2018 IEEE 26th International Conference on Network Protocols

ELI: Empowering LTE with Interference Awareness in Unlicensed Spectrum

Ramanujan K Sheshadri[†], Karthikeyan Sundaresan[†], Eugene Chai[†], Sampath Rangarajan[†], Dimitrios Koutsonikolas[*]

[†] NEC Labs America, Inc.
(ram,karthiks,eugene,sampath)@nec-labs.com

[*] University at Buffalo, SUNY
dimitrio@buffalo.edu

***Abstract*—The advent of LTE into the unlicensed spectrum has necessitated the understanding of its operational efficiency when sharing spectrum with different radio access technologies. Our study reveals that LTE, owing to its inherent transmission characteristics, suffers significant performance degradation in the presence of interference caused by hidden terminals. This motivates the need for interference-awareness in LTE's channel access in unlicensed spectrum.**

To address this problem, we propose ELI. ELI's three-pronged solution equips the LTE base station with novel techniques to: (a) accurately detect and measure interference caused by hidden terminals, (b) collect interference statistics from clients across different channels with affordable overhead, and (c) leverage interference-awareness to improve its channel access performance. Our evaluations show that ELI can achieve 1.5-2x throughput gains over baseline schemes. Finally, ELI is LTE-LAA/MulteFire-standard compliant and can be deployed over the existing LTE-LAA implementation without any modifications.

I. Introduction

To address the increasing demand for additional LTE spectrum, there is growing interest and support to operate LTE in the unlicensed band in conjunction with its operations in the licensed spectrum (3GPP's LTE-LAA/eLAA specifications) [1]. Accordingly, a LTE node operating in the unlicensed spectrum is required to implement an asynchronous channel access mechanism (Listen Before Talk – LBT), that includes both energy sensing (CCA) and back-off to ensure fair co-existence with the incumbent WiFi. Traditional network providers like AT&T and Verizon have already started rolling out LAA technology nationwide [2], [3]. In addition, there is also a significant effort by green-field providers to develop a stand-alone LTE specification (MulteFire [4]) that can operate LTE entirely in the unlicensed band without any assistance from the licensed spectrum for use in private LTE (e.g. IoT) networks. In view of these developments, this work takes an important step towards adapting LTE operations to suit the rigors of a contentious medium like the unlicensed spectrum.

Need for interference awareness. As is the case with any access technology, to operate efficiently in the unlicensed spectrum, LTE needs to intelligently pick a channel and then operate efficiently within the selected channel. While this problem is similar to that in WiFi, LTE's inherent transmission characteristics, makes designing a LTE-specific solution highly challenging.

WiFi employs an asynchronous and distributed access mechanism, where every node is individually responsible to scan the channel in its vicinity and make independent access decisions based on perceived interference. In contrast, while both eNB and clients perform LBT, LTE mandates the eNB to schedule channel access to each of its clients (UEs) both on the downlink (DL) and uplink (UL) (see Fig. 1). Such an eNB-controlled approach delivers optimal performance gains for LTE in licensed spectrum that is free of external interference. However, in unlicensed spectrum, it leads to lack of visibility into channel (interference) state at the clients, resulting in collisions on the DL, and spectrum under-utilization on the UL (Fig.1). This coupled with LTE's longer transmission durations of 2-10 ms [5], [1] (compared to WiFi's transmission duration),along with its lack of interference avoidance mechanisms (e.g. WiFi's RTS/CTS), can degrade its performance significantly when the interfering nodes are hidden from the eNB. As we show later in Section II, the performance degradation for LTE from such hidden terminals can be *twice* as much as that experienced by WiFi.

While introducing interference avoidance techniques (similar to WiFi's RTS/CTS) and reducing transmission durations can alleviate the impact of such interference, this would require changes to the LTE specification – a major road-block to practical realization. In contrast, we focus on the root-cause for this magnified impact of interference, namely LTE's eNB-driven channel access procedure, which lacks information on the interference perceived by its clients. Consequently, any solution to tackle this problem should start by equipping the eNB with the ability to detect and estimate interference at its clients.

Challenges in bringing interference awareness into LTE. (i) *Accurate estimation of interference:* A natural approach to detect hidden terminal interference at the clients would be to measure the failure of eNB's DownLink (DL) transmissions by

978-1-5386-6043-0/18/$31.00 ©2018 IEEE
DOI 10.1109/ICNP.2018.00044

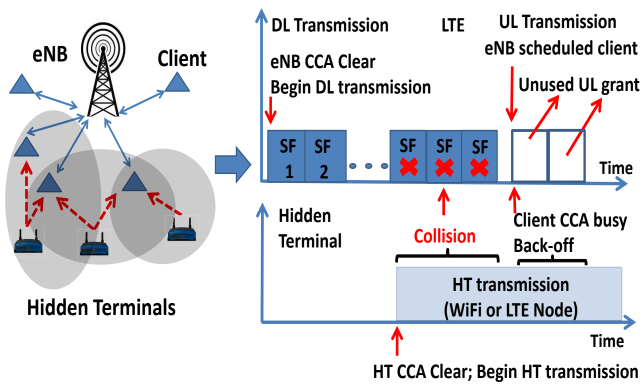

Fig. 1: Illustration of eNB Sub-Frames (SF) colliding on DL; On UL client backs-off due to HT activity causing spectrum under-utilization.

using the HARQ feedback (ACK/NACK) sent by the clients. However, transmission failures can result from collisions (due to hidden terminals), or due to the vagaries of the channel itself (*e.g.,* channel fading). The inability of HARQ to clearly distinguish the reason for failure makes them less useful for estimating interference. Hence, we need an alternate mechanism to accurately estimate interference that is specific to individual clients.

(ii) *Scalable estimation of interference:* Unlike the licensed spectrum, interference in the unlicensed spectrum could vary significantly across different channels (details in Section II). Consequently, for an LTE eNB to make effective channel access decisions, it needs to estimate interference at all its clients on each of the candidate channels that it considers for regular operation. This constitutes an overhead that scales both with the number of clients and channels and can quickly become prohibitive. For e.g., a dense LTE network, potentially hosting hundreds of clients, in 5 GHz and/or CBRS (3.5 GHz) bands that contain multiple channels, could incur substantial overhead that negates the performance gains resulting from such interference awareness. Hence, while it is critical to estimate such interference information, it needs to be done in a highly scalable manner.

(iii) *Leveraging interference-awareness:* Finally, to orchestrate LTE for efficient access in unlicensed spectrum, it is equally important to understand how we can leverage the estimated interference information in LTE's channel access mechanisms in a standards-compliant manner.

Our Proposal – ELI: Towards addressing these challenges and making LTE robust in the presence of interference in unlicensed spectrum, we propose our solution ELI. ELI employs a three-pronged approach, which includes (a) Equipping eNB to accurately detect and measure interference caused by hidden terminals; (b) Collecting interference statistics across channels for all the clients in a scalable manner; and (c) Using collected interference statistics through novel standards-compliant access techniques to deliver improved performance.

(a) While the eNB cannot differentiate between a fading and a (interference) collision loss on the DL through HARQ, ELI leverages their varied manifestation on the UL, where clients employ LBT for access. When hidden terminals are detected by clients during their LBT, they refrain (back-off) from transmitting (both their reference-signals/pilots and data) on their scheduled UL resources. This manifestation is different from a fading scenario, where both pilots and data are transmitted – the data gets corrupted, while pilots can still be recovered due to their use of the most robust rate (MCS). Thus, by scheduling a client on UL resources and observing its UL pilots, ELI is able to accurately estimate the interference statistics (probability of hidden terminal interference) at the client (Section III-B).

(b) ELI employs a two-step solution to minimize the overhead associated with estimating interference. First, it reduces the number of channels that need to be measured by half, while still obtaining the desired interference information on all channels. It does this by intelligently leveraging LTE's OFDMA capability to concurrently transmit on fine-grained spectral resources (LTE resource blocks) within a channel, to measure interference on two channels simultaneously, while incurring the measurement overhead of a single channel. Second, it samples only a fraction of clients for whom interference needs to be estimated on each channel, while still obtaining the desired information for all the clients. Here, it leverages the diversity of interference (hidden terminals) across channels to spatially cluster clients with the same interference statistics (Section III-C).

(c) Finally, ELI leverages the collected interference statistics of all the clients on all channels, to (i) dynamically select an unlicensed channel (at coarse time scales, secs-min) that has minimal impact from hidden terminals on its clients; and (ii) within the chosen channel, transforms the proportionally-fair (PF) scheduler at eNB into an interference-aware PF scheduler that accounts for the true channel state of the clients at fine time scales (millisecs) in its channel access (scheduling) allocations (Section III-D).

Contributions: Through the design of ELI, we make the following contributions in this work.
(a) Motivate the need for interference-awareness when operating LTE in unlicensed spectrum.
(b) Propose accurate, scalable and standards-compliant approach to estimate interference at the clients due to hidden terminals.
(c) Incorporate interference awareness into LTE to improve its access performance at both macro (channel selection) and micro (scheduling) time scales.
(d) Provide detailed evaluations to show that ELI can deliver throughput gains of $\approx$ 1.5x to 2x when compared to existing schemes.

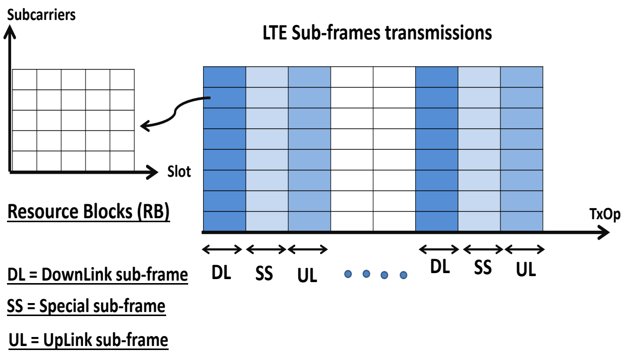

Fig. 2: Illustration showing sub-frame structure in a TDD LTE.

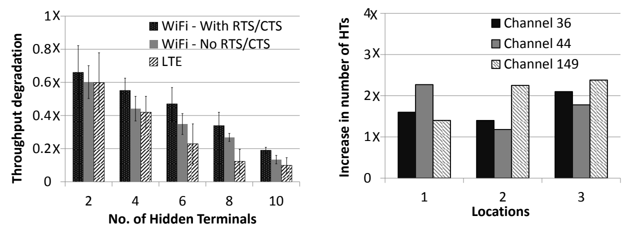

(a) Throughput degradation (b) Increase in HTs for LTE

Fig. 3: LTE in the presence of Hidden Terminals (HT)

II. Motivation

A. Background

LTE overview: LTE is essentially designed to operate in a licensed band, with a synchronous and a scheduled access mechanism, where the eNB schedules clients for both UL and DL traffic. LTE transmissions happen in the granularity of sub-frames, each being 1ms, and consists of two slots (0.5 ms each). Each slot spans both time-frequency resources by comprising of multiple OFDM symbols and sub-carriers as shown in Fig. 2. In a time-divisioned (TDD) system, DL sub-frames (eNB transmissions) are followed by (optional) special sub-frames followed by UL sub-frames (see Fig. 2), where clients transmit in the UL resources allocated to them a priori (based on clients' schedule requests) by the eNB. A special sub-frame is a transition sub-frame consisting of part DL and part UL transmissions along with a guard period. The LTE schedulers allocate blocks of subcarriers to each client, known as resource blocks, which are then used to exchange data with the clients on both DL and UL. The total number of resource blocks available for the eNB depends on the channel width (e.g., 100 RBs in 20 MHz channel bandwidth). Each resource block is orthogonal to each other which enables the eNB to allocate each resource block to a different client. This allows for simultaneous data and control information to be exchanged between the eNB and its clients on the allocated resource blocks using OFDMA principles.

Listen Before Talk (LBT) in LTE: The 3GPP release-14 [1] requires LTE nodes operating in the unlicensed spectrum to implement the LTE-LBT procedure which includes both CCA (Clear Channel Assessment) energy detection, and back-off. If the channel is busy (detected energy is above the prescribed CCA threshold), the eNB implements exponential back-off and waits for the back-off counter to expire before attempting to transmit again on DL. On the other hand, when it is time for the client (UE) to transmit on its scheduled UL resources, due to its LBT procedure it backs-off and simply forgoes its transmission opportunity if the channel is busy and waits for the eNB to schedule its next transmission. TDD, being the dominant flavor covering both license-assisted and stand-alone LTE modes in unlicensed spectrum, forms our focus in this work.

B. Need for interference awareness

Increased impact of interference on LTE: We begin by comparing the impact of a given set of hidden-terminal-interference on an LTE network (single eNB and clients) and a WiFi network (single AP and its clients). To do this, we conduct a small experiment using the NS3 simulator (details of NS3 implementation in sec IV), in which we consider multiple topologies of a WiFi AP/LTE eNB with 25 randomly distributed clients. We then make the AP/eNB transmit DL (SISO) UDP traffic to all its clients in the presence of multiple hidden terminals (WiFi nodes outside transmitting range of the AP/eNB). The impact of interference on the DL throughput in each case is shown in Fig. 3a. In the case of WiFi, we consider two versions, one without RTS/CTS and the other with RTS/CTS. As expected, the decrease in LTE's DL throughput is substantial ($\approx 80\%$ when hidden terminals > 6). But more interestingly, the impact of interference on LTE is much worse even compared to WiFi without RTS/CTS (with 8,10 hidden terminals). This is because, unlike WiFi, whose transmission durations last for a few hundred microseconds [5], LTE transmissions span for a substantially prolonged duration (2-10 ms) [1], making it more susceptible to collisions caused by hidden terminals. The absence of interference avoidance mechanisms in LTE further amplifies its performance degradation compared to WiFi with RTS/CTS.

Increase in the number of interferers for LTE: The LTE-LBT procedure requires the CCA energy threshold to be around -65 dBm to -72 dBm [6]. On the other hand, the WiFi nodes which implement preamble detection along with energy detection have a much lower CCA sensitivity threshold of $\approx$ -85 dBm. This inconsistency in channel sensing threshold (difference of $\approx$ 20 dBm) between LTE and WiFi results in creation of additional hidden terminals to the eNB. To quantify this effect, we collect the WiFi data traces from three locations on the same floor of an academic building, on three different WiFi channels simultaneously by switching the WiFi AP to the monitor mode. We calculate the number of stations that would form hidden terminals to the WiFi AP

(Received RSSI > -85 dBm), and compare it with the number of stations that would form hidden terminals to an LTE eNB, if the AP was replaced by the eNB (RSSI between -65 dBm and -85 dBm). The collected traces indicate a significant increase in the number of hidden terminals (Fig. 3b shows the increase) for a LTE eNB compared to a WiFi AP.
Remark: Lowering CCA energy threshold of LTE to match that of the WiFi is not a viable option as it will lead to severe under-utilization of the channel [6].

Diversity in impact of interference across channels: In addition to the increase in the number of hidden terminals, the graph in Fig. 3b also shows a varied number of hidden terminals across the three channels. This clearly indicates that there is sufficient diversity of interference across channels and also the need for LTE to collect interference statistics from different channels for a more robust channel access.

Lack of existing solutions: The problem of interference due to hidden terminals is not new and has been extensively studied in the past in the context of different wireless access technologies (specially WiFi). However, the solutions prescribed are specific to each individual wireless access technology and cannot be applied to the LTE. For instance, WiFi's RTS/CTS is more suited for a distributed channel-access mechanism. The LTE, however, uses a centralized channel-access mechanism where the eNB alone decides when client's can access the channel. To implement anything like RTS/CTS exchange, the eNB would have to schedule channel resources for the clients to transmit back the CTS messages, making even CTS transmissions vulnerable to interference (also might lead to resource wastage). What we need is a solution specific to LTE.

C. Challenges

The study above highlights the pronounced impact of interference caused by hidden terminals on LTE's performance (compared to WiFi). This clearly necessitates the need to incorporate interference-awareness into LTE's channel access mechanisms. However, as discussed earlier, realizing this in practice requires us to address the following important questions:
(a) How can an eNB accurately detect the interference at its clients, caused by the hidden terminals?
(b) How can an eNB estimate such interference at all clients on all candidate channels in a scalable manner?
(c) Finally, how to incorporate the estimated interference statistics into the LTE operations to benefit its channel access?

III. ELI– Design Components

A. Overview

Towards addressing the above mentioned challenges, we present our solution – ELI. ELI targets a TDD system, where both DL and UL operate sequentially on the same unlicensed channel chosen by the eNB. TDD systems are more generic in scope as they encompass not only LTE-LAA systems but also stand-alone LTE systems, e.g. MulteFire, which are primarily TDD in operation. ELI operates its channel access in epochs as shown in Fig. 4. With LTE being more vulnerable to interference than WiFi, ELI chooses a shorter epoch duration of 1 minute (compared to the typical 5-10 min epochs in WiFi [7]) for channel acccess decisions, to track and adapt to interference at finer time scales. Each epoch consists of two phases – (i) *Interference estimation*: The epoch starts with the estimation of hidden terminal interference to each client on every candidate unlicensed channel; and (ii) *Interference-aware channel access*: Using the estimated information, ELI selects an appropriate channel with the least impact of interference on its clients' rates; then schedules its clients in that channel for actual operation for the rest of the epoch, while accounting for their true (interference-incorporated) channel state. Note that while data is also transferred (clients are scheduled) during the estimation phase, the channel access is not optimized for efficiency but more for measurement and hence constitutes an overhead. Hence, in addition to the actual interference-aware channel access, reducing the time incurred in the estimation phase is a key pillar-stone of ELI's design. In the rest of this section, we describe each of ELI's design components that contribute to both these phases of the epoch.

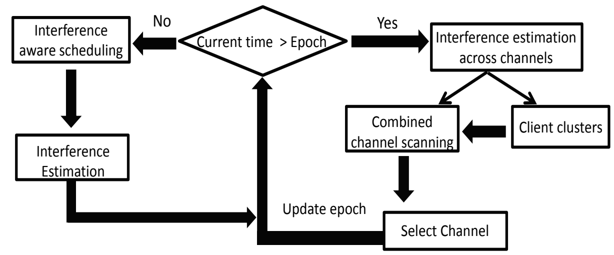

Fig. 4: Design control flow of ELI.

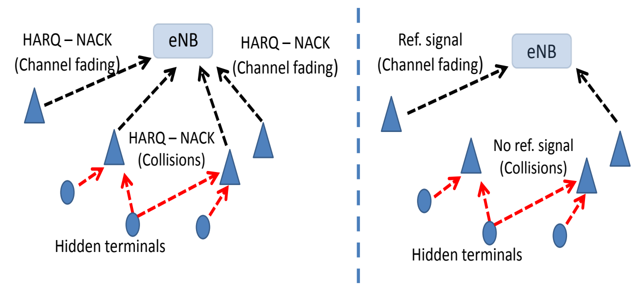

Fig. 5: Left: Inability to distinguish fading from collisions using only HARQ feedbacks; Right: Ability to distinguish between fading and collisions using UL reference signals

B. Accurate Estimation of Hidden Terminal Interference

While eNB is able to directly estimate the interference perceived by it, it does not have information on interference faced by its clients, but hidden from itself. A natural approach to measure such hidden terminal interference is for eNB to use the client's feedback packets like HARQ to verify if its DL transmissions incurred any collisions due to hidden terminal traffic. While the HARQ feedbacks clearly indicate the failure of DL transmissions (through NACKs), they do not reveal the cause, which could be due to collisions or simply due to decreased path loss from channel fading (as shown in Fig. 5). Thus, it is unable to isolate the impact of hidden terminal interference from channel fading. While channel fading plagues all clients and is fleeting in nature, hidden terminal interference is persistent, and varies significantly from one client to another as well as across channels. Hence, ELI needs to isolate and estimate this interference to counteract it.

Leverage interference manifestation on UL during LBT: ELI leverages the combination of LTE's scheduled access along with clients' LBT mechanism on the UL to isolate and detect hidden terminal interference. Recall from Section II that an LTE client performs LBT before its UL channel access, wherein it senses the channel for active transmissions – if it detects energy above a specified (CCA) threshold, it backs-off from its transmission (LBT). Further, unlike WiFi, where every client scans the channel independently before accessing the channel, the LTE eNB schedules the clients for their UL channel access. Hence, the eNB is aware of when and where (in the OFDMA frame) a scheduled client will transmit on in the UL. Hence, when the eNB is able to transmit on the DL (using its LBT), but does not receive any signal (including the reference signal a.k.a pilot tones) from the client on its scheduled UL resources, the eNB can be certain that the client has backed-off from its transmission due to hidden terminal interference (client's LBT). Since the reference signals are sent at the lowest modulation, they are are more resistant to the vagaries of the channel (e.g., path loss due to fading) compared to data. This helps the eNB distinguish between a fading scenario (data corrupted but pilots received) from a hidden terminal (no data and pilots received) one reliably on the UL. Over time, the eNB uses this information to estimate the probability of a hidden terminal blocking the client from transmitting (Q_i), and hence the access probability of a client subject to hidden terminal interference ($P_i = 1 - Q_i$).

Remark: The eNB does not necessarily have to wait for clients to have data to send on the UL to measure the client's access probability. Clients are made to periodically (periodicity between 1 ms to 160 ms) send control packets on the UL, which suffices to estimate their impact from hidden terminals (hereafter captured through client's access probability).

C. Scalable Estimation of Interference

With the above mechanism on the UL set for isolating and detecting hidden terminal interference for a client, ELI's next objective is to measure every client's channel access probability on each of the candidate unlicensed channels for operations. This can potentially constitute a prohibitive overhead that scales with both the number clients and the channels. *e.g.,* A dense network in the 5 GHz ISM band has 9 non-overlapping channels, and can support hundreds of clients. This brings us to ELI's mechanisms that help reduce estimation overhead along both the dimensions of channels and clients.

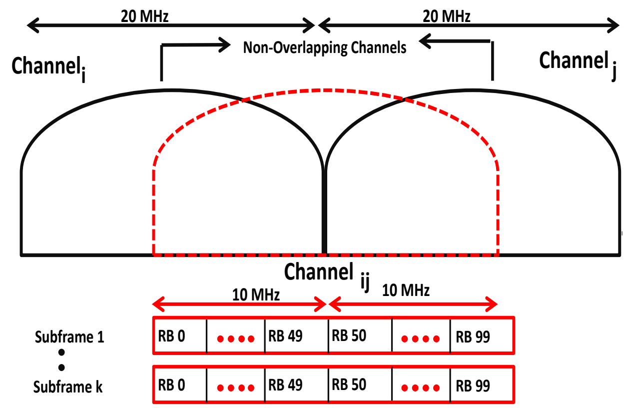

Fig. 6: An overlapping channel (shown in red) and its resource block distribution.

*1) **Reducing Channel Overhead**:* ELI leverages the notion of overlapping channels along with flexible resource allocation (OFDMA) on the DL to reduce the number of channels that need to scanned by *half*. A given spectrum is typically divided into multiple non-overlapping channels (*e.g.,* 9 channels of 20 MHz each in 5 GHz) that incur low cross-channel interference and are hence used for operation. In addition to these non-overlapping channels, there are other overlapping channels(*e.g.,* C_{ij}, Fig. 6) whose central frequencies are exactly in the middle of two adjacent non-overlapping channels (C_i,C_j, in Fig 6). While ELI employs one of the non-overlapping channels during regular operations (channel access phase), it aims to use these overlapping channels intelligently during the measurement phase. This allows it to estimate its clients' access probabilities on two non-overlapping channels (*e.g.,* C_i and C_j) *simultaneously* using a single overlapping channel (C_{ij}), thereby reducing the net overhead by half. Note that when the eNB uses an overlapping channel, its resource blocks (e.g. 100 RBs in 20 MHz channel) are equally distributed between the two non-overlapping channels (C_i, C_j) that it covers partially (*e.g.,* 0-49 RBs span half of C_i, while 50-99 span half of C_j).

While ELI's interference detection is accurate on the UL, the client senses the 20 MHz channel (10 MHz each of the two adjacent non-overlapping channels) as a whole. This makes it difficult for the eNB to detect interference on each of the constituent channels (10 MHz blocks) just based on UL

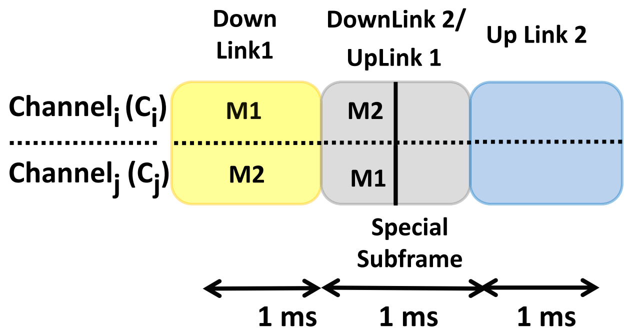

Fig. 7: Sub-frame resource allocation during interference scanning.

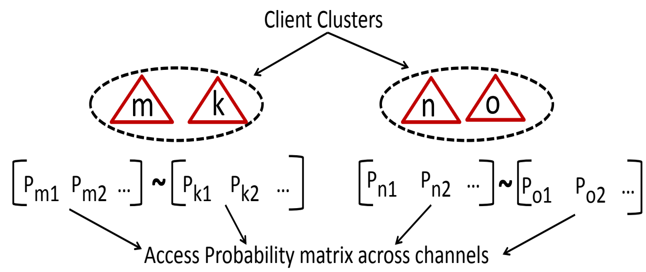

Fig. 8: Illustration showing client clustering based on interference across different channels

access. Hence, ELI intelligently employs *joint* DL and UL access, where it leverages LTE's OFDMA capability on the DL to flexibly schedule a client's DL resource allocation and hence sample interference on both the constituent channels individually, while it uses UL access to accurately detect interference on the immediately preceding DL transmissions on each of the constituent channels. ELI's algorithm proceeds as follows.

Step 1: Schedule Joint DL+UL Measurements – ELI conducts its measurements in a TxOP configuration of 3 ms each, consisting of a DL, a special and a UL sub-frame in that order, as shown in Fig. 7. ELI creates an *alternating* schedule within the DL parts of the first two (DL and special) sub-frames as follows. The eNB transmits to two sets ($\mathcal{M}_1, \mathcal{M}_2$) of $\frac{M}{2}$ clients each on RBs spanning parts of channel C_i and C_j respectively in the first (DL) sub-frame.This is followed by switching the transmissions to $\mathcal{M}_1$ and $\mathcal{M}_2$ on the alternate set of RBs (*i.e.,* C_j and C_i respectively) in the DL part of the second (special) sub-frame, as shown in Fig. 7 . In other words, eNB transmits two transport blocks $T_m(1)$ and $T_m(2)$ to each client $m \in \mathcal{M}_1$ on DL, where $T_m(1)$ is on RBs in C_i and $T_m(2)$ is on RBs in C_j (vice versa for $m \in \mathcal{M}_2$). The UL part of the second (special) and third (UL) sub-frames is primarily to activate UL access (detect interference through back-off) and obtain feedback from the clients in $\mathcal{M}_1$ and $\mathcal{M}_2$ for their DL transmissions, but can also be used to schedule data for them based on their outstanding requests. This alternating DL schedule coupled with immediate UL access allows ELI to sample (on the DL) and detect (on the UL) the interference on both C_i and C_j for each client in tandem.

Step 2: Estimate Interference on Two Channels Simultaneously – To accurately estimate interference on each of the constituent channels C_i and C_j, ELI needs to distinguish between scenarios where there is a hidden terminal on only *one* or *both* of the constituent channels. Using ELI's prior interference estimation procedure on the UL, we know that a missing transmission on the UL, signals hidden terminal interference. When this is immediately preceded by DL transmissions, the result is a *"delayed"* feedback (ACK/NACK) for DL transmissions. This coupled with the alternating DL schedule that samples both parts of the channel (C_{ij}) belonging to C_i and C_j, helps ELI estimate interference accurately on C_i and C_j simultaneously for each client.

(i) when there is a delayed ACK for $T_m(1)$ from the client, the ACK signifies a successful transmission on C_i, while the delay signifies a hidden terminal impact on UL transmission, indicating that the interference is in the other non-overlapping channel C_j.

(ii) when there is a delayed NACK for $T_m(1)$ from the client, the NACK signifies a failed transmission on C_i, while the delay signifies a hidden terminal impact on the UL transmission, indicating that there is interference in the same non-overlapping channel C_i. However, this does not preclude that there is also hidden terminal impact on C_j. To resolve this, ELI does a similar estimation on $T_m(2)$, which then helps accurately identify if there is also simultaneous interference on C_i and C_j.

Step 3: Repeat across Overlapping Channels – The eNB switches across overlapping channels and for each overlapping channel, it schedules multiple such 3 sub-frame TxOPs (Steps 1 and 2), where clients are scheduled in a round-robin fashion. By gathering sufficient samples, it estimates their interference impact on each of the two constituent non-overlapping channels simultaneously, thereby helping reduce the overhead by half.

*2) **Reducing Client Overhead:*** After reducing the channel overhead, we now focus on reducing the channel-dwell time, namely the time required to estimate interference information for each client on a given channel. In dense networks hosting a large number of clients (tens to hundred), the associated overhead can be excessive. However, such dense scenarios also offer an opportunity for ELI, wherein one can expect

Algorithm 1 Clustering clients to reduce channel-dwell time

```
Input: Access probability of all clients in all channels
Output: Client clusters
while N > 1 do                                  ▷ N = No. of clusters
    for i ∈ {1, ..., N} do
        for j ∈ {1, ..., N} do
            if i ≠ j then
                C_i = centroid(i)
                C_j = centroid(j)
                euclid = calc_euclidean(C_i, C_j))
            end if
        end for
    end for
    clusters = do_cluster(min(euclid))
    N = count(cluster)
end while
```

multiple clients that are spatially co-located to share the same set of hidden terminals and hence a similar interference impact statistically. Hence, ELI aims to cluster such spatially co-located clients and measures the interference on only one representative client in each cluster, thereby reducing the channel-dwell time.

Clustering Clients using Interference Diversity: However, the only interference information that ELI estimates is the client's access probability (P_{mi}) on a channel i. With this limited information, it is possible for clients (say m and n, Fig. 8) in spatially distant regions around the eNB to exhibit similar interference impact ($P_{mi} = P_{ni}$), making it a challenge to accurately identify co-located clients. Here, ELI leverages the inherent interference diversity that exists across different channels to solve the problem. Note that the set of hidden terminals that affect a client would vary significantly from one channel to another (§ II). While the interference estimates for two distant clients can be similar on a given channel, when their interference estimates across multiple channels are *jointly* considered ($\mathbb{P}_m = \{P_{m1}, P_{m2}, \cdots\}$), only co-located clients (*e.g.,* clients m and k, Fig. 8) would reveal a higher correlation in their interference estimate vectors ($\mathbb{P}_m$). Distant clients would experience diverse interference across channels and hence automatically reveal lower correlation. Thus, by clustering the clients using their interference estimates across the frequency domain (channels), ELI is able to co-locate clients with high accuracy, where using data from more channels contributes to increasing accuracy.

Clustering technique: To cluster the clients, ELI uses one of the standard clustering techniques called hierarchical clustering. The hierarchical clustering measures the similarity between two vectors (access probabilities in each channel, of two clients – $\mathbb{P}_m$, Fig 8) using the euclidean distance between them. The clustering algorithm uses an agglomerative (bottom-up) approach where each client starts as a single element in its own cluster and in each subsequent iteration two clusters that have the smallest euclidean distance between their respective centroids (calculated as the average access probability of the cluster) are merged together until all the clusters merge together to form one cluster. The advantage of employing hierarchical clustering as opposed to other popular clustering techniques like K-means is that the hierarchical clustering outputs a hierarchy, a structure that is more informative than the unstructured set of clusters returned by other flat clustering techniques. This allows us to further fine-tune the clustering threshold (similarity threshold), and there-by the overhead incurred. The pseudo-code of the clustering process is shown in algorithm 1.

*3) **Benefits of ELI's Overhead Reduction**:* To highlight the potential benefits of ELI's channel scanning mechanisms, consider an instance of an eNB supporting 50 clients with 9 candidate channels, in 5 GHz band. If we employ 50 sub-frames (1 subframe = 1ms, each sub-frame supporting max. 10 clients) per client to estimate its access probability, then the total scan time for 9 channels without ELI's optimizations would incur a total estimation overhead of $\approx$ 2.3 secs (4% of a 1min epoch). With ELI's channel reduction, this overhead reduces to $\approx$ 1.2 secs (2% of epoch); with client clustering, it further reduces to only 300 ms (only 0.5% of epoch) in our evaluations (§ III-A).

D. Leveraging interference awareness

We now explain ELI's schemes to suitably incorporate the interference statistics to enhance LTE's performance at both macro time scale through dynamic channel selection as well as micro time scales through interference-aware access scheduling.

1) Dynamic channel selection: At macro time scales (every minute – § III-A), ELI switches eNB's operating channel to the channel that offers minimal interference (as measured during the estimation phase) as well as a higher average transmission rate to its clients. The concept of dynamic channel switching is practiced in WiFi APs where the AP switches its operating channel periodically to move into a better channel [7]. The advantage that the WiFi AP has over the LTE eNB while selecting the channel is that the WiFi AP does not necessarily have to take into account the interference at the clients' end. Even a sub-optimal channel selection by a WiFi AP may not have a big impact on its performance, because of the WiFi's distributed channel access mechanism (clients access channel independently) and its inherent interference-mitigation techniques (*e.g.,* RTS/CTS mechanism). However, when it comes to LTE, the eNB must take into account the interference in every channel in addition to the client rates, before picking the best channel. ELI periodically uses the access probability statistics gathered from all the available channels (C) to select the best channel ($C_b \in C$), as shown below.

$$C_b = \arg\max_{i \in \mathcal{C}} \{ \sum_{m=1}^{M} P_{mi} \cdot \log(R_{mi}) \}$$

where R_{mi} represents the average channel rate of the client m in channel i, calculated during the measurement phase, and incorporates both the DL and UL average rates as $R_{mi} = \beta R^D_{mi} + (1-\beta)R^U_{mi}$, where $\beta \in (0,1)$ captures the relative importance/load between DL and UL traffic.
Remark: The dynamic channel selection aims to improve performance on a macro time scale (sec III-A). Consequently, it is aimed at static clients and for clients whose mobility changes over a coarse time scale. For more dynamic mobile clients, ELI performs intelligent access scheduling on a micro time scales, as discussed below.

2) Interference-aware PF (IPF) scheduler: To improve LTE's performance at micro time scale (order of milliseconds), ELI transforms the LTE's native interference-agnostic PF scheduler into an interference-aware PF scheduler (IPF). While IPF can be applied to both DL and UL, we focus on the DL here, starting with background on native LTE schedulers.
Native PF Scheduler: Among the many LTE schedulers that exist, the most extensively used is the PF scheduler because of the right balance it achieves between both performance (throughput) and fairness across all the clients. For the PF scheduler the optimal scheduling is obtained by maximizing the logarithmic utility function – $\sum_m log(R_m)$, where R_m is the average channel rate of the client m, calculated over time. Consequently, the PF scheduler prioritizes clients whose instantaneous channel rate (measured at the eNB based on the modulation and coding scheme used) is the highest when normalized to its own average channel rate. Thus, the decision ($S(t)$) of allocating B resource blocks to N clients at each sub-frame (t), is formalized as shown below.

$$S(t) = \arg\max_{x\in S}\{\sum_{b=1}^{B}\sum_{m=1}^{N}\frac{x_{m,b}r_{m,b}(t)}{R_m(t)}\},\ \text{s.t.}\ \sum_{i=1}^{N}x_{m,b}\leq 1, \forall b$$

Where, x is a binary variable capturing the schedule. The $r_{m,b}$ is the instantaneous rate of client m on resource block b in a SISO system. Subsequent to each assignment, the average channel rate of a client m, is then updated as:

$$R_i(t) = \begin{cases} \frac{1}{\alpha}\sum_{b=1}^{B} r_{m,b}(t) + (1-\frac{1}{\alpha})R_m(t-1), & \textit{if } x_{m,b} == 1 \\ (1-\frac{1}{\alpha})\cdot R_m(t-1), & \textit{otherwise} \end{cases}$$

where, α is an exponential weighted constant.

IPF scheduler: One of the pitfalls of the native PF scheduler is that it allocates resources to clients purely based on their channel rates, which works for licensed spectrum. However, the instantaneous channel rates are not reflective of the true channel state of the client in unlicensed spectrum, which is also impacted by hidden terminal interference. Further, such interference also varies from one client to another based on its spatial location. ELI leverages such interference diversity that exists within a cell to convert the native (PF) scheduler into a weighted interference-aware PF scheduler. The latter not only uses channel diversity, but also incorporates interference diversity and its impact to capture the true channel state of the clients. If P_m represents the moving average probability of the client m utilizing its resources (access probability, subject to the HT interference), then the Interference-aware PF scheduler (IPF scheduler) is given by,

$$S(t) = \arg\max_{x\in S}\{\sum_{b=1}^{B}\sum_{m=1}^{N}\frac{P_m\cdot x_{m,b}r_{m,b}(t)}{R_m(t)}\},\ \text{s.t.}\ \sum_{m=1}^{N}x_{m,b}\leq 1, \forall b$$

ELI uses the IPF scheduler to schedule clients within the channel, there-by adapting its resource schedule to account for the impact of interference in clients' channel rate.

Remark: The moving average of clients' channel access probability is calculated over few tens of sub-frames. Since each sub-frame lasts for only 1 ms, the time granularity used to calculate the access probability (and adapt resource scheduling accordingly) is more finer compared to the time granularity of the actual mobility of a client. This allows the IPF scheduler to serve for both mobile and static clients.

IV. Implementation and Evaluation

To evaluate the efficacy of ELI's design, we implement ELI on the NS3 simulator. We extended the current NS3-LAA's [8] basic implementation to fully support the LTE-LBT mechanism [1] on both the eNB and the client. Additionally, we implemented the LTE-TDD module, and combined it with the existing Lena-LTE module of the NS3 [9], to implement a full LTE-stack on both the eNB and the client.

A. Evaluation

We evaluate the efficacy of ELI in a systematic manner. We begin by evaluating the benefits of IPF scheduler alone (no dynamic channel selection) by comparing it with the native (interference-agnostic) PF scheduler. This allows us to estimate the gains resulting by simply making the eNB interference-aware in any given channel. Subsequent to this, we evaluate ELI's dynamic channel selection algorithm combined with the IPF scheduler to show its overall performance gains.

1) IPF scheduler: We consider 20 different topologies where an eNB continuously transmits DL data to its 25 clients. The clients are randomly distributed within the eNB's sensing/transmission range. Additionally, we randomly distribute many Wi-Fi nodes outside the sensing/transmission range of the eNB, that form the hidden terminals. All the hidden terminals are then made to transmit data to each other starting at random times, but for specific duration. The eNB calculates the average access probability of each client over a moving window of 50 sub-frames (50 ms). Each simulation is run for 10 min and the net-throughput in each case (IPF and native PF) is calculated.

Performance: The graphs in Fig.9a shows the average net-UDP-throughput gains of the IPF scheduler over the (native) PF for different number of hidden terminals. As

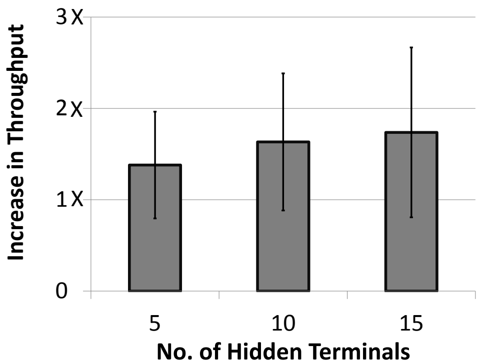

(a) Throughput gains of IPF

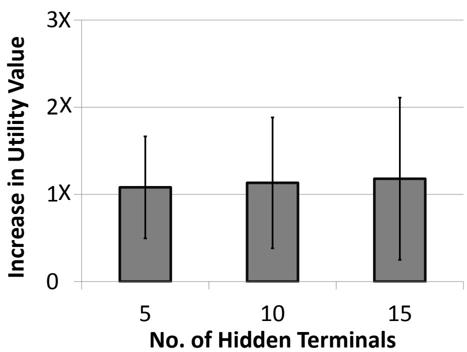

(b) IPF utility value gains

Fig. 9: IPF Vs. PF scheduler performance

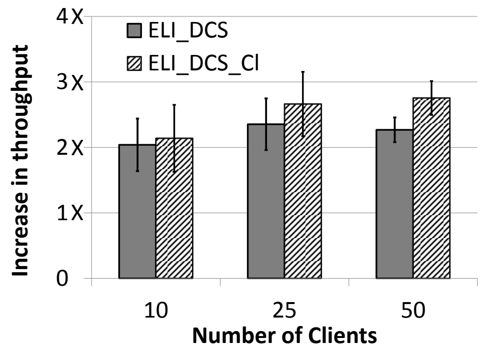

(a) Throughput gains

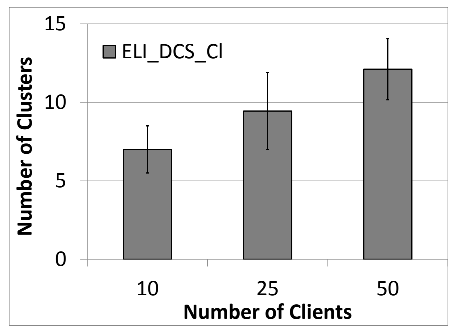

(b) Number of client clusters

Fig. 10: ELI's DCS Vs. Baseline DCS

expected, the IPF schedulers' throughput gains increase with the number of hidden terminals. While the native PF scheduler invariably schedules clients affected by hidden-terminals, IPF makes intelligent scheduling decisions by leveraging the interference diversity among the clients. With 10 and 15 hidden terminals we see that IPF can secure gains of 1.5x and 1.7x respectively, over the native PF scheduler.

Proportional Fairness: The graph in Fig.9b shows the utility($\sum_m log(R_m)$) gains of IPF over native PF (≈1.2X, on an average, when no. of HTs = 15). Since the utility captures both performance (throughput) and fairness, we can safely conclude that the IPF's performance gains do not come at the expense of fairness. IPF gains are because it considers both interference and channel diversity while scheduling access to clients, as against native PF which considers only channel diversity.

2) Dynamic channel selection (DCS): We now evaluate ELI's dynamic channel selection efficiency. The DCS process relies on the periodic interference statistics collected from every client across all the channels. Since ELI employs two optimizations (channel scanning and client clustering) to do this, we evaluate the performance gains achieved by each of the two optimizations in a stepwise manner. We first evaluate ELI's DCS (optimized channel scanning alone) without client-clustering, and then evaluate ELI's DCS with both optimized scanning and client-clustering optimizations. We compare the net-throughputs achieved by ELI each case to the net-throughput achieved when a baseline DCS algorithm (ported from WiFi) is used. The baseline DCS is the state-of-the art algorithm currently used by many different commercial APs [7]. The baseline DCS requires the eNB to scan all the channels periodically and select the channel with the least channel-occupancy. Since the baseline DCS does not require to take into account the interference caused by hidden terminals, the eNB makes channel-selection decisions independently without any client-feedback (akin to the WiFi AP).

Experimental set-up for this evaluation is the same as in IV-A1, except that the number of hidden terminals in each channel is now randomly selected (although in the range [5, 15]). Each simulation is run for 10 min with channel selection epoch set for every 1 min (refer III-A). Baseline DCS's per-channel dwell duration is fixed as 60 ms [7].

Performance: The graph in Fig. 10a shows the net-UDP-throughput gains achieved on the DL using the ELI's DCS algorithm with only optimized scanning (ELI _DCS) and with optimized scanning and client clustering (ELI _DCS_Cl), over the baseline dynamic channel clustering. We see that the ELI's

DCS scheme achieves throughput gains of > 2x (with and without clustering), over the baseline DCS scheme. ELI's client-clustering only helps in furthering the throughput gains over the baseline DCS algorithm, since it reduces the time required to collect interference statistics and make decisions faster. The average number of client-clusters formed across all the topologies when different number of clients were used is as shown in 10b. The time required to collect interference statistics does not proportionally increase with the network density. (10 clusters for 25 clients vs. 12 clusters for 50 clients). In point of fact, ELI's client clustering leverages the close distribution of clients in dense networks. This is further validated by the graph in Fig. 11a which shows that ELI reduces the time required to collect the interference statistics from all channels significantly (≈7x in 50 client cell), when compared to a naive approach of scanning employed by baseline DCS.

3) Interference diversity and client clustering: To show that interference diversity from across channels improves the accuracy of identifying co-located clients (and hence our client-clustering technique), we compare the throughput gains (Fig. 11b) achieved by ELI over baseline DCS (in a 50 client cell) when clients are clustered based on the interference statistics collected from varied number of channels. As evident from the graph, the client clusters' precision (fewer false-positives, leading to higher throughput gains) improves with the interference diversity used for clustering (determined by the number of channels' data used).

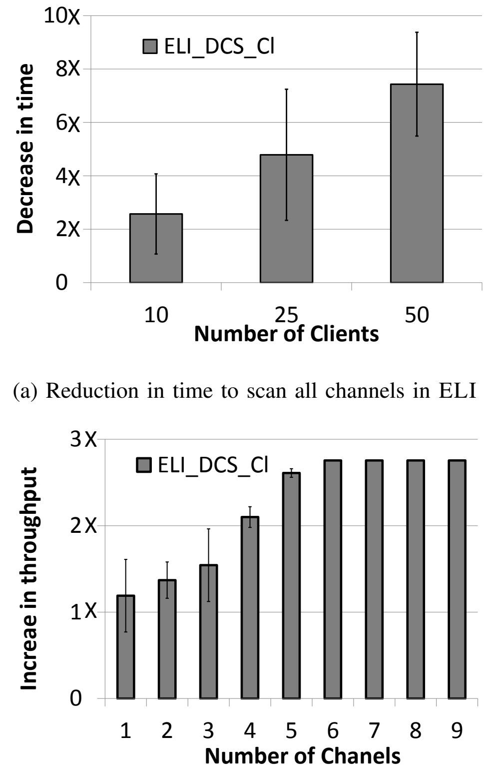

(a) Reduction in time to scan all channels in ELI

(b) Throughput gains Vs. Interference diversity used to cluster

Fig. 11: ELI's performance evaluation

V. Related Work

LTE and WiFi co-existence: While there is a plethora of research on the topic of interference impacting wireless nodes in the unlicensed spectrum, they have all inherently assumed the impact to be on technologies like WiFi and bluetooth that have been traditionally operating in the unlicensed spectrum. The advent of LTE onto the unlicensed band is more recent [10]. Consequently, researchers have initially focused more on ensuring fair co-existence of LTE and WiFi [6], [11], [12], [13], [14], [1]. In this paper, we focus on the next step and study the impact of LTE and WiFi co-existence on LTE's operations, through the lens of interference.

LTE scheduling mechanism: The strategy of manipulating the LTE's scheduling mechanism to achieve an objective has been successfully used the past [15], [16], [17], [18]. However, unlike before ELI leverages the interference diversity in the unlicensed spectrum along with the channel diversity offered by the proportional fair scheduler to schedule client access. [19] takes interference diversity into account to schedule uplink transmissions on clients, but does not deal with dynamic channel selection.

Dynamic channel selection: Dynamic channel selection is a technique that is commonly used to help WiFi APs periodically hop onto a better channel. Both industry [7], [20], [21] and academia [22], [23], [24] have proposed various heuristic based techniques to help APs dynamically select channels. However, all these solutions rely on heuristics or techniques that are intrinsic to WiFi to tackle the interference from hidden terminal in the channel selected. [25] does not take into account interference caused by hidden terminals. On the contrary, ELI argues that when it comes to LTE, considering interference across channels while selecting an operating channel is paramount and advocates using interference itself to implement an efficient dynamic channel selection.

VI. Conclusion

In this work we first showed the importance of handling interference caused by the hidden terminals in an LTE network operating in the unlicensed spectrum. Subsequently, we proposed ELI, a novel system for LTE to counter interference caused by the hidden terminals. ELI's three-pronged solution includes (a) a novel technique that equips LTE eNB to unilaterally, yet accurately measure the interferece at its clients' end; (b) novel and scalable scheme to measure interference at every client across all the available channels with affordable overhead; (c) schemes for LTE to incorporate interference awareness in its operations to improve its performance at both micro (interference aware scheduling) and macro time scales (dynamic channel selection). ELI's evaluation shows significant throughput gains for LTE in the unlicensed spectrum.

REFERENCES

[1] 3GPP Release 14, 36.213, "Channel Access Procedures for LAA," 2016.
[2] Mobile World Live, "LAA deployments," https://www.rcrwireless.com/20170804/carriers/verizon-starts-nationwide-laa-deployment-tag4, 2017.
[3] "AT&T LTE-LAA field trials ," http://about.att.com/story/lte_licensed_assisted_access_field_trials.html, 2017.
[4] Qualcomm White Paper, "MulteFire Specification," 2016.
[5] "WiFi transmission times."
[6] E. Chai, K. Sundaresan, M. Khojastepour, and S. Raagarajan, "Lte in unlicensed spectrum: Are we there yet?" ser. MobiCom '16, 2016.
[7] Cisco White Paper, "Cisco – Dynamic channel selection," http://www.cisco.com/c/en/us/td/docs/wireless/controller/technotes/8-3/b_RRM_White_Paper/b_RRM_White_Paper_chapter_0100.pdf.
[8] "NS3 - LAA laa-wifi-coexistence."
[9] NS3 – Lena project, "The ns-3 LTE module by the LENA project," https://www.nsnam.org/tutorials/consortium13/lte-tutorial.pdf, [Online].
[10] 3GPP Release 12, "Release 12 - 3GPP," 2015.
[11] 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 *GLOBECOM*, Dec 2015, pp. 1–6.
[12] A. Mukherjee, J.-F. Cheng, S. Falahati, L. Falconetti, A. Furuskar, B. Godana, D. H. Kang, H. Koorapaty, D. Larsson, and Y. Yang, "System architecture and coexistence evaluation of licensed-assisted access lte with ieee 802.11," in *Communication Workshop (ICCW), 2015 IEEE International Conference on*, June 2015, pp. 2350–2355.
[13] 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 *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 5083–5088.
[14] S. Yun and L. Qiu, "Supporting wifi and lte co-existence," in *IEEE INFOCOM 2015*, 2015.
[15] S.-B. Lee, S. Choudhury, A. Khoshnevis, S. Xu, and S. Lu, "Downlink mimo with frequency-domain packet scheduling for 3gpp lte," in *INFOCOM 2009, IEEE*. IEEE, 2009, pp. 1269–1277.
[16] S.-B. Lee, I. Pefkianakis, A. Meyerson, S. Xu, and S. Lu, "Proportional fair frequency-domain packet scheduling for 3gpp lte uplink," in *IEEE INFOCOM 2009*.
[17] A. Pokhariyal, G. Monghal, K. I. Pedersen, P. E. Mogensen, I. Z. Kovacs, C. Rosa, and T. E. Kolding, "Frequency domain packet scheduling under fractional load for the utran lte downlink," in *IEEE VTC2007-Spring*.
[18] M. Andrews and L. Zhang, "Scheduling algorithms for multi-carrier wireless data systems," in *ACM Mobicom*.
[19] R. K. Sheshadri, K. Sundaresan, E. Chai, A. Khojastepour, S. Rangarajan, and D. Koutsonikolas, "Blue-printing interference for robust lte access in unlicensed spectrum," in *ACM CoNEXT*, 2017.
[20] Ruckus wireless, "Radio control settings," https://docs.ruckuswireless.com/unleashed/200.4/c-ConfiguringRadioCtrlSettings.html, [Online].
[21] Aruba networks, "Adaptive Radio Management," http://www.arubanetworks.com/assets/tg/TB_ARM.pdf, [Online].
[22] M. Abusubaih, J. Gross, and A. Wolisz, "An inter-access point coordination protocol for dynamic channel selection in ieee802. 11 wireless lans," in *1st IEEE Workshop on Autonomic Communications and Network Management 2007 (ACNM 2007)*, 2007.
[23] D. J. Leith, P. Clifford, V. Badarla, and D. Malone, "Wlan channel selection without communication," *Computer Networks*, vol. 56, no. 4, pp. 1424–1441, 2012.
[24] M. Ihmig and P. Steenkiste, "Distributed dynamic channel selection in chaotic wireless networks," in *13th European Wireless Conference, Paris, France*, 2007.
[25] S. Sen, B. Radunovic, J. Lee, and K.-H. Kim, "Cspy: Finding the best quality channel without probing," in *ACM Mobicom*, 2013.