assigns the WiFi sensor to the channel with the highest WiFi throughput (load), and allocates the LTE traffic according to (2). In contrast, ULTRON determines the ideal sensor placement and traffic allocation jointly. Comparing ULTRON with Max-WiFi highlights the benefits of correctly utilizing the single WiFi sensing interface on a eNodeB.

6.2.1 *Experiment Setup*

Trace collection. We collect 10-minute WiFi traces over five 5GHz ISM bands (to emulate 5 LTE CCs) using the WARPv3 nodes with 3 APs and 3 STAs (one STA per AP). All APs operate concurrently within the same wireless contention region. We also place a WARP node at the location of the LTE eNodeB, so as to obtain a snapshot of all WiFi traffic from the point of view of the eNodeB.

Trace synchronization. We synchronize the clocks on the different WARP devices so that all packet timestamps, even from different WARP boards, are directly comparable to each other. We then simply merge all traces together, while maintaining the global transmit/receive order of all WiFi frames.

LTE transmit opportunities from traces. We use the WiFi traces taken at the eNodeB position to emulate the CCA mechanism. These packet traces also include the RSSI of each received frame. Unlicensed LTE transmit opportunities lie in time intervals with either no WiFi packet, or a WiFi packet with RSSI below the -62dBm CCA threshold. If this transmit opportunity falls between two LTE subframe boundaries, a reservation signal is transmitted until the start of the next subframe.

Sensing epoch, T. The sensing epoch only needs to be long enough for coarse timescale channel statistics to be found. In our experiments, we use $T = 10s$.

Interference Model. We assume that if a WiFi node receives any interference-free LTE transmission, it will freeze its transmissions in response to the NAV field of the embedded CTS packet. This suspended WiFi node will resume, without any additional backoff penalty, at the earliest packet in the WiFi trace after the end of the LTE transmission. The unlicensed LTE eNodeB uses a WiFi-like energy-sensing CCA before transmission. Should a LTE-WiFi collision occur, we assume that both the LTE subframe and WiFi frame involved in the collision are lost. The LTE and WiFi nodes will backoff for an additional duration specified by the WiFi CSMA protocol and resume at the next trace packet (for WiFi) or transmit opportunity (for LTE) thereafter.

This is a simplification of the actual interference behavior — overlapping transmissions may not result in a total loss of all packets/subframes. However, LTE can usually be decoded at a lower SINR than WiFi packets. Hence, the impact of this simplification is greater on ULTRON than WiFi.

6.2.2 ULTRON *WiFi Throughput*

We study the impact of ULTRON on WiFi under two different LTE offered loads: a light load of 5Mbps ($L = 1$), and a heavier load of 15Mbps ($L = 3$). Fig. 21 shows the normalized throughput of the WiFi network under $L = 1$. The output is normalized w.r.t. to the upper-bound throughput achieved under (A). We evaluate ULTRON alongside the three other protocols under two different WiFi network conditions: one where more than 75% of the

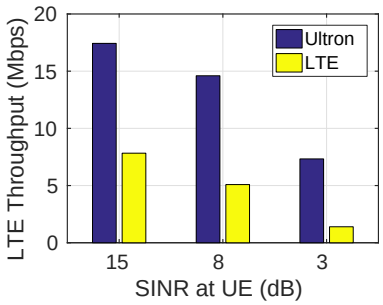

Figure 19: LTE throughput at MCS 12 under different UE SINR

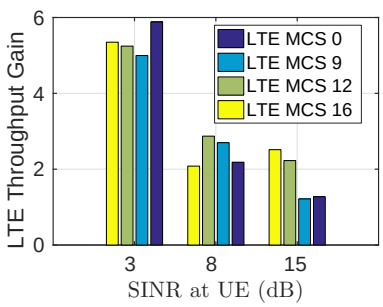

Figure 20: LTE throughput gain under different SIR

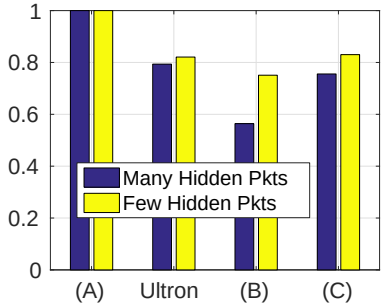

Figure 21: Normalized WiFi throughput ($L = 1$).

WiFi packets fall outside the energy-sensing threshold of ULTRON ("Many Hidden Pkts") and the other where only less than 19% of the WiFi packets are outside the CCA threshold of ULTRON ("Few Hidden Pkts"). In both of these trace sets, the aggregate throughput over the WiFi network is the same.

Most notably, observe that ULTRON *achieves substantial gains in WiFi throughput over (B) in the "Many Hidden Pkts" scenario* ((C) is a variation of ULTRON). Note that the performance under (A) is optimal, as (A) provides sufficient resources for complete and concurrent monitoring of all LTE CCs. As we will see in §6.2.3, this WiFi performance of ULTRON is reached while simultaneously achieving the highest LTE throughput. Note that this near-optimal performance is attained under the constraint of only one WiFi sensing interface. This is because ULTRON selects both the WiFi sensor placement and traffic allocation jointly, thereby effectively minimizing LTE-WiFi interference without penalizing either LTE or WiFi throughput. These gains are also observed under relatively heavy LTE offered loads ($L = 3$), as seen in Fig. 22, although the lack of a WiFi sensor on all the channels in real-time increases sub-optimality.

6.2.3 ULTRON *LTE Throughput*

Unless otherwise stated, we fix the offered load by the LTE eNodeB at $L = 1$. For clarity, we normalize all LTE results w.r.t. the upper-bound throughput obtained by protocol (A), where each LTE CC has a WiFi sensing interface.

ULTRON achieves near-optimal performance with one WiFi sensor. Fig. 23 shows the normalized LTE throughput achieved by ULTRON, and the other three protocols(labelled (A), (B) and (C)). Note that in (A), 3 out of the 5 CCs are used to carry the LTE load of $L = 1$. Observe that ULTRON achieves a throughput of 0.94, that is close to this upper-bound, in a network with many hidden WiFi transmissions (i.e. outside the energy-sensing range of ULTRON). More importantly, observe that ULTRON outperforms all other protocols that rely only on one WiFi sensor, i.e. (B) and (C). In fact, ULTRON maintains a significant fraction of the optimal throughput even under heavier LTE load ($L = 3$). This can be seen in Fig. 24 where ULTRON outperforms (B) and (C).

ULTRON traffic allocation is essential to minimize LTE-WiFi interference. Fig. 23 also shows that ULTRON achieves higher throughput then (B), where LTE load is equally allocated to all CCs. This is because unlike (B), ULTRON takes into account the WiFi traffic load and LTE capacity, and successfully packs the LTE load into as few CCs as possible (in this case, only 3 out of the 5 CCs are used). This minimizes the impact of LTE interference on WiFi. We note that when there are fewer hidden WiFi packets, the eNodeB sees a higher WiFi traffic load and thus, will obtain a smaller share of the channel due to longer unlicensed LTE backoff durations.

ULTRON WiFi sensor placement is critical to LTE performance. When we compare ULTRON to (C), we can see that even under low LTE loads, correct placement of the WiFi sensing interface can improve the LTE throughput by over 24% (0.76 with (B) to 0.94 with ULTRON). The WiFi sensor placement becomes even more critical when the LTE load increases. When $L = 3$, ULTRON requires all CCs to deliver the offered LTE load. Hence, traffic allocation has little impact on throughput. However, observe from Fig. 23 that ULTRON still outperforms (C) due to an optimal choice of the CC with the WiFi sensor.

7. DISCUSSION: APPLYING ULTRON TO LTE-LAA

ULTRON is applicable to both LTE-U and LTE-LAA. In this section, we give an overview of the differences between LTE-U and LTE-LAA, and describe how ULTRON can be applied to LTE-LAA.

7.1 LTE-LAA vs LTE-U

The main differences between LTE-LAA and LTE-U lies in its subframe format and channel access mechanisms.

Listen-Before-Talk. LTE-LAA uses Listen-Before-Talk (LBT), which operates similarly to the CSMA mechanism in WiFi: LTE-LAA performs a energy-sensing CCA before each unlicensed transmission. LBT is still plagued by the asymmetric sensing problem described in this paper because its CCA mechanism only relies on energy-sensing. Hence, a CTS-to-Self is still needed to ensure fair co-existence with WiFi. Furthermore, since LTE-LAA does not recognize and decode WiFi transmissions, a WiFi sensor is still necessary to obtain statistics on the channel utilization and behavior of neighboring WiFi devices.

Subframe Slot Boundary. An LTE-LAA channel access opportunity can arise between subframe boundaries. The LTE-LAA specification addresses this challenge by allowing LTE-LAA transmissions to begin on a *slot* boundary instead (recall that each 1ms LTE subframe consists of two 0.5ms slots). LTE-LAA transmissions will thus begin more quickly after a transmit opportunity arises. A LTE-LAA eNodeB will also transmit a "reservation signal" between the channel access opportunity and the next nearest slot boundary to reduce the probability of WiFi transmissions before the start of the LTE-LAA slot.

Shorter Transmit Bursts. LTE-LAA only allows transmissions of up to 10ms, as compared to 20ms for LTE-U. In some regions such as Japan, this upper limit can be as low as 4ms. The use of LBT also permits fine-grained channel access in response to WiFi traffic. LTE-U only uses CSAT for dynamically adapting its on and off periods over longer periods (up to 20ms) of time. LTE-U does not support LBT and thus, cannot react to fine-grained WiFi channel access characteristics.

Our measurements in §3.2 shows that these shorter transmit bursts

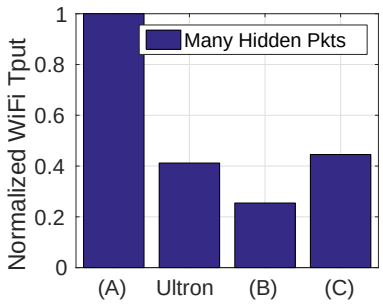

Figure 22: Normalized WiFi throughput ($L = 3$).

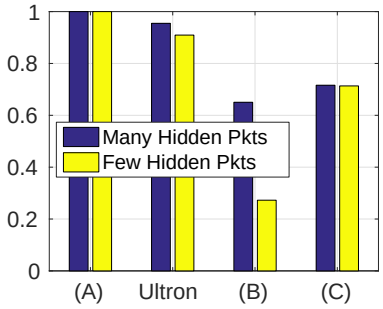

Figure 23: Normalized LTE throughput with low LTE load ($L = 1$)

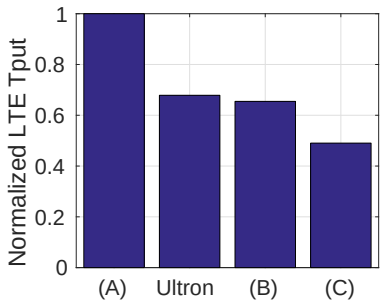

Figure 24: Normalized LTE throughput with high LTE load ($L = 3$)

of LTE-LAA can potentially reduce WiFi throughput and increase its latency even more severely. Hence, CTS embedding and WiFi sensing is even more important in LTE-LAA networks to ensure fair co-existence between LTE and WiFi.

7.2 ULTRON in LTE-LAA

ULTRON, with its CTS-embedding and WiFi sensing components, can be applied to LTE-LAA with very little (if any) modifications. In this section, we will briefly discuss how these two key components can be used together with LTE-LAA.

7.2.1 Integrating ULTRON into LAA

CTS-Embedding. The CTS embedding feature is designed to interoperate seamlessly with the LTE PHY encoding process. Since LTE-LAA (as defined in the LTE Release 13 specification) largely shares the same subframe format with LTE, the CTS embedding feature in ULTRON can be directly applied to LTE-LAA. In fact, maximum effectiveness of CTS embedding can be better realized when it is operated together with the fine-timescale LBT access mechanism of LTE-LAA. Such an integration between LTE-LAA and ULTRON will further the industry goal for a more "WiFi-like" channel access mechanism in unlicensed LTE [16].

WiFi Sensing. The LTE-LAA specification uses a energy-sensing CCA, rather than a WiFi-like preamble-sensing one. While having LBT is definitely useful in LTE-LAA, ULTRON advocates the need for WiFi "carrier" sensing in addition to LBT. ULTRON employs a middle-ground solution, without having to change the LTE PHY or employ multiple WiFi interfaces. This WiFi-sensing component is even more critical in LTE-LAA given that the goal of LTE-LAA is to achieve fine-grained coordination with WiFi.

7.2.2 CTS Transmit Latency

One key challenge in integrating ULTRON into LTE-LAA stems from the CTS transmission latency — since the embedded CTS can only be transmitted at slot boundaries, there is a delay between the time of channel availability (as determined by the energy-sensing CCA) and the CTS transmission.

ULTRON will utilize the reservation-signal feature of LTE-LAA to interoperate with WiFi. Under ULTRON, the energy-sensing asymmetry will exist only for the duration of this short reservation signal. The LTE-LAA transmission thereafter will be protected with the CTS-to-Self reservation. This duration is a significant reduction as compared to LTE-LAA without ULTRON where the sensing asymmetry will exist for the entire LTE-LAA transmission.

There is an alternative solution where we extend the carrier sensing duration such that it will always end at a slot boundary. This CCA duration will always be longer than that mandated by the LTE-LAA specification. Such a CCA extension will ensure that when CTS-to-Self is transmitted, the channel will always be clear according to LBT. This will disadvantage LTE-LAA, at the cost of ensuring the WiFi transmissions are always protected.

8. RELATED WORK

PHY Techniques: PHY-layer techniques are necessary to achieve LTE-WiFi co-existence and have been confirmed by multiple simulation studies [17, 18]. Frame-blanking [19, 20] exploits an existing inter-cell coordination feature to enable co-existence, while power control [21] is another direct approach to reduce LTE interference to WiFi. [22] distributes LTE and WiFi across the degrees of freedom of a MIMO system, where each MIMO stream can transmit either LTE or WiFi. Reinforcement learning [23] can also be used to determine the optimal duty-cycles for LTE. [24] studies the CCA sensitivity control algorithm in 802.11ax, and shows performance gains in dense WiFi networks. **Traffic Allocation:** [25] uses a proportional-fair allocation algorithm for LTE-U (not LTE-LAA) and WiFi co-existence. [26] uses an SDN architecture to enable co-existence between multi-operator WiFi and LTE networks. [27] further enhances LTE-WiFi coexistence concept by dynamically selecting between LTE-WiFi resource sharing and simple WiFi offloading. **WiFi Co-existence:** ULTRON is also related to a large body of research to ensure that WiFi co-exists efficiently with other ISM protocols such as Zigbee and Bluetooth [28, 29, 30, 31]. **Software-Defined PHY:** An alternative approach to co-existence is to employ software-defined WiFi and LTE PHY layers and frameworks [32, 33] together with flexible multi-protocol RF platforms [34, 35, 36, 37]. Such platforms are open to non-standards-compliant PHY modifications that can efficiently achieve LTE/WiFi co-existence. ULTRON differs from these works in that it addresses the fundamental channel access asymmetry between LTE and WiFi in a standards-compliant manner.

9. CONCLUSION

In this paper, we demonstrate that LTE in the unlicensed spectrum still faces significant hurdles. In particular, the asymmetric channel sensing of LTE (energy-sensing) and WiFi (preamble/carrier-sensing) will significantly degrade the performance of WiFi in the presence of LTE transmissions. ULTRON is a 3GPP standards-compliant solution to this co-existence problem — it utilizes *WiFi embedding* to transmit WiFi CTS frames over a unmodified LTE PHY, and scalable *WiFi sensing* to efficiently deploy only a single WiFi sensing interface to jointly optimize both LTE and WiFi performance. Our evaluations show that ULTRON can improve LTE and WiFi throughput by up to 5× and 6× respectively.

10. REFERENCES

[1] "U-LTE: Unlicensed spectrum utilization of LTE," *Huawei Whitepaper*, 2014.

[2] "LTE License Assisted Access," *Ericsson*, 2015.

[3] "Verizon to trial spidercloud lte-u scalable in-building system for enterprises and venues," http://www.spidercloud.com/news/press-release/verizon-trial-spidercloud-lte-u-scalable-building-system-enterprises-and-venues, 2016, [Online; Accessed March 14 2016].

[4] "Google's fcc letter," http://apps.fcc.gov/ecfs/document/view?id=60001078145, 2015, [Online; Accessed March 14 2016].

[5] "Sercom multi-mode small cell," http://www.sercomm.com/contpage.aspx?langid=1&type=prod3&L1id=2&L2id=1&L3id=1&Prodid=65, [Online; Accessed March 14 2016].

[6] "Cisco universal small cell 8718 and 8818 data sheet," http://www.cisco.com/c/en/us/products/collateral/wireless/universal-small-cell-8000-series/datasheet-c78-733945.html, [Online; Accessed March 14 2016].

[7] Xiang Zhang, Wenbo Wang, and Y Yang, "Carrier aggregation for lte-advanced mobile communication systems," *IEEE Communications Magazine*, p. 89, 2010.

[8] Z. Shen, A. Papasakellariou, J. Montojo, D. Gerstenberger, and F. Xu, "Overview of 3gpp lte-advanced carrier aggregation for 4g wireless communications," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 122–130, February 2012.

[9] "Evolved universal terrestrial radio access (e-utra); physical channels and modulation (release 13), 3GPP TS 36.211 V13.1.0," 2016.

[10] Iyappan Ramachandran and Sumit Roy, "On the impact of clear channel assessment on mac performance," in *Global Telecommunications Conference, 2006. GLOBECOM'06. IEEE*. IEEE, 2006, pp. 1–5.

[11] "LTE in unlicensed spectrum: Harmonious coexistence with Wi-Fi," *Qualcomm Research*, June 2014.

[12] P. Robertson, "Illuminating the structure of code and decoder of parallel concatenated recursive systematic (turbo) codes," in *Global Telecommunications Conference, 1994. GLOBECOM '94. Communications: The Global Bridge., IEEE*, Nov 1994, vol. 3, pp. 1298–1303 vol.3.

[13] "Multiplexing and channel coding (release 13), 3GPP TS 36.212 V13.0.0," 2015.

[14] "LTE System Toolbox," http://www.mathworks.com/products/lte-system/.

[15] "WARP 802.11 reference design," http://warpproject.org/trac/wiki/802.11, [Online; Accessed March 14 2016].

[16] R. Kwan, R. Pazhyannur, J. Seymour, V. Chandrasekhar, S. R. Saunders, D. Bevan, H. Osman, J. Bradford, J. Robson, and K. Konstantinou, "Fair co-existence of licensed assisted access lte (laa-lte) and wi-fi in unlicensed spectrum," in *Computer Science and Electronic Engineering Conference (CEEC), 2015 7th*, Sept 2015, pp. 13–18.

[17] N. Rupasinghe and Äř GÃijvenÃğ, "Licensed-assisted access for wifi-lte coexistence in the unlicensed spectrum," in *Globecom Workshops (GC Wkshps), 2014*, Dec 2014, pp. 894–899.

[18] A. M. Cavalcante, E. Almeida, R. D. Vieira, F. Chaves, R. C. D. Paiva, F. Abinader, S. Choudhury, E. Tuomaala, and K. Doppler, "Performance evaluation of lte and wi-fi coexistence in unlicensed bands," in *IEEE 77th Vehicular Technology Conference (VTC Spring)*, June 2013, pp. 1–6.

[19] E. Almeida, A. M. Cavalcante, R. C. D. Paiva, F. S. Chaves, F. M. Abinader, R. D. Vieira, S. Choudhury, E. Tuomaala, and K. Doppler, "Enabling lte/wifi coexistence by lte blank subframe allocation," in *IEEE International Conference on Communications (ICC)*, June 2013, pp. 5083–5088.

[20] T. NihtilÃd', V. Tykhomyrov, O. Alanen, M. A. Uusitalo, A. Sorri, M. Moisio, S. Iraji, R. Ratasuk, and N. Mangalvedhe, "System performance of lte and ieee 802.11 coexisting on a shared frequency band," in *Wireless Communications and Networking Conference (WCNC), 2013 IEEE*, April 2013, pp. 1038–1043.

[21] F. S. Chaves, E. P. L. Almeida, R. D. Vieira, A. M. Cavalcante, F. M. Abinader, S. Choudhury, and K. Doppler, "LTE UL Power Control for the Improvement of LTE/Wi-Fi Coexistence," in *Vehicular Technology Conference (VTC Fall), 2013 IEEE 78th*, Sept 2013, pp. 1–6.

[22] Sangki Yun and Lili Qiu, "Supporting wifi and lte co-existence," in *Computer Communications (INFOCOM), 2015 IEEE Conference on*, April 2015, pp. 810–818.

[23] N. Rupasinghe and Äř GÃijvenÃğ, "Reinforcement learning for licensed-assisted access of lte in the unlicensed spectrum," in *Wireless Communications and Networking Conference (WCNC), 2015 IEEE*, March 2015, pp. 1279–1284.

[24] M Shahwaiz Afaqui, Eduard Garcia-Villegas, Elena Lopez-Aguilera, Graham Smith, and Daniel Camps, "Evaluation of dynamic sensitivity control algorithm for ieee 802.11 ax," in *Wireless Communications and Networking Conference (WCNC), 2015 IEEE*. IEEE, 2015, pp. 1060–1065.

[25] C. Cano and D. J. Leith, "Coexistence of wifi and lte in unlicensed bands: A proportional fair allocation scheme," in *Communication Workshop (ICCW), 2015 IEEE International Conference on*, June 2015, pp. 2288–2293.

[26] Shweta Sagari, Samuel Baysting, Dola Saha, Ivan Seskar, Wade Trappe, and Dipankar Raychaudhuri, "Coordinated dynamic spectrum management of lte-u and wi-fi networks," in *Dynamic Spectrum Access Networks (DySPAN), 2015 IEEE International Symposium on*. IEEE, 2015, pp. 209–220.

[27] Q. Chen, G. Yu, H. Shan, A. Maaref, G. Y. Li, and A. Huang, "An opportunistic unlicensed spectrum utilization method for lte and wifi coexistence system," in *2015 IEEE Global Communications Conference (GLOBECOM)*, Dec 2015, pp. 1–6.

[28] Chieh-Jan Mike Liang, Nissanka Bodhi Priyantha, Jie Liu, and Andreas Terzis, "Surviving wi-fi interference in low power zigbee networks," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, New York, NY, USA, 2010, SenSys '10, pp. 309–322, ACM.

[29] Song Min Kim and Tian He, "Freebee: Cross-technology communication via free side-channel," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, New York, NY, USA, 2015, MobiCom '15, pp. 317–330, ACM.

[30] Yan Yubo, Yang Panlong, Li Xiangyang, Tao Yue, Zhang Lan, and You Lizhao, "Zimo: Building cross-technology mimo to harmonize zigbee smog with wifi flash without intervention," in *Proceedings of the 19th Annual International Conference on Mobile Computing &*

Networking, New York, NY, USA, 2013, MobiCom '13, pp. 465–476, ACM.

[31] Shyamnath Gollakota, Fadel Adib, Dina Katabi, and Srinivasan Seshan, "Clearing the rf smog: Making 802.11n robust to cross-technology interference," in *Proceedings of the ACM SIGCOMM 2011 Conference*, New York, NY, USA, 2011, SIGCOMM '11, pp. 170–181, ACM.

[32] Navid Nikaein, Mahesh K. Marina, Saravana Manickam, Alex Dawson, Raymond Knopp, and Christian Bonnet, "Openairinterface: A flexible platform for 5g research," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 5, pp. 33–38, Oct. 2014.

[33] Manu Bansal, Aaron Schulman, and Sachin Katti, "Atomix: A framework for deploying signal processing applications on wireless infrastructure," in *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*, Berkeley, CA, USA, 2015, NSDI'15, pp. 173–188, USENIX Association.

[34] Steven S. Hong, Jeffrey Mehlman, and Sachin Katti, "Picasso: Flexible rf and spectrum slicing," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, New York, NY, USA, 2012, SIGCOMM '12, pp. 37–48, ACM.

[35] Bo Chen, Vivek Yenamandra, and Kannan Srinivasan, "Flexradio: Fully flexible radios and networks," in *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*, Berkeley, CA, USA, 2015, NSDI'15, pp. 205–218, USENIX Association.

[36] Kun Tan, Haichen Shen, Jiansong Zhang, and Yongguang Zhang, "Enable flexible spectrum access with spectrum virtualization," in *Dynamic Spectrum Access Networks (DYSPAN), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 47–58.

[37] Sangki Yun, Daehyeok Kim, and Lili Qiu, "Fine-grained spectrum adaptation in wifi networks," in *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, New York, NY, USA, 2013, MobiCom '13, pp. 327–338, ACM.