Fall 2019

Leveraging Geographical and Spectral Information for Efficient Cellular Systems

Matthew Tonnemacher
Southern Methodist University, mtonnemach@smu.edu

LEVERAGING GEOGRAPHICAL AND SPECTRAL INFORMATION FOR EFFICIENT CELLULAR SYSTEMS

Approved by:

Dr. Joseph Camp
Dissertation Committee Chairperson

Dr. Dinesh Rajan

Dr. Carlos Davila

Dr. Ping Gui

Dr. James Dunham

LEVERAGING GEOGRAPHICAL AND SPECTRAL INFORMATION FOR EFFICIENT CELLULAR SYSTEMS

A Dissertation Presented to the Graduate Faculty of the

Lyle School of Engineering

Southern Methodist University

in

Partial Fulfillment of the Requirements

for the degree of

Doctor of Philosophy

with a

Major in Electrical Engineering

by

Matthew Tonnemacher

M.S., Electrical Engineering, Southern Methodist University
B.S., Electrical Engineering, Southern Methodist University

December 21, 2019

ACKNOWLEDGMENTS

I have so much thanks to give to the friends, family members, professors, and peers, without which this achievement would not have been possible.

My advisor, Dr. Joseph Camp, has guided my growth as a graduate student through the entirety of this research. He has been both a role model and a friend. Likewise, I would like to thank Dr. Dinesh Rajan for his beyond invaluable insight and advice. I would also like to thank my colleague Rita Enami for the helpful discussions related to the work in Chapter 2.

Chance Tarver, my colleague and frequent co-author, has also helped me tremendously in my work. Specifically, I would like to thank him for his help with the hardware implementation in Chapter 2, editing and figure generation for Chapters 3 and 4, and frequent back-and-forth brainstorming sessions that helped inspire many other aspects of this work.

My family, especially my mother, have been a major source of support throughout my entire college career thus far and have my utmost appreciation. Most of all, I would like to thank my wife, Michelle, the love of my life who has stood by me through the good and bad alike.

Tonnemacher, Matthew M.S., Electrical Engineering, Southern Methodist University, 2013
B.S., Electrical Engineering, Southern Methodist University, 2011

Leveraging Geographical and Spectral Information for Efficient Cellular Systems

Advisor: Dr. Joseph Camp

Doctor of Philosophy conferred December 21, 2019

Dissertation completed January 9, 2020

With the unprecedented increase in mobile data demand and limited usable spectrum to provide for it, a paradigm shift towards spectrum sharing is a promising solution. However, there are many challenges that limit current spectrum sharing practices. One challenge is that proper spectrum sharing requires engaging devices to have an understanding of the impact they have on the ecosystem while transmitting in terms of spacial interference footprint and the implications to devices in their interference range. Another is that operators, especially licensed ones, have strict quality of service requirements for their subscribers, discouraging them from allowing unlicensed access of their purchased spectrum unless the sharing scheme can guarantee minimal impact to their systems. This dissertation seeks to address these challenges across three distinct works.

First, we used geographical features of a region to reduce in-field propagation experimentation by predicting the number of measurements required to accurately characterize its path loss, which can then be used to model coverage of arbitrarily positioned base stations. By exploiting the relationship between terrain feature complexity and measurement requirements, we found that the number of measurements collected to achieve a certain path loss accuracy over the entire region can be reduced by up to 58% in a high density drive testing scenario.

Next, we looked at applying Listen-Before-Talk (LBT) schemes in Citizens Broadband Radio Service (CBRS) networks for increasing the spatial reuse at secondary users while minimizing the interference footprint on incumbent and primary users. We used a novel Q-learning scheme to adapt the contention EDT to the changing network topology and traffic conditions, providing up to 350% gains in average secondary node user perceived throughput (UPT) in certain difficult

topologies with merely a 4% reduction in primary node UPT.

Finally, we studied channel selection in unlicensed Long-Term Evolution (LTE) cellular systems. We propose a mechanism for unlicensed LTE channel selection that not only takes into account interference to and from Wi-Fi access points but also considers other LTE operators in the unlicensed band. By collecting channel utilization statistics and sharing this information periodically with other unlicensed LTE base stations (eNBs), each eNB can improve their channel selection given their limited knowledge of the full topology via a proposed statistical and machine learning approach. We simulate operation in the unlicensed band using our channel selection algorithm and show how Wi-Fi load and inter-cell interference estimation can jointly be used to select transmission channels for all eNBs in the network.

TABLE OF CONTENTS

LIST OF FIGURES

Figure Page

LIST OF TABLES

This dissertation is dedicated to my wife, Michelle.

Chapter 1
Introduction

Mobile communication networks today are all enabled by a single, vital resource: spectrum. Like any resource, proper management of the accessible electromagnetic spectrum is necessary to enable coexistence between the myriad of technologies that seek to use it. In the United States, spectrum is managed by the Federal Communications Commission (FCC) to ensure that every commercial, military, medical, and all other applications have some place in the spectrum where they can operate with assurance to meet their operating requirements. Historically, spectrum has been a plentiful resource, able to meet the demands of all applications that require use of the medium. In the past couple of decades, however, the amount of practical spectrum has been stretched thin due, primarily due to the ever increasing media demands of mobile device users and the arrival of Internet of Things (IoT) applications. To keep up with this demand, an increasing volume of spectrum is being made available for mobile data, spanning each generation of communications technology (3G, 4G, and now 5G).

This scaling, however, is unsustainable. While spectrum is infinite in theory, the reality is that not all spectrum is created equal. Physical properties such as propagation, material penetration, and transceiver hardware requirements vary wildly at different frequencies, leading to bands of spectrum being selectively useful for different applications. In general, for large-scale mobile communications, frequencies below 6 GHz are significantly more useful than those over 6 GHz, thus making the practical use spectrum limited. With the resource limited and bandwidth demands on the rise, the generally accepted solution is to enable smarter spectrum sharing policies to improve the efficiency of its use. This thought is further driven by the fact that large chunks of licensed and unlicensed spectrum remain underutilized, blocked off from other technologies' access. The general idea is to have spectrum starved technologies access these spectrum holes in such a way that the impact to incumbent systems is limited.

Despite the elegance of the shared spectrum solution, there has yet to be widespread practice

of spectrum sharing policies. There are many reasons for this gap between idea and implementation. One major reason is that proper spectrum sharing requires engaging devices to have an understanding of the impact they have on the ecosystem while transmitting in terms of spacial interference footprint and the implications to devices in their interference range. Another is that operators, especially licensed ones, have strict quality of service requirements for their subscribers, discouraging them from allowing unlicensed access of their purchased spectrum unless the sharing scheme can guarantee minimal impact to their systems. In this dissertation, these challenges in spectrum sharing are addressed across three separate works.

Ensuring cellular coverage and spacial frequency isolation is an important consideration for carriers when trying to enable efficient spectrum sharing. Frequency boundaries of cells, especially in higher density, lower power deployments, is of utmost importance to manage interference footprints in a network. However, landscapes are complex and can change over time, making maintaining accurate mappings a costly concern for carriers due to the expense of in-field experimentation (i.e., drive testing). With the ubiquity of smartphones, apps, and social media, there has been an explosion of crowdsourcing to understand a vast array of trends and topics at a minimal cost to the organization. While cellular carriers might seek to replace the expensive act of drive testing with the nearly cost-free crowdsourcing, questions remain as to: *(i)* the accuracy of crowdsourcing, considering the lack of user control, *(ii)* the detection of when drive testing might still be required, and *(iii)* the quantification of how many additional in-field measurements to perform for a certain accuracy level. In the first chapter of the dissertation, we use geographical features of a region to reduce in-field propagation experimentation by predicting the number of measurements required to accurately characterize its path loss. In particular, we study the path loss prediction accuracy of drive testing and crowdsourcing by taking millions of measurements in a suburban and downtown region. We then use statistical learning to build a relationship between these geographical features and the measurements required. In doing so, we find that the number of measurements collected to achieve a certain path loss accuracy over the entire region can be reduced by up to 58% in a high density drive testing scenario.

The upcoming deployments of devices on the new 3.5 GHz, Citizens Broadband Radio Service (CBRS) is expected to enable spectrum sharing innovation by lowering the barrier to entry into LTE and other technologies. With a three-tiered spectrum-sharing solution, the CBRS band promises

to allow coexistence of federal incumbent users, priority licensees, and general users. While there have been many works studying cellular traffic offloading to unlicensed bands (e.g., Licensed Assisted Access) or minimizing interference in Cognitive Radio Networks, there has been comparatively little work on maximizing spatial reuse of spectrum in a shared spectrum CBRS network. In the second chapter of this dissertation, listen-before-talk (LBT) schemes are used in a CBRS network framework for increasing the spatial reuse at secondary (general) users while minimizing the interference footprint on higher-tier (incumbent) users. We propose LBT schemes that allow opportunistic access to licensed, CBRS spectrum and test our LBT schemes on a custom testbed with multiple software-defined radios and a real-time signal analyzer. We find that by allowing LBT spectrum sharing in a two carrier, two eNB scenario, we see upwards of 50% user perceived throughput (UPT) gains for *both* eNBs. Furthermore, we examine the use of Q-learning to adapt the energy-detection threshold (EDT), combating problematic topologies such as hidden and exposed nodes. When adapting the EDT of opportunistically transmitting nodes, we see up to 350% gains in average secondary node UPT in certain difficult topologies with merely a 4% reduction in primary node UPT.

As mentioned earlier, due to the ever increasing mobile data demand, cellular providers have been investing into high-density, low coverage small cell deployments, enabling greater spatial frequency reuse. Additionally, unlicensed industrial, scientific, and medical (ISM) bands are seen as a potential bandwidth extension to increase capacity. The desire to utilize unlicensed bandwidth for cellular data has prompted several standardization efforts for LTE operation in the unlicensed band such as LTE-U and Licensed Assisted Access. The major concern for these efforts has been coexistence with existing devices, however we found the research into inter-cell interference between these unlicensed LTE small cells, especially in the presence of heterogeneous Wi-Fi contention between nearby cells, to be lacking. In the third chapter, We propose a mechanism for unlicensed LTE channel selection that not only takes into account interference to and from Wi-Fi access points but also considers other LTE operators in the unlicensed band. By collecting channel utilization statistics and sharing this information periodically with other unlicensed LTE eNBs, each eNB can improve their channel selection given their limited knowledge of the full topology. Moreover, we expand on our statistical channel selection formulation to include reinforcement learning and show that the addition of reinforcement learning can augment channel selection in cases where

there is high similarity in channel utilization for neighboring nodes. We simulate operation in the unlicensed band using our channel selection algorithm and show how Wi-Fi load and inter-cell interference estimation can jointly be used to select transmission channels for all small cells in the network. Our approaches lead to an increase in user-perceived throughput and spectral efficiency across the entire band when compared to the greedy channel selection.

This dissertation is organized as follows. In the first chapter, we present GeoRIPE, a framework under which we can reduce the cost of cellular network drive testing. In the second chapter, we propose a novel spectrum sharing methodology for CBRS networks based off of popular listen before talk solutions. In the third and final chapter, we study channel selection in unlicensed LTE networks, specifically considering heterogeneous Wi-Fi interference between cells.

Chapter 2

GeoRIPE: Efficiently Harvesting Field Measurements for Map-Based Path Loss Modeling

2.1 Introduction

To address multi-fold increases in cellular demand, carrier cell sizes are shifting downwards to maximize network capacity. In doing so, the accurate and fine-grained coverage estimation of coverage becomes a critical issue for spatial reuse, inter-cell interference, and smooth handoffs between cells. Historically, in-field experimentation (i.e., drive testing) has been used to estimate the cellular propagation of a given region, which is costly for network operators due to the man-power and equipment required. The emergence of smartphones and their apps have offered a far cheaper alternative, recording in-field network observations directly from the cellular users (i.e., crowdsourcing). Crowdsourcing can allow network observations to be recorded in areas to which in-field testers may not have access.

Despite the availability of crowdsourced measurements, network providers continue to use extensive drive-testing to validate network coverage and quality of service metrics. One of the main goals in supplementing crowdsourced measurements with drive testing measurements is to fill in the gaps of crowdsourcing. However, the drive testing process can be costly and time consuming, creating a market for network drive testing reports. Thus, collecting these measurements in an efficient manner is a high priority. The most direct approach would be to minimize the number of measurements that needs to be collected to achieve the accuracy requirements set. In this work, we attempt to reduce the number of measurements required for accurate path loss characterization throughout a given region by understanding and exploiting the variation in geographical features. In order for this to happen, several questions need to be answered:

1. What are the relationships between geographical features and signal strength measurement requirements?

2. Can the number of measurements required to achieve a particular accuracy level be deter-

mined for a specific area?

3. How can these spatial differences be exploited to reduce the total number of measurements required to meet a particular regional path loss characterization accuracy?

In this work, we use geographical features of a region to more efficiently collect signal strength measurements, thereby reducing the amount of time spent on in-field propagation experimentation. To do this, we first introduce GeoRIPE, a statistical learning framework to situationally predict the number of measurements required to meet a specified path loss characterization precision. With this framework, geographical feature distribution input is used to suggest measurement collection requirements in a grid-like fashion over the target region. Then, we developed and deployed an Android application to gather signal strength measurements from real users throughout the world. We use a specially modified version of this application to collect high-density drive testing measurements from two distinct region types in a major metropolitan area. By using commercially available smartphones, measurements gathered are comparable to those gathered using purely crowdsourcing. Next, we explore the effect of land use on path loss characterization, showing how geographical feature diversity plays a large role in determining regional measurement requirements. We show especially strong correlation between the number of measurements required to accurately characterize the path loss in a region and the geographical feature ratio of small, medium, and large buildings, foliage and free space in an area. Finally, we validate the framework by comparing GeoRIPE to uniform measurement collection approach. We show that when fixing the accuracy requirement in path loss evaluation over the entire region, using the GeoRIPE framework can significantly reduce the required number of measurements that need to be taken to meet it.

2.2 GeoRIPE Framework

To illustrate the GeoRIPE[1] framework of using geographical features to reduce in-field propagation experimentation, we have depicted an aerial view of a region's terrain in the left-most image in Figure 2.1 with north on the top of the image. The terrain is classified into the following geographical features: buildings, dense foliage, and free space. Since in-field testing is expensive, our

[1] As a mnemonic for this work, consider that fruit should be in the field the appropriate amount of time before harvested (i.e., to be *ripe*). Similarly, we seek to find the minimum amount of time necessary for in-field experimentation to accurately predict the path loss of a region.

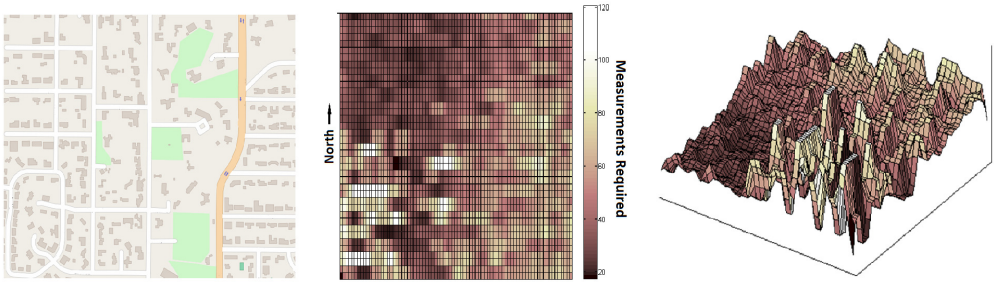

Figure 2.1: GeoRIPE Framework: using geographical features of a region (left) to infer the number of measurements required to predict path loss with a given accuracy level. The middle figure shows minimum measurements required (light is high, dark is low) of the same spatial distribution as the left-most figure. The right figure is a 3-D version of the middle figure.

goal is to predict the least number of propagation measurements required to characterize the path loss in the region according to a specified level of accuracy. For example, the path loss of a region which has entirely free space (e.g., a desert) could be characterized with very few measurements. However, a diverse metropolitan region would require far more measurements to characterize.

The middle and right images of Figure 2.1 depict the end result of the GeoRIPE framework. The middle image depicts a two-dimensional overlay of the measurement density required to characterize the region shown in the left-most image. The southwest portion of the region is the lightest color, which means that it requires the greatest number of measurements to characterize due to the high concentration of buildings. In contrast, the northwest portion of the region is the darkest, which means that it requires the least number of measurements to characterize due to sparse building placement and less foliage. A three-dimensional version of the same figure can be seen in the right-most image to show the quantity of measurements required in each portion of the region.

A similar analysis could be done for a given region by an excessive amount of in-field testing and finding when each portion of the terrain converged to a particular level of path loss prediction accuracy. However, such an approach, by definition precludes any in-field testing reduction. While we take a very large number of measurements in certain regions in our work, we do so to train a statistical learning decision structure to infer the number of in-field measurements required. As a result, for any mix of such terrain features, we can avoid the two in-field testing extremes of: *(i)* too few measurements, resulting in an inaccurate path loss estimate, or *(ii)* too many measurements, resulting in excessive experimentation costs. We can then evaluate the viability of using crowdsourcing to lower the drive testing cost.

The GeoRIPE framework's measurement distribution prediction is made with path loss accuracy in mind. So, before we evaluate the framework itself, we first need to give some background on path loss models in general as well as what model we use for our analysis. Path loss models attempt to predict the electromagnetic propagation as a function of distance. Many of these models rely on *a priori* information, using environmental details, a theoretical foundation, empirical findings, or some combination of the three for their prediction [4–6]. Other techniques operate under the assumption that *a priori* information is insufficient. These models supplement an existing model with a correction factor or factors based on measurements collected throughout a region to be modeled and tend to be more accurate than their *a priori* counterparts [7]. These active measurement models consist of two fundamental components: *(i)* how the measurements are gathered, and *(ii)* how they are incorporated into the model.

2.2.1 Path Loss Measurement and Supplemented Models

W. C. Lee studied the initial theoretical methodology of gathering active measurements for modeling path loss [8]. Lee proposed arced measurements at incremental distances from the transmitter while averaging measurements that fall within 20 to 40 wavelengths of each other, a claim corroborated by Shin using IEEE 802.11b measurements some years later [9]. In practice, it is often difficult to collect measurements strictly following the theoretically ideal guidelines due to environmental inaccessibility. This can be due to permission limitations, such as access restricted buildings or construction sites, or infrastructure limitations, where equipment setups are subject to the same mobility freedoms as the vehicles in which they operate. With a crowdsourced approach, a greater access diversity can be achieved with the limitations of a lack of control over data validity and input distribution. Due to these practical considerations, our work considers geographical complexity and its role in characterizing a region, both with vehicle-based drive testing and app-enabled crowdsourcing.

One of the more recent path loss models utilizing collected measurements is one proposed by Robinson *et al.* [10]. Using the Technology For All (TFA) network in Houston, TX, they utilized a modified Flexible Path Loss Exponent model with a terrain correction factor derived iteratively from collected measurements. The model is an extension of Friis' fundamental study [4] and can

be written as:

$$\begin{aligned} P_{rx} =& P_{tx} + 10\alpha \log(d) - 20 \log(f) \\ & - 20 \log \left(\frac{4\pi}{c} \right) \end{aligned} \tag{2.1}$$

Here, P_{tx} and P_{rx} are the transmitted and received signal powers, respectively, α is the path loss exponent, f is the transmit frequency, d is the distance from the transmitter, and c is the speed of light. In their work, the authors use existing wireless mesh nodes and detailed terrain information to determine sections that are likely to share a similar path loss exponent. They then incrementally gather measurements around the borders of these sections in a push-pull algorithm to refine the coverage estimate of the mesh node.

2.2.2 Obtaining a Path Loss Exponent

In our statistical learning approach, it is necessary to train a classifier with path loss exponent observations derived from existing measurements to motivate predictions in areas that lack those same measurements. We borrow the idea of a spatially-dependent path loss exponent from Robinson *et al.* without the push-pull measurement adjustment algorithm, a reference node, and detailed terrain information (including material loss estimations). Instead, we use (2.1) in a square-shaped moving window over the region, using linear regression to obtain a path loss exponent for each window. Since the measurements are obtained from many different towers distributed over the area, each using potentially different transmit powers at different heights, we rely on a larger quantity of data to average out these inconsistencies. However, the accuracy (which we define as inversely proportional the standard deviation of obtained path loss exponents over several calculations using orthogonal measurements) is increased, which we rely on more heavily for our statistical learning framework. To calculate the metric of standard deviation on the path loss exponent, we divide the data considered into several independent sets, calculate path loss exponent for each independent set, and compute the standard deviation of the exponents derived. Again, this gives us a solid metric for path loss precision, even if the exponents themselves are biased by the data collection limitations.

2.3 In-Field Wireless and Geographical Data

In this section, we present our Android-based measurement gathering platform, which will be leveraged locally by us to gather a dense measurement set of wireless signal strengths in both a downtown region and suburban region. We also introduce the geographical feature data set that we use from the drive tested regions to establish a relationship between geodata and the attenuation of wireless signals. By using a smartphone based collection platform, we can gather Received Signal Strength Indicator (RSSI) measurements that relate more directly to user experience than measurements collected with traditional network analyzing hardware.

2.3.1 Local Measurement Collection

Over the span of two weeks (over 30 in-car hours), we collected 6.7 million drive testing measurements by placing LG Nexus 4 smartphones in a vehicle and thoroughly driving throughout two regions in a snake like pattern, covering all available roads in each region. Since we are using the measurements for studying region-based path loss characteristics, the specific cellular technology used is less important. Therefore, the measurements were collected on GSM networks as they are still the most prevalent. The measurements were obtained at a relatively constant speed of 30 mph in two different areas of the Dallas metropolitan area. The first area is a suburban region several miles north of the city center with lush greenery prevalent throughout and is predominantly residential. The second area is in downtown, where there is far less vegetation, and the buildings are far taller than the suburban structures with non-uniform heights. Our goal is to use these two distinct regions to examine how differences in feature distribution affect the number of required measurements to characterize path loss to a certain degree of accuracy in each region.

2.3.2 Received Signal Strength in Android API

Each cellular measurement contains an RSSI field for each visible cellular tower, a GPS location, an accuracy reading, and physical speed of the device. While we now obtain RSSI readings in terms of dBm, most of our measurements were taken when the API reported RSSI in terms of Arbitrary Strength Units (ASU), an android specific quantized signal strength metric, which quantizes obtained RSSI values for GSM to 32 different levels shown in the equation below from the

Android API [11].

$$P_{rx}(dBm) = 2 * P_{rx}(ASU) - 113 \tag{2.2}$$

$$P_{rx}(ASU) = [0, 31] \tag{2.3}$$

We consider $P_{rx}(ASU) = 0$ and 31 unusable since they correspond to SNR in an unlimited range; an ASU value of 31 includes any RSSI value above -51 dBm. Not including these measurements, however, clips the natural distribution of RSSI readings at locations with measurements near the quantization limits. The lower and upper bounds set by omitting measurements where $P_{rx}(ASU) = 31$ and $P_{rx}(ASU) = 0$, respectively, move the average RSSI at certain distances from the tower. Distances closer to the tower that generally have higher RSSI measurements near the upper bound may have a lowered average RSSI. Conversely, distances farther from that tower that generally have lower RSSI measurements near the lower bound may have a heightened average RSSI. The bias in the movement of average RSSI near the boundaries could end up changing (likely reducing) the value of the obtained path loss exponent. While the exact values of RSSI and path loss exponents are likely affected by the quantization error, we are not evaluating absolute path loss accuracy, only relative accuracy in our experiments, so the bias does not affect our results.

2.3.3 Geographical Feature Data

In order to obtain geographical feature information, we utilized an open-access online resource, Open Street Maps (OSM) [12], to identify, outline, and label specific regional features and output them to an easily accessible data structure for parsing. To this end, we mapped hundreds of offices, parks, houses, and other features in both the suburban and downtown regions and grouped them into the feature category classes for our statistical learning system. With statistical learning, the number of training observations necessary for accurate divisions scales up proportionally with the number of features used in the training. Due to this so-called curse of dimensionality, as well as the limited number of possible features to label in each region considered, it is necessary to divide all possible geographical features into relatively few feature categories for processing. With this in

mind, we selected five feature categories under which all features were labeled: small buildings, medium buildings, large buildings, high foliage, and open space. In this system, we define small buildings consist of buildings that are under 5 stories tall (ground footprint is not considered for the category, but is implicitly considered when calculating feature distributions in a region). Similarly, we define the range of medium buildings as being between 5 and 15 stories tall and large buildings as being over 15 stories tall. These building height tiers were chosen to give each feature type non-trivial representation in the learning algorithm. Finally, we consider high foliage areas in the regions are areas with a large number of trees, and open space is the area defined by the complement to the set of all other features combined and includes structures such as roads, parking lots, etc. It is important to note that the feature set we consider is far from ideal; with more detailed geo-spatial feature data that is currently unavailable to us (such as exact building and foliage canopy heights), the GeoRIPE framework's accuracy will only improve.

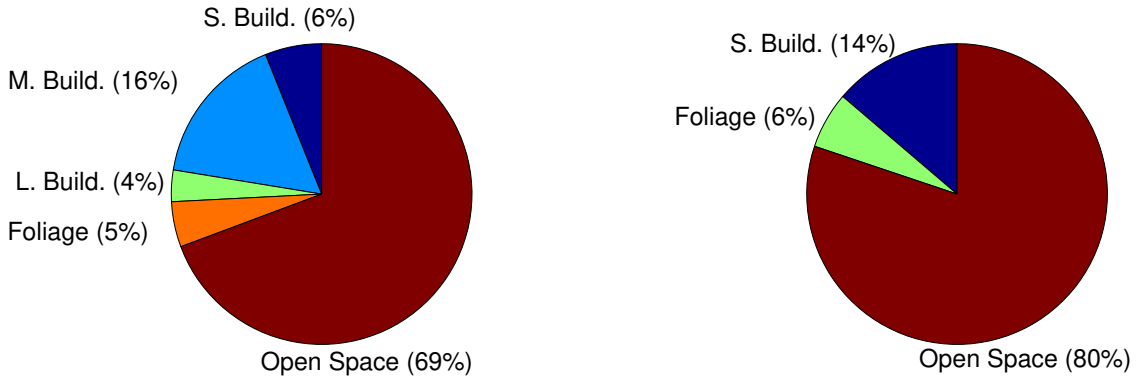

(a) Downtown

(b) Suburban

Figure 2.2: Regional Feature Distributions

The overall ratio of features in the downtown and suburban regions we examine are shown in Figures 2.2a and 2.2b, respectively. These ratios represent the relative space occupied by each

feature according the following equations:

$$s = [f_1, f_2, ..., f_5] \tag{2.4}$$

$$\sum_i s_i = 1 \tag{2.5}$$

where s is a weighted vector for the normalized occupancy of each of the 5 features in the full region in terms of total feature area. From this figure, we can see that the suburban area lacks medium and large buildings and has a higher percentage of open space than the downtown region, as anticipated. Ideally, we will be able to further differentiate and parse members of the open space set to derive additional feature categories in the future.

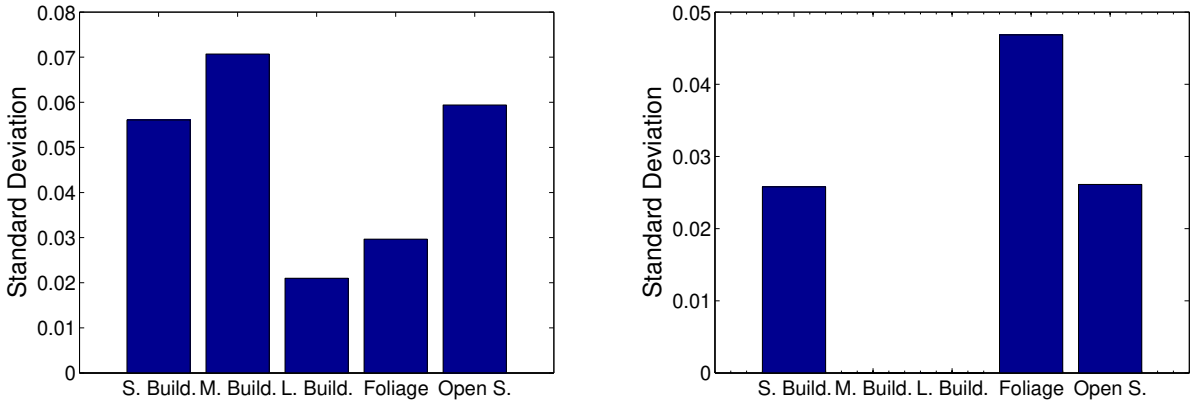

(a) Downtown

(b) Suburban

Figure 2.3: Regional Feature Distribution Deviation

Primarily, we want to examine how the features change over each region. To do this, we calculated the ratio of features in a moving window over each region. Treating the percentage of each feature in the windows as a random variable, we then compute the standard deviation, giving us a picture of the feature variability as we move across each region. The standard deviation of features can be seen for the downtown and suburban regions in Figures 2.3a and 2.3b, respectively.

From these figures, we can see that the variance of features in the suburban region is, overall, significantly smaller than in the downtown region. This is because, in the suburban region, the grouping of features are polarized (ex: houses in half the region, foliage in the other half), while the in the downtown region, neighboring areas have a higher diversity in their feature composition.

2.4 Density-Dependent Tiling of In-Situ Data

While we have examined the differences in geographical feature distribution of the two areas, we have yet to explore the impact of changing the size of the subregion, or tile, used to group measurements spatially. In this section, we explore the differences in path loss exponent changes between these regions, their relative sub-regions, and the trade off between tile size, measurement density, and measurement error in evaluating path loss.

2.4.1 Extreme 1: Highly-Sparse Crowdsourced Data

The first scenario is one in which the data set has very few measurements. In such a situation, we need all the measurements we have to assign a single path loss exponent to characterize an entire region, similar to the traditional approach. In other words, dividing the region into smaller areas to have more path loss precision cannot occur because there is a lack of a sufficient number of measurements to compose a path loss exponent estimate. The result here can be considered a rough average of path loss over the entire region; however, accuracy at any given area depends on the variability of the region itself. While a single exponent over an entire region may create a simpler coverage calculation, it may not be accurate, especially in more diverse region types such as large cities. When enough measurements are available, we can divide the region into independent tiles for characterization based on measurement density and region type instead. Figure 2.4 shows the suburban region is divided into 6 and 24 tiles for path loss evaluation, demonstrating the disparity in derived path loss exponents for the same area using different tile sizes. While some smaller tiles match the their large tile counterparts, others are different, alluding to diverse environmental characteristics.

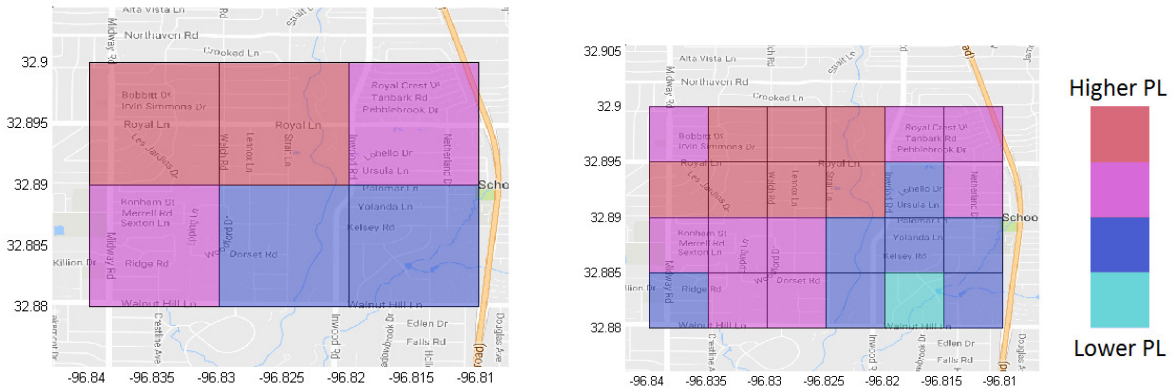

(a) 6 Tiles

(b) 24 Tiles

Figure 2.4: Suburban region path loss divided into a) 6 and b) 24 tiles.

2.4.2 Extreme 2: Highly-Dense Crowdsourced Data

We now examine the other extreme, when a very large number of measurements are available. In this case, the tile size is not limited by measurement density or acceptable error. With highly-dense measurements (*e.g.*, as the number of measurements approach infinity), tile size is virtually unbounded, and the standard deviation of path loss approaches zero (orthogonal subsets measurements would regress to the same exponent when evaluating as the subset size approaches infinity). Instead, the variability of the terrain determines the effective lower bound on the terrain characterization resolution, preventing the tile size from going to zero. In other words, decreasing the title size resolution after a certain point does not provide any additional information about path loss in the region.

Region	Tiles	Diff. Mean	Diff. Variance
Downtown	6	0.0668	0.0040
	24	0.0944	0.0059
Suburban	6	0.0742	0.0029
	24	0.0464	0.0014

Table 2.1: Mean and variance of differences between neighboring tiles' path loss.

Depending on the nature of the region being analyzed, the mean and variance of the difference between neighboring tiles changes with the tile size. Table 2.1 shows the mean and the variance in path loss exponent calculation differences between neighboring tiles for both the downtown and suburban regions of different tile sizes. As the tile size decreases, we observe different behavior from the two region types. In the downtown region, the differential mean and variance increase with a smaller tile size, while in the suburban region, the opposite is true. For more diverse regions like downtown, using smaller tile sizes has a larger benefit in characterizing the spatial diversity of path loss. In less diverse regions like homogeneous neighborhoods, the differential path loss throughout the region does not require as high of a resolution; the path loss variability seen from smaller tile sizes is below the noise floor in generating the path loss exponents. Thus, the measurement density available along with the region type's path loss variability must be jointly considered in determining a minimum tile size for characterization.

2.4.3 General Case: In-Situ Tile Size Adaptation

In most cellular networks, it is likely that the set of available measurements is neither highly sparse nor infinitely dense. Instead, the system is generally in a state between these two extremes. Hence, choosing the tile size of the region becomes a critical issue since it is not initially clear if the measurement density or the terrain heterogeneity will drive the tile size.

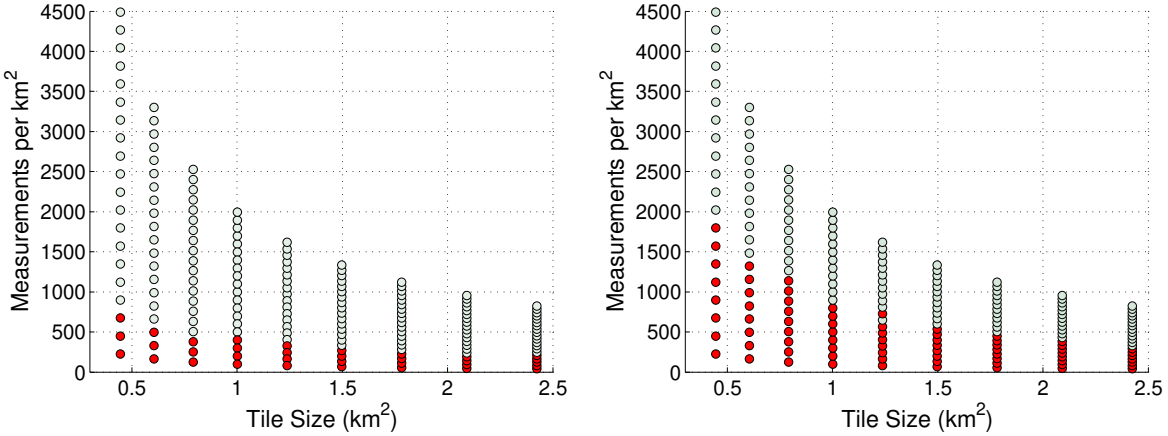

(a) Suburban (b) Downtown

Figure 2.5: Examining tile size versus measurement density over different land uses for a given acceptable error.

Given a set of measurements with a specific measurement density, a minimum tile size exists that remains below the reliability threshold. In Figures 2.5a and 2.5b we set a standard deviation threshold of 0.0125 for our path loss evaluation and examined the minimum tile size for different measurement densities available. In these figures the lighter points are below the threshold and darker, red dots are above the threshold. For both land use types, we see that as the measurement density increases, the minimum tile size achievable under the threshold decreases, enabling a finer grain resolution while maintaining the reliability we desired. However, the suburban region consistently requires a lower measurement density to be below the error threshold than the downtown region because the suburban region has a lower terrain variability. From these figures, we can see that to achieve a tile size of 1 km^2 under the error threshold, the suburban region requires a measurement density of 500 measurements per km^2 while the downtown region requires 800. This relationship between the different land uses holds for each other tile size as well. The measurement density required for a certain resolution of path loss increases with the heterogeneity of the region.

2.5 Experimentally Evaluating Map-Based Modeling

Despite that fact that there have been several works that suggest measurement distribution and geographical features play an enormous role in the resulting path loss characterization of a region, there has not been a study showing how these metrics can be used to quantify the number of measurements required to characterize an area. In this section, we take a critical look at the impact of measurement distribution and geographical feature components on path loss precision. More specifically, we compare measurement distributions obtained from crowdsourcing versus drive testing measurements, examine geographical feature components of our two metropolitan region types, and correlate these feature distributions with both path loss exponents as well as the number of measurements required to obtain a certain precision in characterization. Our goal is to use regional geographical features to learn how to properly collect measurements, ensuring a predetermined precision in path loss characterization.

2.5.1 Path Loss Metric and Geographical Feature Correlation

Using geographical features as a region specific identifier, we want to understand how specific geographical features can be used to characterize path loss throughout a region. We now explore four different path loss related metrics to determine which had the closest relationship, and there-

fore the highest suitability, to be used as the target for our geographical feature based statistical learning approach. The four metrics we examine are path loss exponent (PLE), differential path loss exponent (DPLE), number of measurements required (MR) for path loss convergence, and the differential number of measurements required (DMR) for convergence. The MR and PLE metrics are calculated for a given region using Algorithm 1, which is initialized with parameters listed in Table 2.2. Algorithm 1 can be visualized as a sliding window filter moving across the region as illustrated in Figure 2.6. In this algorithm, the first two loops control the moving window as it shifts vertically and horizontally, respectively. For a given window at position v, h, we compute the path loss exponent directly with all available data, giving the PLE metric. Following that, we divide the data into G separate groups, calculate the path loss exponent in each group, and take the standard deviation over all exponents. We increase the number of measurements in each of these groups by S until the standard deviation is under a certain threshold (chosen to be whatever accuracy is acceptable, we chose 0.03 because that was about the point that an a linear increase in the number of measurements started to have diminishing returns). Additionally, When the standard deviation falls under this threshold, we record the measurements in each group as the MR metric.

We use this algorithm to determine a map of the measurements required and path loss exponent metrics over a region. With our feature data for the region, we can derive a similar map of feature distributions using the same windowing method. The differential metrics, differential path loss exponent and differential measurements required, can be easily derived from column and row differentiation of the PLE and MR matrices, respectively. A corresponding differential feature distribution map can be derived in the same manner. With matching metric and feature maps, we can correlate each metric with the corresponding feature map to obtain a sample Pearson correlation

Table 2.2: Spatial feature and path loss metric algorithm parameters.

Parameter	Setting	Description
W	$1km^2$	Moving window area
V	20	Number of vertical shifts
H	40	Number of horizontal shifts
σ	0.03	Desired std. dev.
S	20	Measurement step size
G	30	Number of orthogonal groups

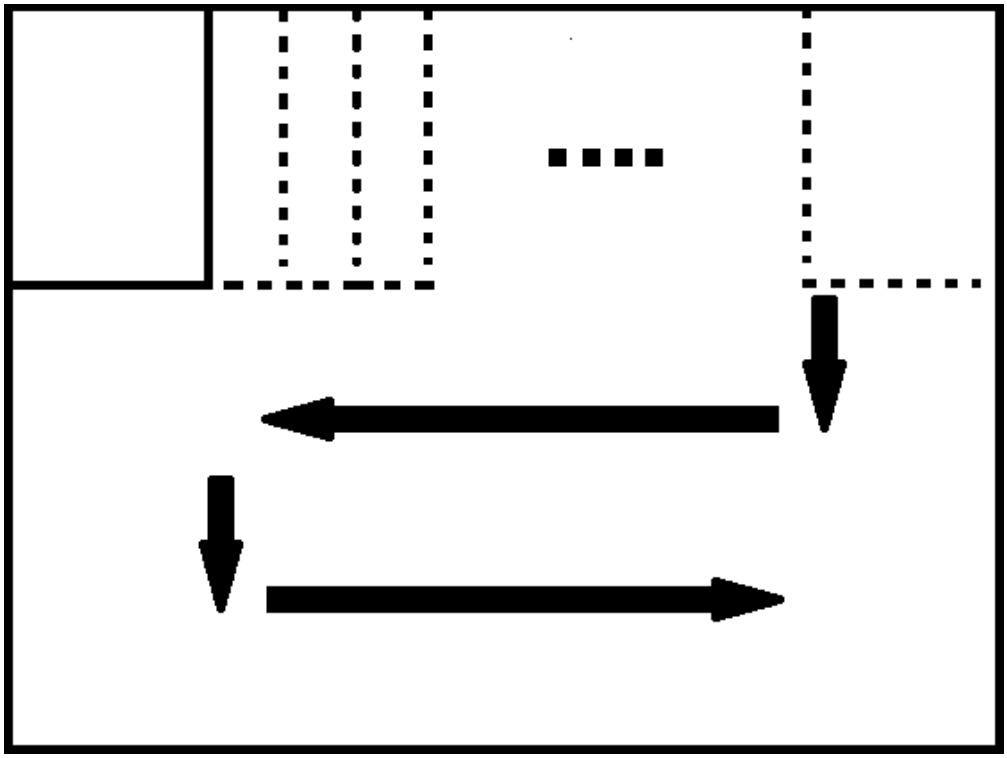

Figure 2.6: Visualization of Algorithm 1.

Data: measurements (M)
Result: V by H PLE and MR matrices
Initialize Parameters;
for $v \leftarrow 1$ **to** V **do**
  for $h \leftarrow 1$ **to** H **do**
    PLE$(v, h) \leftarrow ComputeExponent(\forall M \in W)$;
    $group(1..G) \leftarrow \forall M \in W$ split into G sets;
    $P_{size} \leftarrow 0$;
    while $\sigma_{temp} \geq \sigma$ **do**
      $\sigma_{temp} \leftarrow \infty$;
      $P_{size} \leftarrow P_{size} + S$;
      for $g \leftarrow 1$ **to** G **do**
        $P \leftarrow P_{size}$ elements $\in group(g)$;
        $exponent(g) \leftarrow ComputeExponent(P)$;
      end
      $\sigma_{temp} \leftarrow ComputeStdDev(exponent(1..G))$;
    end
    MR$(v, h) \leftarrow P_{size}$;
    $W \leftarrow W$ horizontally shifted by 1;
  end
  $W \leftarrow W$ vertically shifted by 1;
end

Algorithm 1: Algorithm for computing PLE and MR metrics.

coefficient (the standard equation for correlating discrete groups) using Equation 2.6.

$$r_{ij} = \frac{\sum_{k=1}^{n}(x_{ik} - \bar{x_i})(y_{jk} - \bar{y_j})}{\sqrt{\sum_{k=1}^{n}(x_{ik} - \bar{x_i})^2 \sum_{k=1}^{n}(y_{jk} - \bar{y_j})^2}} \tag{2.6}$$

In this equation, n is the number of samples, x_{ik} is sample k of feature i, y_{jk} is the sample k of path loss metric j, and $\bar{x_i}$ and $\bar{y_j}$ are the average distribution of feature i and the average of path loss metric j, respectively.

We want to select a path loss metric to use as a training class for the statistical learning framework that has the highest correlation coefficients with the feature set to provide clear decision boundaries. The correlation coefficients for each of the path loss metrics in the downtown and suburban regions are shown Tables 2.3 and 2.4, respectively.

We can see that the MR metric has the highest overall correlation coefficient magnitude and is likely the best contender for a simple decision tree based learning algorithm. Interestingly, we see that for the suburban region in particular, the correlation coefficient for the MR metric are negative with small buildings and positive with foliage, while positive with both for the PLE metric. This suggests that while increased buildings and foliage contribute to a larger path loss exponent (as expected), the number of measurements required to drop below the 0.03 path loss exponent standard deviation increases only with the percentage of foliage.

From this result, we can see that while the small buildings feature increases the path loss exponent, it decreases received power variability, while the foliage feature increases received power variability. This trend is visualized for the downtown and suburban areas in Figures 2.7a and 2.7b.

To understand how each terrain feature affects the measurement requirements individually, we

Table 2.3: Downtown metric-feature correlation coefficients.

Metric	S. Building	M. Building	L. Building	Foliage	Open Space
PLE	-0.23	0.02	-0.05	0.17	0.18
DPLE	-0.05	-0.03	0.02	-0.04	0.07
MR	**-0.32**	**0.34**	**0.49**	-0.10	-0.23
DMR	-0.06	0.05	0.03	0.01	-0.02

Table 2.4: Suburban metric-feature correlation coefficients.

Metric	S. Building	M. Building	L. Building	Foliage	Open Space
PLE	**0.36**	NA	NA	**0.31**	**-0.38**
DPLE	-0.15	NA	NA	0.05	0.05
MR	**-0.53**	NA	NA	**0.44**	-0.27
DMR	-0.06	NA	NA	0.01	0.04

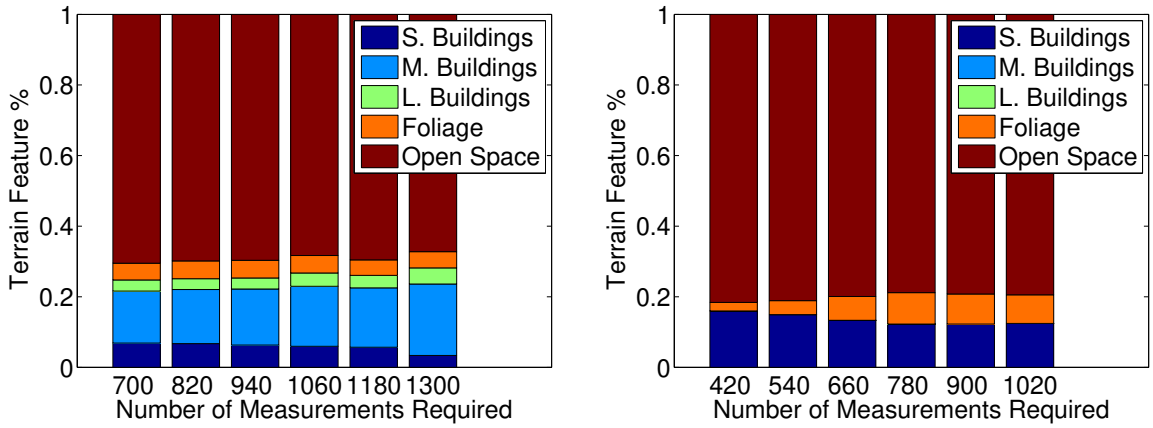

(a) Downtown (b) Suburban

Figure 2.7: Average feature distributions for different MR tiers.

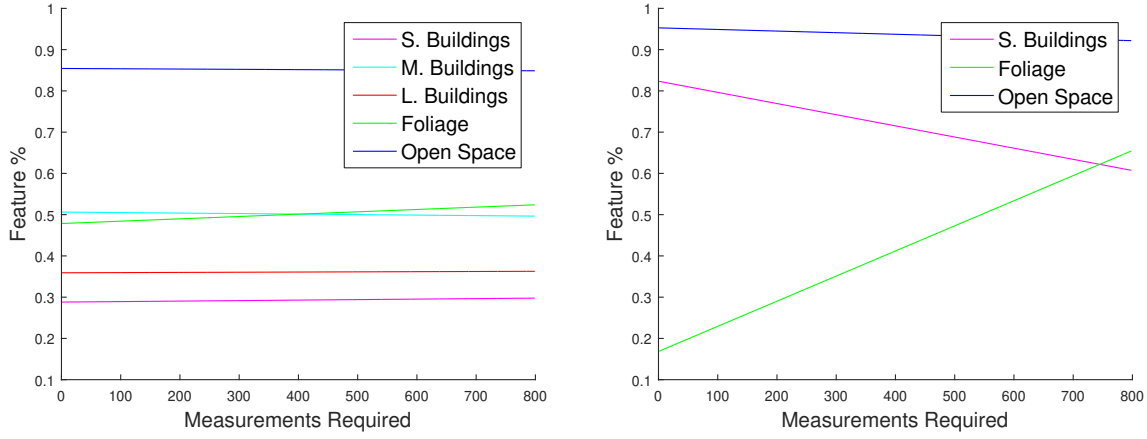

(a) Downtown (b) Suburban

Figure 2.8: Feature trends for increasing measurement requirements.

examined the trends of each feature distribution as the number of required measurements increases in Figure 2.8. From this figure, we can see that in the suburban region, the individual feature impact is quite clear; increases in the percentage of foliage and decreases in the percentage of small buildings increases the number of measurements required, while the open space component doesn't fluctuate much at all. Conversely, we cannot induce distinct trends from the downtown region. We see in the suburban region, there are only two features driving the increase in measurement requirements, thus trends can be easily seen. In the downtown region, however, each feature apart from open space has an effect on the measurements, thus trends from individual factors cannot be so easily derived.

2.5.2 Classifier Training for MR Prediction

To validate the GeoRIPE framework, we divide the MR results for the downtown and suburban regions into 6 same-sized class bins. As seen in Figure 2.9, the class groupings are not homogeneous for either of the regions. Unsurprisingly, the downtown region class distribution has a higher mean number of required measurements than the suburban region due to its higher geographical complexity.

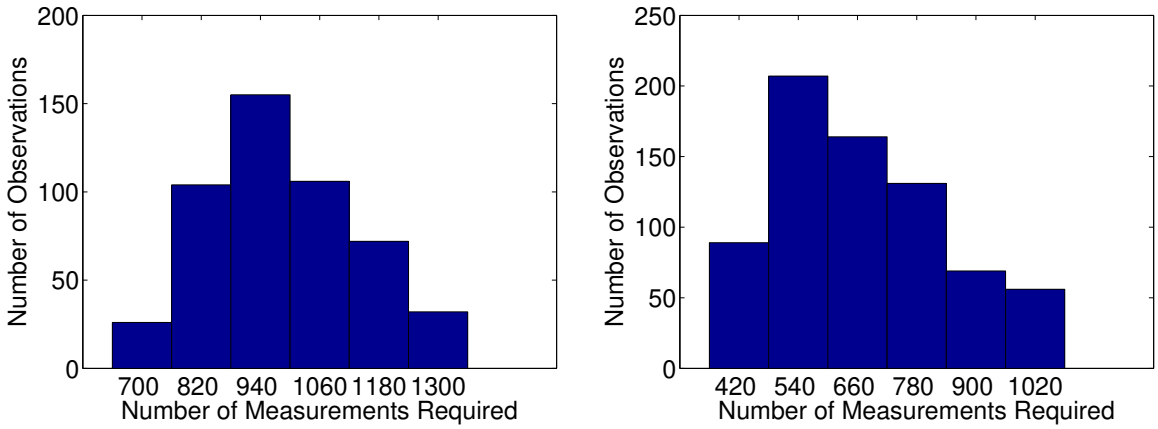

(a) Downtown (b) Suburban

Figure 2.9: Regional MR Class Distributions

Under these class groupings, the input terrain feature distributions used each measurement class grouping are shown in Figure 2.7. In reference to Figure 2.9, we see that the majority of the regional features fall into groups centered around 940 and 540 measurements for the downtown and suburban regions, respectively. Thus, these bins will have a higher weight under the learning framework.

To train each the decision tree classifier, it is important to allow equal training weights per class as much as possible to balance the tree and not over-fit the data. For this, we randomly selected an equal number of observations for resulting in each class to balance the observations per bin. We further divided this set of observations into two separate training and validation observation sets, again being sure to have equal class representation in each set. We then trained the decision tree classifier with the training set and predicted MR classes with the validation set.

Due to the linear relationship between the MR metric and class, it is important to look at not only the prediction accuracy in choosing the correct class, but also the distribution of predicted class offsets (how many classes away from the correct class) when the correct class is not chosen. This is because a lower average offset between the predicted and correct MR class is nearly as important as the accuracy in choosing the correct class. For example, predicting the adjacent class is not as detrimental to the measurement number estimation as predicting multiple classes away.

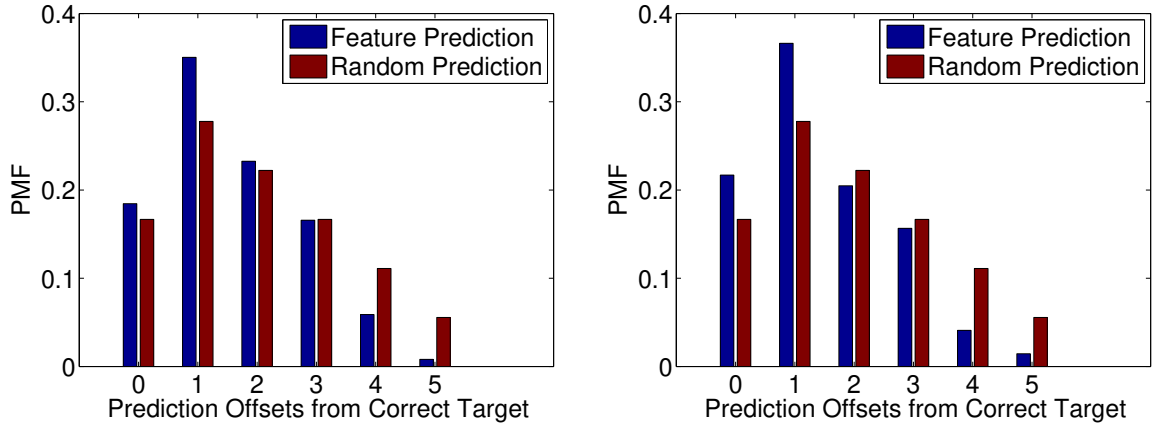

(a) Downtown

(b) Suburban

Figure 2.10: Regional Feature Versus Random Prediction Offsets

Figure 2.10 shows the class prediction offset magnitudes for using feature prediction to choose a class and choosing a random class (according to the frequency of occurrence). This result shows that for both the downtown and suburban region, the average predicted class offset is significantly lower using the feature prediction than choosing a class at random, even if there may not be a very high accuracy in actually predicting the correct class.

2.5.3 Uniform Drive Testing Comparison

Using our trained and validated decision tree classifier, we wanted to compare GeoRIPE to uniform drive testing in two scenarios. First is a dense uniform drive testing scenario, in which measurements are gathered according to the requirements of the subregion (window from Algorithm 1) with the highest geographical variability. More specifically, we experimentally found the number of measurements required in the worst case subregion to meet the accuracy threshold and uniformly take that number of measurements over every subregion. The other is a sparse uniform drive testing scenario, in which measurements are gathered according to the measurement requirements of the subregion with the lowest geographical variability. In this experiment, the goal is to stay under a predetermined path loss exponent standard deviation (corresponding to an accuracy that a network provider would require) while using the lowest amount of measurements. To do this, we divided the regions into several uniformly sized tiles and gathered several orthogonal sets of measurements from each tile according to the sparse, dense, and GeoRIPE predicted number of measurements. For each orthogonal set in each scheme, we calculated the path loss exponent and took the standard deviation over all exponents for each of the three techniques. By doing this, we can compare the path loss exponent accuracy and the number of measurements required for each technique. For this experiment, we trained the GeoRIPE classifier to predict the measurements required for the standard deviation of 0.03 using half of the region, and predicted the number of measurements required for the for the other half. We repeated this experiment for both the downtown and suburban regions, and the results can be seen in Table 2.5.

From this table, we can see that the sparse drive testing does not meet the required standard deviation of below 0.03 that we set at the start of the experiment. The dense drive testing does stay under the standard deviation requirement, using the minimum number of measurements to do so over all areas. Using GeoRIPE, the standard deviation requirement is also met, but it requires 58%

Table 2.5: GeoRIPE Standard Deviation and Measurements Compared to Sparse and Dense Uniform Drive Testing Scenarios

Technique	Region	Average σ	Avg. Meas. per km^2
GeoRIPE	Downtown	0.0286	194
GeoRIPE	Suburban	0.0284	186
Sparse	Downtown	0.0631	40
Sparse	Suburban	0.0440	80
Dense	Downtown	0.0205	400
Dense	Suburban	0.0188	440

fewer measurements than uniform dense drive testing to get all subregions below the threshold. From the GeoRIPE results, the effect of geographical complexity can be clearly seen; a lower standard deviation of path loss exponents is obtained using fewer measurements in the suburban region than the geographically more complex downtown region.

In addition to meeting the accuracy requirements using the least number of measurements, we wanted to evaluate the benefits of using the GeoRIPE framework over a uniform distribution that uses the same number of total measurements. To do this, we used measurements from the GeoRIPE distribution given by Equation 2.7.

$$p_x = \frac{M_x}{\sum_{\hat{x}}^{X} M_{\hat{x}}} \tag{2.7}$$

Here, x is a single section in the set of all tiles X, p_x is the fraction of measurements to be collected in section x, and M_x is the set of predicted MR values of tile x. We collected several orthogonal subsets of measurements in each tile for an increasing number of total measurements in each region and compared the accuracy of the two techniques. For each orthogonal subset in each tile, we calculated a path loss exponent and computed the standard deviation of the path loss exponents in each tile. The standard deviation for all tiles was averaged at each number of total measurements and the results were organized by standard deviation. For selected standard deviation, each of the techniques required a different number of measurements per km^2, resulting in Table 2.6.

From this table, we see that as the threshold for standard deviation is lowered, measurements required increases approximately 10% 'faster' using uniform drive testing. So, while the true value in using the GeoRIPE framework is predicting the number of measurements required over

Table 2.6: GeoRIPE versus Uniform Drive Testing Measurements to Achieve a Fixed σ

σ	GeoRIPE # Meas.	Uniform # Meas.
0.050	64	64
0.045	80	84
0.040	96	108
0.035	128	138
0.030	178	196
0.025	252	276
0.020	400	436

a region to meet a certain path loss exponent accuracy, the normalized GeoRIPE distribution also achieves the desired path loss accuracy with proportionally fewer measurements than the uniform counterpart. This result, however, only analyzes the average path loss over the entire region. Using a similar windowing method previously described, we wanted to see how many measurements it took to bring the standard deviation of the path loss exponent in all the windows to fall below these thresholds. We found that while the number of measurements for the GeoRIPE framework to accomplish this is similar to the numbers in Table 2.6, uniform drive testing required an average of 20% more measurements than the listed numbers. The biggest difference in this experiment is alluded to in Table 2.5, wherein GeoRIPE requires 58% fewer measurements to go below the standard deviation threshold than uniform drive testing.

2.6 Related Work

Measurement Collection Approaches. Due to the low cost of crowdsourcing from smartphones, the technique has been used by many other groups to collect data about wireless networks. In a study by Huang *et al.* [13], LTE performance data was collected by creating an Android application named 4GTest. This application gained 3,000 users during 2 months of data collection and collected data that focused on media streaming by mobile clients. With this data, [13] was able to show that with the download speed increase seen with LTE networks, the traffic bottleneck shifted from the network to the processing power of the mobile devices. In [14], an Android application was again used to capture network speed data. This study focused on comparing the speeds of 802.11 networks to the speeds of LTE networks in major cities around the globe. Neidhardt *et al.* used a crowdsourced infrastructure to provide an open source and more accurate base station

location and coverage estimation system [15]. While they had promising results on the base station localization aspect, they concede that cellular coverage estimation was lacking with their purely crowdsourced measurements, especially in urban environments with diverse terrain features. Our work focuses on the minimum measurement requirements according to different geographical features of a given region.

Measurement-Driven Path Loss Evaluation. There have been several measurement studies that strive to more accurately characterize path loss in specific region types. Hata *et al.* [6] and Okumura *et al.* [16] specifically focus on accurate characterization in urban regions. Using measurements gathered by [16] in Tokyo, Japan, Hata *et al.* empirically derived a path loss prediction formula with correction factors for various region types such as large-city urban, small-city urban, suburban, and open areas. Additionally, the Hata model considers base station transmitter height. Similar to the path loss prediction curves found by [16] in Japan, Allsebrook *et al.* [17] evaluated path loss prediction curves for three British cities: Birmingham, Bath, and Bradford. Akimoto *et al.* [18] derived a model based on gathered measurements in a rural area using the 2 and 5 GHz bands. Similarly, [19] studied measurements collected in a suburban neighborhood at 5.7 GHz as did [20] with measurements taken in Istanbul in the GSM-900 band. More recently, Robinson *et al.* sought to minimize the number of measurements necessary to accurately characterize mesh node coverage in the TFA network in Houston [10]. Their work uses an online push-pull measurement gathering approach, taking very few active measurements on an existing deployment based on terrain features in the area. Additionally, Sayrac *et al.* [21] and several others [22–24] try to reduce the number of drive testing measurements required for coverage evaluation via Baysian kriging, showing how their techniques can be used to detect coverage holes. However, their analysis replies on the spatial correlation between the measurements themselves to detect coverage holes from existing transmission infrastructure. In contrast, our approach aims to analyze geographical features of a region and predict the number of measurements required to obtain an accurate estimate of path loss throughout, including from transmission sources that do not yet exist, by tying the measurement requirements to the terrain itself.

2.7 Conclusion

In this chapter, we built the GeoRIPE framework which predicts the minimum number of in-field measurements required to accurately characterize the path loss of a region according to that region's geographical features. To find if such measurements would be sufficient for a given area, we gathered millions of signal strength measurements along with geographical feature ratios in both a downtown and suburban region. Using this data, we correlated several distinct geographical features with different metrics for path loss evaluation complexity. We found that, together, these features are correlated with the number of measurements required to achieve a fixed path loss accuracy. We also evaluated the merit of using area bounded path loss metrics. By abstracting propagation loss parameters away from specific paths and binding them to a specific area, we are able to evaluate path loss for arbitrary paths through the area. We found that the size of the individual path loss evaluation areas should be selected based on the complexity of the terrain features residing in each area. In general, the more complex the area, the smaller the evaluation area should be. Finally, to validate our work, we compared drive testing using our GeoRIPE framework to uniform drive testing in each region. We found that our technique, as opposed to spatially uniform drive testing, required fewer measurements to achieve a similar path loss characterization accuracy.

Chapter 3

Enabling a "Use-or-Share" Framework for PAL–GAA Sharing in CBRS Networks via Reinforcement Learning

3.1 Introduction

Additional spectrum availability and increased efficiency in the use of existing resources are needed to accommodate the rapidly increasing density and subsequent data demands of wireless devices around the world. The United States government recognized this need in 2010 and put into motion what would ultimately become the Citizens Broadband Radio Service (CBRS) [2].

In the standard CBRS architecture, there is a three-tiered system managed by a dynamic database called the Spectrum Access System (SAS), illustrated in Fig. 3.1. The top tier consists of incumbent users, the second tier consists of Priority Access Licenses (PALs), and the bottom tier is for General Authorized Access (GAA). Many users will likely operate in a pseudo-unlicensed fashion on the GAA tier, and in competitive markets, some carriers may choose to purchase a PAL license to ensure a minimum QoS. For more background on the CBRS band, see the Appendix.

At the same time that the federal government is opening up new bands for shared use, there is increasing congestion on the unlicensed bands. For example, 802.11 Wi-Fi and Bluetooth devices densely occupy the unlicensed 2.4 GHz and 5 GHz industrial, scientific and medical (ISM) radio bands. Cell providers are increasingly interested in using free, unlicensed spectrum to supplement their licensed networks, exacerbating this crowding. The idea of using unlicensed bands to supplement licensed networks has been pushed to multiple standards such as 3rd Generation Partnership Project (3GPP) Licensed Assisted Access (LAA) [25], LTE-U [26], and MulteFire [27].

It is clear that spectrum sharing will become more prevalent in future medium access policies. The cornerstone of these policies is the requirement for a way to manage the sharing. Spectrum sharing is traditionally done through contention-based protocols such as a listen-before-talk (LBT) scheme like Carrier-Sense Multiple Access (CSMA) for 802.11 devices. In these schemes, the

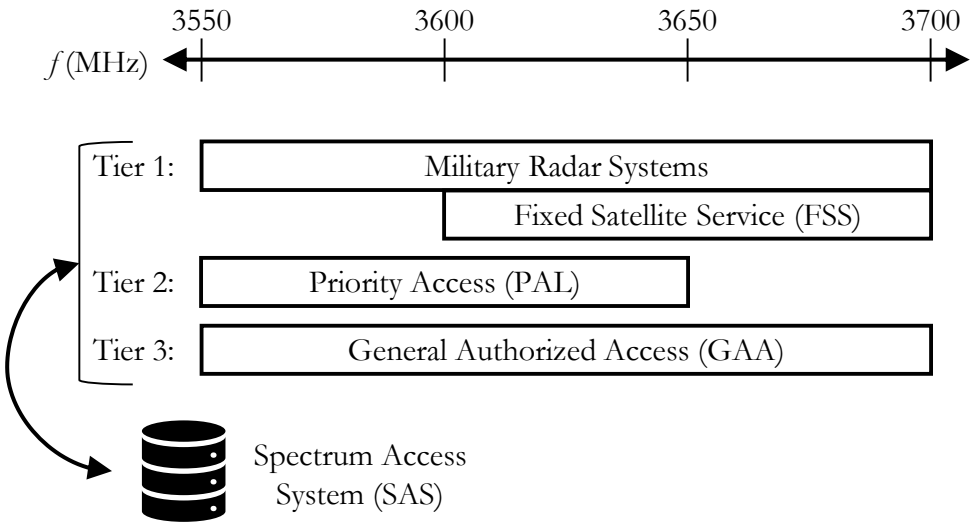

Figure 3.1: Structure of the CBRS band. Three tiers of different priorities share the band, and a central database called the SAS dynamically manages the users. PAL users are licensed through auction while GAA users may use any spectrum not reserved for a higher tier [2].

channel must be determined to be unoccupied for a certain amount of time before it is accessed.

In the current release of the CBRS standard, there is no formal mechanism for GAA users to access PAL allocated channels. Instead, it is framed such that "use-or-share" [28] policies should be adopted to prevent spectrum warehousing by allowing PAL–GAA sharing but without any explanation for how it should be done. We propose a method to address this with the following contributions.

3.1.1 Contributions

In this paper, we present and compare two LTE-based LBT schemes for use in the CBRS infrastructure to allow the GAA users to access licensed PAL spectrum when available. Specifically, we analyze the trade-off between GAA user gain and PAL user interference through simulations where we aggregate results over many random user topologies and packet arrival rates. Moreover, we test our LBT schemes on a custom testbed with multiple software-defined radios and a real-time signal analyzer to show the feasibility on real hardware. In doing so, we make the following five contributions:

1. We design and evaluate two LBT schemes to be used in CBRS networks for PAL–GAA

spectrum sharing.

2. We show that while one scheme has higher performance, both schemes significantly improve GAA UPT with a minor decrease in PAL UPT.

3. We find that the decreased PAL UPT is a function of PAL traffic load and problematic network topologies between PAL–GAA users.

4. To reduce the negative consequences of spectrum sharing on the PAL, we formulate a novel Q-learning algorithm that adjusts GAA opportunistic access via learning an improved energy-detection threshold (EDT) for carrier sensing.

5. By using average and differential PAL buffer occupancy as the environmental observations, we find the detriment to PAL UPT from spectrum sharing can be greatly reduced.

3.1.2 Related Work

Many works in the cognitive radio and dynamic spectrum access literature deal with sharing between a higher priority Primary Node (PN) and a lower priority Secondary Node (SN). Some recent investigations also consider this while adopting machine learning with promising results. In [29], the authors use Q-learning for dynamically choosing the channel for cells as opposed to static assignments. In [30], a decentralized Q-learning scheme is used for reducing the interference seen by 802.22 PN users. Although Q-learning has been considered for power allocations and channel assignments, it has not been used, to the best of the authors' knowledge, for adapting a dynamic EDT for SNs in a shared spectrum environment.

Overall, CBRS is still an emerging standard, but we can learn from similar experiences on other bands. When considering the coexistence of GAA users, there are many similarities to unlicensed bands which have been studied extensively. Notably, there is substantial work that has been done for the coexistence of LTE and Wi-Fi nodes in unlicensed bands for LAA [31–34] with [35] using reinforcement learning to alter the duty cycle of the LTE nodes and [36] using Q-learning to adjust the channel occupancy time of the LAA nodes.

Given that CBRS is a new band, the case of LTE nodes coexisting with another LTE node that has a higher priority is a new area. Much of the previously mentioned LTE–LAA works are not

sufficient as the higher-tiered node does not necessarily engage in LBT. There have not yet been extensive studies for CBRS-specific performance improvements; so far, there are only initial proof of concept demonstrations reported. In [37], a field trial of CBRS devices (CBSDs) working with the SAS was shown where they suggest improvements to the SAS protocol based on their results, and recently Verizon has deployed a CBRS network in Florida [38].

The rest of the paper is organized as follows. We present our new listen-before-talk schemes for sharing the spectrum amongst LTE devices on the CBRS band in Section II. In Section III, we give an overview of the reinforcement learning scheme we apply. We additionally provide Q-learning simulation results in Section III and then conclude the paper in Section IV. Following the conclusion is an appendix devoted to the state of the CBRS band.

3.2 Listen Before Talk

Although the CBRS band allows for spectrum sharing between multiple tiers and encourages such “use-or-share” frameworks, there is currently no mechanism to facilitate this sharing, a deficiency that we address here.

Throughout the analysis in the remainder of the paper, we adopt the conventional spectrum sharing terminology of primary and secondary nodes. In most scenarios, the PAL acts as a PN which will transmit whenever it has traffic for a connected user in what is called “On/Off" mode. The GAA acts as an SN which will need to contend for access to the PAL’s channel via LBT. However, we also present a mutual sharing scenario in Section 3.2.3 where two PAL operators may use each other’s spectrum as GAA users. In this case, an operator would be a PN on its licensed PAL spectrum and an SN on other operators’ spectrum.

When it comes to opportunistic random access, LBT schemes are a proven method that can be used to allow GAA users to access PAL spectrum opportunistically. Wi-Fi, one of the most popular random access schemes available, has used LBT in the form of CSMA for sharing the spectrum between multiple users to great success. A version of LBT is essential in any shared-spectrum environment and is legally required in the European Union and Japan for operation on an unlicensed band. Ideally, under such a scheme, the secondary GAA nodes (SNs) gain additional spectrum resources, increasing throughput, while the primary PAL nodes (PNs) are unaffected. While it is impossible to achieve such perfect coordination due to the inability for secondary nodes

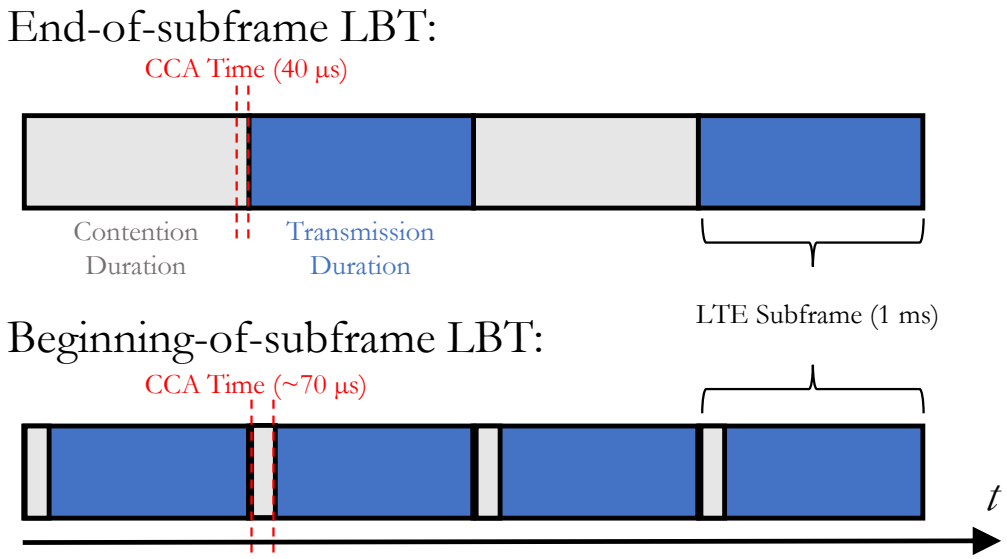

Figure 3.2: Comparison of two proposed LBT schemes for CBRS. In End-of-subframe LBT, the SN uses the last 40 μs of a subframe for a CCA, leading to a 50% duty cycle but relatively easier implementation. In Beginning-of-subframe LBT, the 1st symbol of a subframe acts as the CCA leading to a higher, 93% duty cycle at the cost of a more complex implementation.

to predict future PN packet arrival, it is possible for LBT to adequately allow LTE nodes to coexist with marginal harm to the PN.

LBT has already been adopted in some LTE standards, for example, 3GPP's LAA specification. The LBT scheme used in LAA is as follows. Whenever a device needs to transmit, it needs an initial clear-channel assessment (CCA). It must sense that the channel is idle for at least 34 μs. If so, it can transmit for the length of one transmit opportunity (TXOP). If there is additional traffic to send, an exponential random backoff mechanism is used, similar to Wi-Fi [33].

However, in the context of CBRS Alliance LTE devices, it is possible to further tailor the LBT scheme specifically for LTE devices. Although there could be many ways for performing LBT, we develop and compare the performance of two schemes that seem to be the natural choices: sensing at the end or sensing at the beginning of a subframe. These are shown in Fig. 3.2, and each is evaluated below. Similar schemes have also been considered for LTE/Wi-Fi coexistence in [34]. For the developed LBT schemes, we consider only the downlink, assuming that devices are operating similarly to LAA in Release 13 of LTE where the CBRS carrier is considered as a supplemental downlink carrier [33]. For more information on the typical structure of an LTE

frame, see [39] and the references therein.

3.2.1 End-of-subframe LBT

In end-of-subframe LBT, an entire subframe functions as a contention window for the SN with the last 40 μs in a subframe used for the CCA. If the channel is determined to be idle, the SN uses the next subframe for transmission. Using this scheme has an advantage of not altering the structure of a subframe. However, it results in, at most, a 50% transmission duty cycle. Moreover, the scheme could measure the channel to be idle during the contention window only for the primary node to start transmitting on the next subframe, leading to a collision. As this is a tiered access system, collisions with the PN need to be avoided at all costs.

3.2.2 Beginning-of-subframe LBT

In beginning-of-subframe LBT, the potential transmitter always senses in the first symbol of every subframe. If successful, the SN transmits in the remaining 13 symbols of the subframe. This LBT scheme has the advantage of sensing at the time that a PN would start a transmission, reducing the likelihood of a collision, given synchronization of subframe boundaries between PN and SN. Moreover, this scheme sacrifices a single symbol out of each subframe, a 93% duty cycle. However, LTE uses the first symbol for a control channel, so this scheme may require altering the subframe structure, though this omission of the control channel may be inconsequential in cases where SN's use cross-carrier scheduling.

3.2.3 LBT Scheme Comparison

There are apparent differences between the LBT schemes by construction. Given that the best-case duty cycle for end-of-subframe LBT and beginning-of-subframe LBT are 50% and 93%, respectively, beginning-of-subframe LBT is preferable. However, there is a tradeoff between performance and implementation complexity between the two schemes. With end-of-subframe LBT, we do not require any modification to the subframe structure, but beginning-of-subframe LBT would require proposing a change to the relevant standards. This change to the standards could be done through the CBRS Alliance, but since it may modify the subframe this change may also require 3GPP efforts, resulting in higher implementation complexity. Such a proposal would be a lengthy effort. Ultimately, after further study, there would need to be a mechanism to allow the

control channel information to be sent on any symbol, or there would need to be a requirement of cross-carrier scheduling for LBT Scheme 2. However, standardization efforts are beyond the scope of the paper. Instead, we merely present our interpretation of the most likely scenario for the benefits that these LBT schemes could have for coexistence purposes.

To evaluate the performance of these LBT schemes in a CBRS-like framework, we first simulate their operation in various scenarios. In Fig. 3.3, we show each scenario considered to help understand the effect of the LBT scheme on both the PN and SN. For a baseline, we consider in Fig. 3.3a, Scenario 1, where two operators are operating on their carriers without any sharing. They operate in an "On/Off" mode without contention. In Fig. 3.3b, Scenario 2, we test the result when both operators engage in mutual sharing onto each other's carrier. In Fig. 3.3c, Scenario 3, we consider the case that a single operator is on its carrier and performs LBT on another carrier that is entirely unoccupied. This case represents an upper bound on gains for an SN. In Fig. 3.3d, Scenario 4, sharing is performed on a single component carrier to see the realistic gains for operator 2 (Op. 2) when sharing and the effect it has on the PN, operator 1 (Op. 1). Due to the nature of the band, PAL–GAA sharing will likely require an agreement between two operators on a secondary market. Considering this case, more complicated topologies with many operators and carriers are not likely to be realistic, and hence, the simulation scenarios in Fig. 3.3 give a realistic insight to potential sharing scenarios.

For the sake of comparing the performance of each LBT scheme, we considered the scenarios with no sharing (Scenario 1) and sharing with a PN (Scenario 4) from Figs. 3.3a and 3.3d as before and after cases. We report the change in the user-perceived throughput (UPT) for Op. 1 and Op. 2. In this figure, UPT is given by

$$UPT = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P_{total}}\left[\sum_{j=1}^{P_{served}}\frac{M \cdot r_{ij}}{t_{ij}} + \frac{b_i}{t_{serving,i}}\right] \tag{3.1}$$

where N is the number of UEs served by the eNB, and i indexes the UEs. P_{total} is the total number of packets, elaborated by $P_{total} = P_{served} + P_{serving}$, where P_{served} and $P_{serving}$ are the number of packets served and being served, respectively. M is the number of bits per packet, r_{ij} is the ratio of successfully transmitted bits over all bits in the packet to UE i for packet j, and t_{ij} is the time taken to send the same packet. b_i is the number of bits sent to UE i as a partial packet still in flight,

and $t_{serving,i}$ is the time spent by the packet.

We simulated our spectrum-sharing scheme using MATLAB by reusing the 3GPP LAA evaluation assumptions for an indoor scenario [3]. Fig. 3.4 shows this topology, and the rest of the simulation settings are as follows:

- Two operators with four small cells each in a single floor building (Fig. 3.4)
- 18 dBm TX power
- 10 randomly distributed UEs per operator
- -72 dBm EDT
- 20 MHz system bandwidth
- 10 independent simulations with random positioning of the UEs
- 20,000 subframes per simulation.

Fig. 3.5 shows the simulation results. For each test, we show the mean, 5th percentile, median, and 95th percentile to illustrate the variance across all simulations. In Fig. 3.5a, we see a maximum increase in UPT for the SN of 40%. However, for the same case, there can be a 10% reduction in UPT for the PN. In Fig. 3.5b, we see a nearly 80% increase in performance for the SN with a similar drop in performance for the PN. Overall, beginning-of-subframe LBT performed significantly better for both the PN and SN. So, in the next subsection, we select this scheme for use in additional simulations to determine the possible spectrum-sharing gain. Although throughout the paper we show

3.2.4 Simulations with Static EDT

Fig. 3.6 shows the results for simulating beginning-of-subframe LBT across different spectrum sharing scenarios for two different traffic arrival rates to see the effect on UPT. In Fig. 3.6a, the average traffic arrival rate was 0.5 MB/s for PNs and SNs. The first cluster of results is the Scenario 1 baseline from Fig. 3.3a, where both operators are on their carriers without sharing the spectrum. The second cluster of results shows Scenario 2 from Fig. 3.3b, where each operator mutually shares

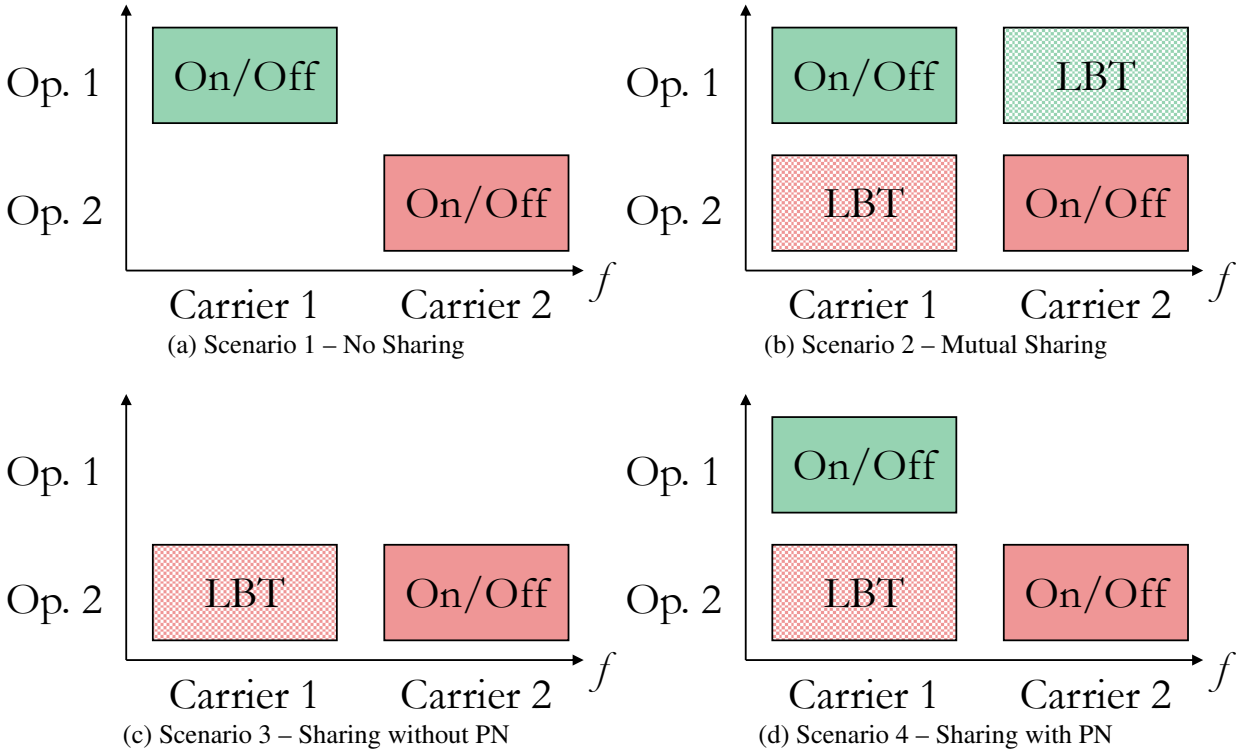

(a) Scenario 1 – No Sharing (b) Scenario 2 – Mutual Sharing

(c) Scenario 3 – Sharing without PN (d) Scenario 4 – Sharing with PN

Figure 3.3: Simulation scenarios considered throughout the paper. Simulations using these scenarios give us a performance baseline, a possible gain when two carriers mutually share their spectrum, an approximate best case SN gain when sharing, and an expected PN loss.

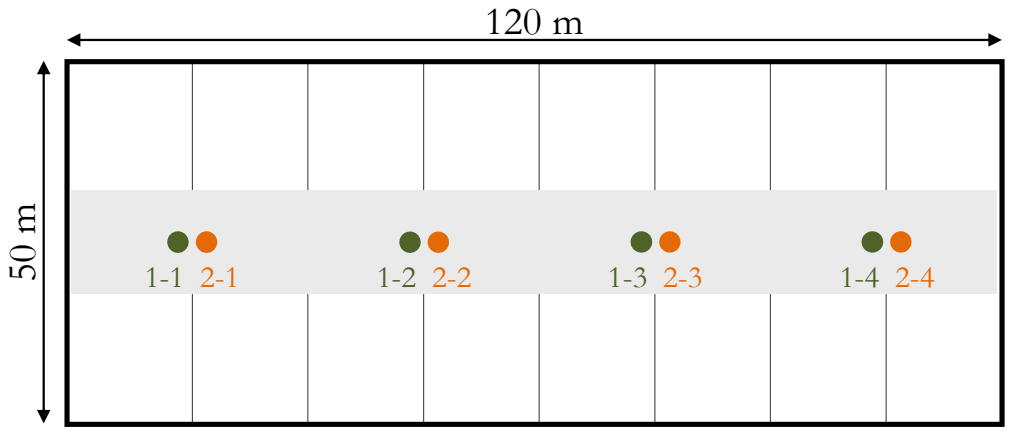

Figure 3.4: 3GPP indoor scenario for LAA coexistence evaluations with two operators and four nodes per operator [3]. This standard scenario provides an industry agreed-upon simulation scenario for our LBT evaluations.

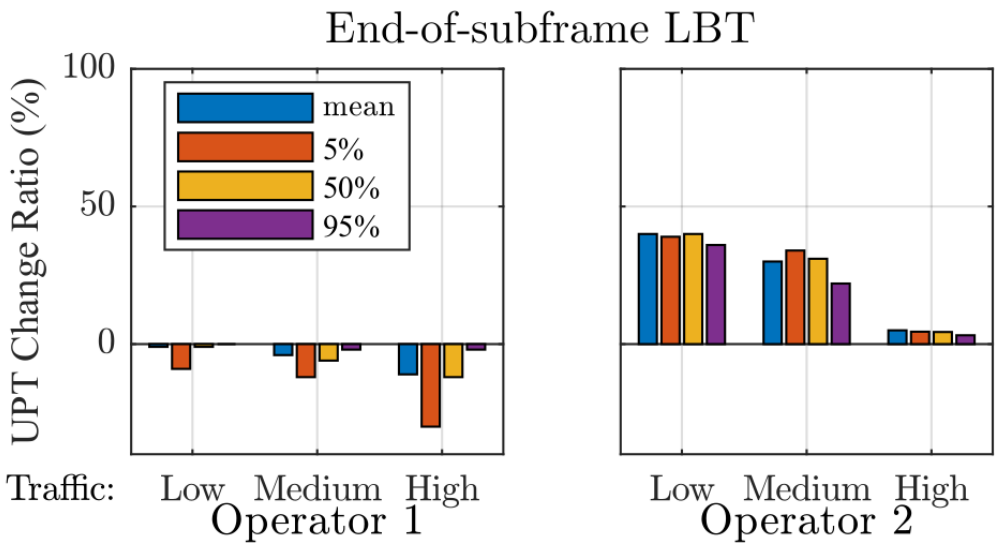

(a) UPT change when secondary node uses end-of-subframe LBT.

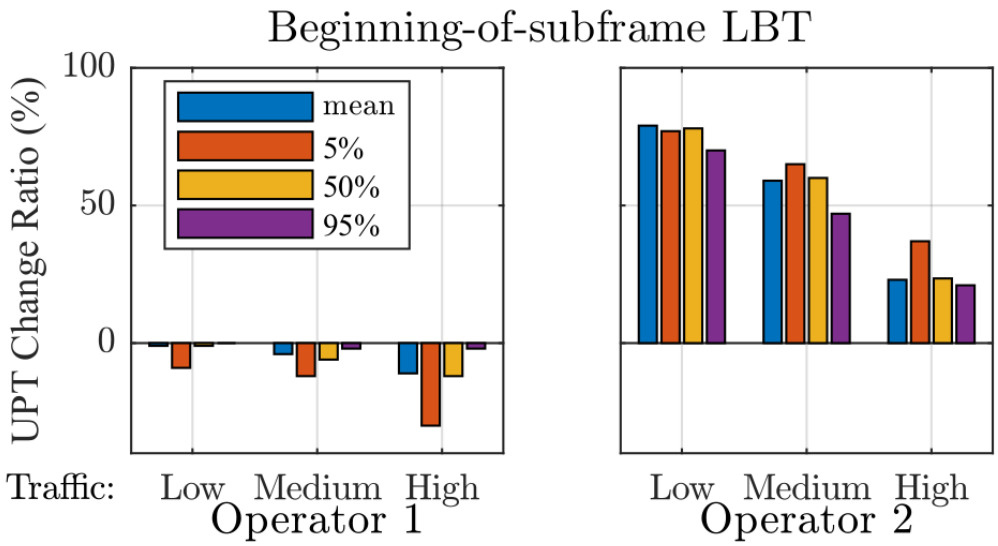

(b) UPT change when secondary node uses beginning-of-subframe LBT

Figure 3.5: Performance of each LBT scheme for different traffic loads. In the results, a low-traffic load corresponds to a traffic arrival rate per UE of 0.5MB/s, the medium is 0.75 MB/s, and the high is 1.05 MB/s.

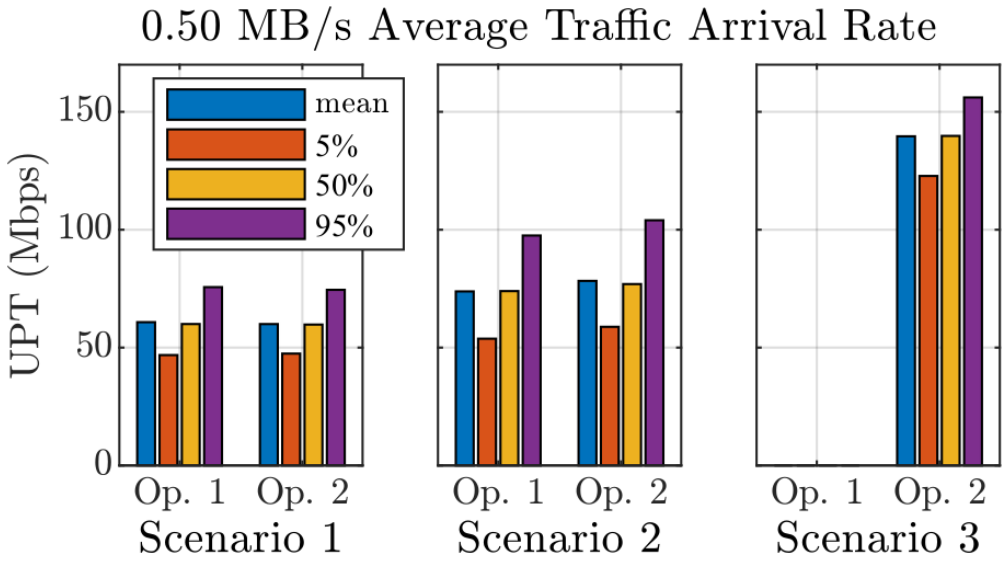

(a) UPT for spectrum-sharing scenarios when both operators use a relatively slow average traffic arrival rate of 0.50 MB/s.

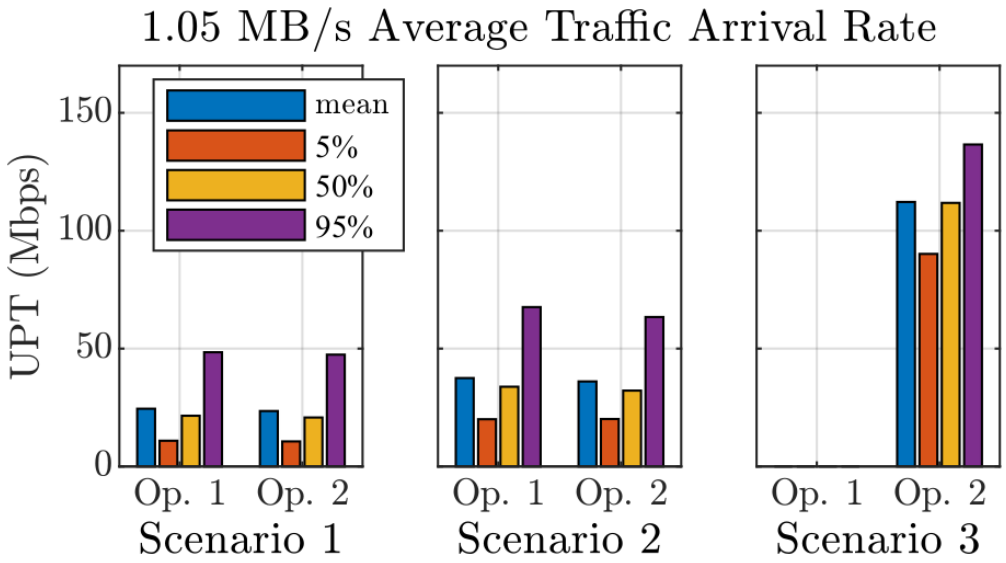

(b) UPT for spectrum-sharing scenarios when both operators use a relatively fast average traffic arrival rate of 1.05 MB/s, similar to each user streaming a 1080p video.

Figure 3.6: Performance of LBT for two different average traffic arrival rates.

its primary spectrum with the other operator. Here, we can see that each operator experiences an increase in the mean UPT by about 25%.

The third cluster shows Scenario 3 from Fig. 3.3c, which provides a (coarse) upper bound on the maximum achievable spectrum sharing gain. Here, we can see that an operator can achieve a maximum of 133% gain when adopting a spectrum sharing scheme.

The simulation is repeated for the case of a higher-traffic arrival rate of 1.05 MB/s in Fig. 3.6b, which is similar to each user streaming a 1080p video using the H.264 codec. Similar to the results for the slower traffic rate, when each operator engages in a mutual sharing as in Scenario 2, each operator sees an improvement in UPT. For the higher traffic case, the gains are doubled with approximately 50% increase in UPT for both operators.

We then use beginning-of-subframe LBT and consider the performance for various EDTs. In Fig. 3.7a, we show the results when we consider different EDTs for Scenario 2 from Fig. 3.3b. Here, we see that there is an "optimal" EDT around -52 dBm. These results highlight that for different scenarios, there may be different "ideal" EDTs.

In Fig. 3.7b, we plot the UPT vs. EDT for Scenario 4 from Fig. 3.3d. In this figure, we can see the effect of a higher EDT. Here, a higher EDT at Op. 2 implies more frequency of channel access at the expense of increased downlink interference at Op. 1. As the EDT increases, the UPT of Op. 1 decreases, and the UPT of Op. 2 increases.

It is worth noting that the UPT decreases for the PN are significantly smaller in the low-traffic load case and more significant in the high-traffic load case for both schemes. Ideally, in situations where the PN has a high-traffic load, the SN would behave more passively when on the PN's carrier. In Section 3.3, we will explore the use of machine learning in adjusting SN EDT to improve LBT performance in hidden and traffic-heavy node scenarios. Using our algorithm, we show that scenario-specific, poor-LBT performance can be significantly reduced.

3.2.5 Multiple Secondary Nodes

In many cases where PAL–GAA sharing may occur, it is likely that a sharing arrangement between the two will need to be formed on a secondary market [28]. In this case, there would be only two operators, and the operator on the GAA tier would perform its cell planning so that it would not experience significant self interference.

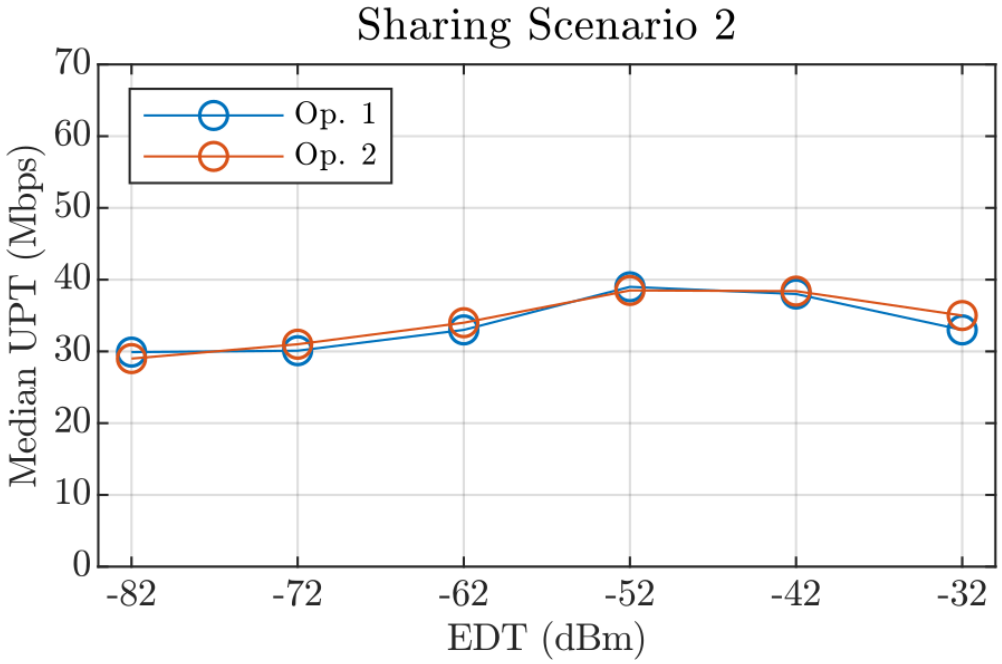

(a) Mutual sharing on two carriers, Scenario 2.

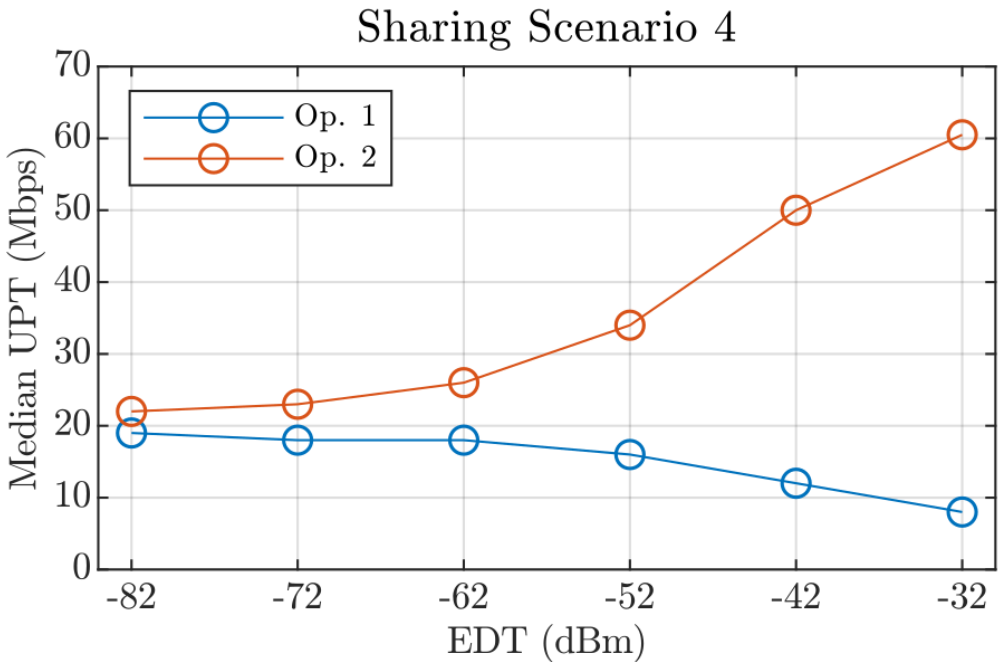

(b) Sharing on one carrier, Scenario 4.

Figure 3.7: Median UPT vs. EDT for tests with a 20 MHz system bandwidth and a traffic arrival rate of 1.05 MB/s.

In the case where multiple GAA tier operators attempt to use the same PAL spectrum, the LBT scheme could be modified to accommodate this. One simple way would be to apply random back-off in the unit of subframes whenever an SN experiences a collision. Another possibility is to enforce a Maximum Channel Occupancy Time (MCOT) to limit the number of subframes an SN can transmit on consecutively.

3.2.6 Shared-Spectrum Testbed

To further evaluate the LBT schemes outlined so far, we developed a shared-spectrum testbed shown in Fig. 3.8. The testbed consists of the following:

- 4 USRP SDRs with the possibility of including more for more extensive tests with many nodes and UEs.
- Ideal ethernet backhaul via python/UDP for statistics
- Real-time spectrum analyzer.

The nodes of the testbed can be arranged to emulate various topologies such as hidden/exposed nodes with additional possibilities of including mobility, allowing us to see how well the LBT schemes behave under real channel conditions where there may be a rich, multipath environment that changes quickly.

We modified the National Instruments LTE Application Framework to implement our LBT schemes. Fig. 3.9 shows an example result from the real-time signal analyzer spectrogram. Here, we use end-of-subframe LBT where the SN uses a subframe for contention and a subframe for transmission. We show Scenario 2, where each operator can engage in spectrum sharing. In this case, Op. 2 is under heavy load while Op. 1 is not. Op. 2 augments its services by aggregating onto Op. 1's spectrum while Op. 1 has no traffic. For this demo, we restrict each operator to occupy only half of each 20-MHz carrier so that we can easily distinguish between the operators on the spectrogram. In this figure, red represents high measured power for the corresponding time/frequency unit, and green represents low power. We can see that Op. 2 is using Carrier 1 with a 50% duty cycle as it is the SN on this carrier and must perform end-of-subframe LBT. We can also see that when Op. 1 begins using its primary spectrum, Op. 2 detects the presence of the PN and waits for the carrier to become available again before transmitting.

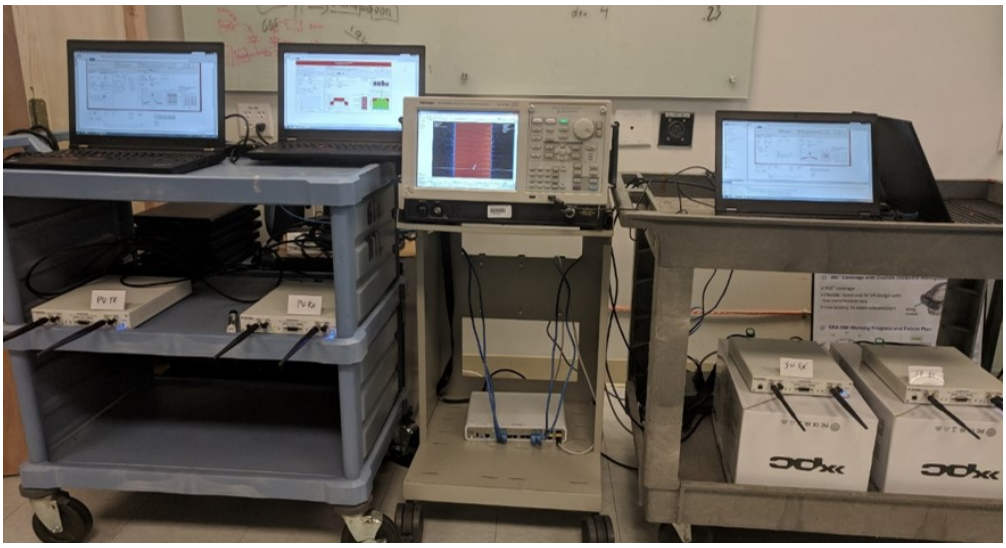

Figure 3.8: Photograph of the shared-spectrum testbed. Four USRPs connected to host PCs running LabVIEW Communications with a real-time signal analyzer.

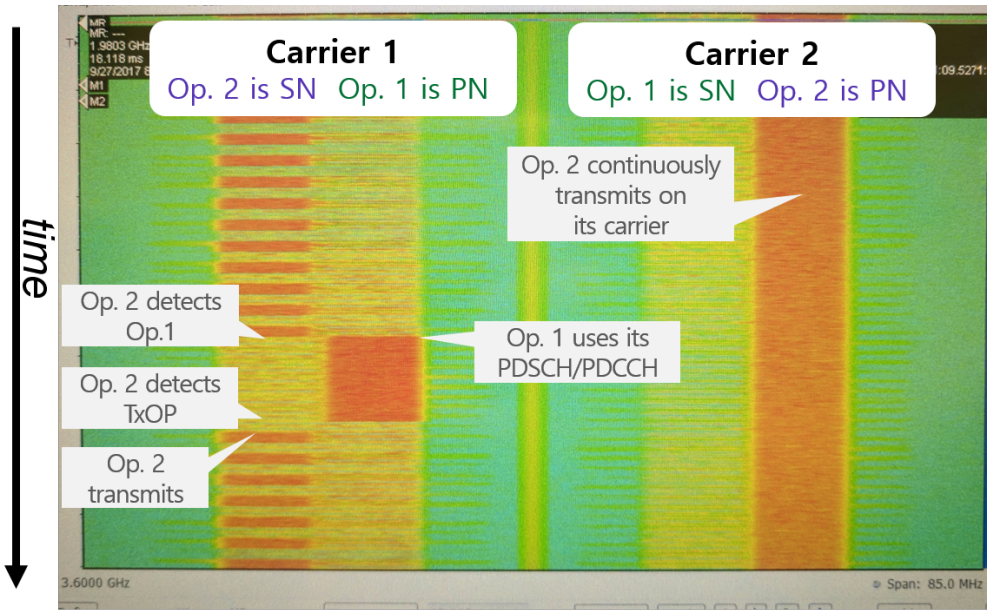

Figure 3.9: Real-time signal analyzer spectrogram for end-of-subframe LBT, Scenario 2 from the testbed. Each operator is restricted to half of the full carrier to allow for visual distinction on the spectrum analyzer — Op. 2 transmits on the outer resource blocks, and Op. 1 transmits on the inner resource blocks. Here, Op. 2 has high traffic while Op. 1 has low traffic. On Carrier 2 where Op. 2 is the PN, Op. 2 continuously transmits as indicated by the orange color representing higher measured power. It can aggregate onto Carrier 1 while Op. 1 has no traffic. Whenever Op. 1 does transmit, Op. 2 can sense and avoid until the channel is free again.

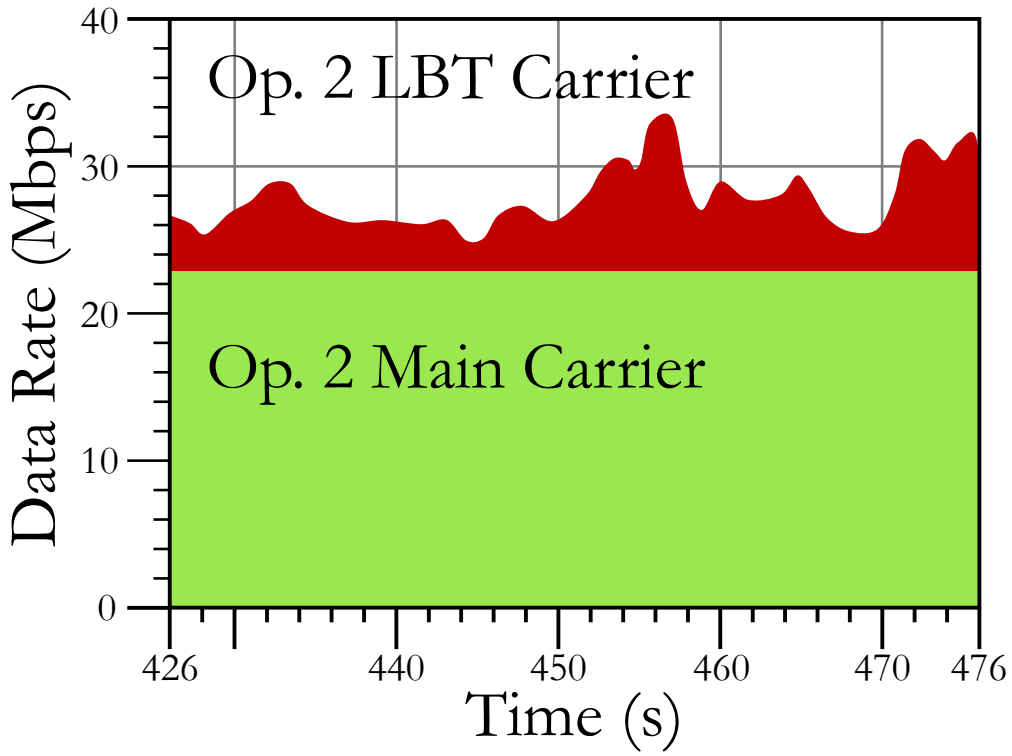

Figure 3.10: Downlink throughput result from the testbed. Here, the green represents the throughput achieved on Op. 2's main carrier where it does not need to perform LBT. The red represents the throughput on its secondary carrier where it performs end-of-subframe LBT to avoid collisions with Op. 1. Periodically, Op. 2 is able to get an additional 50% throughput by utilizing the other carrier opportunistically.

Fig. 3.10, which shows an example output from the GUI corresponding to the custom receiver design, shows the downlink performance for this example. Here, Op. 2 is under high traffic. It fully utilizes its main carrier to achieve approximately 22 Mbps UPT. To enhance performance, it performs end-of-subframe LBT on Op. 1's carrier. Throughout the experiment, Op. 2 is able to aggregate opportunistically onto Op. 1's primary spectrum. At times where Op. 1 has nearly no traffic, Op. 2 can get an additional 50% throughput.

In this testbed, the SN synchronizes with the timing of the PN of a channel. Using the existing LTE synchronization signals, the SN detects the subframe boundaries, measures the energy in the channel at the appropriate time, and then if the SN determines the PN not to be transmitting, the SN transmits to its users during the available TXOP. This result highlights the feasibility for the SN to sync to a PN for performing LBT in real-time. Moreover, Fig. 3.10 highlights what the throughput gain for the SN may realistically be as the PN transmits with a low traffic-arrival rate.

In this testbed, the SN synchronizes with the timing of the PN of a channel. The SN detects the primary synchronization signal (PSS) of the PN every 5 ms. It then uses this to estimate the subframe boundaries so that it can measure the energy in the channel at the appropriate time. If the SN determines the PN not to be transmitting, the SN transmits to its users during the available TXOP. This result highlights the feasibility for the SN to sync to a PN for performing LBT in real-time. Moreover, Fig. 3.10 highlights what the throughput gain for the SN may realistically be as the PN transmits with a low traffic-arrival rate.

3.3 Reinforcement Learning

While LBT schemes have been successfully implemented in 802.11 standards to great effect, certain network situations can result in poor performance. While other applications of LBT may have ways of reducing network problems in topologies that include hidden and exposed terminals, shown in Fig. 3.11 [40], such as the collision avoidance in CSMA, similar schemes are not applicable when applying LBT to the CBRS tiered architecture as the primary node does not necessarily engage in LBT behavior. In this section, we mitigate this issue by proposing a novel reinforcement Q-learning technique to adapt an energy-detection threshold (EDT) for secondary nodes in a shared spectrum environment. We show that by using machine learning, we can increase both SN and PN gains over LBT schemes that use a static EDT.

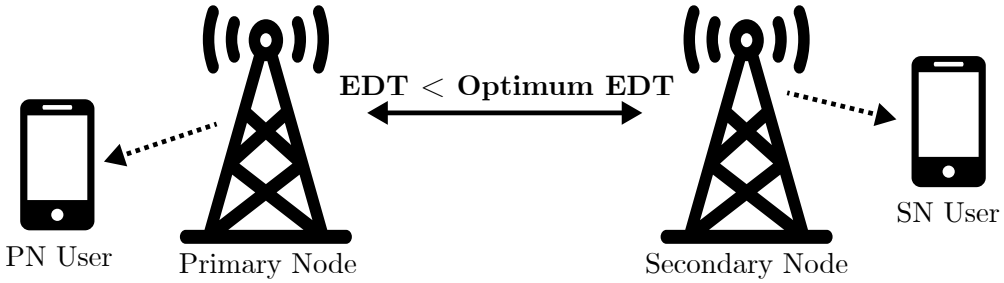

(a) Exposed node case with a Primary Node (PN) inside the sensing range of a Secondary Node (SN). Here, the energy detection threshold (EDT) is too sensitive, so whenever the PN transmits, the SN defers to avoid possible collisions with the PN's users. However, given the positions of the users, the SN could transmit without interference. This wasted opportunity could be remedied with a higher EDT.

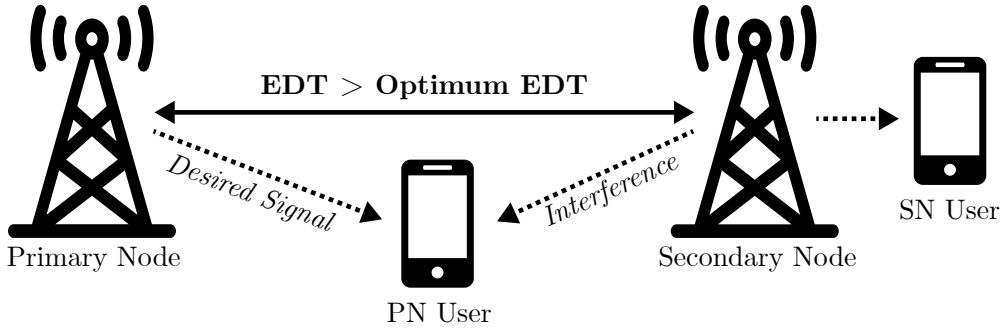

(b) Hidden terminal case where a PN user is impacted by SN interference, despite PN being outside of SN sensing range. This collision could be avoided with a lower EDT.

Figure 3.11: Example primary node (PN)–secondary node (SN) topologies where an adaptive EDT could benefit the network.

3.3.1 Reinforcement Learning Primer

We first review the general reinforcement-learning strategy. Roughly speaking, reinforcement learning addresses the general problem of learning from interactions to achieve a goal [41]. The learner and decision maker is called the agent. Everything outside the agent that the agent interacts with is called the environment.

Agents interact with the environment via actions. For each action, α, that the agent executes, it influences the state of the environment and receives an evaluative feedback, or reward, r. This reward is used to learn/adapt its subsequent actions, should it encounter the same state in a subsequent time slot. We define periodic time intervals, $t = 0, T, 2T, ...$, in which each agent represents its observation, o, of the surrounding environment at time t as a state $s \in S$, where S designates a finite set of environmental states. In summary, at each step t:

The agent:

- Executes action α_t
- Receives observation α_t of s_t
- Receives reward r_t

The environment:

- Receives action α_t
- Emits observation o_{t+1} of s_{t+1}
- Emits scalar reward r_{t+1}

At each time step, the agent implements a mapping from states to probabilities of selecting each possible action. This mapping is called the agent's policy π_t, where $\pi_t(s, \alpha)$ is the probability that $\alpha_t = \alpha$ if $s_t = s$. Reinforcement learning methods specify how the agent changes its policy as a result of its experience. The agent's goal is to maximize the total amount of reward it receives over the long term.

One popular reinforcement-learning algorithm is Q-learning [41]. This model-free learning strategy can be used to learn an optimal decision policy for any Markov decision process. We adopt

Q-learning with the objective to minimize interference at an incumbent or PN due to spectrum sharing with an opportunistic SN. In this scenario, each SN acts as an agent adapting its action in response to the reward obtained for its previous action.

3.3.2 Target Improvement Areas

Our objective is to use reinforcement learning to assist the SN in harvesting unused bandwidth from the PN in an efficient fashion. Specifically, we leverage Q-learning to dynamically adjust the SN's EDT to maximize network UPT while subsequently minimizing the impact on PN UPT. We identify two scenarios in which EDT adjustment can mitigate poor LBT performance:

Hidden/Exposed Terminals In hidden or exposed terminal topologies illustrated in Fig. 3.11 an adaptive EDT can benefit network performance. For example, in the hidden node case, a UE served by the primary node potentially sees a significant interference if a secondary node transmits at the same time. If the queue size at the PN increases, a possible reason is because of interference from a (hidden) SN. In this scenario, the network would benefit if the SN had a more conservative EDT.

Load Adaptation The SN node opportunistically adapts to fluctuations in offered traffic at PNs. If the PN traffic load is low, it may be able to use lower modulation and coding schemes (MCS) while maintaining a similar quality of service. By using more robust coding, higher interference can be tolerated without an increase in packet loss. Thus, the SN EDT can be reduced, allowing for more aggressive SN behavior, depending on the distance of the PN. Alternatively, if the PN traffic load is high, the SN EDT should be increased to prevent interference, even if the PN is further away, allowing higher PN MCS schemes to be used.

With these scenarios in mind, we now present our Q-learning algorithm followed by scenario-specific results.

3.3.3 Q-Learning Algorithm Description

In designing the reinforcement-learning algorithm, our objective is to determine a policy (sequence of state/action pairs) by which the agent (SN eNB) adapts its EDT based on observations taken during the latest epoch to maximize long-term rewards. In our setup, we assume that the PN shares transmit buffer occupancy/queue length information with the overarching CBRS network architecture, making this information available to the SN. In turn, the SN uses this information as

the environmental observation for the reinforcement learning.

Our assumption of the PAL sharing its buffer occupancy and queue lengths is not currently part of the CBRS Alliance standard. However, the topic of PAL–GAA coexistence is also not addressed in the current standard. PAL–GAA coexistence is expected to be addressed in a future version. The standard is actively evolving, and such information could be added to the measurement reports that currently exist between CBSDs, SASs, and the CBRS Alliance Coexistence Manager (CxM) [42]. In fact, the governing bodies are actively considering new measurements to include in these reports. We believe that buffer occupancy and queue length should be considered as they could be, as we demonstrate in this paper, useful tools to improve coexistence in future standards.

The buffer occupancy and queue length are helpful in facilitating coexistence because they, in part, provide a snapshot of the interference observed by the PN. The buffer occupancy is a percentage of the time in the previous epoch that the eNB had data in the buffer waiting to be sent. The queue length is the instantaneous amount of data in bytes to be sent at the end of each training epoch. The choice of buffer occupancy is a result of it being able to capture partial observations of three key factors in the environment: physical interference topology of the entire network, traffic load at the PN, and the virtual topology (sensing topology based on the EDT). Effectively, the buffer occupancy can be seen as a function of both the PN traffic arrivals and the inconsistencies between the interference topology and virtual topology (collisions). Since the goal of the EDT adjustment is to alter the transmission aggressiveness of the SN based on the PN traffic load and to make the virtual topology match the interference topology, this makes buffer occupancy the most appropriate metric to use.

The basic Q-learning implementation is as follows. Let epoch m, with duration T, refer to time interval $mT \leq t < (m+1)T$. The epoch duration, T, needs to be long enough (e.g., 10s - 100s of sub-frames) to avoid adapting to short-lived flows. At time $t = mT$, the agent chooses an action which maximizes its Q-table. At time $t = (m+1)T$, the agent receives the observation of the environment state from its last action, receives the associated reward, updates the Q-table, and then chooses an action α_{m+1} for epoch $m+1$. Fig. 3.12 depicts this iterative process.

Given state space S, the environment lies in one of two states $s \in \{1, 2\}$ defined in Table 3.1. Here, L_m is the instantaneous PN transmit queue size at the end of epoch m, and γ_1 is a threshold used to differentiate high and low traffic loads. γ_1 selection can be used to adjust the relative

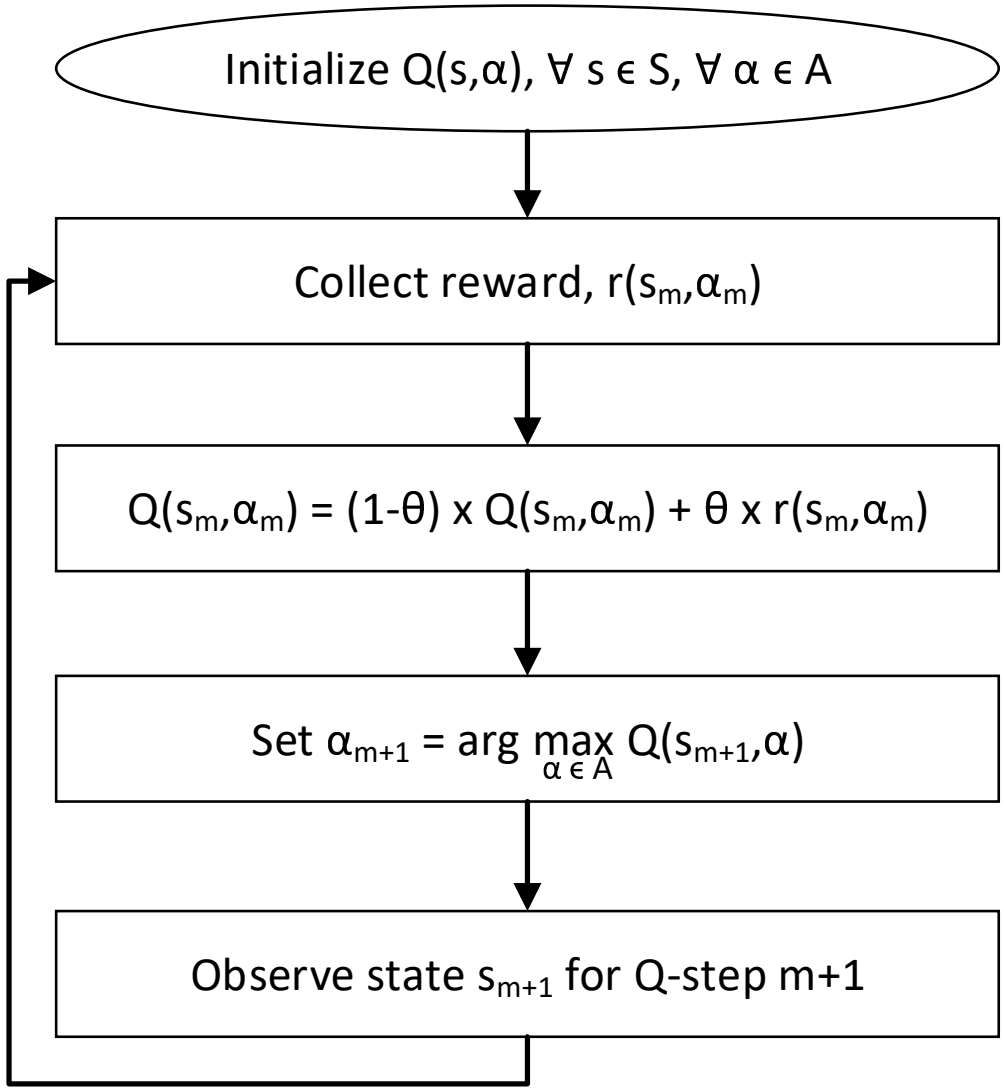

Figure 3.12: Overview of Q-learning algorithm. In each epoch, an action is taken based on the state and learned "quality" of each state/action pair. Based on the change in environment from our action, the quality of the previous action is updated.

weight between PN and SN users. We define these states such that there are binary light/heavy traffic conditions to reduce variables in our performance evaluation. However, this definition can easily be extended to multiple states if it is necessary to define more nuanced packet load conditions by defining multiple thresholds.

The agent is rewarded or punished according to the intuitive guidelines listed in Table 3.2. Table 3.3 elaborates on this by showing specifically provide rewards in our Q-learning algorithm. In general, a positive reward is given if the PN's state improves or if a higher EDT threshold is chosen at the SN without a negative impact on the PN. Consequently, a negative reward is given in cases where the PN transitions to a worse state or if the SN chooses a low EDT value without any benefits. More specifically, we define numerical rewards according to the state transition and the average buffer occupancy over the previous epoch, B_m. Each γ in Table 3.3 is a tunable parameter outlined in Table 3.4. Each can help control the Q-learning to better tailor it for various goals and constraints. Z_m can be considered as a soft reward when outcomes are between actionable thresholds.

As mentioned earlier, the selection of γ_1 plays a critical role in how the learning algorithm behaves between the thresholded values, with higher values allowing the SN to be more responsive to changes in the PN queue length. These threshold values need to be tuned experimentally for different deployments, as there is no absolute rule for how they should be set. In general, each threshold contributes in one way or another to how quickly the algorithm adapts to changes. Depending on the specific situation, more or less rapid responses could be advantageous.

The agent updates the Q-table after each action according to (3.2). Here, $\theta \in (0, 1)$ is a discount factor that is used to control the importance of the current reward, r_m, in terms of updating the Q-function. Larger θ values will prioritize longer-term reward, while a lower θ applies more weight to the next iteration reward. The learning rate is controlled by $\beta \in (0, 1)$ which is a weighted

Table 3.1: States for the Q-Learning

State	Average Primary Node Queue Size	Comment
1	$0 \leq L_m < \gamma_1$	Primary node traffic load is light
2	$L_m \geq \gamma_1$	Primary node traffic load is heavy

Table 3.2: Reward intuition

Reward	Conditions
Positive	• The average primary node queue length is low (state 1), and the last action was to choose a high EDT value (e.g. -62 dBm). • The average primary node queue length goes from high to low (state 2 to 1) following the last epoch.
Negative	• The average primary node queue length is low (state 1), but the last action was to choose a low EDT value (e.g. -77 dBm). • The average primary node queue length goes from low to high (state 1 to 2) following the last epoch, and the buffer occupancy is large.

Table 3.3: Rewards for the Q-Learning

(s_m, s_{m+1})	Reward r_m
$(1,1)$	• γ_4, if $B_m \leq \gamma_2$ and $\alpha_m \geq \gamma_3$ • $-\gamma_4$, if $B_m \leq \gamma_2$ and $\alpha_m < \gamma_3$ • Z_m, otherwise
$(1,2)$	• γ_4, if $B_m \leq \gamma_2$ and $\alpha_m \geq \gamma_3$ • $-\gamma_4$, if $B_m \leq \gamma_2$ and $\alpha_m < \gamma_3$ • $-\gamma_4$, if $B_m > \gamma_2$
$(2,1)$	• 0, if $B_m \leq \gamma_2$ • γ_4, otherwise
$(2,2)$	• 0, if $B_m \leq \gamma_2$ • Z_m, otherwise

Table 3.4: Tunable Parameters for the Q-Learning

Parameter	Description	Units
γ_1	Low/high traffic state boundary	Bytes
γ_2	Buffer occupancy threshold	$\in (0,1)$
γ_3	EDT setting threshold	dBm
γ_4	Standard reward value	Scalar
ϵ	Exploration parameter	$\in$(0,1)
θ	Discount factor	$\in$(0,1)
β	Learning Rate	$\in$(0,1)

average between the previous value of the Q-table for a state/action pair. This affects the tradeoff between convergence time and the relative stability of the entries.

$$Q(s_m, \alpha_m) \leftarrow (1-\beta)Q(s_m, \alpha_m) + \beta\left(r_m + \theta \max_{\alpha} Q(s_{m+1}, \alpha)\right) \tag{3.2}$$

The next action is chosen at each epoch according to the probability distribution in (3.3).

$$P(\alpha_{m+1}) = \begin{cases} 1-\epsilon, & \arg\max_{\alpha \in A} Q(s_{m+1}, \alpha) \\ \epsilon, & \text{rand}(\alpha \in A) \end{cases} \tag{3.3}$$

Here, $\epsilon \in [0,1]$ is an exploration parameter, allowing for occasional random actions to be taken. In general, allowing for exploration prevents the learning algorithm from getting locked into sub-optimal operation by filling in more of the Q-table than would occur otherwise. Furthermore, the exploration probability can be reduced over time as more iterations of the algorithm have occurred.

Although it is possible that PN queue length could increase due to factors besides SN interference such as increased MAC-layer contention, the Q-learning can be resilient to these factors as the Q-table is updated after an epoch consisting of 100s of subframes. Moreover, to further minimize the effects of non-interference caused changes to the PN queue length, the learning rate of the Q-learning algorithm, β, can be chosen to be small. The combination of large epochs and a low learning rate has the effect of "averaging out" most non-interference caused changes to the queue.

3.3.4 Simulations with Adaptive EDT

To evaluate the performance of the adaptive EDT, we perform system simulation in MATLAB. We examine several distinct scenarios to examine how the performance of LBT compares with and without the Q-learning based adaptive EDT.

3.3.4.1 Hidden Node – Mitigating Interference

For the first simulation, we consider the topology shown in Fig. 3.13. All of the PN UEs are placed equidistant from the PN and SN so that if the SN is transmitting, the SINR that they would receive would be approximately 0 dB. This simulation emulates a hidden-node case, where the distance between the PN and SN is far greater than the distance between the SN and the PN UEs. In Fig. 3.14, we compare PN UPT with a fixed, -62 dBm EDT to the adaptive EDT using Q-learning. In this figure, the upper bound on PN transmission is the situation in which there is no secondary user; thus, the PN can transmit interference free. We can see that when using a fixed EDT, the PN UPT drops drastically as expected in a hidden node scenario. However, by allowing the EDT to increase in response to the detection of increasing buffer occupancy at the PN, the penalty received by the PN is greatly reduced.

3.3.4.2 Exposed Node – Exploiting Spatial Reuse

For this simulation, we consider the topology shown in Fig. 3.15. All of the PN UEs are placed at the cell edge opposite of the SN so that they are not susceptible to interference from the SN. This simulation emulates an exposed-node case, where the distance between the PN and SN is far less than the distance between the SN and the PN UEs. In Fig. 3.16, we compare PN UPT with two fixed EDTs to the adaptive EDT using Q-learning. In this figure, the upper bound on PN transmission is the situation in which there is no secondary user; thus, the PN can transmit interference free. We can see that no matter what the EDT of the SN is, the PN UPT remains nearly constant.

For the SN, we see in Fig. 3.17 that by allowing an adaptive EDT, the SN UPT is 2.8 times greater than the case in which the SN uses a static -72 dBm EDT. However, when the SN uses a static EDT of -62 dBm, the average UPT is 7.4 Mbps. Although the SN with an adaptive EDT achieves 76% of the performance of the static -62 dBm case, we see an impressive gain in the UPT over the -72 dBm EDT. It is important to choose parameters such as γ_1 to balance the possible

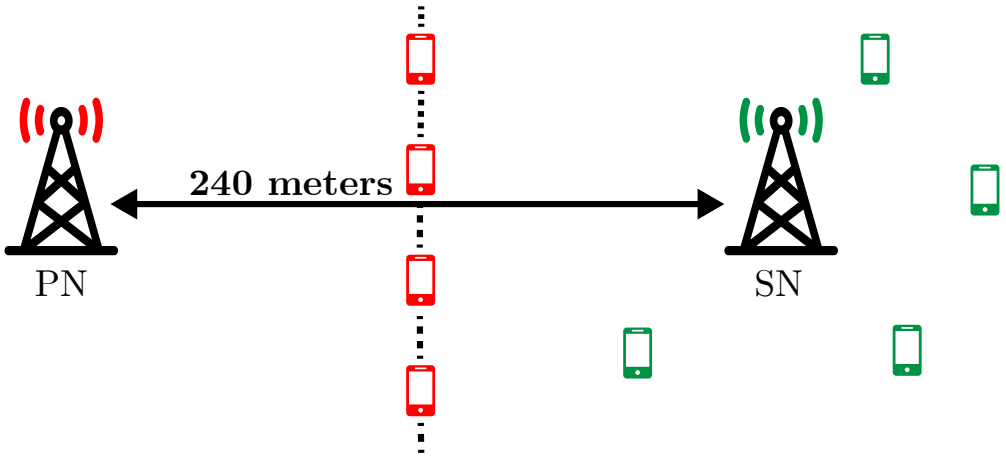

Figure 3.13: Hidden node test topology where the PN UEs are equidistant from the PN and SN. Here, the PN users experience 0 dB SINR and are susceptible to collisions from the SN. The SN needs to adjust its EDT to be more sensitive to PN transmissions.

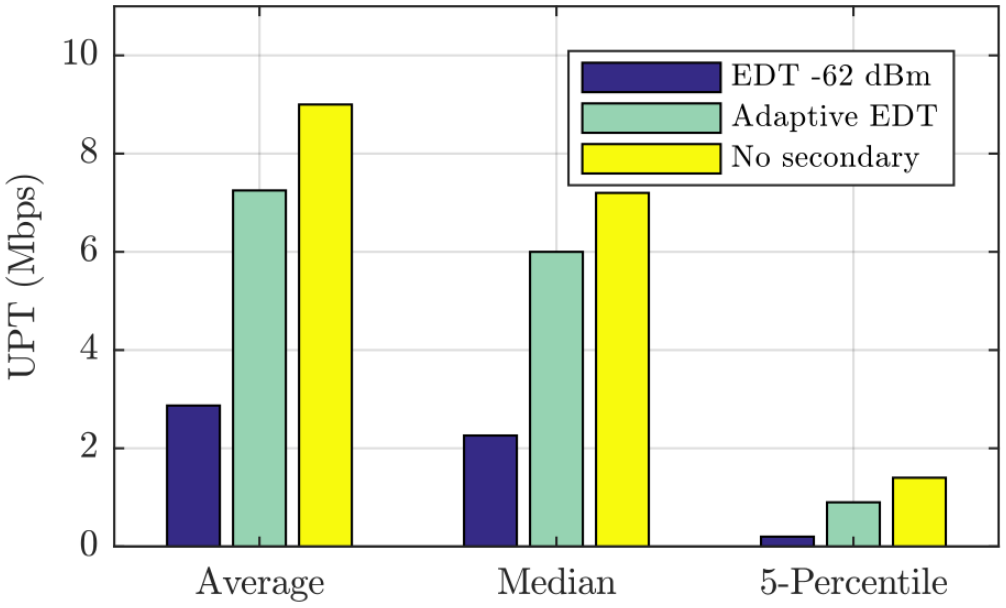

Figure 3.14: PN performance when its UEs are hidden terminals to the SN. An adaptive EDT in the SN allows the SN to reduce its interference to the PN.

losses in hidden node cases with the possible gains in exposed node cases. In this particular simulation, the PN is in a high traffic state with an average queue length of 38,294 bytes. The γ_1 parameter for the Q-learning is chosen to be 75,000 bytes. By choosing a higher γ_1, it would be possible to achieve a greater UPT for the Adaptive EDT in this example at the possible expense of generality for the algorithm, tailoring it too much for one specific scenario.

3.3.4.3 Adapting to PN Load

In the next set of simulations, we have four nodes with two operators in a shared-carrier topology, as shown in Fig. 3.18. In this scenario, the PN load is effectively doubled, as each SN needs to defer to two different PNs. The UE distribution for each node is randomized in proximity around each node. We present the simulation results in Fig. 3.19. We can see that there is not a significant change in the PN UPT for any scheme, as they all approach the upper bound. For the case with no secondary node and the case with a -72 dBm fixed EDT, this is to be expected, as the SNs will be able to sense the PN and defer for nearly every PN transmission. However, when using an adaptive EDT, the SN UPT is increased by a factor of four. This is because the reinforcement learning can adaptively shrink when buffer occupancy remains low at the PN, by taking advantage of momentarily light traffic loads and/or transmissions to UEs located further away from the SN.

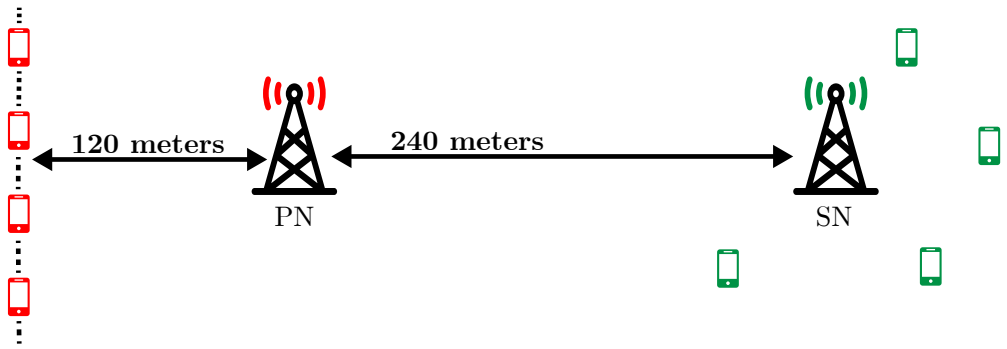

Figure 3.15: Exposed node test topology where the PN UEs are at the cell edge opposite of the SN. Here, the PN users experience a low SINR but are not susceptible to collisions from the SN. The SN needs to adjust its EDT to be less sensitive to PN transmissions.

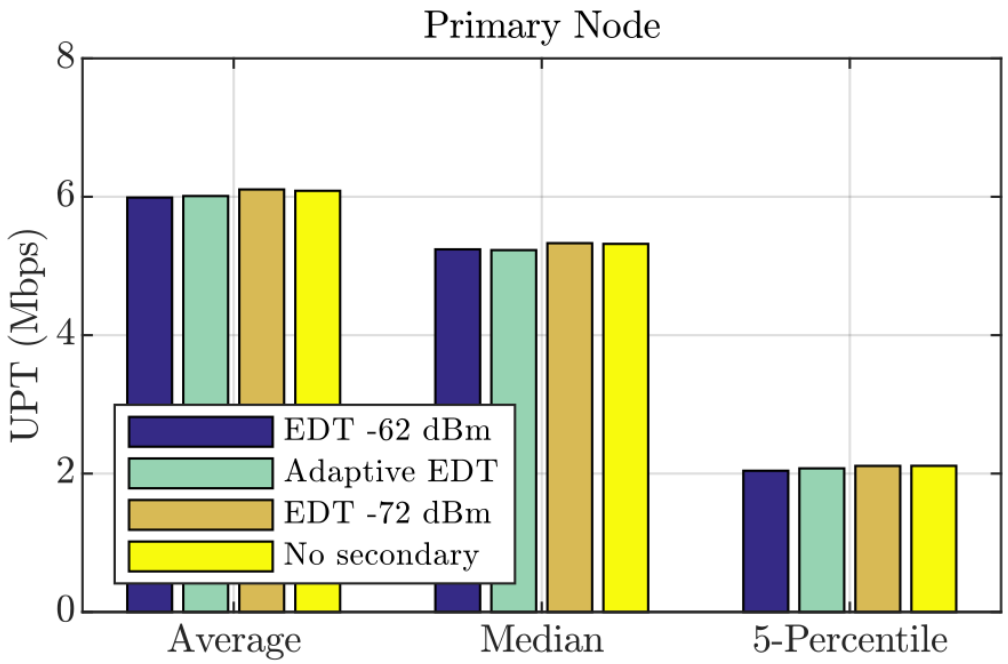

Figure 3.16: PN performance when the SN is an exposed node. An adaptive EDT in the SN allows the SN to reduce its interference to the PN.

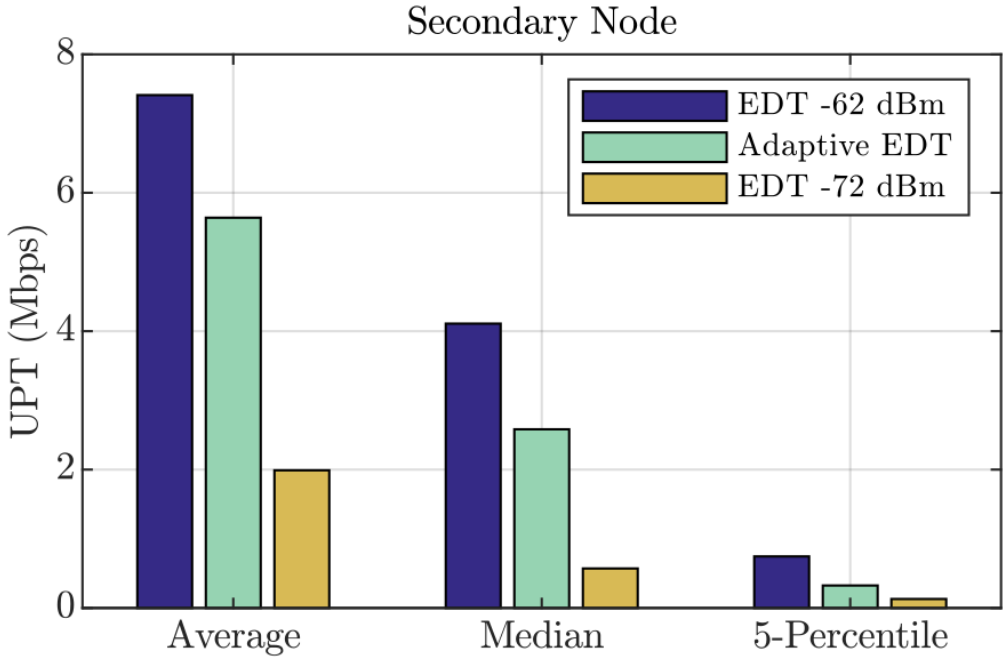

Figure 3.17: SN performance for the exposed node topology. An adaptive EDT allows the node to learn that a higher EDT is acceptable with negligible effect on the PN.

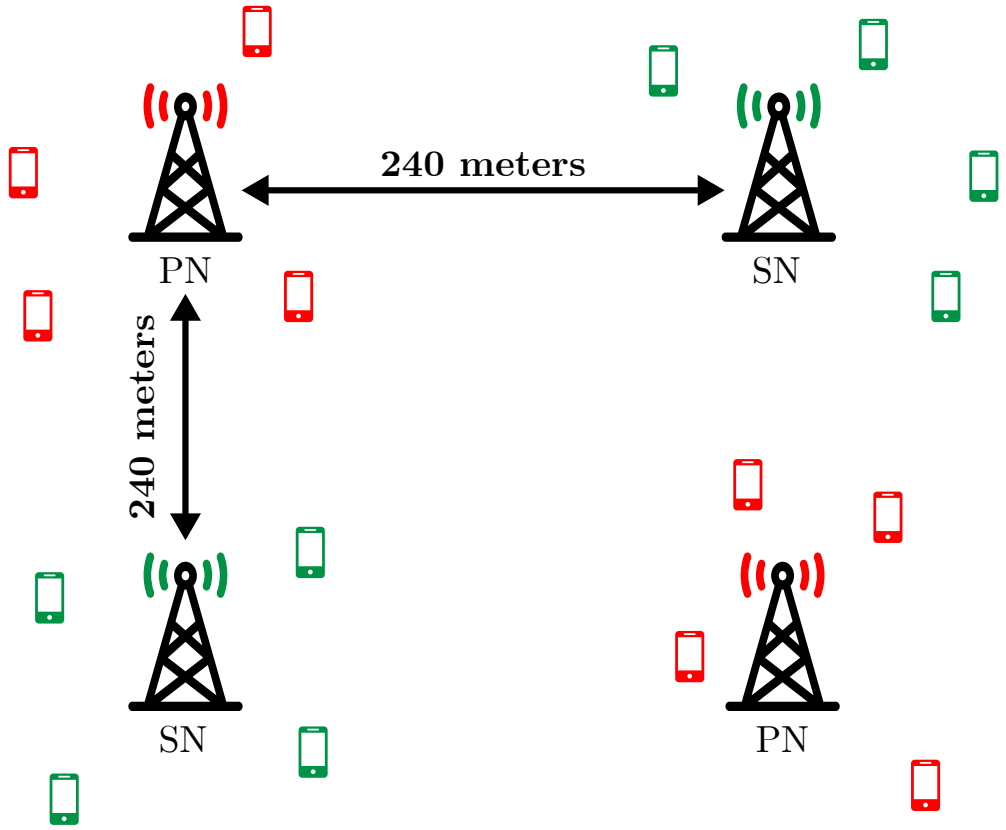

Figure 3.18: Four-node test topology with two SNs and two PNs.

3.3.4.4 Delayed Feedback Scenario

One concern when using PN network statistics to drive the adaptive EDT algorithm is the effect of non-real-time buffer occupancy reports on Q-learning training. In any real system, sharing of data cannot be instantaneous, so we investigate the effect of system delay in PN buffer occupancy reporting in this simulation and measure the performance degradation. Fig. 3.20 shows changes in primary and secondary node UPT as system delay increases for both hidden nodes and normally-distributed UE topologies. As the delay in the reporting of PN statistics increases, there is only a minor drop in performance of the adaptive EDT algorithm showing the resilience of the Q-learning to latency in PN buffer occupancy and queue length reports. Even with significant network delays, the PN performance with an adaptive EDT is shown to be significantly better than the static -62 dBm EDT. We do not evaluate delays greater than 200 ms, as we consider this to be a rough upper bound of the intra-network latency experienced by a real deployment. However, even when

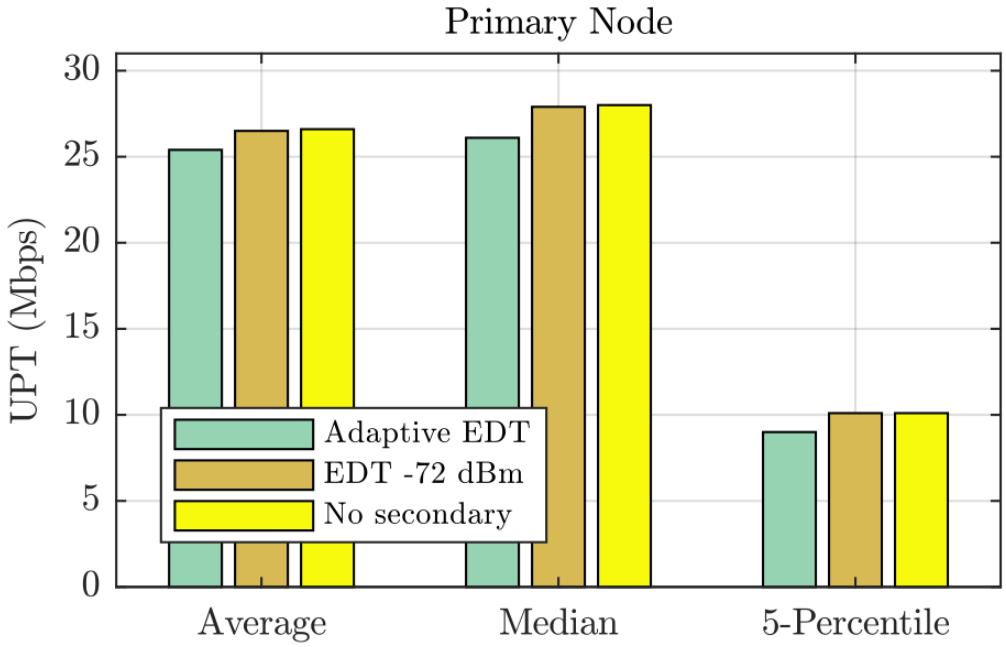

(a) PN UPT for the case where the SN adapts to PN load.

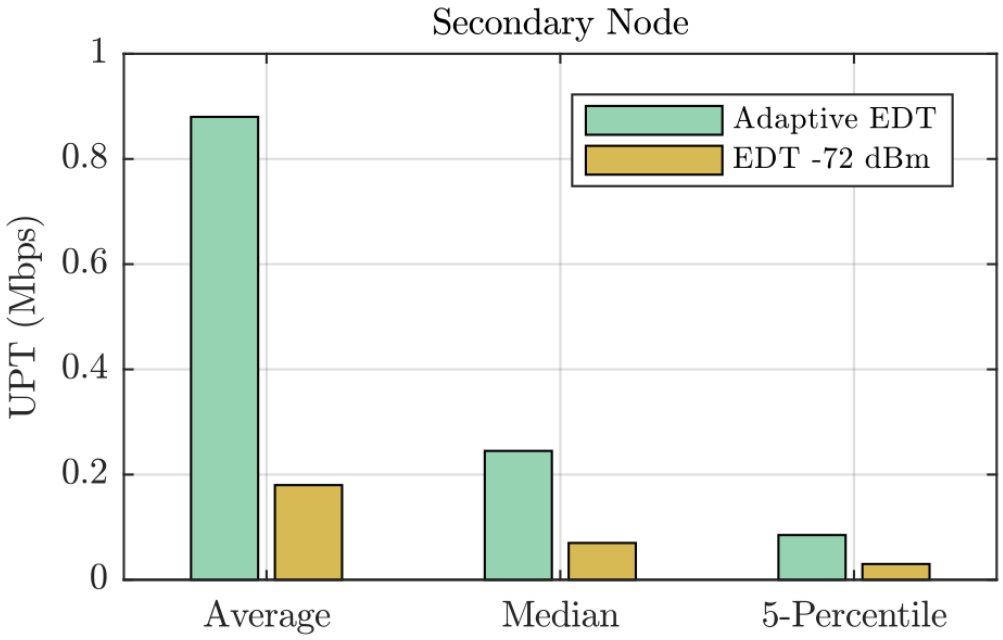

(b) SN UPT for the case where the SN adapts to PN load.

Figure 3.19: Performance of PN and SN with adaptive versus fixed EDTs.

reaching the upper end of this range, standard user mobility would not drive considerable changes to the network topology or traffic load in such a short time window. Therefore, these results are more a reflection of the increased convergence time to a static scenario rather than measuring adaptability to change. With that in mind, we initialized this simulation with a more aggressive EDT of -62 dBm and limited the experiment to 5 seconds to amplify the impact of the initial convergence time.

In the results, we can see that the normally-distributed topology converges to acceptable EDT values quickly even with delay, resulting in no perceived UPT loss for SN or PN. In the hidden node topology, however, the primary UPT experiences a slight inverse relationship with the system delay as a result of the extended convergence time. These results show that the Q-learning based adaptive EDT is resilient to cases where the PN buffer occupancy reports may not be exactly real time.

3.3.4.5 Effect of the State Transition Boundary, γ_1

In the next set of simulations, we have two nodes with two operators, a subset of the scenario in Fig. 3.18. We perform the simulation with two different γ_1 settings for the Q-learning, 25,000 bytes and 75,000 bytes. The choice in γ_1 is notable in that it governs the division between states 1 and 2 in the Q-learning algorithm. A transition from state 1 to 2 causes a negative reward for the action that caused the state transition.

Fig. 3.21 shows these simulation results. In Fig. 3.21a, we show the UPT of the PN. In the case of a fixed EDT of -62 dBm, our highest considered EDT, we see the lowest performance for the PN as the SN will not as readily defer to the PN. Inversely, for a fixed EDT of -72 dBm, our lowest considered EDT, the SN will defer heavily to the PN, resulting in a performance similar to the case where there is no secondary node. When using an adaptive EDT, by changing the value of γ_1, we can balance the performance of the SN and PN. In Fig. 3.21b, we can see the complementary performance of the SN. This result shows that tuning of the γ_1 parameter can be used to balance the tradeoff between PN interference and SN channel access.

3.3.4.6 Scaling State Space

To better understand the impact of state space selection, we examined UPT performance when using different numbers of Q-learning states. We repeated this experiment for normal, hidden

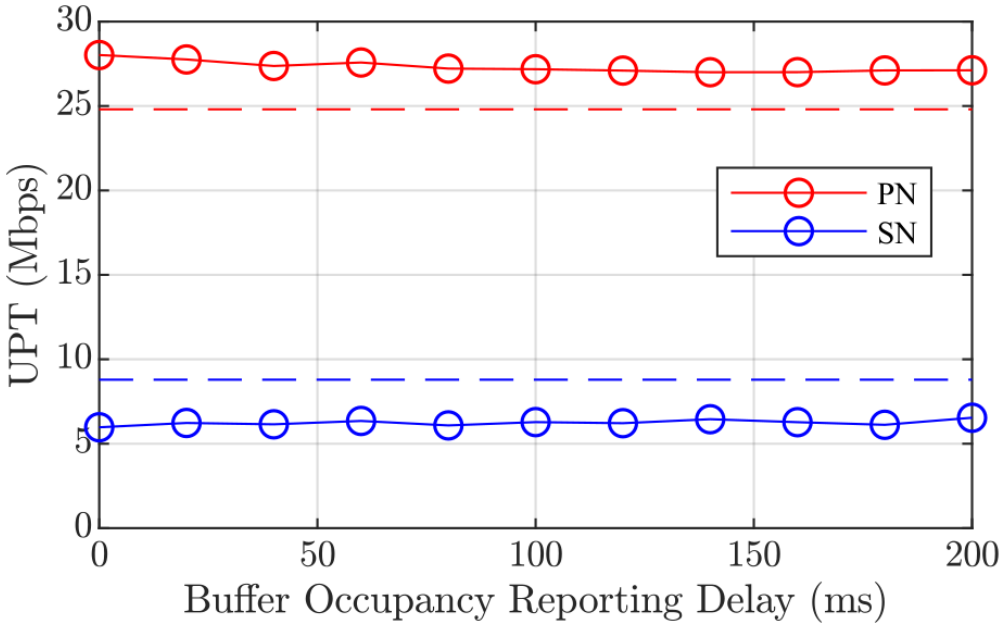

(a) Normal distribution of UEs simulation.

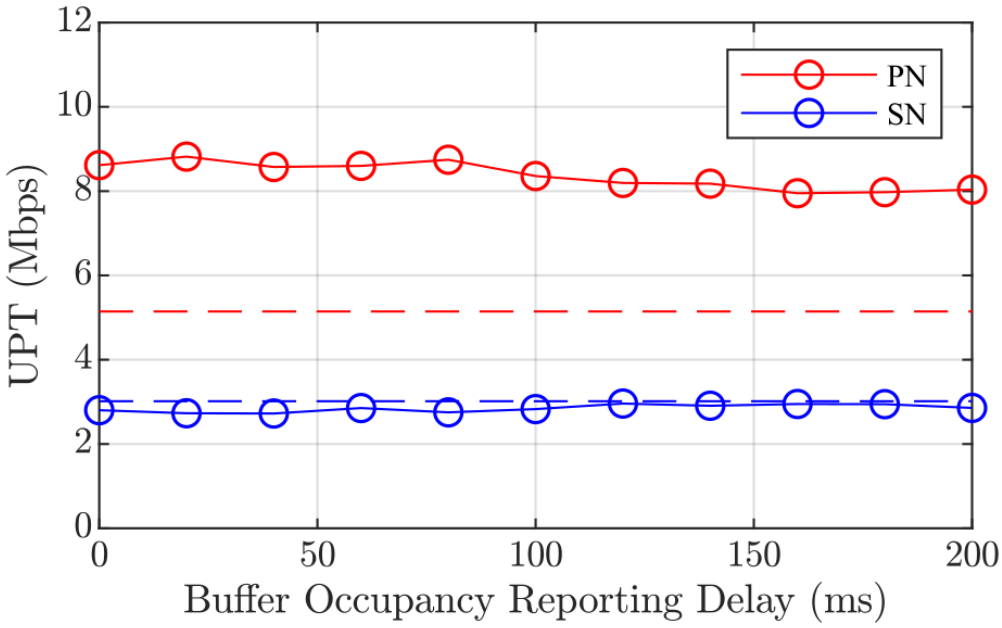

(b) Hidden node simulation.

Figure 3.20: Spectrum sharing with a simulated network latency in PN reports. The solid line with markers represents the results with the adaptive EDT using Q-learning, and the dashed line represents the baseline, static -62 dBm EDT without Q-learning.

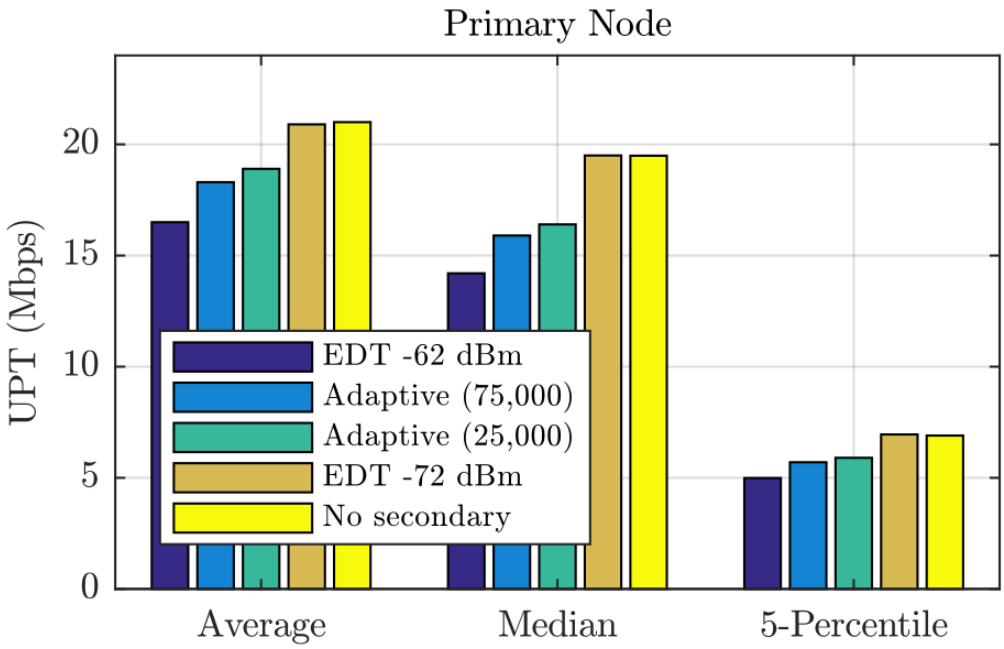

(a) PN UPT with differing γ_1 values.

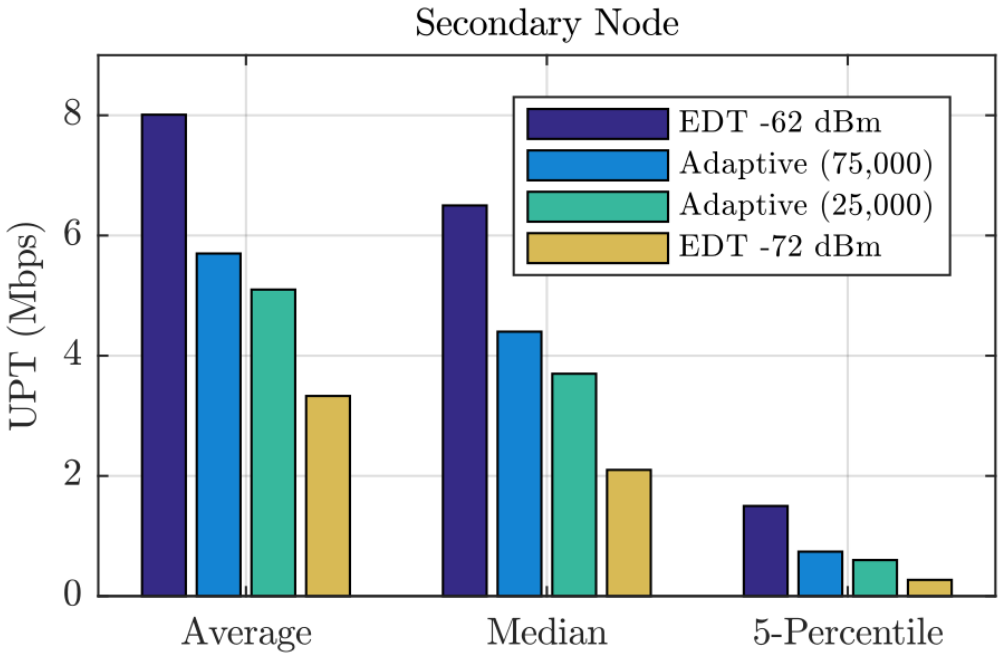

(b) SN UPT with differing γ_1 values.

Figure 3.21: Spectrum sharing with different L_m threshold values.

node, and exposed node topologies to see if perhaps different state space sizes would be preferable in specific situations. Our results are catalogued in Table 3.5. From this table, we can see that there is not a large difference between state size and performance for PN traffic, however there is a slight advantage for using only 2 states in terms of SN UPT in the hidden and exposed node scenarios and using more than 2 states in the normally distributed topology. Ideally, a more nuanced state space could help better select an EDT for SN use, however we found that in the hidden and exposed node cases, the buffer tends to fill and empty very quickly, resulting in most of the state representation lying at the boundaries, with only brief occurrences in the load states between. Additionally, in those problematic topologies, the ideal EDT is min-maxed regardless of current UE transmissions. In the normal distribution, different combinations of UE traffic can result in different instantaneous optimal EDTs, causing convergence time to increase. It is also important to note that if the simulation relies on longer, averaged buffer occupancy's for state assignment, this result could be slightly different.

3.3.4.7 Alternative Sharing Strategies

In this set of simulations, we explore two alternative spectrum sharing strategies, comparing performance to our proposed solution.

For the first alternative spectrum sharing strategy, we consider a simple control system in which the SN directly adapts its EDT according to changes in the PN buffer occupancy. When the PN buffer occupancy increases, the SN EDT decreases by one stage. Similarly, when the PN buffer occupancy decreases, the SN EDT increases by one stage. If the PN buffer occupancy remains the same or the EDT is at the boundary of the action set and cannot be increased/decreased, no EDT

Table 3.5: Parameter Tuning - Number of States

	Normal		Hidden		Exposed	
States	PN	SN	PN	SN	PN	SN
2	29.88	6.23	17.38	2.17	17.84	2.05
3	30.28	6.88	17.39	1.81	17.56	1.84
4	30.15	6.88	17.70	1.77	17.61	1.91
5	29.76	6.69	17.53	1.80	17.93	1.87

change is made.

The second alternative spectrum sharing strategy we consider to be a "perfect" sharing strategy. We want to see what the prospective performance could be if the SN could access the spectrum with complete noninterference with the PN. To achieve this, the SN is provided perfect knowledge of the PN's upcoming transmission information and channel conditions. The SN only transmits when it will not result in a PN CRC failure, resulting in no PN performance loss despite SN channel access. This algorithm is meant to provide an upper bound on sharing performance and is not feasible for implementation.

The results can be seen in Figures 3.22 and 3.23, where we show PN UPT relative to the case where no PN is present and SN UPT relative to the minimum EDT (-62 dBm) in both hidden and exposed node scenarios. We find that the adaptive Q-learning approach consistently outperforms the simple control system across all scenarios.

3.3.4.8 Multi-Node Scenario

In our final set of simulations, we use the LAA indoor scenario as outlined previously (shown in Fig. 3.4). Fig. 3.24 shows the benefit of Q-learning for the case where Op. 1 has at first a low, 0.125 MB/s average traffic arrival rate and then a high, 1.05 MB/s one while the SN traffic is kept high at 1.05 MB/s. For the low-traffic case, there is only about a 5% gain in UPT when using Q-learning because the SNs were already exploiting the many spectrum holes created by the limited traffic activity at PNs. However, when Op. 1 has a higher traffic load, the SN significantly benefits from Q-learning where the adaptive EDT leads to a median UPT improvement of over 30%.

3.4 Conclusions

In this chapter, we examine the challenges of using LBT for PAL–GAA spectrum sharing in CBRS networks by evaluating two different LBT schemes and showing that they can be used to improve significantly the SN UPT with a minor decrease in PN UPT. To reduce the negative consequences of spectrum sharing on the PN, we presented a novel, Q-learning algorithm that adjusts SN opportunistic access via learning an EDT for carrier sensing. We showed that by using average and differential PN buffer occupancy as the environmental observations, the SN can improve their throughput by up to 350% with only marginal losses to the PN UPT (4%). In future work, we can extend the intelligence globally from the local learning framework presented in this work, to

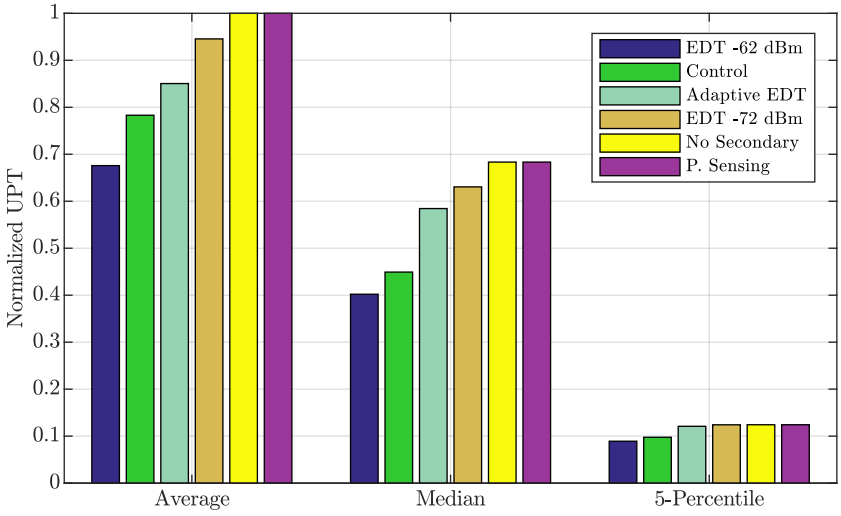

(a) PN UPT relative to performance with no SN.

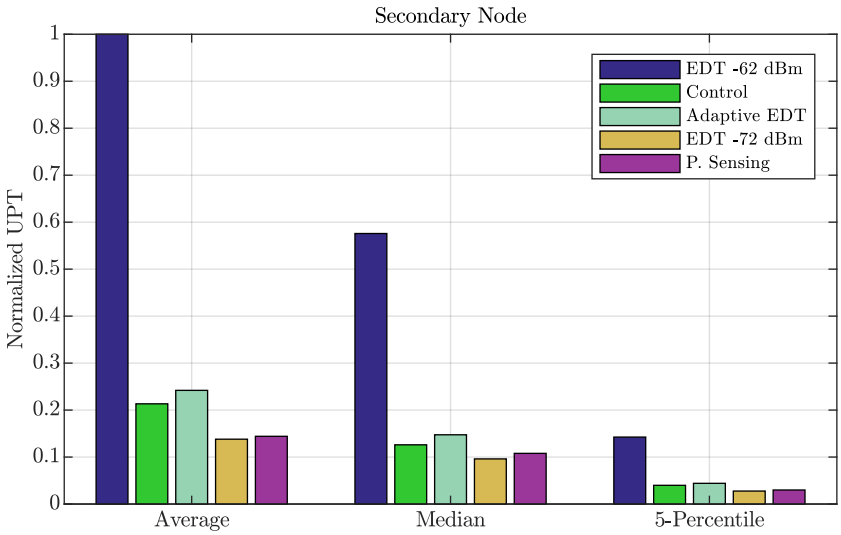

(b) SN UPT relative to performance using highest available EDT.

Figure 3.22: EDT selection algorithm comparison in a hidden node scenario.

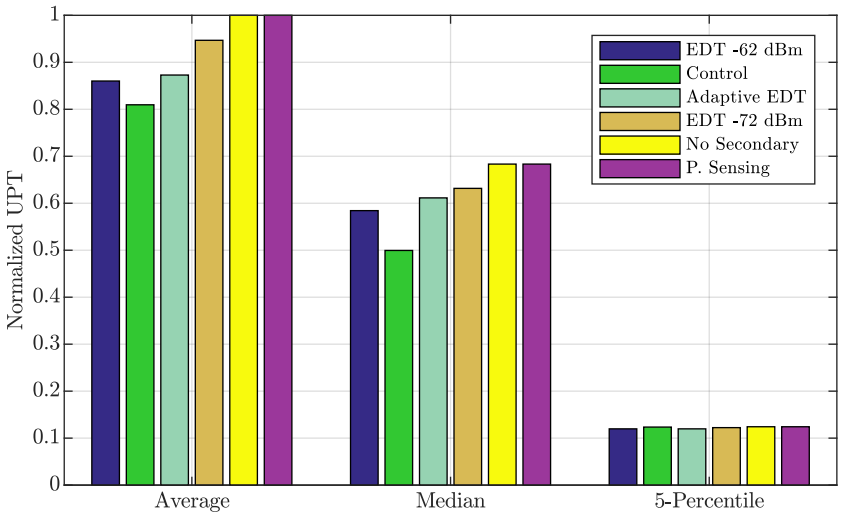

(a) PN UPT relative to performance with no SN.

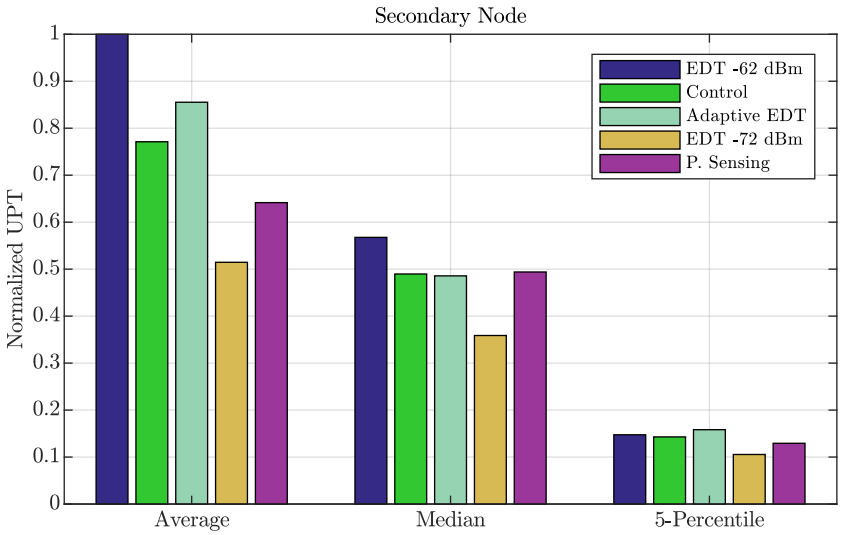

(b) SN UPT relative to performance using highest available EDT.

Figure 3.23: EDT selection algorithm comparison in an exposed node scenario.

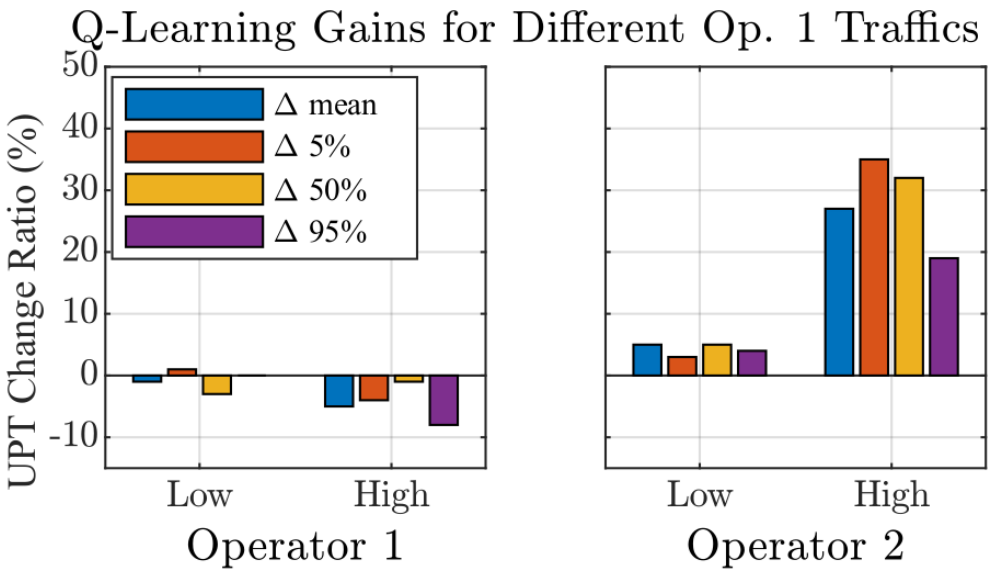

Figure 3.24: Change in UPT for Op. 1 and Op. 2 when Op. 1 has a 0.125 MB/s (Low) and a 1.05 MB/s (High) average traffic arrival rates while Op. 2 always has a 1.05 MB/s traffic arrival rate. Here, each bar shows the change in the mean or corresponding percentile when Q-learning is used to alter the EDT of the SN.

jointly optimize within and across different shared-spectrum deployments, and examine how this work can scale to situations with multiple SNs.

Chapter 4

Machine Learning Enhanced Channel Selection for Unlicensed LTE

4.1 Introduction

Deployment of LTE technology on unlicensed bands allows new spectrum opportunities for cellular-based mobile broadband. This new spectrum helps to improve downlink data rates, increase network capacity, and, being free for use, allows for new use cases such as neutral host. Moreover, the low-power small/femto cells promise low spatial radio footprints that enable multiplicative capacity gains via frequency reuse. However, being unlicensed, it is possible to face significant interference from incumbent Wi-Fi devices and other unlicensed LTE operators.

Many other works have studied Wi-Fi coexistence with specific implementations of unlicensed LTE (generically referred to as uLTE hereafter), such as LTE-U and its sister technology, Licensed Assisted Access (LAA), and found that LTE can coexist with Wi-Fi better than Wi-Fi can coexist with itself [3, 43, 44]. As operator deployment on the unlicensed band ramps up [45, 46], it is necessary to turn our attention to studying coexistence between LTE cells operating in the unlicensed band, which is largely neglected by past works.

Figure 4.1 shows a simple illustration of the challenges we face on unlicensed spectrum. In dense deployments, there may be many uLTE eNBs from many operators, uLTE user equipments (UEs) communicating with uLTE Evolved Node Bs (eNBs), Wi-Fi access points (APs), and Wi-Fi clients operating on the same or adjacent channels potentially causing interference.

There are two obvious ways to attempt to improve spectrum sharing: improve sharing in the time-domain via improved channel contention mechanisms, or improve sharing in the frequency domain via improved channel selection. There has been no shortage of papers on the issue of improving time-domain sharing via contention and Listen Before Talk (LBT) schemes [47]. However, the LBT mechanism is difficult to optimize for since there are many competing regulatory requirements. Alternatively, channel selection is not as burdened by such issues, and improvements to the

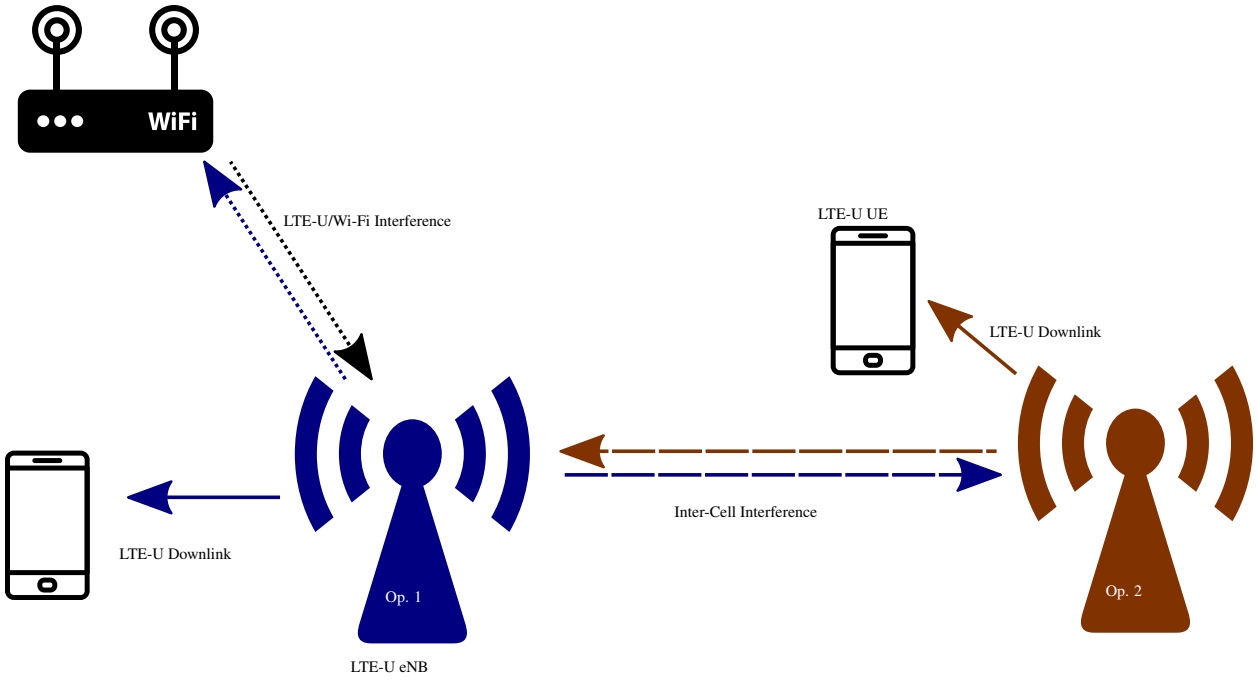

Figure 4.1: Unlicensed band deployment scenario. In practice, there may be many incumbent Wi-Fi devices, new uLTE eNBs from multiple operators (Ops), and their clients all possibly causing interference with each other. Some sort of interference/collision avoidance mechanism will be necessary to maintain a minimum QoS on the band.

channel selection procedure can yield concrete gains, agnostic to any specific contention protocol. Therefore, in this chapter, we seek to improve channel selection. To do this, we utilize statistical modeling of channel utilization over time.

We make the following observation and seek to exploit it in this work. *If the full unlicensed band channel utilization at each small cell can be estimated and shared with other small cells, a statistical model can be built for each cell selecting the best channel to transmit such that expected network capacity is maximized.* This has the effect of not only improving uLTE/Wi-Fi coexistence, but also improving uLTE inter-operator coexistence. We first present a framework that allows uLTE small cells to collect and share their estimated local channel utilizations with neighboring cells. Next, we propose a novel reinforcement algorithm in which each small cell can leverage channel utilization information provided by neighboring cells to select the best operating channel. Finally, we implement, simulate, and compare the performance of our proposed algorithm against the few existing state-of-the-art schemes. We find that by leveraging shared information and reinforcement

learning techniques for proper channel selection, aggregate eNB performance is improved for both random scenarios and ones with a high degree of spectral congestion from Wi-Fi devices.

The rest of the chapter is organized as follows. In Section 4.2, we present a mathematical model of the problem and then develop our channel selection algorithm. We simulate and compare our channel selection algorithm to existing and other proposed solutions in Section 4.3. Finally, we review related works in Section 4.4 and conclude this work in Section 4.5.

4.2 Channel Selection Algorithm Overview

In this section, we outline various algorithms that could be used for channel selection strategies. However, before doing so, we first provide details relating to our system model.

We consider a field of operation in which there are three types of entities deployed: eNBs operating in the unlicensed band, Wi-Fi APs, and UE devices connected to the eNBs. We lump Wi-Fi device behavior into that of the Wi-Fi AP, as we are only looking at the downlink channel and are not directly considering Wi-Fi performance in this dissertation, as it is a well-studied topic. Most unlicensed LTE protocols require or at least have some form of LBT available. In our formulation, we consider an abstracted form of LBT between cells and Wi-Fi APs, in which transmission opportunities are slotted. In each slot, eNB transmission requires a clear channel assessment (CCA) be performed to determine whether or not the medium is idle. Since this work focuses primarily on inter-cell coexistence, we do not specifically model Wi-Fi MAC layer operation, instead assuming that Wi-Fi APs capture the channel according to Poisson arrival process with independent uniformly distributed random access duration bounded by the CSMA protocol. Each eNB only contends for the channel in slots where Wi-Fi traffic is absent; channel access for the slot is then randomized between all contending eNBs for each slot. By modeling our setup in this way, we can abstract away from specific protocol implementations (allowing for this work to extend towards future spectrum sharing protocol design) and reduce overall variability between simulations. We define eNBs that exist within contention range of any given eNB as its neighbors. The neighbor set will be used to simplify much of our analysis.

The major objective of this work is to provide an improved methodology for selecting an unlicensed channel for transmission, such that each eNB can the best quality of service for their connected UEs. In order to do so, we modify the fixed-channel operation of LTE and adapt it

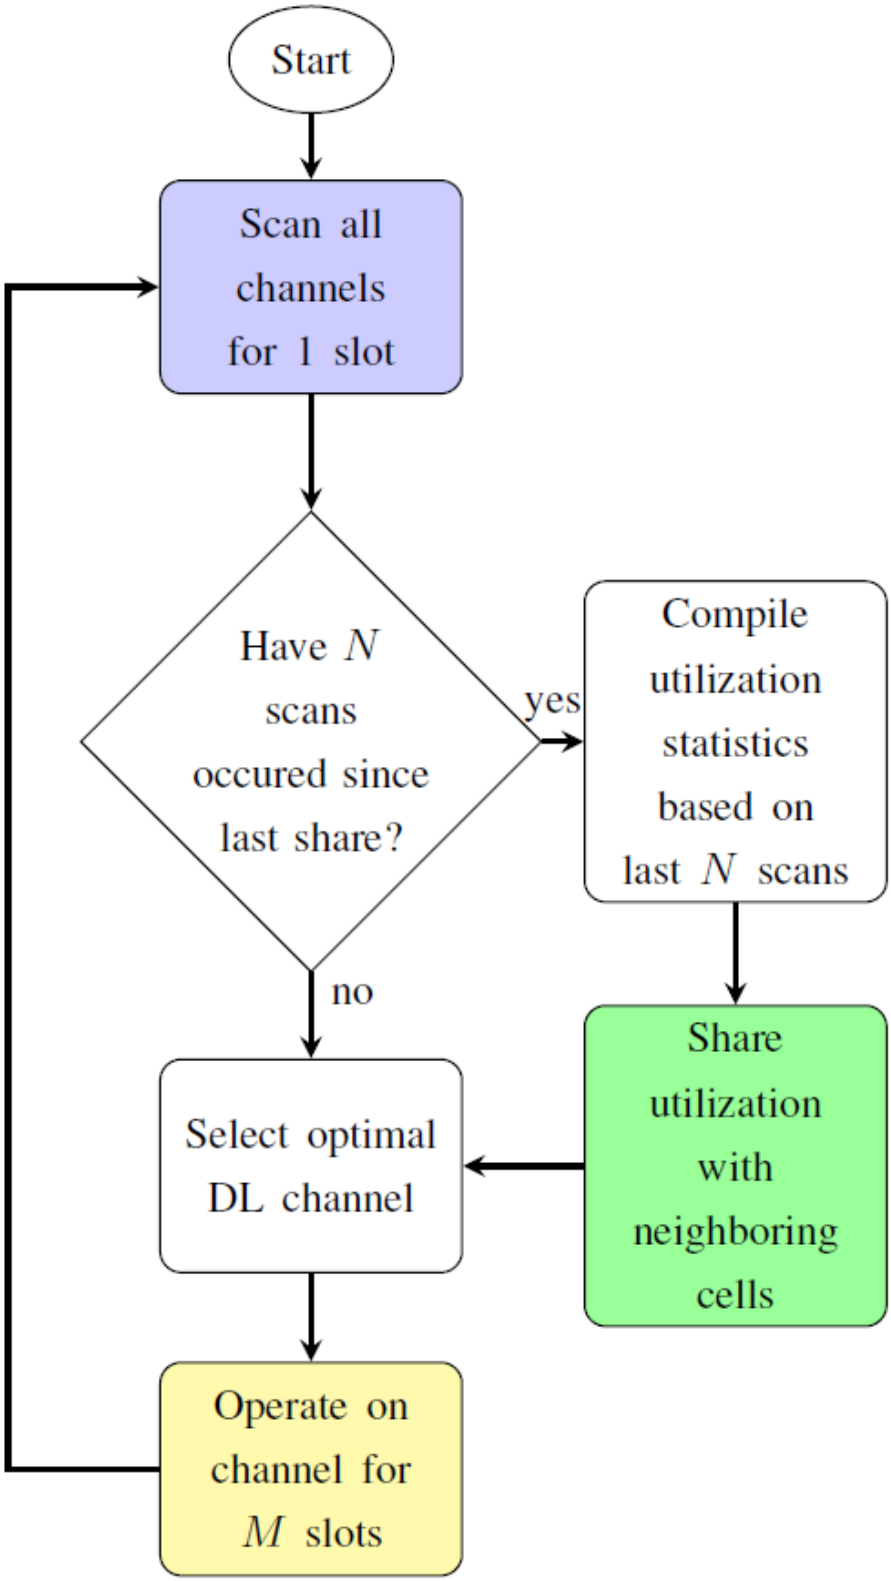

Figure 4.2: MAC system overview. With a period of every M slots (or 1 epoch), the eNB will operate on a DL channel. The eNB will scan the channel and average utilization statistics over a window of the last N scans.

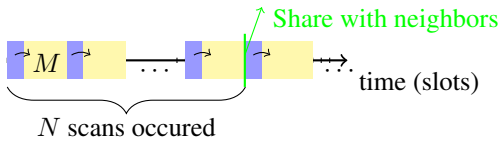

Figure 4.3: Time domain representation of eNB behavior. An eNB scans all channels for one slot to measure the energy in each channel (shown in blue). Using this info, the eNB chooses a DL channel to operate on for the next M slots (yellow). This process repeats for a total of N scans before the utilization statistics are reported to neighboring cells.

to our slotted transmission system. The channel selection is done periodically according to the following steps.

1. Each eNB scans the band for 1 full slot to measure the power in each channel.

2. Based on this, it will choose its DL channel to operate on for the next M slots based on the channel selection algorithm implemented.

3. After M slots have passed, the eNB will re-scan the set of channels. We call this duration of M an epoch.

4. After N scan cycles have completed, the eNB will share its current channel utilization statistics with neighboring eNBs before the channel selection.

A flow chart of the overall system algorithm is shown in Figure 4.2, and a time domain depiction of basic operation is shown in Figure 4.3. We note that the 1 slot scanning duration is considered a silent period for the eNB in that it is unable to transmit throughout. Additionally, we assume silent periods are done simultaneously between all eNBs. The primary reason for this is to allow each eNB to capture the Wi-Fi utilization without the addition of inter-cell interference. While this is not strictly necessary due to the ease in which an eNB could differentiate between Wi-Fi and LTE transmissions, it is convenient for notation and simulation implementation.

Since each scanning period is limited in duration by design, we track the utilization over time via an exponentially weighted moving average as:

$$\mathbf{u_i} \leftarrow \alpha \hat{\mathbf{u}}_\mathbf{i} + (1-\alpha)\mathbf{u_i}. \tag{4.1}$$

Here, the moving average for eNB i, $\mathbf{u_i} \in [0,1]^{|\mathcal{K}|}$, is updated by using its past value and most recent measurement $\hat{\mathbf{u}}_\mathbf{i}$, both being vectors representing the utilization on each channel, k, in the set of channels, $\mathcal{K}$. The scalar $\alpha \in [0,1]$ is a weight that balances recent measurements versus past measurements. If, for example, eNB i measures channel k to be fully utilized by Wi-Fi devices in its most recent measurement, then $u_{ik} = 1$. The moving-average formulation strikes a reasonable balance of prioritizing recent measurements while maintaining a long-term history and having low implementation complexity.

4.2.1 Greedy channel selection

With a greedy channel selection algorithm, each eNB seeks to optimize its throughput by choosing the channel with the lowest utilization in the previous scan. This is similar to how many uLTE base stations currently operate [48–50]. The idea is to pick the least occupied channel possible, thus minimizing the impact on existing Wi-Fi networks in the area and maximizing channel access for the eNB. This is shown in Equation 4.2 where k_j is the channel for eNB j, t is the current epoch, and $u_{j,k}$ is the most recent measured utilization by eNB j for channel k.

$$k_j(t+1) = \arg\min_{k \in \mathcal{K}} u_{j,k} \tag{4.2}$$

The greedy channel selection algorithm has the benefit of being completely decentralized with each eNB making decisions based only on local measurements. However, this simple algorithm may suffer in scenarios where there are many eNBs in close proximity as all eNBs would collide on the channel with the lowest utilization.

4.2.2 SOPI — Stochastic Optimization with Partial Information

To improve on the greedy approach, we present two different algorithms that each consider not only local channel utilization measurements but also neighboring measurements. In both of these algorithms, each eNB locally calculates the expected capacity for the network given its local sensing of the unlicensed channels and neighboring reports of their sensing. For the first algorithm, we use a statistical approach to leverage local and neighboring channel utilization measurements to optimize sum capacity. We consider this an optimization with partial information, as the only information shared from neighbors is the Wi-Fi utilization. This algorithm is performed indepen-

dently at each eNB. The process of each eNB is to first estimate the expected network capacity for each channel, then select the channel that maximizes this estimated capacity. To properly calculate the expected capacity, we first estimate the joint probability of channel selection between all neighboring eNBs and each combination's respective capacity.

We begin by estimating the probability that an eNB is able to capture the channel for transmission. Each eNB, i, measures the utilization of channel k, u_{ik} given from Eq. 4.1. For each eNB, we can derive the probability of access, P_a, to channel k as:

$$P_a(i,k) = 1 - u_{ik}. \tag{4.3}$$

Let eNB i have a set of neighbors, $\mathcal{N}$. We define a subset of neighbors simultaneously operating on the same channel as $J \subseteq \mathcal{N}$. We also define binary indicator vector, $\mathbf{y}$, of length $|J|$ where each element represents the instantaneous channel availability for eNB $j \in J$. All combinations of $J \subseteq \mathcal{N}$ must be iterated through to calculate the expected network capacity. We can then define the probability of eNB i capturing channel k for a slot given the subset of neighbors J also choose k as:

$$P_{\text{capt.}}(i,k,J) = \sum_{\mathbf{y} \in Y} \frac{\prod_{j \in J} \left((1 - P_a(j,k))^{\bar{y}_j} \cdot (P_a(j,k))^{y_j} \right)}{w(\mathbf{y}) + 1} \tag{4.4}$$

Here, $w(\mathbf{y})$ is the Hamming weight of binary indicator vector $\mathbf{y}$, $\bar{\mathbf{y}}$ is the binary complement of $\mathbf{y}$, and $Y = \mathbb{Z}_2^{|J|}$ is a set containing all realizations of $\mathbf{y}$. This channel capture probability is based on the premise that the channel can only be captured when the same channel neighbors either do not have access to the channel for the slot, or they do have access to the channel, and eNB i wins contention against them. With the channel capture probability found, we calculate the channel capacity for eNB i with contending set J on channel k as:

$$C_i(k,J) = P_a(i,k) P_{\text{capt.}}(i,k,J) C_{i,\text{max}} \tag{4.5}$$

Here, the probability of transmission is scaled by both the access probability, P_a, for node i from Eq. 4.3 and the capture probability, $P_{\text{capt.}}$ from Eq. 4.4, since the channel needs to be available before it can be contended for. We define $C_{i,\text{max}}$ as the unimpeded channel capacity estimated by

i based on the average SNR experienced over all connected UEs as $W \log_2(1 + \text{SNR}_i)$. Using a similar method, eNB i estimates the expected capacity of neighboring eNB j by assuming that j's only neighbor is i and that average SNR for j is identical to its own.

$$C_j(k, J) = P_a(j, k) P_{\text{cap}}(j, k, i) C_{i,\text{max}} \tag{4.6}$$

Now that we have found both the local and neighboring expected capacities for all k and J, we can begin calculating the overall expected network capacity given eNB i selects channel k.

For eNB i, we can estimate the probability that a neighboring eNB $j \in \mathcal{N}$ will have access to channel k by converting the channel availability of each eNB to a probability distribution given as:

$$P_s(j, k) = \frac{1 - u_{jk}}{\sum_{\gamma \in \mathcal{K}} (1 - u_{j\gamma})} \tag{4.7}$$

While ideally, we would perform a joint optimization over all channels, our model assumes eNB i only has local Wi-Fi utilization of its neighbors and nothing else. Thus, this estimate is subject to error based on discrepancies in actual access probabilities, even if all neighboring eNBs are using identical channel selection strategies. We address this problem in more detail later on.

Next, we derive the probability that multiple neighbors, denoted by the set $J \subseteq \mathcal{N}$, will access the single channel k. To calculate this, we multiply the individual probabilities that a neighbor in J accesses the channel by the probabilities that other neighbors access any other channel. This is given as:

$$P_s(J, k) = \prod_{j \in J} P_s(j, k) \cdot \prod_{j \in \mathcal{N} \setminus J} (1 - P_s(j, k)) \tag{4.8}$$

The total capacity experienced by the network when eNB i with the subset of neighbors J transmits on channel k is given by:

$$\begin{aligned} C_{i,\text{net}}(J, k) =& (1 - u_{ik}) C_i(J, k) \\ &+ \sum_{\gamma \in \mathcal{N}} (1 - u_{\gamma k}) \left(x_\gamma C_\gamma(J, k) + \bar{x}_\gamma C_{i,\text{max}} \right) \end{aligned} \tag{4.9}$$

Here, x_γ is an indicator. Namely, 1 if $\gamma \in J$ and 0 for all other cases. The variable $\bar{x}_\gamma$ is an

indicator that is the complement to x_γ.

The expected network capacity if eNB i chooses channel k is given by summing the probability that a subset of neighboring eNBs, J, transmit on the channel multiplied by the network capacity in such a scenario over all possible combinations of eNBs in the neighboring set, $\mathcal{N}$, which is given by the power set, $\mathbb{P}(\mathcal{N})$.

$$\hat{C}_{i,\text{net}}(k) = \mathbb{E}\left[C_{i,\text{net}}(J,k)\right] \tag{4.10}$$

$$= \sum_{J \in \mathbb{P}(\mathcal{N})} P_s(J,k) C_{i,\text{net}}(J,k) \tag{4.11}$$

The channel chosen by eNB i for epoch $t+1$ should therefore be selected such that the expected network capacity is maximized according to:

$$k_i(t+1) = \arg\max_{k \in \mathcal{K}} \hat{C}_{i,\text{net}}(k) \tag{4.12}$$

The result of this algorithm, then, is to select a channel that maximizes the expected sum capacity in the network. However, because the second-order neighbors (neighbors of each neighbor to a given eNB i) are unknown, the only second order neighbor we can assume exists is eNB i itself. Without a full neighbor graph, the neighbor channel access probabilities estimated by $P_s(J,k)$ can differ from reality. Unfortunately, in the case where channel utilization statistics are similar across all eNBs for each channel, each eNB will have a high probability of selecting the same channel despite any attempted avoidance due to all eNBs operating under the same channel selection policy. An eNB following the same policy with the same information will necessarily make the same decision in a hard-decision maximization.

In such a situation, the issue of channel selection is a game-theoretical problem. The algorithm performance can be expected to be worse than the greedy selection algorithm under the premise that at least the greedy selection algorithm will lead to mass collisions on the channel with the lowest measured Wi-Fi occupation, whereas the partial information in the SOPI algorithm can drive eNBs to all simultaneously select a busier channel, degrading performance further.

We can partially mitigate the damage in this situation by converting the vector of expected network capacities for each channel, $\hat{C}_{i,\text{net}}(k)$, to a probability distribution, where $P_c(i,k)$ represents

the probability of eNB i selecting channel k given by:

$$P_c(i,k) = \frac{\hat{C}_{i,\text{net}}(k)}{\sum_{\gamma \in \mathcal{K}} \left(\hat{C}_{i,\text{net}}(\gamma) \right)} \tag{4.13}$$

However, such a solution is not ideal as the result-per-game is memory-less and does not converge to a Nash equilibrium over time, resulting in a pseudo-random collision behavior.

4.2.3 Romero-Q — Reinforcement learning using throughput from [1]

In [1] the authors adapt reinforcement learning for uLTE channel selection. We detail their algorithm along with a brief introduction to reinforcement learning below.

Reinforcement learning is an iterative interaction between the agent (which hosts the algorithm) and the environment used in many fields. Agents interact with the environment via actions. Each action performed by the agent impacts the environment in some way, causing the environment to emit an observation and reward. The observation and reward are used to refine the decision-making function, or policy, of the agent. Reinforcement learning methods specify how the agent changes its policy as a result of its experience. The agent's goal is to maximize the total amount of reward it receives over the long term.

In [1], a subcategory of reinforcement learning called Q-learning is used. The quality $Q_i(k)$ of eNB i taking the action to select channel k based on the number of available channels, $k \in K$, is defined by:

$$Q_i(k) \leftarrow (1 - \alpha_L) \cdot Q_i(k) + \alpha_L \cdot r_i(k) \tag{4.14}$$

Here, α_L is the learning rate and $r_i(k)$ is the reward the agent received from the environment for selecting channel k. They derive the reward is determined based on the achieved throughput seen on channel k and the final decision is probabilistic based on the soft-max policy described in [51]. This decision is given as:

$$P_{Qi}(k) = \frac{\exp\left(\frac{Q_i(k)}{\tau}\right)}{\sum_{\gamma \in \mathcal{K}} \exp\left(\frac{Q_i(\gamma)}{\tau}\right)} \tag{4.15}$$

4.2.4 RLPI — Reinforcement Learning with Partial Information

In order to properly manage situations with similar channel utilizations across neighboring eNBs, we propose a novel reinforcement learning algorithm that combines prior work for LTE-U channel selection from a purely game-theoretical perspective [1] with our analytical model that includes both local and neighboring channel utilization information.

Rather than allocate reward based on achieved throughput as done in [1], we prefer to measure the average SINR over the past epoch and calculate the reward directly using how successful the eNB was at transmitting on the chosen channel. Directly using a notion of contention is a natural choice as throughput is a function of many parameters such as UE placement, modulation and coding scheme (MCS), and inter-cell interference. The reward used in this work is given as:

$$r_i(k) = \frac{m_{\text{cap}}(i)}{M}, \tag{4.16}$$

Here, $m_{\text{cap}}(i)$ is the number of slots successfully transmitted on by eNB i in the previous epoch, and M is the number of slots in each epoch. By using a unit-less reward, we can abstract away many of the dynamic aspects of the system, such as interference, channel fluctuations, UE scheduling, and modulation/coding scheme selection.

We define a metric for similarity, S, based on the Jensen-Shannon divergence between all normalized channel utilizations. This normalization is given by:

$$P_u(j) = \frac{u_{jk}}{\sum_{k' \in \mathcal{K}} u_{jk'}}. \tag{4.17}$$

Using these normalized channel utilization distributions, we define the channel utilization similarity for eNB i as:

$$S_i = \frac{1 - \sum_{j \in \mathcal{J}} JSD(P_u(i)|P_u(j))^{1/\omega}}{|J|} \tag{4.18}$$

Here, $JSD(P_u(i)|P_u(j))$ is the Jensen-Shannon divergence between the normalized channel utilization experienced by eNB i and neighbor j. We normalize by the cardinality of J to bound the similarity between $[0, 1]$, where a value of 1 indicates the channel utilization between the eNB i and all of its neighbors is identical, and a value of 0 indicates all neighboring channel utilizations di-

verge from the locally measured channel utilization. We add an ω exponent as a tun-able parameter that functions to scale how aggressively the algorithm relies on the net capacity estimates over the Q-value. Unless specified otherwise, we set $\omega = 2$, which commonly used as the Jensen-Shannon distance [52].

Similar to Romero-Q, RLPI performs channel selection through a soft-decision function based on the soft-max policy [51] combined with the estimated channel capacities. The probability for eNB i to select channel k via the modified soft-max decision is given by:

$$P_{Qi}(k) = \frac{\exp\left(\frac{S_i \tilde{Q_i}(k) + \bar{S}_i \tilde{C}_{i,\text{net}}(k)}{\tau}\right)}{\sum_{\gamma \in \mathcal{K}} \exp\left(\frac{S_i \tilde{Q_i}(\gamma) + \bar{S}_i \tilde{C}_{i,\text{net}}(\gamma)}{\tau}\right)} \tag{4.19}$$

Here, $\bar{S}_i$ is the complement to S_i given by $\bar{S}_i = 1 - S_i$, $\tilde{Q_i}(k)$ and $\tilde{C}_{i,\text{net}}(k)$ are the scaled sets of eNB i's Q values and expected network capacities for each channel $k \in K$, such that the maximum value in the set is 1 and other members of the set are scaled relative to it. This scaling to unit values allows the Q values and expected network capacities to be compared on the same relative scale. τ is the temperature, which is a function of the number of epochs experienced by the algorithm as $\tau = \frac{\tau_0}{\log(1+t)}$. The temperature is used as a way to increasingly polarize the selection probabilities the longer the algorithm runs [1]. Consequently, more exploration will occur during the earlier epochs, decreasing as time goes on. The value of τ_0, or the initial temperature, should be set to tune the exploration. Equation 4.19 aims to strike a balance between past performance represented by the Q value and the expected network capacity for a given channel selection. When the similarity is high, the selection will be primarily based on past behaviors to avoid group-think between all eNBs by using the experience-driven Q-value as the primary selection metric. When the similarity is low, the selection will be primarily based on the relative expected network capacity between selected channels, exploiting the shared information between eNBs.

4.2.5 SOFI — Stochastic Optimization with Full Information

To provide a reasonable upper bound for performance, we evaluate a fully-informed algorithm in which the channel selection for each eNB is performed jointly. This algorithm operates as an exhaustive search over all possible combinations assuming the decision is made for all jointly. While this is not possible in a distributed cellular network, it would be possible with centralized

base-band processing across all nodes in a cloud-RAN architecture. The SOFI algorithm operates as follows.

Let $\mathbf{k} \in \{1, \ldots, |\mathcal{K}|\}^{|\mathcal{X}|}$ be a vector where each element denotes the channel selected for the eNB with the corresponding index. This vector encapsulates the channel selection for all eNBs in the system, which is denoted by the set $\mathcal{X}$. The idea here is to choose the $\hat{\mathbf{k}}$ that maximizes the net capacity, i.e., we globally assign each eNB a channel instead of having them independently select their channel. This optimization is given by:

$$\mathbf{k} = \arg\max_{\hat{\mathbf{k}}} C_{\text{net}}(\hat{\mathbf{k}}) \tag{4.20}$$

To accomplish this, we calculate the resulting expected network capacity for all channel selection vectors $\mathbf{k}$. $P_a(j, k_j)$ is the probability that eNB j can access the channel assigned in the scenario (Eq. 4.3). $L_j \subseteq \mathcal{X}$ is the set of eNBs on the same channel as eNB j.

$$C_{\text{net}}(\mathbf{k}) = \sum_{j=1}^{|\mathcal{X}|} \frac{P_a(j, k_j)}{1 + |\mathcal{N}_j \cap L_j|} C_{j,\text{max}} \tag{4.21}$$

For a channel selection scenario, this equation sums the product of the probability that a node will access the channel, the probability it will win contention, and its max capacity if it were to win contention. The probability of winning contention is given by assuming that all neighboring nodes on the same channel are equally likely to win contention.

This global channel selection is done by explicitly calculating the expected network capacity over all possible vectors $\mathbf{k}$, which may be a large search space in cases when there are many channels and many nodes. This model is a reasonable approximation of an upper bound. However, it is not a true upper bound in the sense that we are unaware of future UE and Wi-Fi traffic.

4.3 Simulation Results

In this section, we compare the relative performance of each algorithm in a system level simulator we developed to model generalized uLTE operation. We first present details of our simulation environment. Using our simulator, we compare the performance of the five presented channel selection algorithms in a randomly generated topology with fixed parameters. We then evaluate algorithm performance for both high and low average similarity, S, across all deployed eNBs. We

further split the high average similarity result into low and high overall Wi-Fi utilization to see the impact of abundant and sparse transmit opportunity availability, respectively, for each algorithm. Finally, we examine the channel selection performance relative to topology scaling, including the number of eNBs, channels, and WiFi APs.

4.3.1 Simulation Environment

We test the algorithm in MATLAB where we design a simulator that randomly places eNBs and UEs, models interference using the ITU-INH channel model, and unique traffic being requested by UEs at each TTI according to a fixed traffic rate. TTIs can be equated to slots from our previous formulation, and have a duration of a single LTE subframe, 1 ms. Based on user positions, an appropriate MCS is chosen. If the UE see severe interference, the cyclic redundancy check (CRC) on their transmission may fail, prompting a re-transmission. We evaluate and update performance metrics upon successful packet reception, which can take several TTIs. In this simulator, we only model the downlink traffic using a 20 MHz, FDD LTE-based signal. As the scope of this work is focused on uLTE performance, the Wi-Fi traffic is modeled as transmissions that occur according to a regular traffic rate without modeling the individual Wi-Fi clients, their detailed MAC behavior, or their throughputs.

The optimization from Equation 4.12 is performed through a brute-force search. This is not unreasonable given the relatively small, discrete space that can explicitly be calculated. However, this space does grow exponentially with the number of available channels and the number of neighbors. In this case, heuristics or advanced optimization techniques can be explored, however such analysis is outside the scope of our current work.

To quantify performance, we use the user perceived throughput (UPT). The UPT is a metric of the throughput experienced by each individual user after accounting for the total time, including re-transmissions, taken for the complete packet to be received by the user. Using this metric has several beneficial consequences for analysis. UPT is significantly impacted by temporary outages in service. For our application, this amplifies the impact of poor channel selection, as the buffered data for transmission grows. From a UE perspective, UPT is also impacted by fairness and latency. This allows it to be used as a general quality of service metric, something unachievable when

comparing throughput alone. We formally define UPT in Equation 4.22.

$$UPT = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P_{total}}\left[\sum_{j=1}^{P_{served}}\frac{M \cdot r_{ij}}{t_{ij}} + \frac{b_i}{t_{serving,i}}\right] \tag{4.22}$$

Here, N is the number of UEs served by the eNB, and i indexes the UEs. P_{total} is the total number of packets, elaborated by $P_{total} = P_{served} + P_{serving}$, where P_{served} and $P_{serving}$ are the number of packets served and being served, respectively. M is the number of bits per packet, r_{ij} is the ratio of successfully transmitted bits over all bits in the packet to UE i for packet j, and t_{ij} is the time taken to send the same packet. b_i is the number of bits sent to UE i as a partial packet still in flight, and $t_{serving,i}$ is the time spent by the packet.

4.3.2 Overall Performance

Using our simulation environment, we first evaluate performance over many different random topologies. By doing so, we can gain an understanding of how each algorithm can be expected to perform in a generalized scenario. For this evaluation, we employ Monte Carlo simulation over many different “drops", or realizations of random simulation parameters. These drop-randomized parameters include eNB, Wi-Fi AP, and UE deployment locations, initial channel assignments, and data traffic realizations.

In Figure 4.4a, we show an example random realization of the network topology or “drop." Here, the four eNBs and eight Wi-Fi APs are uniformly randomly distributed in a 240×240 meter space. Each eNB has ten connected UEs uniformly distributed in a sixty meter radius around it.

For this simulation, we set the parameters as shown in Table 4.1 and then perform a Monte Carlo simulation. We simulate 100 unique drops with random topologies which are run for 5,000 TTI, or 5 seconds. The relatively short simulation duration functions as a snapshot of performance within the physical network topology, and traffic sources are unlikely to change significantly. Our further assumption is that the performance of these algorithms over this duration is extendable to multiple subsequent durations with incremental time-dependent changes to the topology. We plot the average UPT for each algorithm in Figure 4.5.

We find, over an extensive series of random topologies, an example of which is shown in Figure 4.4a, the Greedy algorithm has the worst performance, while RLPI performs the best. It is

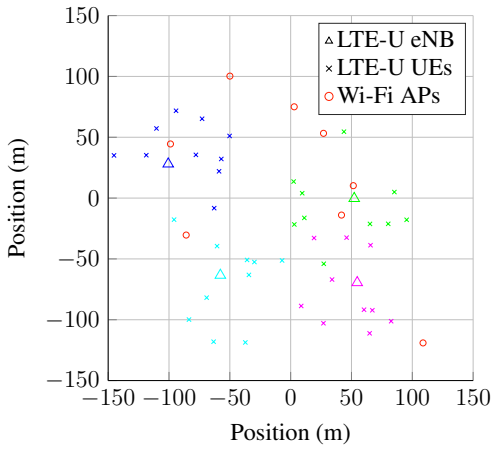

(a) Example realization of the random topology used in Monte Carlo simulation.

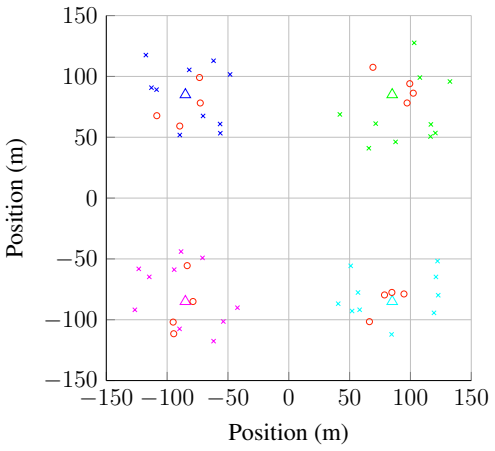

(b) Example realization of the fixed topology used for generating a low dependence scenario.

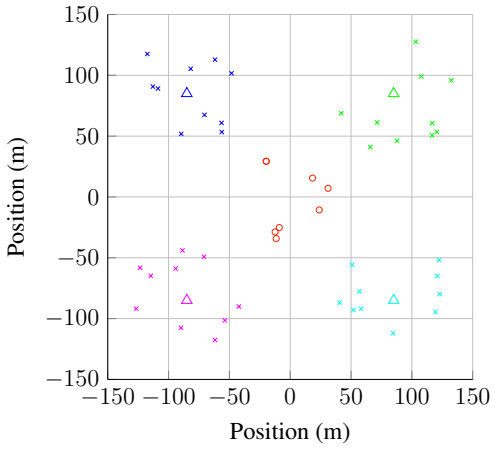

(c) Example realization of the fixed topology used for generating a high dependence scenario.

Figure 4.4: Example experimental topologies with four separate eNBs each with ten clients and eight Wi-Fi APs. The colors (blue, green, cyan, and magenta) represent an eNB and its clients. The eNBs and the Wi-Fi APs are uniformly randomly distributed in a room that is 240×240 meters. The UEs are distributed randomly in a 60-meter fixed radius around their eNB.

Table 4.1: Parameters used for overall performance simulation.

Parameter	Value in Simulation
Number of Channels, $\|\mathcal{K}\|$	4
Number of uLTE eNBs, $\|\mathcal{X}\|$	4
Number of Wi-Fi APs	8
Scan Period, M	20 TTI
Update Period, N	100 TTI
Deployment Area	240m x 240m
Drops	100
eNB TX Power	18 dBm
Wi-Fi AP TX Power	14 dBm
ED Threshold	-62 dBm

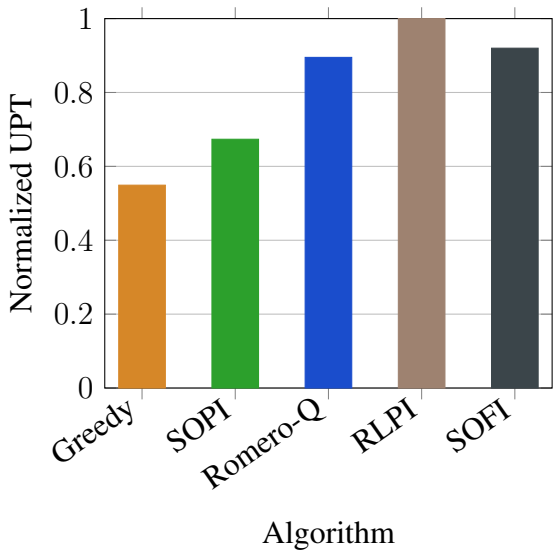

Figure 4.5: Overall algorithm performance, averaged across 100 random drops over 5,000 TTI to gauge the relative performance of each algorithm in generalized circumstances.

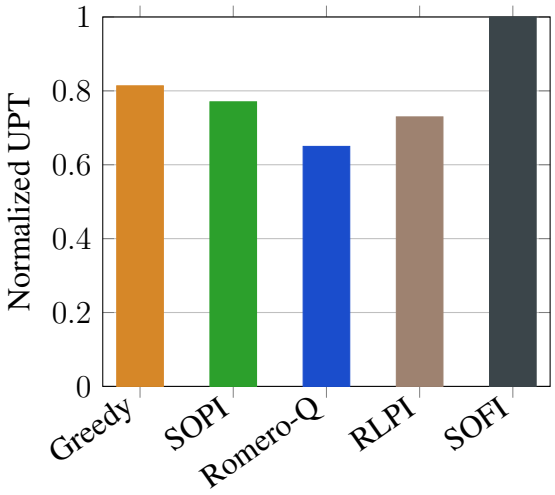

Figure 4.6: Algorithm performance in a low similarity scenario.

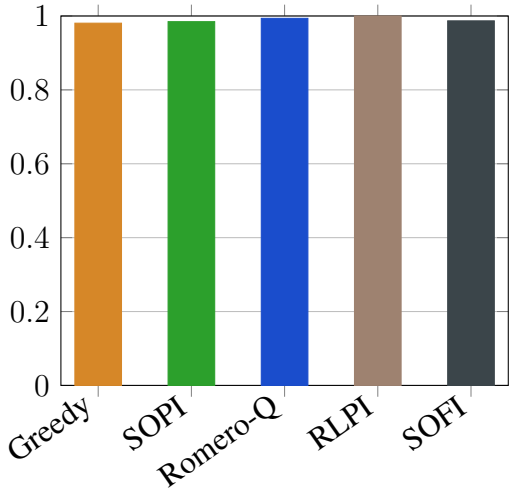

Figure 4.7: Algorithm performance in a high similarity scenario with low average Wi-Fi utilization.

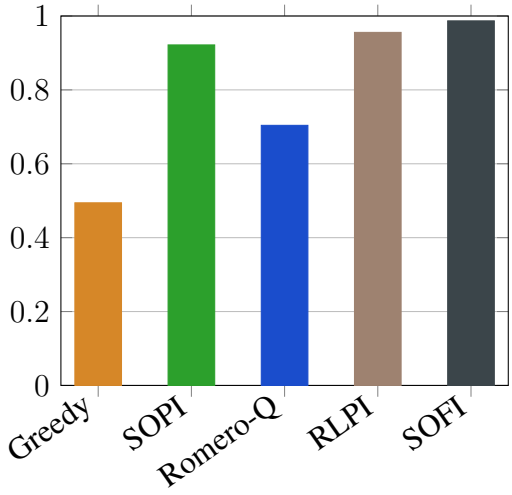

Figure 4.8: Algorithm performance in a high similarity scenario with high average Wi-Fi utilization: saturated channel conditions where improvements are critical

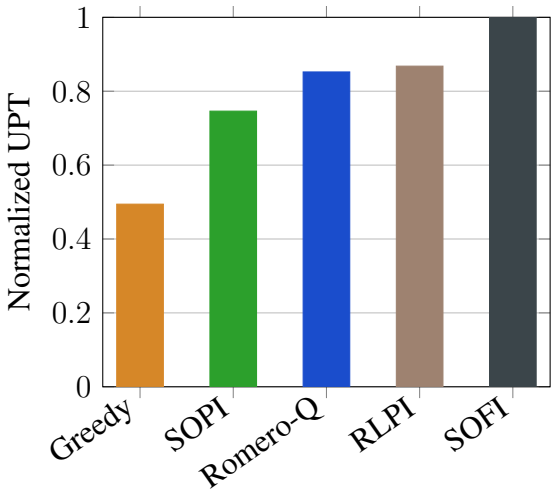

Figure 4.9: Algorithm performance in a high similarity scenario with independent channel access probabilities.

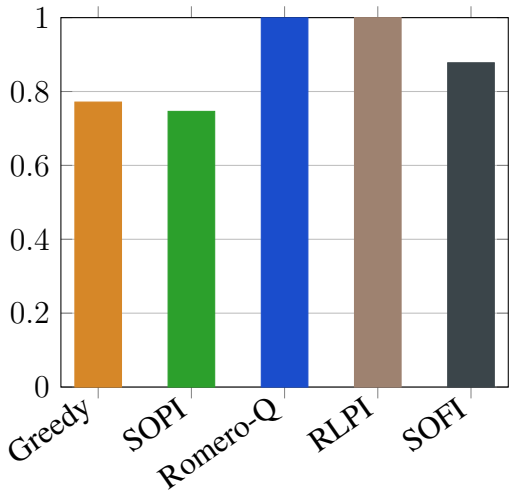

Figure 4.10: Algorithm performance in a high similarity scenario with fully dependent channel access probabilities.

interesting to note that both of the reinforcement algorithms outperform both SOPI and SOFI algorithms. One explanation for this is the partial and full optimization algorithms in SOPI and SOFI only considers channel selection with regards to channel and Wi-Fi utilization, while remaining oblivious to well known hidden factors such as exposed/hidden terminals between eNBs. Reinforcement learning performs exceptionally well in environments with hidden influences, allowing it to out-perform others in the general case. However, these factors alone may not account for such a discrepancy in performance over a large number of simulation iterations. In the next set of experiments, we isolate the utilization similarity metric, dependence between channel access probabilities, and Wi-Fi traffic load to better explain the underlying causes of performance differences between these algorithms.

4.3.3 Utilization Similarity, Dependence, and Load

For our RLPI algorithm, we introduced a similarity metric, S, to balance the Q-value impact and expected network capacity when calculating channel selection probabilities. This similarity metric can be thought of as the relative difference in channel availability an eNB and its neighbors. In other words, a high similarity indicates the Wi-Fi occupancy is roughly the same between an eNB and its neighbors for each channel, and a low similarity indicates the opposite.

However, we can also think of similarity as a metric of how correlated the channel access probabilities between each eNB. For example, if two eNBs share a common Wi-Fi AP nearby on channel k, the access probabilities for the two eNBs on channel k as defined in Equation 4.3 are not strictly independent, distorting the calculation for capture probabilities in Equation 4.4. This interaction between instantaneous channel availability correlation and similarity is not straightforward. For example, while a high dependence between channel access probabilities implies a high similarity metric, a high similarity can be realized with completely independent channel access probabilities between neighbors. This relationship strictly depends on whether an eNB neighborhood shares common Wi-Fi APs (high dependence) or if they do not share APs, but the utilization happens to be the same (low dependence). Unfortunately, the eNBs and, by extension, the channel selection algorithm they employ are unable to differentiate the access probability dependence and must rely on the similarity metric alone.

Thus, we need to explicitly quantify the impact of similarity on the performance for each

algorithm. To do so, we first define an average similarity metric, S_{avg} as the average similarity for a simulation as:

$$S_{\text{avg}} = \sum_{i \in \mathcal{X}} \frac{S_i}{|\mathcal{X}|} \tag{4.23}$$

For our analysis, we categorize similarities into two extreme scenarios: low average similarity, $S_{\text{avg}} < 0.5$, and high average similarity, $S_{\text{avg}} > 0.9$. The results for similarity values in-between the high and low can be considered the general case, which represents the majority of drop outcomes in the Monte Carlo simulation. To elucidate these thresholds, we have designed representative scenarios for each.

Low similarity To properly analyze low similarity, we modify our simulation environment to provide a more deterministic topology for study. We fix the positions of 4 eNBs in a square pattern such that each eNB is considered a neighbor to the two closest eNBs, creating a unique neighbor set for each eNB. Four Wi-Fi APs are distributed near each eNB such that other eNBs are not impacted. This creates completely independent Wi-Fi interference sources for each of the eNBs. An example deployment under this fixed topology constraint is shown in Figure 4.4b.

To achieve the desired similarity for simulation, we randomly select channels and generate traffic for the Wi-Fi APs constrained such that $S_{\text{avg}} < 0.5$. We show the resulting algorithm performance for the low similarity scenario in Figure 4.6.

We find that in a low similarity scenario, the greedy selection algorithm and SOFI end up performing particularly well. This is because a low similarity metric is generally only achievable when the Wi-Fi channel occupancy tends to be vastly different between neighboring eNBs across channels and the channel with the lowest Wi-Fi occupancy in each eNB will tend to be different from its neighbors. In such a situation, the best channel for each eNB ends up being the one with the lowest measured Wi-Fi occupancy locally, which is ideal for the greedy algorithm and easily discovered by SOFI as well. However, SOPI ends up under-performing due to the probabilistic selection, since SOPI lacks sufficient information to guarantee neighboring eNBs will select the channel with the most local availability. Romero-Q's performance is below greedy simply because of the learning ramp-up.

High similarity For the high similarity scenario, we modify the fixed topology from the low similarity analysis such that all Wi-Fi APs are deployed in a small radius around the environment center. This placement results in a high dependence between channel access probabilities, as each Wi-Fi AP placed has a high probability of impacting multiple eNBs. We reduce the total number of Wi-Fi APs to better control average channel utilization while still randomizing traffic and Wi-Fi AP channel assignments. An example deployment under this fixed topology constraint is shown in Figure 4.4c.

Additionally, we simulate the high similarity scenario with two different sub-conditions: low and high average Wi-Fi utilization. We label the simulation realization as having a low average Wi-Fi utilization when the average utilization experienced at each eNB over all channels is under 10%. Conversely, we label the simulation realization as having a high Wi-Fi utilization when the average utilization is over 50%. We show the resulting algorithm performance for both low and high Wi-Fi utilization settings in Figures 4.7 and 4.8.

When the Wi-Fi utilization is not a factor due to low occupancy, none of the eNBs have trouble offloading all their traffic. This is true even if there is some frequency overlap between them. This situation can be considered an abundance of spectral resources, and thus the UPT is similarly high across all channel selection algorithms. However, as licensed cellular providers and a myriad of new Internet of Things (IoT) devices pop up in unlicensed spectrum, the more likely (and important) situation is one in which the spectrum is saturated.

When spectral resources are scarce in the high Wi-Fi utilization case, proper channel selection becomes vital to achieving serviceable performance. In this situation, slight differences in channel availability between channels can have a significant impact on expected capacity. The SOPI, RLPI, and SOFI algorithms can evaluate and exploit these differences to great effect. Our results show that these algorithms significantly outperform Romero-Q and Greedy, demonstrating one of the major benefits of sharing the Wi-Fi utilization information between eNBs.

When considering the impact of Wi-Fi utilization similarities between nodes, we can further differentiate between those that are similar due to interference from the same Wi-Fi APs and those that have different Wi-Fi APs, but similar traffic statistics. The impact of these differences is most noticeable in the channel access probability and the resulting channel capacity calculation of the SOPI algorithm in Equations 4.4 and 4.5. The presented equation assumes independent channel

access probabilities between each eNB and its neighbors, which is achievable when different Wi-Fi APs within the sensing range of each eNB. However, when the Wi-Fi APs are shared, the neighboring channel access probabilities should instead be conditioned on i. This error leads to a differential in capture probability and subsequent capacity calculation depending on the dependence between neighboring eNB access probabilities. The impact of this error can be seen in Figures 4.9 and 4.10, where we compare the performance of each channel access algorithm given fully dependent or fully independent channel access probabilities between eNBs. We generate the independent access probability scenario using the topology shown in Figure 4.4b and dependent access probability scenario using the topology in Figure 4.4c. Both scenarios are crafted with random Wi-Fi traffic generation.

In the fully dependent scenario, the SOPI and SOFI algorithms overestimate their channel capture probabilities due to the assumption of independence and end up being overshadowed by the reinforcement based algorithms. Conversely, the SOFI algorithm ends up outperforming all other algorithms when the access probabilities are independent, and the joint optimization between eNBs can accurately calculate expected capacities.

Overall, we see that while the context-reliant algorithms are all impacted by the access probability dependencies, the reinforcement learning algorithms remain unaffected. Alternatively, context-reliant algorithms perform better when spectral resources are scarce. Our proposed RLPI algorithm balances these two approaches and achieves high performance across each of the scenarios studied, and results in improved performance in the general case as well.

4.3.4 Scalability

In this series of experiments, we examine the scalability by testing each algorithm for various numbers of channels, eNBs, and Wi-Fi APs.

4.3.4.1 Number of Channels

We consider the cases of 2, 4, 6, and 8 unlicensed channels available in the band. The results are plotted in Figure 4.11a. We average the UPT across all the connected UEs in a drop. Across all ten drops, we plot the average, median, and fifth-percentile for the UPT seen versus the number of unlicensed channels available in the system. For scenarios with a fewer number of channels, the UPT is lower as the many eNBs, UEs, and Wi-Fi APs all must contend across the available chan-

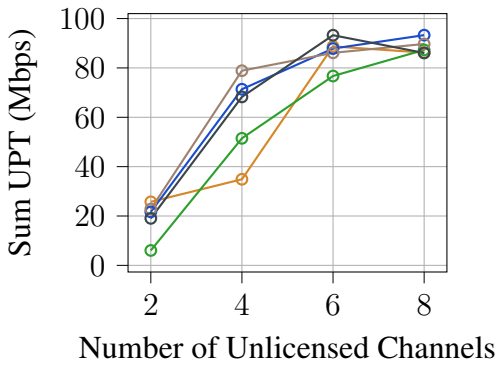

(a) Varying the number of unlicensed channels while the number of Wi-Fi nodes is fixed to 8 and the number of eNBs is fixed to 4.

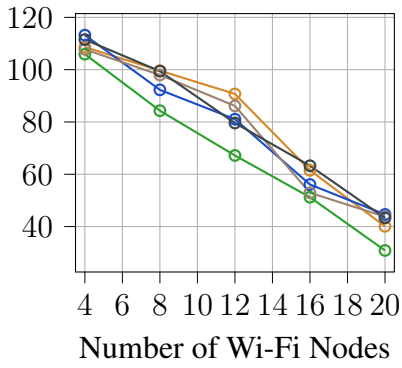

(b) Varying the number of Wi-Fi nodes while the number of channels is fixed to 4 and the number of eNBs is fixed to 4.

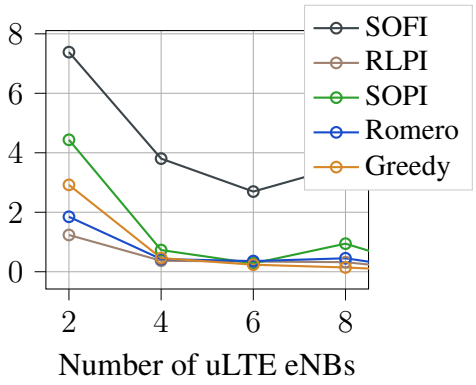

(c) Varying the number of uLTE eNBs while the number of of channels is fixed to 4 and the number of Wi-Fi nodes is fixed to 8.

Figure 4.11: Scalability experiments. For each algorithm, we test performance as we vary the number of unlicensed channels, Wi-Fi nodes, and uLTE eNBs.

nels. As the number of channels increases, we see that the UPT increases as the traffic is able to be divided across more channels leading to fewer collisions. However, the greedy algorithm consistently is less able to take advantage of the new channels as eNBs that measure similar utilizations will likely collide. SOPI is able to perform better than greedy throughout. The two reinforcement learning algorithms behave similarly throughout the experiment. SOFI also performs comparably to the Q-learning is able to

4.3.4.2 Number of Wi-Fi APs

We consider the case of 4, 8, 12, 16, and 20 Wi-Fi APs operating across 4 channels. The results are plotted in Figure 4.11b. For this result we see that UPT consistently falls as more Wi-Fi APs are added to the topology.

4.3.4.3 Number of uLTE eNBs

We consider the case of 2, 4, 6, and 8 eNBs operating across 4 channels with 16 Wi-Fi APs. The results are plotted in Figure 4.11c. For this result we see that UPT consistently falls as more eNBs are added to the topology.

4.4 Related Work

Many works in the literature are concerned with the general topic of LTE on unlicensed spectrum [1, 3, 43–50, 53–59]. However, comparatively few are concerned with channel selection or LTE to LTE coexistence on unlicensed spectrum. In this section, we discuss other works that investigate this problem.

In [1] and [56], the authors focus on channel selection and propose a Q-learning, reinforcement algorithm with [1] originally formulating the algorithm and [56] expanding on it for non-stationary environments. However, their Q-learning is based entirely on past throughput that each small cell achieves on the channel, neglecting the utilization of the channel. Moreover, throughput is realistically a function of more than interference levels, i.e., traffic rates, MCS, number of users, etc. This could cause the Q-Learning to take unnecessary actions depending on the instantaneous status of the other settings, especially in a dynamic environment. It will also take a long time to explore the entire Q-table and will not be able to adapt to rapid fluctuations in utilizations across the channels. Moreover, the authors do not place Wi-Fi devices in any of their simulations. The

authors later also considered a game theory based approach and found it to converge faster than their Q-learning method [55]. One notable strength in all of their work is that their methods are entirely decentralized with each small cell making decisions based on their own throughput.

An alternative approach is taken in [57] where the authors instead consider a joint optimization of the channel selection and the frame scheduling to improve coexistence in the frequency and time domains. Fairness to Wi-Fi devices is included in the optimization as a constraint while maximizing LTE throughput. However, this work relies on a complicated cloud radio access network (RAN) scenario which would be unlikely to account for multiple operators.

Deep reinforcement learning was used in [58]. This work is also entirely distributed. Although their results are promising, the training data may be prohibitive as multiple days of Wi-Fi utilizations and neighboring eNBs traffic loads are required. Moreover, online training to continuously update traffic models may be too computationally complex to do for each small cell in a network. The Wi-Fi loads dataset used in this work for training and evaluating their method is from a 2003 dataset that polls for AP activity every five minutes. It is unclear how their method would perform in modern and likely less predictable traffic patterns where there are many more Wi-Fi APs and clients in a region.

4.5 Conclusions

In this chapter, we examine uLTE coexistence with other uLTE eNBs in the presence of spatially heterogeneous Wi-Fi utilization. While other similar works consider uniform Wi-Fi behavior or ignore Wi-Fi presence entirely, our work uses knowledge of Wi-Fi incumbency to outperform existing solutions. To do so, we first formulate a framework under which neighboring eNBs can share Wi-Fi utilization estimates with each other. Next, we present a unique algorithm in which deployed eNBs exploit the shared utilization estimates to select an optimal transmission channel. We then combine this analytical model with state-of-the-art reinforcement learning techniques and present a novel reinforcement learning approach that leverages the shared utilization statistics. Finally, we simulate five relevant channel selection algorithms, showing significant improvements in performance in both general random deployments and scenarios with high spectral congestion.

For future work, we would like to study the convergence characteristics of our presented reinforcement learning algorithm, especially in the presence of dynamic changes to the simulation

environment.

Chapter 5

Future Research

Wireless spectrum sharing is a considerably large topic that changes drastically over time with new wireless technology generations, deployments, and use-cases. In this dissertation, we have primarily focused on simulation-based scenarios in regards to unlicensed spectrum access between similar technologies.

One future direction would be to test these algorithms in real systems, in which both technology and user behavior will likely diverge from presumed models. Currently, live CBRS deployments are up in several locations around the United States and could potentially be used to implement and fine-tune using buffer occupancy statistics to better adjust GAA behavior. Similarly, private uLTE nodes operating with LTE-U or LAA could be used in evaluating the channel selection algorithms presented in this work. Alternatively, software defined radios such as the ones we used in the CBRS evaluation can be utilized to build up larger-scale test networks for a more controlled evaluation using real hardware.

Another direction would be to expand and refine existing simulation models to include more sophisticated wireless techniques such as directional beamforming, MIMO configurations, and dynamic mobility models. Each of these additions would add a new dimension to the evaluation of these algorithms, and brings them closer to behavior in a real deployment.

Chapter 6

Conclusion

In this dissertation, three distinct works were presented in the area of spectrum sharing in cellular networks, with the aim of moving the state of the art closer to realizing efficient spectrum usage throughout sub-6 GHz frequency bands.

First, we used geographical features of a region to reduce in-field propagation experimentation by predicting the number of measurements required to accurately characterize its path loss, which can then be used to model coverage of arbitrarily positioned base stations. By exploiting the relationship between terrain feature complexity and measurement requirements, we found that the number of measurements collected to achieve a certain path loss accuracy over the entire region can be reduced by up to 58% in a high density drive testing scenario.

Next, we looked at applying LBT schemes in CBRS networks for increasing the spatial reuse at secondary users while minimizing the interference footprint on incumbent and primary users. We used a novel Q-learning scheme to adapt the contention EDT to the changing network topology and traffic conditions, providing up to 350% gains in average secondary node UPT in certain difficult topologies with merely a 4% reduction in primary node UPT.

Finally, we studied channel selection in unlicensed LTE systems, proposing a mechanism for unlicensed eNBs to share statistical channel occupancy information with neighboring cells. By collecting channel utilization statistics and sharing this information periodically with other unlicensed LTE eNBs, each eNB can improve their channel selection given their limited knowledge of the full topology via a proposed statistical and machine learning approach. We simulate operation in the unlicensed band using our channel selection algorithm and show how Wi-Fi load and inter-cell interference estimation can jointly be used to select transmission channels for all eNBs in the network.

BIBLIOGRAPHY

[1] O. Sallent, J. Pérez-Romero, R. Ferrús, and R. Agustí, "Learning-based coexistence for lte operation in unlicensed bands," in *2015 IEEE International Conference on Communication Workshop (ICCW)*, June 2015, pp. 2307–2313. ix, 77, 78, 79, 92

[2] *FCC 15-47*. Federal Communications Commission, Apr. 2015. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-47A1.pdf x, 29, 30

[3] 3GPP, "Study on Licensed-Assisted Access to Unlicensed Spectrum," 3GPP, Technical Report (TR) 36.889, June 2015, version 13.0.0. [Online]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2579 xi, 36, 37, 68, 92

[4] H. T. Friis, "A note on a simple transmission formula," *proc. IRE*, vol. 34, no. 5, pp. 254–256, 1946. 8

[5] B. Sklar, "Rayleigh fading channels in mobile digital communication systems. i. characterization," *Communications Magazine, IEEE*, vol. 35, no. 7, pp. 90–100, 1997. 8

[6] M. Hata, "Empirical formula for propagation loss in land mobile radio services," *Vehicular Technology, IEEE Transactions on*, vol. 29, no. 3, pp. 317–325, 1980. 8, 27

[7] C. Phillips, D. Sicker, and D. Grunwald, "A survey of wireless path loss prediction and coverage mapping methods," *Communications Surveys & Tutorials, IEEE*, vol. 15, no. 1, pp. 255–270, 2013. 8

[8] W. C. Lee, "Estimate of local average power of a mobile radio signal," *Vehicular Technology, IEEE Transactions on*, vol. 34, no. 1, pp. 22–27, 1985. 8

[9] H. Shin, "Measurements and models of 802.11 b signal strength variation over small distances," Ph.D. dissertation, University of Delaware, 2010. 8

[10] J. Robinson, R. Swaminathan, and E. W. Knightly, "Assessment of urban-scale wireless networks with a small number of measurements," in *Proc. of ACM MobiCom*, 2008. 8, 27

[11] Google, "Android api description," 2017. [Online]. Available: https://developer.android.com/reference/android/telephony/CellSignalStrength.html 11

[12] "Open street maps," www.openstreetmap.org, license: http://www.openstreetmap.org/copyright. 11

[13] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4g lte networks," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 2012, pp. 225–238. 26

[14] S. Deng, R. Netravali, A. Sivaraman, and H. Balakrishnan, "Wifi, lte, or both?: Measuring multi-homed wireless internet performance," in *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 2014, pp. 181–194. 26

[15] E. Neidhardt, A. Uzun, U. Bareth, and A. Kupper, "Estimating locations and coverage areas of mobile network cells based on crowdsourced data," in *Wireless and Mobile Networking Conference (WMNC), 2013 6th Joint IFIP*. IEEE, 2013, pp. 1–8. 27

[16] Y. Okumura, E. Ohmori, T. Kawano, and K. Fukuda, "Field strength and its variability in VHF and UHF land-mobile radio service," *Rev. Elec. Commun. Lab*, vol. 16, no. 9, pp. 825–73, 1968. 27

[17] K. Allsebrook and J. Parsons, "Mobile radio propagation in british cities at frequencies in the VHF and UHF bands," *Vehicular Technology, IEEE Transactions on*, vol. 26, no. 4, pp. 313–323, 1977. 27

[18] M. Akimoto, T. Shimizu, and M. Nakatsugawa, "Path loss estimation of 2 GHz and 5 GHz band FWA within 20 km in rural area," *Proc. of ISAP*, 2006. 27

[19] S. Ghassemzadeh, H. Worstell, and R. Miller, "Wireless neighborhood area network path loss characterization at 5.7 GHz," in *Proc. of IEEE VTC-Fall*, 2010. 27

[20] B. Hanci and I. Cavdar, "Mobile radio propagation measurements and tuning the path loss model in urban areas at GSM-900 band in istanbul-turkey," in *Proc. of IEEE VTC-Fall*, 2004. 27

[21] B. Sayrac, J. Riihijärvi, P. Mähönen, S. Ben Jemaa, E. Moulines, and S. Grimoud, "Improving coverage estimation for cellular networks with spatial bayesian prediction based on measurements," in *Proceedings of the 2012 ACM SIGCOMM workshop on Cellular networks: operations, challenges, and future design*. ACM, 2012, pp. 43–48. 27

[22] H. Braham, S. B. Jemaa, G. Fort, E. Moulines, and B. Sayrac, "Fixed rank kriging for cellular coverage analysis," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 5, pp. 4212–4222, 2017. 27

[23] A. Galindo-Serrano, B. Sayrac, S. B. Jemaa, J. Riihijärvi, and P. Mähönen, "Automated coverage hole detection for cellular networks using radio environment maps," in *Modeling & Optimization in Mobile, Ad Hoc & Wireless Networks (WiOpt), 2013 11th International Symposium on*. IEEE, 2013, pp. 35–40. 27

[24] H.-W. Liang, C.-H. Ho, L.-S. Chen, W.-H. Chung, S.-Y. Yuan, and S.-Y. Kuo, "Coverage hole detection in cellular networks with deterministic propagation model," in *Intelligent Green Building and Smart Grid (IGBSG), 2016 2nd International Conference on*. IEEE, 2016, pp. 1–6. 27

[25] D. Flore, "Laa standardization: coexistence is the key," 3GPP, July 2016. [Online]. Available: http://www.3gpp.org/news-events/3gpp-news/1789-laa_update 29

[26] M. Chmaytelli, "LTE-U Forum: Ensuring LTE and Wi-Fi fairly coexist in unlicensed spectrum," Qualcomm, Mar. 2016. [Online]. Available: https://www.qualcomm.com/news/onq/2015/03/03/lte-u-forum-ensuring-lte-and-wi-fi-fairly-coexist-unlicensed-spectrum 29

[27] "MulteFire Release 1.0 Technical Paper," MulteFire Alliance, July 2017. [Online]. Available: https://www.multefire.org/wp-content/uploads/MulteFire-Release-1.0-whitepaper_FINAL.pdf 29

[28] *FCC 18-149*. Federal Communications Commission, Oct. 2018. [Online]. Available: https://docs.fcc.gov/public/attachments/FCC-18-149A1.pdf 30, 40

[29] J. Nie and S. Haykin, "A Q-learning-based Dynamic Channel Assignment Technique for Mobile Communication Systems," *IEEE Trans. Veh. Technol.*, vol. 48, no. 5, pp. 1676–1687, Sept. 1999. 31

[30] A. Galindo-Serrano and L. Giupponi, "Distributed Q-Learning for Aggregated Interference Control in Cognitive Radio Networks," *IEEE Trans. Veh. Technol.*, vol. 59, no. 4, pp. 1823–1834, May 2010. 31

[31] C. Chen, R. Ratasuk, and A. Ghosh, "Downlink Performance Analysis of LTE and WiFi Coexistence in Unlicensed Bands with a Simple Listen-Before-Talk Scheme," in *2015 IEEE 81st Veh. Technol. Conf.*, May 2015, pp. 1–5. 31

[32] N. Rupasinghe and I. Güvenç, "Licensed-assisted access for WiFi-LTE coexistence in the unlicensed spectrum," in *2014 IEEE GC Wkshps*, Dec. 2014, pp. 894–899. 31

[33] H. J. Kwon *et al.*, "Licensed-Assisted Access to Unlicensed Spectrum in LTE Release 13," *IEEE Commun. Mag.*, vol. 55, no. 2, pp. 201–207, Feb. 2017. 31, 33

[34] B. Jia and M. Tao, "A Channel Sensing Based Design for LTE in Unlicensed Bands," in *2015 IEEE International Conference on Communication Workshop (ICCW)*, June 2015, pp. 2332–2337. 31, 33

[35] N. Rupasinghe and I. Güvenç, "Reinforcement learning for licensed-assisted access of LTE in the unlicensed spectrum," in *Proc. of 2015 IEEE Wireless Commun. and Netw. Conf.*, Mar. 2015, pp. 1279–1284. 31

[36] A. Galanopoulos, F. Foukalas, and T. A. Tsiftsis, "Efficient Coexistence of LTE With WiFi in the Licensed and Unlicensed Spectrum Aggregation," *IEEE Trans. on Cogn. Commun. Netw.*, vol. 2, no. 2, pp. 129–140, June 2016. 31

[37] M. Palola *et al.*, "Field trial of the 3.5 GHz citizens broadband radio service governed by a spectrum access system (SAS)," in *Proc. of DySPAN 2017*, Mar. 2017, pp. 1–9. 32

[38] "You don't need high grade Navy radar systems to spot which companies just achieved another industry milestone for customers," May 2018. [Online]. Available: http://www.verizon.com/about/news/you-dont-need-high-grade-navy-radar-systems-spot-which-companies-just-achieved-another 32

[39] A. Ghosh, R. Ratasuk, B. Mondal, N. Mangalvedhe, and T. Thomas, "LTE-advanced: next-generation wireless broadband technology [Invited Paper]," *IEEE Wireless Commun.*, vol. 17, no. 3, pp. 10–22, June 2010. 34

[40] L. Wang, K. Wu, and M. Hamdi, "Combating Hidden and Exposed Terminal Problems in Wireless Networks," *IEEE Trans. Wireless Commun.*, vol. 11, no. 11, pp. 4204–4213, Nov. 2012. 45

[41] C. J. C. H. Watkins, "Learning from Delayed Rewards," Ph.D. dissertation, 1989. 47

[42] *CBRS Coexistence Technical Specification*, CBRS Alliance, Feb. 2018. [Online]. Available: https://www.cbrsalliance.org/wp-content/uploads/2018/04/CBRS-Coexistence-Technical-Specification.pdf 49

[43] Qualcomm Incorporated, "Extending LTE Advanced to unlicensed spectrum," Tech. Rep., Dec 2013. 68, 92

[44] T. Novlan, B. L. Ng, H. Si, and J. C. Zhang, "Overview and Evaluation of Licensed Assisted Access for LTE-Advanced," in *2015 49th Asilomar Conference on Signals, Systems and Computers*, Nov 2015, pp. 1031–1035. 68, 92

[45] T-Mobile, "T-Mobile Completes Nation's First Live Commercial Network Test of License Assisted Access (LAA)," Jun 2017. [Online]. Available: https://www.t-mobile.com/news/lte-u 68, 92

[46] K. Hill, "GSA Sees Momentum for CBRS Market and use of LTE in Unlicensed Spectrum," May 2019. [Online]. Available: https://www.rcrwireless.com/20190509/network-infrastructure/gsa-sees-momentum-for-cbrs-market-and-use-of-lte-in-unlicensed-spectrum 68, 92

[47] M. Haider and M. Erol-Kantarci, "Enhanced LBT Mechanism for LTE-Unlicensed Using Reinforcement Learning," in *2018 IEEE Canadian Conference on Electrical Computer Engineering (CCECE)*, May 2018, pp. 1–4. 68, 92

[48] Qualcomm Technologies Inc., "Qualcomm Research LTE in Unlicensed Spectrum: Harmonious Coexistence with Wi-Fi," Jun 2014. [Online]. Available: https://www.qualcomm.com/media/documents/files/lte-unlicensed-coexistence-whitepaper.pdf 73, 92

[49] H. Li, Y. Chang, F. Hao, A. Men, J. Zhang, and W. Quan, "Study on dynamic channel switch in license-assisted-access based on listen-before-talk," in *2016 International Symposium on Wireless Communication Systems (ISWCS)*. IEEE, 2016, pp. 506–510. 73, 92

[50] S. F. Thomas Cheng and D. Larsson, "Licensed Assisted Access: Practical Coexistence Solutions," Feb 2015. [Online]. Available: https://www.ericsson.com/en/blog/2015/2/licensed-assisted-access-practical-coexistence-solutions 73, 92

[51] R. S. Sutton, A. G. Barto *et al.*, *Introduction to reinforcement learning*. MIT press Cambridge, 1998, vol. 135. 77, 79

[52] S.-H. Cha, "Comprehensive survey on distance/similarity measures between probability density functions," *City*, vol. 1, no. 2, p. 1, 2007. 79

[53] *LTE-U CSAT Procedure TS V1.0*, LTE-U Forum, 10 2015. 92

[54] 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 *2013 IEEE International Conference on Communications (ICC)*, June 2013, pp. 5083–5088. 92

[55] A. Castañé, J. Pérez-Romero, and O. Sallent, "On the Implementation of Channel Selection for LTE in Unlicensed Bands using Q-learning and Game Theory Algorithms," in *2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC)*, June 2017, pp. 1096–1101. 92, 93

[56] J. Perez-Romero, O. Sallent, R. Ferrus, and R. Agusti, "A robustness analysis of learning-based coexistence mechanisms for lte-u operation in non-stationary conditions," in *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*, Sep. 2015, pp. 1–5. 92

[57] H. Ko, J. Lee, and S. Pack, "Joint optimization of channel selection and frame scheduling for coexistence of lte and wlan," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 7, pp. 6481–6491, July 2018. 92, 93

[58] U. Challita, L. Dong, and W. Saad, "Proactive resource management for lte in unlicensed spectrum: A deep learning perspective," *IEEE Transactions on Wireless Communications*, vol. 17, no. 7, pp. 4674–4689, July 2018. 92, 93

[59] F. Liu, E. Bala, E. Erkip, M. C. Beluri, and R. Yang, "Small-cell traffic balancing over licensed and unlicensed bands," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 12, pp. 5850–5865, Dec 2015. 92

[60] A. Balasubramanian, R. Mahajan, and A. Venkataramani, "Augmenting mobile 3G using WiFi," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 209–222.

[61] B. Han, P. Hui, and A. Srinivasan, "Mobile data offloading in metropolitan area networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 14, no. 4, pp. 28–30, 2011.

[62] S. Wiethölter, M. Emmelmann, R. Andersson, and A. Wolisz, "Performance evaluation of selection schemes for offloading traffic to ieee 802.11 hotspots," in *2012 IEEE International Conference on Communications (ICC)*. IEEE, 2012, pp. 5423–5428.

[63] S. Dimatteo, P. Hui, B. Han, and V. O. Li, "Cellular traffic offloading through wifi networks," in *2011 IEEE Eighth International Conference on Mobile Ad-Hoc and Sensor Systems*. IEEE, 2011, pp. 192–201.

[64] N. Rupasinghe and I. Güvenç, "Licensed-assisted access for WiFi-LTE coexistence in the unlicensed spectrum," in *2014 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2014, pp. 894–899.

[65] R. Ratasuk, N. Mangalvedhe, and A. Ghosh, "LTE in unlicensed spectrum using licensed-assisted access," in *2014 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2014, pp. 746–751.

[66] A. Mukherjee, J.-F. Cheng, S. Falahati, L. Falconetti, A. Furuskär, B. Godana, H. Koorapaty, D. Larsson, Y. Yang *et al.*, "System architecture and coexistence evaluation of licensed-assisted access LTE with ieee 802.11," in *2015 IEEE International Conference on Communication Workshop (ICCW)*. IEEE, 2015, pp. 2350–2355.

[67] N. Rupasinghe and İ. Güvenç, "Reinforcement learning for licensed-assisted access of LTE in the unlicensed spectrum," in *2015 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2015, pp. 1279–1284.

[68] Q. Chen, G. Yu, R. Yin, A. Maaref, G. Y. Li, and A. Huang, "Energy efficiency optimization in licensed-assisted access," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 4, pp. 723–734, 2016.

[69] M. Bennis, M. Simsek, A. Czylwik, W. Saad, S. Valentin, and M. Debbah, "When cellular meets wifi in wireless small cell networks," *IEEE Communications Magazine*, vol. 51, no. 6, pp. 44–50, 2013.

[70] T. Novlan, B. L. Ng, H. Si, and J. C. Zhang, "Overview and evaluation of licensed assisted access for LTE-advanced," in *2015 49th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2015, pp. 1031–1035.

[71] M. Gerasimenko, N. Himayat, S.-p. Yeh, S. Talwar, S. Andreev, and Y. Koucheryavy, "Characterizing performance of load-aware network selection in multi-radio (WiFi/LTE) heterogeneous networks," in *2013 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2013, pp. 397–402.

[72] Q. Chen, G. Yu, H. Shan, A. Maaref, G. Y. Li, and A. Huang, "Cellular meets WiFi: Traffic offloading or resource sharing?" *IEEE Transactions on Wireless Communications*, vol. 15, no. 5, pp. 3354–3367, 2016.

[73] S. Yun and L. Qiu, "Supporting WiFi and LTE co-existence," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 810–818.

[74] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, "Mobile data offloading: how much can WiFi deliver?" in *Proceedings of the 6th International COnference*. ACM, 2010, p. 26.

[75] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[76] I. . W. Group *et al.*, "Ieee 802.11-2007: Wireless lan medium access control (mac) and physical layer (phy) specifications," *IEEE 802.11 LAN Standards 2007*, 2007.

[77] T. S. Rappaport, *Wireless communications: principles and practice*, 2nd ed. Prentice Hall PTR New Jersey, 2001.

[78] "Study on licensed-assisted access using LTE," *3GPP Study Item: RP-141397*, 2014.

[79] F. Liu, E. Bala, E. Erkip, M. C. Beluri, and R. Yang, "Small-cell traffic balancing over licensed and unlicensed bands," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 12, pp. 5850–5865, 2015.

[80] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking (ToN)*, vol. 8, no. 5, pp. 556–567, 2000.

[81] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *New frontiers in dynamic spectrum access networks, 2008. DySPAN 2008. 3rd IEEE symposium on*. IEEE, 2008, pp. 1–11.

[82] X. Li and S. A. Zekavat, "Traffic pattern prediction and performance investigation for cognitive radio systems," in *2008 IEEE Wireless Communications and Networking Conference*. IEEE, 2008, pp. 894–899.

[83] "Extended paper," https://www.dropbox.com/sh/okgikrtekkug41e/AACD-TPk85wrwGFQIC3qzj2aa?dl=0.

[84] J. Johansson, W. Hapsari, S. Kelley, G. Bodog *et al.*, "Minimization of drive tests in 3gpp release 11," *Communications Magazine, IEEE*, vol. 50, no. 11, pp. 36–43, 2012.

[85] K. Hodges, "Crowdsourcing early warnings of natural disasters," *Science*, vol. 348, no. 6231, pp. 196–198, 2015.

[86] M. Demirbas, M. A. Bayir, C. G. Akcora, Y. S. Yilmaz, and H. Ferhatosmanoglu, "Crowd-sourced sensing and collaboration using twitter," in *World of Wireless Mobile and Multimedia Networks (WoWMoM), 2010 IEEE International Symposium on a*. IEEE, 2010, pp. 1–9.

[87] S. Hu, L. Su, H. Liu, H. Wang, and T. F. Abdelzaher, "Smartroad: a crowd-sourced traffic regulator detection and identification system," in *Information Processing in Sensor Networks (IPSN), 2013 ACM/IEEE International Conference on*. IEEE, 2013, pp. 331–332.

[88] "Smartphone users worldwide will total 1.75 billion in 2014," http://www.emarketer.com/Article/Smartphone-Users-Worldwide-Will-Total-175-Billion-2014/1010536.

[89] J. De Vriendt, P. Laine, C. Lerouge, and X. Xu, "Mobile network evolution: a revolution on the move," *IEEE Communications magazine*, p. 104, 2002.

[90] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "Characterizing radio resource allocation for 3g networks," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 2010, pp. 137–150.

[91] R. Meikle and J. Camp, "A global measurement study of context-based propagation and user mobility," in *Proceedings of the 4th ACM international workshop on Hot tomath.pics in planet-scale measurement*. ACM, 2012, pp. 21–26.

[92] M. A. Hoque, M. Siekkinen, and J. K. Nurminen, "Using crowd-sourced viewing statistics to save energy in wireless video streaming," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 377–388.

[93] C. B. Nelson, B. D. Steckler, and J. A. Stamberger, "The evolution of hastily formed networks for disaster response: technologies, case studies, and future trends," in *Proc. of IEEE GHTC*. IEEE, 2011.

[94] M. Portmann and A. A. Pirzada, "Wireless mesh networks for public safety and crisis management applications," *Internet Computing, IEEE*, vol. 12, no. 1, pp. 18–25, 2008.

[95] C. Maiers, M. Reynolds, and M. Haselkorn, "Challenges to effective information and communication systems in humanitarian relief organizations," in *Proc. of IEEE IPCC*, 2005.

[96] O. D. Antillon, "Hastily formed networks (hfn) as an enabler for the emergency response community," DTIC Document, Tech. Rep., 2012.

[97] P. C. Smith and D. M. Simpson, "Technology and communications in an urban crisis: The role of mobile communications systems in disasters," *Journal of Urban Technology*, vol. 16, no. 1, pp. 133–149, 2009.

[98] D. Abusch-Magder, P. Bosch, T. E. Klein, P. A. Polakos, L. G. Samuel, and H. Viswanathan, "911-now: A network on wheels for emergency response and disaster recovery operations," *Bell Labs Technical Journal*, vol. 11, no. 4, pp. 113–133, 2007.

[99] S. F. Midkiff and C. W. Bostian, "Rapidly-deployable broadband wireless networks for disaster and emergency response," in *IEEE Workshop on Disaster Recovery Networks (DIREN)*, 2002.

[100] A. Nelson, I. Sigal, D. Zambrano, and S. John, *Media, information systems and communities: lessons from Haiti*. John S. and James L. Knight Foundation, 2010.

[101] M. R. Souryal, J. Geissbuehler, L. E. Miller, and N. Moayeri, "Real-time deployment of multihop relays for range extension," in *Proc. of ACM MobiSys*, 2007.

[102] S. Seidel and T. Rappaport, "A ray tracing technique to predict path loss and delay spread inside buildings," in *Proc. of IEEE GLOBECOM*, 1992.

[103] M. Feuerstein, K. Blackard, T. Rappaport, S. Seidel, and H. Xia, "Path loss, delay spread, and outage models as functions of antenna height for microcellular system design," *Vehicular Technology, IEEE Transactions on*, vol. 43, no. 3, pp. 487–498, 1994.

[104] L. C. Liechty, "Path loss measurements and model analysis of a 2.4 GHz wireless network in an outdoor environment," 2007.

[105] K. Rizk, J. Wagen, and F. Gardiol, "Two-dimensional ray-tracing modeling for propagation prediction in microcellular environments," *Vehicular Technology, IEEE Transactions on*, vol. 46, no. 2, pp. 508–518, 1997.

[106] T. Zhou, H. Sharif, M. Hempel, P. Mahasukhon, W. Wang, and T. Ma, "A deterministic approach to evaluate path loss exponents in large-scale outdoor 802.11 wlans," in *Local Computer Networks, 2009. LCN 2009. IEEE 34th Conference on*, 2009, pp. 348–351.

[107] T. Rao, S. Rao, M. V. S. N. Prasad, M. Sain, A. Iqbal, and D. R. Lakshmi, "Mobile radio propagation path loss studies at VHF/UHF bands in southern india," *Broadcasting, IEEE Transactions on*, vol. 46, no. 2, pp. 158–164, 2000.

[108] P. K. Sharma and R. Singh, "Comparative analysis of propagation path loss models with field measured data," *International Journal of Engineering Science and Technology on*, vol. 2, no. 6, pp. 2008–2013, 2010.

[109] M. Iskander and Z. Yun, "Propagation prediction models for wireless communication systems," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 50, no. 3, pp. 662–673, 2002.

[110] C. Phillips, D. Sicker, and D. Grunwald, "Bounding the error of path loss models," in *Proc. of IEEE DySPAN*, 2011.

[111] M. A. Weissberger, "An initial critical summary of models for predicting the attenuation of radio waves by trees," DTIC Document, Tech. Rep., 1982.

[112] S. Bosio, A. Capone, and M. Cesana, "Radio planning of wireless local area networks," *Networking, IEEE/ACM Transactions on*, vol. 15, no. 6, pp. 1414–1427, 2007.

[113] E. Amaldi, A. Capone, M. Cesana, and F. Malucelli, "Optimizing wlan radio coverage," in *Communications, 2004 IEEE International Conference on*, vol. 1. IEEE, 2004, pp. 180–184.

[114] R. Meikle and J. Camp, "A global measurement study of context-based propagation and user mobility," in *Proc. of ACM HotPlanet*, 2012.

[115] A. Sridharan and J. Bolot, "Location patterns of mobile users: A large-scale study," in *INFOCOM, 2013 Proceedings IEEE*. IEEE, 2013, pp. 1007–1015.

[116] Y. Zhang, "User mobility from the view of cellular data networks," in *Proc. of IEEE INFOCOM*, 2014.

[117] H. Braham, S. B. Jemaa, B. Sayrac, G. Fort, and E. Moulines, "Low complexity spatial interpolation for cellular coverage analysis," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (Wiopt), 2014 12Th International Symposium On*. IEEE, 2014, pp. 188–195.

[118] B. Obama, "Presidential Memorandum: Unleashing the Wireless Broadband Revolution," June 2010. [Online]. Available: https://obamawhitehouse.archives.gov/the-press-office/presidential-memorandum-unleashing-wireless-broadband-revolution

[119] President's Council of Advisors on Science and Technology, "Realizing the Full Potential of Government-Held Spectrum to Spur Economic Growth," Executive Office of the President, July 2012. [Online]. Available: https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/pcast_spectrum_report_final_july_20_2012.pdf

[120] *FCC 08-260*. Federal Communications Commission, Nov. 2008. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf

[121] K. Mun, "OnGo: New Shared Spectrum Enables Flexible Indoor and Outdoor Mobile Solutions and New Business Models," May 2018. [Online]. Available: https://www.cbrsalliance.org/wp-content/uploads/2018/04/Mobile-Experts-OnGo.pdf

[122] A. Sahoo, "Fair resource allocation in the citizens broadband radio service band," in *Proc. of DySPAN 2017*, Mar. 2017, pp. 1–2.

[123] "Snapdragon x20 modem with category 18 gigabit class lte," Mar 2018. [Online]. Available: https://www.qualcomm.com/products/snapdragon/modems/x20

[124] M. Tonnemacher *et al.*, "Opportunistic Channel Access Using Reinforcement Learning in Tiered CBRS Networks," in *Proc. of DySPAN*, Sept. 2018.