SPECTRUM MANAGEMENT – PHYSICAL LAYER EFFICIENCY, COLLABORATIVE SPECTRUM ACCESS COORDINATION AND POLICY

by

DRAGOSLAV STOJADINOVIC

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Electrical and Computer Engineering

Written under the direction of

Dr. Wade Trappe

and approved by

_____

_____

_____

New Brunswick, New Jersey

May, 2022

ABSTRACT OF THE DISSERTATION

Spectrum Management – Physical Layer Efficiency, Collaborative Spectrum Access Coordination and Policy

By Dragoslav Stojadinovic

Dissertation Director:

Dr. Wade Trappe

This work presents the results of multiple studies focused around spectrum usage efficiency improvements. The specified topic has been approached in several different ways, and considered approaches include coexistence of heterogeneous networks in the same spectrum, improving the efficiency of a network in the available spectrum, as well as dynamic spectrum allocation.

The growing shortage in available wireless spectrum has resulted in opening of new frequency bands by the FCC. Some bands that have previously been underutilized are now opened for public usage under the 3-tier spectrum sharing framework. Spectrum access is regulated by a Spectrum Access System (SAS). The three user tiers, in order of priority, are incumbent users with guaranteed spectrum access, Priority Access License (PAL) holders with exclusive access to bands unused by incumbents, and General Authorized Access (GAA) users who access the spectrum opportunistically. The long duration of Priority licenses of 10 years, and the large geographical area the licenses cover, can lead to sub-optimal spectrum utilization. This work presents a market model for spectrum sub-leasing in underutilized areas and spectrum bands with short-term license allocation in smaller, focused geographical regions, which PAL holders can use

to maximize profits. The paper published based on this work showed a basic sub-license allocation model along with two improved hypergraph based models which allow sub-license region adjustments or spectrum sharing and coexistence. The improved models resulted in the financial gain increase between 15 and 35% over the base model, depending on the geographical distribution of the sub-lease bid regions.

The increasing need for more spectrum access prompted DARPA (Defense Advanced Research Projects Agency of the US Department of Defense) to start the The Spectrum Collaboration Challenge (SC2) in 2016 to further expand research on spectrum usage efficiency, and mitigate the ever-growing problem of spectrum scarcity. Teams that participated in SC2 designed and developed wireless networks, called Collaborative Intelligent Radio networks (CIRNs), which used Artificial Intelligence and collaboration with other networks to coordinate spectrum usage schedules. To facilitate this collaboration, DARPA has established the CIRN Interaction Language (CIL) – a language CIRNs can use to communicate with other networks and exchange relevant information to establish common spectrum goals and ways to achieve them. One of CIL's main functionalities was to enable teams to announce their intended spectrum usage and provide information other teams can use to adapt their own channel selection. While potentially a beneficial concept, CIL's effect on ensemble throughput of all networks was never evaluated. The paper published at DySPAN 2019 showed a simplified simulation of the spectrum usage announcement functionality of the CIL, explained the experiments with varying numbers of transmitter/receiver pairs which were run to evaluate CILs gains, and showcased an overall increase in spectrum utilization of 4 to 12%, depending on the number of present spectrum users.

Team SCATTER has been participating in the SC2 since its beginning in 2016. SCATTER's open-source software-defined physical layer (SCATTER PHY) has been developed as a standalone application, in cooperation with imec and the University of Ghent, with the ability to communicate with higher layers of SCATTER's system via ZeroMQ, and used USRP X310 software-defined radio devices to send and receive wireless signals. SCATTER PHY relied on USRP's ability to schedule timed commands, used both physical interfaces of the radio devices, utilized the radio's internal FPGA

board to implement custom high-performance filtering blocks in order to increase its spectral efficiency as well as enable reliable usage of neighboring spectrum bands. The thesis describes the design and main features of SCATTER PHY and showcases the experiments performed to verify the achieved benefits. The experiments show the significant reduction of out-of-band transmission leakage achieved by FPGA-based FIR filtering, the achieved throughput of the implemented SCATTER PHY as a function of SNR and the bandwidth of the utilized channel, the attained computation complexity reduction of the improved frequency offset estimation, as well as the achieved packet reception rate.

With many cognitive radio based spectrum access methods naturally applying reactive methods for detecting spectrum availability and their spectrum utilization, it is shown how a proactive approach which uses short-term forecasting to predict spectrum occupancy could help avoid the shortcomings of reactive methods, such as reaction delays which are dominant in scenarios with short bursts of a large amount of data traffic. As channel occupancy can be represented as a time series, the chapter affirms the emergence of Long Short-Term Memory (LSTM) neural network architecture as an efficient prediction mechanism compared to several traditional statistics-based predictions, describes the testing framework built for signal generation, data collection, prediction and evaluation, and finds the benefits of LSTM-based spectrum occupancy prediction, while explaining the scenarios in which different methods (such as collaborative methods based on CIL) would be advantageous, and how the two methods could be used in tandem to gain the benefits of both while avoiding all the shortcomings.

With many available spectrum access and utilization efficiency approaches, operating on various layers of spectrum access methods, the thesis concludes by reiterating the benefits of each of the specified methods, and explains the gains of using them collectively and how this could shape future research.

Acknowledgements

First and foremost, I would like to thank my advisor, Prof. Wade Trappe, for all the support, help, guidance and patience over all these years. I would also like to extend my gratitude to the committee members, Prof. Predrag Spasojevic, Prof. Richard Martin and Dr. Milind Buddhikot, who have all been a delight to work with and were instrumental in both my personal and professional development.

Special thanks to Ivan Seskar, who spent countless hours in numerous discussions with me, provided a limitless amount of support, feedback, input and all possible kinds of help. You were a mentor, a coworker, and a friend at times when I needed it, and were irreplaceable in the years it took to complete this thesis.

Additionally, I would like to thank Prof. Carlos Caicedo for extremely useful discussions and collaboration which helped form the chapter on SCM-based spectrum access coordination.

I would also like to thank Kryptowire, my new company, and Dr. Angelos Stavrou, whose understanding and support was crucial in the last steps made to finish the thesis.

It is impossible to list all my friends who helped me, personally or professionally, on my journey to the end, but I will do my best. Thank you Vlada, Ivan, Milan, Maja, Nenad, Djordje, Ivan, Francesco, Karthik and all your families, for making a far away place feel like home. Thank you to all my WINLAB friends and colleagues – James, Jake, Ali, Mohsen, Jenny, Mike, Elisa, Shreyasee, Shubham, Prasad, Ratnesh, Neel, Mohammed... There's many of you, and I apologize to anyone I forgot. You made work fun.

I would like to thank my ex-wife Mila, we will not celebrate this together but your love and support all these years were instrumental in leading me all the way to the end, and for that I will always be grateful.

Thank you Marko, Petar, Miroslav, Sanja, Marina, Marko, Katarina, Bogdan, for being an indivisible part of who I am.

Finally, thank you to my family, you made me what I am and inspire me daily to do everything that I do. Thank you Deke i Bake, Tata, Mama, Natasa, Bole, Dusan, Vuk and uja.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The ever-increasing shortage of spectrum availability has attracted the research community for decades [1, 2, 3, 4]. Early research in this area was focused on improving the spectrum usage efficiency. This has resulted in major breakthroughs over the years, with significant improvements to the existing modulations and waveforms. However, these efforts have recently started providing diminishing returns, and the focus of the researchers has shifted accordingly. In recent years, other approaches to the problem of spectrum availability have become more attractive topics. Two of these topics that will be addressed in this thesis are spectrum access policy, specifically in 3-tier spectrum sharing frameworks, as well as using collaboration between independent networks to mutually coordinate spectrum access. These issues are of particular importance in the wireless world with an increasing number of heterogeneous, independent networks, all operating in the same bands. Coordinating their access to the spectrum, either centrally or using an interaction language between the networks, becomes an increasingly attractive idea as shown on figure 1.1.

1.1 Spectrum Micro-licensing and License Sub-leasing in 3-tier Spectrum Sharing

The abundance of effort put into solving the wireless throughput problems in the past several decades has led to this problem being approached in many different ways, and multiple improvements found on all communication layers. A lot of research has been focused on the physical layer, specifically on spectral efficiency improvements. As the spectral efficiency gain for a given amount of effort and time has diminished, other ideas have spawned. Instead of further improving spectral efficiency, researchers have turned

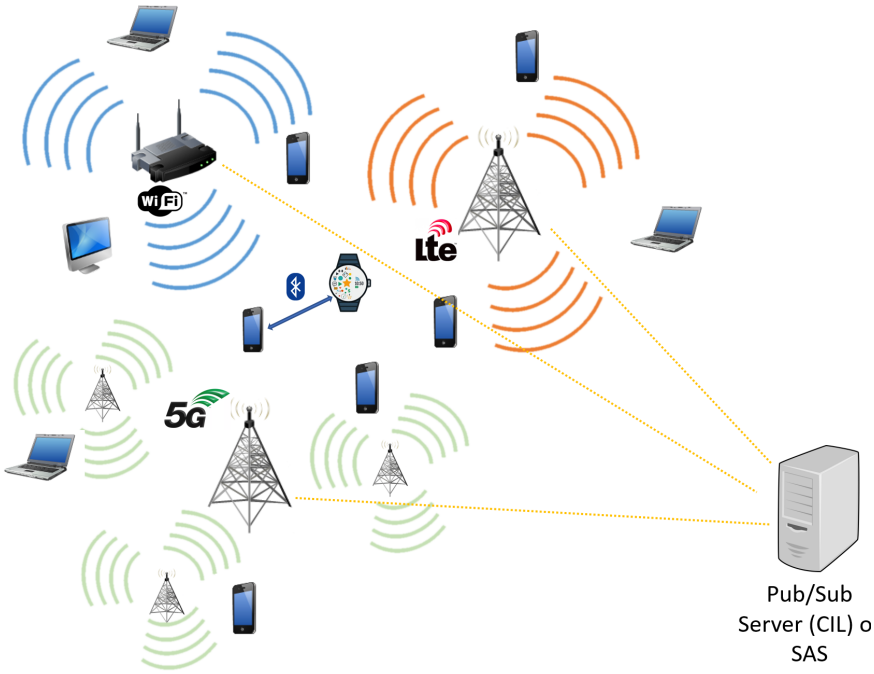

Figure (1.1) With an increasing number of heterogeneous, independent wireless protocols and networks, all operating in the same spectrum, it is important to coordinate their access to the spectrum, either centrally (using SAS) or by enabling the networks to interact and collaboratively determine their access.

their effort onto searching for ways to increase the amount of available spectrum. In some cases, when licensed bands of spectrum had not been used at all times by their primary users (license holders), secondary users were allowed to utilize available channels. This was enabled by using cognitive radio.

However, with the growing requirement for more spectrum availability, some other solutions were recently suggested. Spectrum bands that were until recently reserved exclusively for federal operators (such as navy or satellite communications), are now being opened to secondary users. One such example is the Citizens Broadband Radio Service band (CBRS), spanning from 3550 to 3700 MHz. Cognitive radio's margin of error when ensuring that a channel is available to the secondary user invalidates it as a feasible solution in this case, as primary users must have guaranteed access to the spectrum with a predetermined guaranteed maximum interference from secondary users.

The proposed solution is a 3-tier user structure and a centralized entity that will collect requests for channel access, and assign channels according to availability to each

tier of users [5]. The entity responsible for channel allocation is called the Spectrum Access System (SAS). The first and highest tier provides protection to users of highest priority, incumbent federal users such as naval radar operators. These users must always have access to the required number of channels upon request. Second tier provides exclusive licenses, Priority Access Licenses (PAL) to highest bidding operators, for a given channel within a defined geographical area. PAL holders will be given access by the SAS to the channels they acquired the license for, whenever they are not being used by the incumbent operators. The third tier, called General Authorized Access (GAA), is intended for opportunistic users. The SAS may allocate a channel to them when it is not being used by any of the operators from the two higher tiers [6].

The described tiered spectrum access significantly improves spectrum utilization, while at the same time ensuring that the incumbent users and license holders have unobstructed access to channels when they need it. However, the framework's design with relatively large licensed areas and 10 year long licenses motivates operators to acquire a PAL as soon as it is available, potentially even before they have clear plans on using those channels in their entire county in which the license is assigned. Any unused parts of the county will inevitably end up being utilized by GAA users without any benefit to the license holder. To mitigate the cost of obtaining the license, the operators can sub-lease the unused portions of the licensed county to potential bidders. Indeed, the idea of spectrum leasing in other spectrum sharing scenarios has already existed before the introduction of the 3-tier spectrum sharing framework [7].

A naive way to allocate sub-leasing licenses would be on a first-come first-served basis. A submitted request for a license sub-lease that interferes with any currently assigned license is automatically denied. If the sub-leasing market is designed so that submitted requests are all evaluated at a given deadline, the efficiency, i.e., the number of assigned sub-leasing licenses, can be increased. Depending on the exact shapes and sizes of areas inside the county that the bidders may be interested in, the problem of finding the most efficient way to sub-lease the unused portions of the county may be complex. A graph can be built based on all submitted bids and their regions, such that a bid is a node of the graph and interference with another bid is represented as

an edge between the corresponding vertices. This makes the problem equivalent to the well known graph problem - maximum weighted independent set (MWIS), which is the first proposed improvement to the naive algorithm.

Allowing interference with a limited number of other bidders in small parts of their defined regions may significantly increase the number of assigned sub-leased licenses, and thus bring more profit to the PAL holder. At the same time, this method increases the complexity of the problem, as such geographical distribution of sub-licensed regions can not be adequately represented with regular graphs any more. Thus, section 2.2 of this thesis shows two hypergraph based bidding models the SAS can utilize to improve the sub-leasing allocation, of the licensed channels while maximizing profit and/or licensed area coverage, and shows the increased efficiency of between 15 and 35% with these models compared to the simple interference avoidance method.

1.2 Spectrum Scarcity and the DARPA Spectrum Collaboration Challenge

DARPA has established the second spectrum challenge, this time named Spectrum Collaboration Challenge (SC2), with the aim to boost research on innovative technologies that enable and facilitate coexistence amongst various heterogeneous wireless networks sharing the same wireless spectrum. The new name indicated new goals – instead of simply beating the wireless throughput achieved by other teams, the new challenge was focused on driving the independent networks to collaborate in order to maximize the ensemble throughput.

SC2 was structured around multiple teams that each built their own wireless communication system, a Collaborative Intelligent Radio Network (CIRN). The name CIRN was coined to outline some of the main goals of the entire challenge – to encourage teams to include Artificial Intelligence elements in their networks to help with increasing spectral efficiency and coexistence with other teams and their radios, as well as to promote collaboration between teams, mainly in terms of spectrum usage and occupancy, to achieve the common goal of increasing the combined throughput of all teams.

To enable and facilitate this collaboration, DARPA established a common language for all teams, to help them announce their intentions in spectrum usage, request channel availability, etc. The language was named CIRN Interaction Language (CIL). CIL is based on a publish-subscribe system, where each team can receive messages sent by any other team, and was designed to operate using a designated protocol (Ethernet) in order to bypass the challenges present in the main wireless communication channels being used competition-related throughput, and to guarantee successful delivery of the messages to all teams.

Over the course of the challenge, CIL evolved with the needs and requirements of the teams and has included various functions. This thesis shows an evaluation of the spectrum usage announcement functionality of the CIL. In particular, it attempts to measure the impact that correct spectrum voxel announcements have on the combined throughput of the participating teams. In other words, it shows a comparison between the achieved throughput when teams have information related to the intended spectrum usage of other teams in the near future with the scenario where teams have no information related to the spectral behavior and intentions of other teams.

Within the same challenge, one of the participating teams was SCATTER, a team with the name motivated by their usage of available spectrum as well as the geographical distribution of its members (imec, Ghent University and University of Antwerp in Belgium, and Rutgers University in New Jersey). Team SCATTER has designed and built a multi-layered system, with physical, MAC and other layers being completely autonomous but interconnected with a publish-subscribe messaging system, in particular ZeroMQ [8]. Thus, SCATTER's physical layer (PHY) was designed as an independent system, able to transmit and receive data packets over the air using USRP software-defined radio (SDR) devices. The design of SCATTER PHY uses the existing srsLTE software library [9] as the basis for its communication, and builds on top of it with various features and modifications targeted for the SC2, such as the ability to communicate with higher layers of the system (MAC, AI, etc.), and using custom-built FPGA filters and other blocks to reduce out-of-band emissions and maximize channel utilization. This thesis describes the design and features of SCATTER PHY and experimentally

verifies the benefits achieved with the outlined design decisions.

The rest of the thesis is organized as follows: chapter 2 describes the work done on developing spectrum sub-leasing algorithms in three-tier spectrum sharing scenarios. Section 2.1 explains the problems that motivated the research, 2.2 shows the design of the sub-leasing framework and the models for determining the spectrum band allocation. Section 2.3 describes the implementation of the models, and the simulation utilized to showcase the results of the developed models. Section 2.4 concludes chapter 2. Chapter 3 represents a report on the development of the physical layer by members of team SCATTER for the DARPA Spectrum Collaboration Challenge, as well as an evaluation of the effects of the innovative network interaction language used by the participating teams to coordinate spectrum access. Sections 3.1 and 3.2 describe the design and features of SCATTER's physical layer, and show experimental verification of its achievements and benefits. Section 3.3 shows the work that preceded the CIRN Interaction Language (CIL) used in DARPA SC2. Section 3.4 describes the design of DARPA's CIL, while sections 3.5 and 3.6 describe how CIL's effects were evaluated and show the evaluation results. Section 3.7 briefly touches on the outcomes of the DARPA SC2 and concludes chapter 3. Chapter 4 describes an experimental expansion on CIL. The collaboration is enhanced by implementing the CIL functionality with Spectrum Consumption Model (SCM) based messages, which were selected as an IEEE standardized method of describing transmitter and receiver spectrum utilization characteristics. After motivating the design of SCMs, the chapter describes an experimental setup where wireless nodes utilize CIL to exchange SCM messages and adapt their transmission parameters based on current spectrum availability. Additionally, the chapter describes the architecture and design of a decentralized Dynamic Spectrum Management system, where participating wireless networks collaborate by sharing the spectrum sensing resources and avoid interference by utilizing the SCM based standardized compatibility calculation to precisely determine the available spectrum bands at any given location. Chapter 5 compares the efficiency of various spectrum utilization approaches decribed in previous chapters as well as a predictive approach, provides an analysis of advantages and shortcomings of each approach, and ultimately proposes methods on how different

approaches can be used collectively to gain benefits of various designs while avoiding all the drawbacks. Finally, chapter 6 summarizes the results from the described work, and proposes the problems to be addressed in the future.

Chapter 2

Three-tier Spectrum Sharing and Spectrum Sub-leasing

One of the proposed solutions for spectrum scarcity in bands that were underutilized by their incumbent users, briefly described in chapter 1, is the 3-tier user structure and the Spectrum Access System (SAS) as the centralized entity that performs the channel allocation to interested second-tier users [5]. As shown on figure 2.1, the highest tier provides protection to incumbent users, such as federal and naval radar operators, who lose exclusive access to these spectrum bands but retain the highest priority and can access the spectrum whenever needed. Operators with highest bids obtain Priority Access Licenses (PAL), which guarantee access to the licensed number of channels not utilized by the incumbent users, within well defined geographical areas – counties. General Authorized Access (GAA) users, can only access the spectrum opportunistically, as determined by the SAS [6]. The published work can be found in [10].

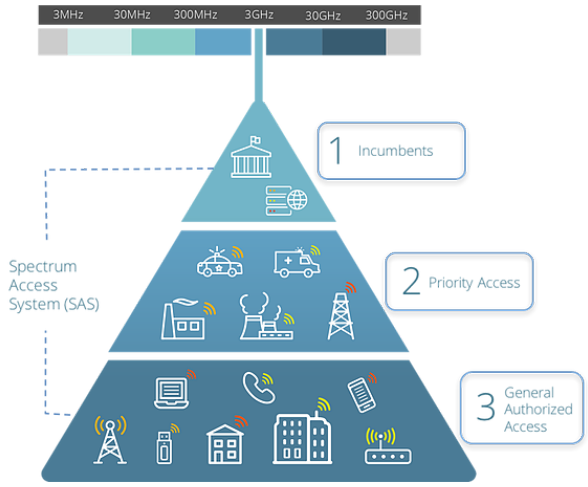

Figure (2.1) Three-tier spectrum sharing infrastructure. The first tier users such as incumbent federal operators can access the spectrum whenever needed; the second tier are licensed users who may use the channels not utilized by incumbent users. The lowest tier users only access the spectrum opportunistically.

As mentioned in chapter 1, the described method of spectrum access significantly improves utilization and ensures interference-free spectrum access to parties with higher priority. One of the important characteristics of the proposed framework remain the relatively large areas the spectrum licenses are assigned for (counties), as well as long license duration (10 years). These features of the described framework may present an issue to operators who are motivated to obtain spectrum access licenses as soon as they become available, but do not plan network deployment in the immediate future. In such cases, the unused portions of their licensed areas present a financial loss.

This chapter describes a spectrum sub-leasing model, which allows the operators to alleviate the high license costs. The idea of spectrum leasing, which already existed in other spectrum sharing scenarios [7], is being applied to the 3-tier spectrum sharing framework in this thesis. First, a brief discussion is made on several aspects of the sub-leasing framework. Then, a general sub-leasing model is defined in the beginning of chapter 2.2. The model is simple, and serves as a base for developing more complex models that allow for more efficient sub-leasing of the unused county area. The efficiency of the sub-lease bidding model can be defined as either the highest achieved profit, largest area covered by secondary PAL users (license sub-lease holders), the highest number of satisfied bidders, or a combination thereof.

Afterwards, this chapter shows the efficiency of the proposed models in each of the defined criteria. The simulation was run with synthetically generated data, with two different geographical distributions of sub-lease region offers, uniform and clustered. Both of the distributions are motivated by realistic base station distribution requirements in rural or urban areas with various degrees of population density.

2.1 Related Work and Motivation

The described 3-tier spectrum sharing framework has been widely accepted as an expected prospect for channel allocation for the future of wireless networks. The subject has garnered a significant amount of research, and various related topics have attracted attention.

The work presented in [11] started the research on this topic by introducing the 3-tier shared spectrum framework, along with the Spectrum Access System, incumbents, Priority Access License and General Authorized Access users, and showcased a prototype end-to-end 3.5GHz LTE/WiFi testbed.

In [12], the authors have shown a comparison between various shared spectrum scenarios and confirmed that the 3-tier shared spectrum framework provides the most combined benefits to both consumers and service providers.

Usage of 3rd tier spectrum access for PAL holders to increase the number of customers that can be serviced is proposed in [13], and then presented in more detail in [14]. The paper explains that this may come at the cost of their customers' perceived Quality of Service, and then suggests a low-complexity algorithm for PAL holders to determine when such 3rd tier access can be profitable, without any prior knowledge of customer demand or channel availability.

An adaptation of video streaming, a high throughput demanding and popular service today, in a 3-tiered framework has been presented in [15]. The algorithm shown in the paper adapts the video streaming service for service providers using opportunistic access under the GAA license with unreliable future network conditions.

A method for GAA users to improve the capacity of their networks while remaining below the required interference limits by optimizing their transmission power is proposed in [16].

In [17], the authors showcase a live field trial of CBRS by utilizing existing network systems, operational base stations and commercial 3.5 GHz equipment to implement the SAS and spectrum users from different access tiers.

A sub-licensing model similar in purpose to the models presented in this thesis is explored in [18]. The entire licensed region is divided into a grid in order to simplify spatial division and reduce interference between adjacent sub-license holders. The paper focuses mainly on the fineness of the grid to achieve best interference reduction.

Even more published scientific work related to the 3 tier spectrum sharing framework topic can be found. The described papers demonstrate the direction in which the scientific community is headed with this topic. The increasing value of undisturbed

spectrum access implicates the benefits a service provider can have by obtaining a PAL. However, for smaller providers, the prices of spectrum access and the fact that licenses are awarded for 3 year periods may seem daunting. This chapter describes the design of a spectrum renting based approach the PAL holders can utilize to profit by sub-leasing unused portions of their counties to other users in an auction. At the same time, this provides an opportunity for smaller service providers to acquire a license for a shorter time period in a smaller area, as much as they need to deploy their network, at a significantly lower cost than obtaining a long-term license for the entire county.

2.2 License Sub-leasing Framework and Models

The proposed license sub-leasing framework is auction based. A PAL holder which determines that the entire county does not need to be serviced in a certain time frame sub-leases the license in the unused area of the county to other users. Bids for the license sub-lease are assumed to provide the defined region within the unused area the bidder is interested in, as well as the financial offer, i.e. the price the bidder is willing to pay for the sub-license. In cases when two or more bids' regions overlap, the SAS needs to determine the "most efficient" way to award sub-licenses, where the discussion about defining "efficiency" is shown in chapter 2.1.

The framework warrants the use of graphs to model the auction scenarios. There are other aspects that PAL holder needs to consider. The duration of sub-lease needs to be short enough to accommodate the dynamics of the market, but at the same time long enough to warrant obtaining a sub-license over simply using a channel with the GAA license when available. Another important aspect is the precise definition of rules of the bidding process itself. An example of these rules can be whether region modifications are allowed, i.e. are the bidders willing to sacrifice smaller parts of their defined regions in order to receive a sub-license, or are they only interested in all-or-nothing? For example, it may be more profitable for the PAL holder to not assign a sub-license to a bidder A, because other bidders with regions that overlap with the region specified by A have made more lucrative offers. However, if A is willing to reduce its region and give up on the parts intersecting with other bidders' regions, it may still receive a sub-license.

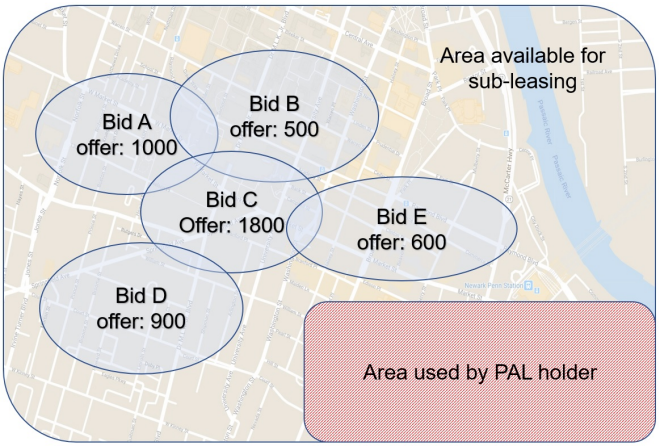

Figure (2.2) The complete geographic area available to the PAL holder can be partitioned into smaller regions which can then be sub-leased. Red shaded area is used by the license holder and is unavailable for sub-leasing. Five light blue regions represent submitted bids for license sub-leasing.

Some of these aspects will be involved in the presented sub-leasing models. There are many possible ways in which the presented models can be improved. This thesis presents two models with the aim to show the general idea behind the sub-leasing framework, as well as benefits that can be gained by slightly adapting the basic general model.

2.2.1 Base Model

This section describes the simple base sub-leasing model. The model is based on an intersection graph. The model assumes that each bidder defines the geographical region it is interested in, as well as the financial compensation it is willing to offer in order to obtain a temporary license in the defined region. Regions are defined by the bidders at the time of making the offer and can not be altered. Furthermore, the model assumes that there is only one channel licensed by the PAL holder and available for sub-leasing. Figure 2.2 shows a scheme of a county, surrounded by the black border. The red shaded area in the lower right corner of the county represents the area intended to be used by the original PAL holder, which is therefore not available for sub-leasing. In the remaining part of the county, there are five light blue regions which represent offers for license sub-leasing.

The bidding offers can be represented by a weighted intersection graph of bids. Figure 2.3 shows an intersection graph which corresponds to the scheme example shown

in Figure 2.2. The weights are not shown in the Figure. The nodes of the graph represent the offers, and an edge between two nodes means that there is an overlap between the regions defined by the two bids corresponding to two adjacent nodes. The nodes are assigned weights which correspond either to the financial offer made by the bidder, or to the size of the region defined by the bidder, depending on whether the PAL holder aims to maximize profit or the size of the covered area. The "best" license sub-lease assignment is then equivalent to finding the Maximum Weighted Independent Set (MWIS) of the graph of bids. Finding the MWIS is an NP hard problem itself, but the solution only needs to be found once when the entire graph is provided. Extensive work has been done in graph theory and several algorithms have been found for finding approximate solutions to this problem, with varying degrees of efficiency, optimization and guarantees for proximity to the actual solution, for various classes of graphs. Recent work on this topic includes [19], [20], [21], [22].

The described model is intentionally simplistic. It has some obvious shortcomings and provides a lot of room for improvement, but it also provides a common base for

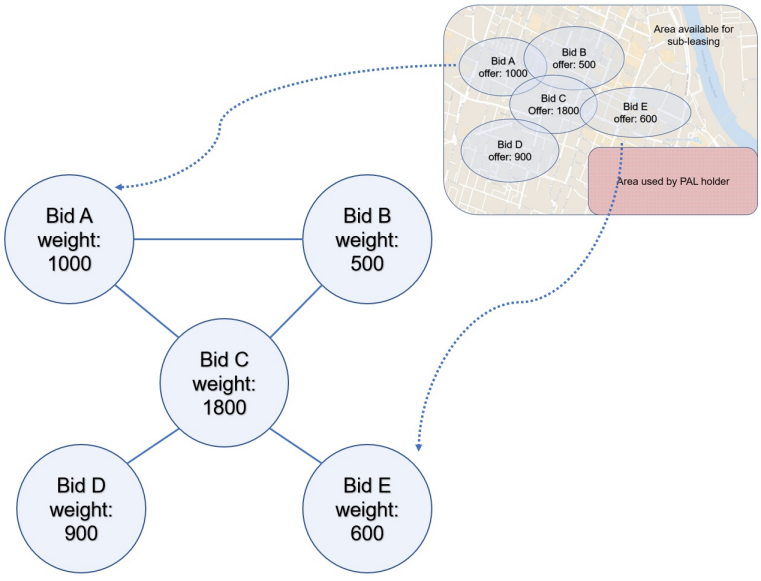

Figure (2.3) The relationships between the interdependencies of the partitioned sub-areas can be captured in a graph structure. Each node corresponds to a submitted bid. Edges between two nodes represent overlap between the defined regions of the corresponding bids. Weights of the nodes are equal to the financial offers made by the corresponding bidders.

various improved and refined sub-leasing models. The base model does not allow bidders to redefine their request regions in order to improve their chances of being awarded a license sub-lease, it doesn't consider the possibility of coexistence with other bidders in parts of their defined regions, and it can not provide optimization in both profit and covered area, except in special cases where financial offers are approximately proportional to the requested region sizes, and in those cases such optimization of both of the mentioned criteria is purely incidental. An example of a simple improvement that would allow optimizing both of these criteria could be to impose a limit to the covered area percentage, and find the most profitable sub-leasing assignment within that limit.

2.2.2 Region Reduction Model

The first suggested improvement to the base model is to assume some bidders are willing to reduce the sub-lease regions they requested, in order to improve their chances of obtaining a license. The idea is to reduce the region just enough to avoid overlap with regions of neighboring bidders. In the graph representation, avoiding this overlap is equivalent to removing the edge between two corresponding nodes. The model proposes a "fair" reduction, i.e., when both adjacent bidders are willing to allow the reduction of their regions, the node that will first be selected for reduction will be the one with the larger requested region, as the "damage" the bid would sustain in this manner would be proportionally lower. To clarify, let's assume two bidders, A and B, made offers with overlapping regions, as shown on Figure 2.4. Bidder A requested a region of size of 13 units, the size of the region of bidder B is 10 units, and the size of the overlapping area is 2 units. In this case, region requested by A is reduced by removing the overlapping area, so that the resulting region's size is now 11 units. Algorithmically, this is presented by a directed graph. Instead of having an undirected edge between nodes A and B, there are now two directed edges, both $A \rightarrow B$ and $B \rightarrow A$. Edges are weighted, and the weight of each edge is equal to the size of overlapping area divided by the size of the region of its source node. In the mentioned case, the weight of edge $A \rightarrow B$ is equal to $2/13$, and the weight of edge $B \rightarrow A$ is $2/10$. This model provides a simple, yet effective improvement, and increases the chances of willing bidders to obtain a sub-license, and

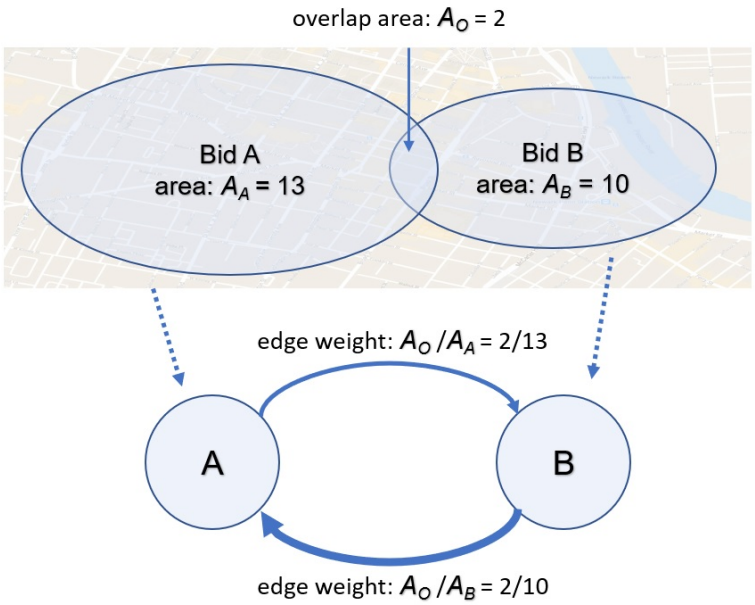

Figure (2.4) Directed graph representation of overlapping bid regions which facilitates bid selection for region reduction. Overlap is represented by two edges, each edge's weight equal to the ratio of the overlap area size and the source node bid area size. Lower edge weight means that its source node has a larger area and is thus the better candidate to give up the area and avoid interference.

ultimately provides a more profitable opportunity for the PAL holder.

2.2.3 Limited Coexistence Model

The second improvement to the base model is to allow coexistence between license sub-lease bidders, with limitations on the number of coexisting operators, as well as the area size in which coexistence is allowed. For example, operator A may be willing to allow coexistence with other operators (and thus potentially lower throughput) in up to 20% of its given region, and may be willing to coexist with 2 but not 3 other operators in the same region.

In order to simplify the explanation of the model, Figure 2.5 shows a scheme of overlapping regions of several license sub-leasing bids, from the point of view of bidder A. The region requested by bidder A overlaps with regions requested by bidders B, C, D and E. The smaller areas where regions overlap are marked as P_i, where $i = 1, ..., 5$. Note that P_4 stands for the area where regions A and D overlap, but excludes the part where this sub-area also overlaps with the region E. Similarly, P_3 stands only for the

area where A and E overlap, excluding the part where they also overlap with D. The area where all A, D and E overlap is marked with P_5.

From the point of view of bidder A, the limitation imposed by this model refers to the entire area where A may be required to coexist with other bidders. Let S_A be the size of the entire region defined by A, and $P = \sum_{i=1}^{5} P_i$. The bidder A may then allow coexistence with other sub-lease holders in the same channels, but it may also limit the size of the area where this coexistence is allowed to $P_{max}(A)$. This limit may be given as a percentage p_A of the entire area S_A, so that $p_A = \frac{P_{max}(A)}{S_A}$. For example, by setting a smaller $P_{max}(A)$, the bidder A can ensure that a larger part of its assigned region is coexistence-free and available exclusively to A. However, if other bidders are more willing to coexist, this may reduce A's chances of receiving a sub-lease region at all. If the licensed channel in the defined region was allocated to A, this would now mean that potentially all other competing bidders B, C, D and E would need to be excluded, but it is likely that this kind of assignment would be sub-optimal in profit and/or covered area, so the optimal assignment now may end up not including A.

The financial offer made by A also needs to be adjusted accordingly. For the sake of explaining this adjustment, let us assume all of the bids A, B, C, D and E are accepted, and they all need to coexist with each other. Let k_i be the number of coexisting operators in an overlapping area P_i. The amount that will now be charged to A per area unit for the area P_i is fractionally smaller than the amount charged for the area exclusive to A. Let us denote this fraction with α_i. One possible value for α_i would be $\alpha_i = \frac{1}{k_i}$. If the original financial offer made by A was O_A^{init}, the final amount charged to A will be

$$O_A = (1-p)O_A^{init} + \sum_{i=1}^{5} p_i \alpha_i O_A^{init},$$

where $p_i = \frac{P_i}{S_A}$.

This sub-leasing model can not be represented by a graph similar to the region reduction model graph. Similarly to that model, in order to represent the overlapping areas and their sizes, the new graph will have directed edges. There will be a pair of directed edges between each two nodes that represent overlapping bids. Edges will

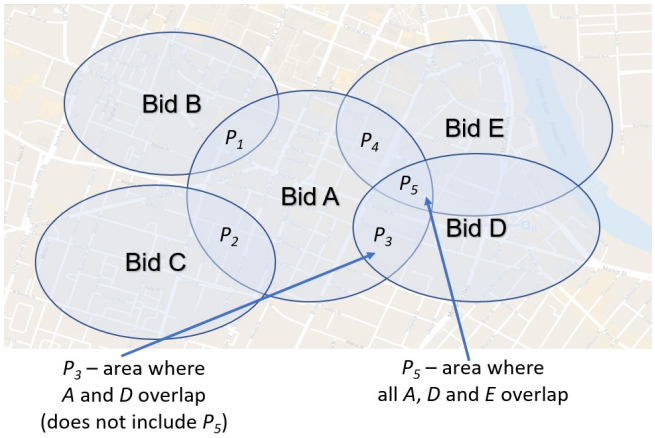

Figure (2.5) Limited number of allowed coexisting operators makes it important to distinguish areas where two, three or more bid regions overlap.

be weighted, such that the weight of the edge is equal to the size of the overlapping area between two corresponding regions. However, a regular graph doesn't make it straightforward to represent areas where 3 or more of the regions overlap, which is crucial for this sub-leasing model.

A more natural graph representation of such scenarios is a hypergraph. A hypergraph is a generalization of graph, such that an edge is not limited to connecting only 2 nodes.

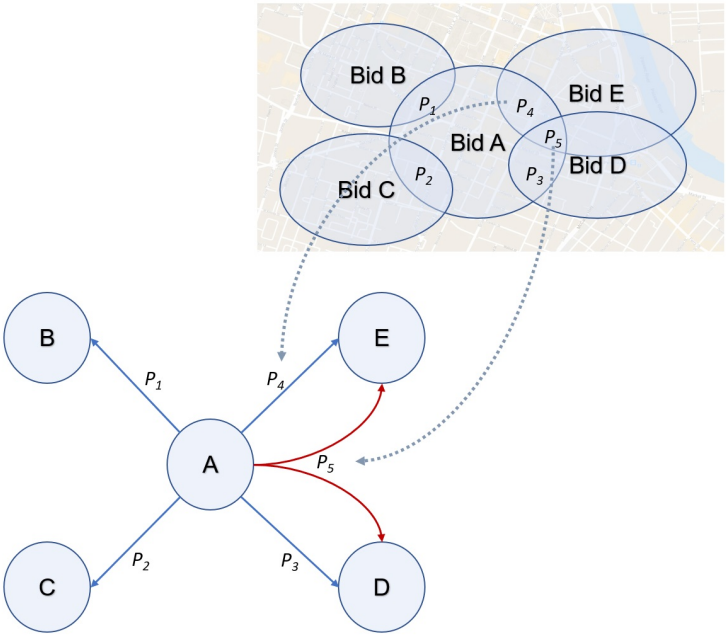

Figure (2.6) Hyperedge which corresponds to P_5, area where three regions overlap, with source at A and destination set $\{D, E\}$, is shown in red.

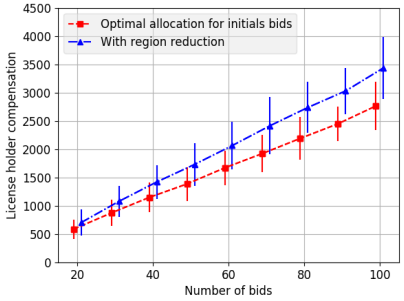

(a) Average compensation for uniform distribution

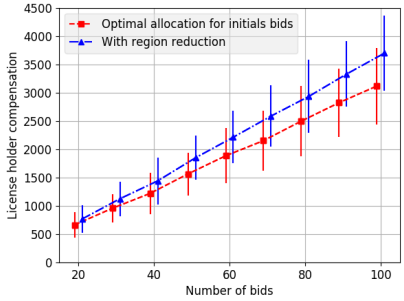

(b) Average compensation for distribution with 3 clusters

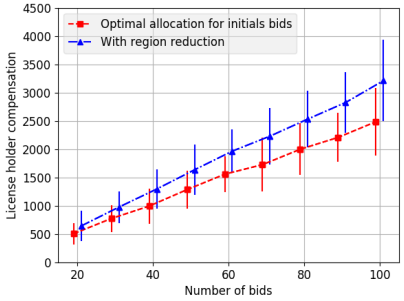

(c) Average compensation for distribution with 6 clusters

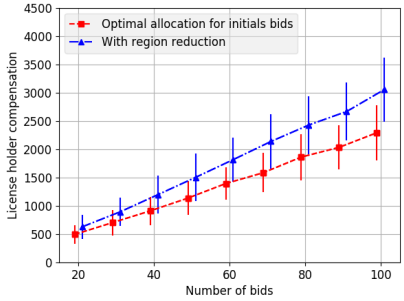

(d) Average compensation for distribution with 8 clusters

Figure (2.7) Profit gained by the license holder. Blue line shows the profit after the necessary bid region reduction is applied and limited coexistence is allowed, red line shows the profit gained by using the base model. Different plots show gain increases for different geographical bid distributions. All plots include the confidence interval of 0.95.

A hypergraph corresponding to the region scenario from Figure 2.5 is shown in Figure 2.6. Edges are shown as light blue arrows, and the hyperedge that connects node A to nodes D and E is represented by the divided light blue arrow. In general, defining the direction of an edge in a hypergraph is arbitrary. We will define the direction of an edge that connects more than 2 nodes only by the source node, the order of destination nodes will not matter.

With the hypergraph with hyperedges representing the overlapping regions, it is now straightforward to assign weights to the edges of the graph. The weight of edge $A \rightarrow B$ will be P_1/S_A, and other weights of 2-node edges can be assigned similarly. The weight of edge $A, \{D, E\}$ is $p_5 = \frac{P_5}{S_A}$. In order to find the optimal sub-leasing assignment to bidders, one needs to find a maximum weighted set of the hypergraph, such that for each

node X, the sum of its overlapping areas with other nodes does not exceed a maximum $P_{max}(X)$ defined by X's bid itself, which is equivalent to the sum of its outgoing edge weights not exceeding $\frac{P_{max}(X)}{S_X}$.

It is important to note that the shortcoming of this model is its complexity. As was already mentioned, finding a maximum weighted independent set is an NP hard problem. Finding a maximum weighted set of nodes as defined for the previous models is even more complex than simply finding MWIS, due to the imposed additional limitations. Finding such set for a hypergraph further increases the complexity of the problem. However, approximations to the solution can be found using heuristic approaches. Additionally, the algorithm efficiency is not critical in this case, as such algorithms do not need to run in real time, but need only to provide results once in a longer period of time with given input data.

2.3 Implementation and Simulation

In order to show the performance of the proposed models, the base model and the region reduction were implemented and tested. To make it more general, the implementation was based on hypergraphs. This makes the implementation easily extensible for the third proposed model, as well as the potential new models that could be devised in the future. However, the problem of such implementation is the overall complexity. Not much work has been done by the research community on developing efficient algorithms for hypegraphs, so the implementation used for this thesis was done using brute force. The consequence of this approach is that the implementation does not scale well for larger input sizes. However, it is once again important to note that in the proposed market design, these algorithms are not meant to perform in real time, but only once in a given time period.

As input data, a given number of nodes is generated at random spots in an imaginary county. Tests were run with the bid numbers ranging from 20 to 100. Two different classes of geographical distributions of the bids were used to simulate different realistic scenarios. The distribution of requested sub-license regions can reasonably be expected

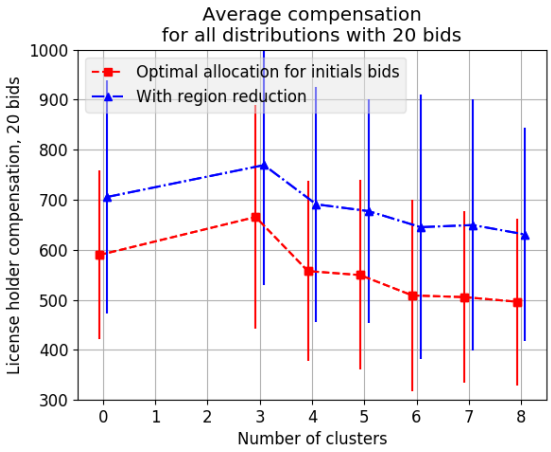

Figure (2.8) Profit increase trends across the distributions for 20 bids. X axis represents all used distributions (0 clusters corresponds to the uniform distribution), Y axis shows the profit increases between the red and blue lines.

to be similar to the distribution of the population in the same area. Uniform approach represents a scenario where the population is uniformly distributed, regardless of the density. This can either be a densely populated part of a large city, such as Manhattan, or a scarcely populated area in a rural region. In this approach, the nodes are uniformly scattered across the entire area. The other type of approach is using the clustered distribution, which can, for example, represent an area with several smaller towns scattered across the area. In this approach the nodes are gathered around a given number of clusters (towns). For the clustered approach, testing was performed with different number of clusters, ranging from 3 to 8. The sizes of regions requested by the bidders, i.e. the nodes of the hypergraph, are also randomly distributed. Normal distribution was used to represent these sizes, such that mean and standard deviation are again selected randomly, to accommodate scenarios with both higher and lower population density, i.e. cases with high amount of interference between nodes, as well as the ones with lower amounts of interference.

Figure 2.7, and its subfigures 2.7a, 2.7b, 2.7c and 2.7d show the simulation results for the uniform distribution, and distributions with 3, 6 and 8 clusters respectively. In each plot, the x axes represents the number of nodes (bids), and the y axes shows the corresponding profit gained by the license holder, averaged over 50 different runs. Note that the profit units do not have any assigned value, as the potential range of absolute

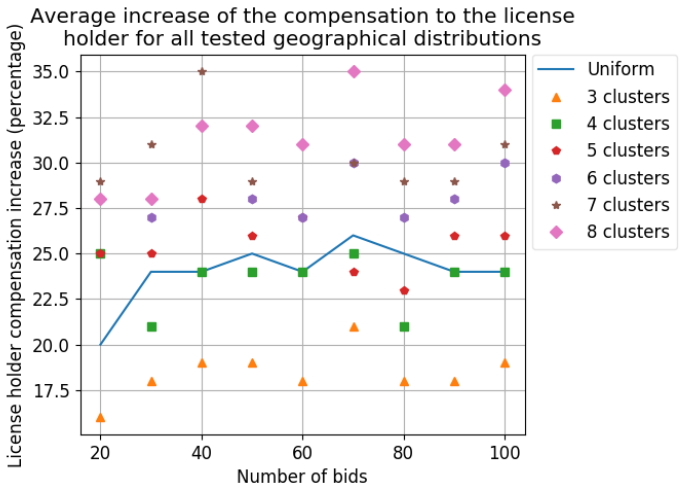

Figure (2.9) Profit increase percentage for all geographical distributions and all number of bids. Highest profit increase is gained for the geographical distribution with the most clusters, lowest profit is gained for the distribution with the fewest clusters, while uniform distribution (indicated by the line) achieves medium increase.

values financial offers of the bids are not in the scope of the presented work. Rather, of interest are relative values between different bids, which are also randomized to allow for different realistic scenarios. Red line shows the profit achieved for the first base model, where no interference or region adjustments are allowed. Blue line shows the values for the second model, which allows a limited region reduction (25% of reduction was allowed for the sake of these tests). All cases show a steady improvement in the increased gain ranging between 15 and 35%.

As was already mentioned, all distributions show a steady increase of profits achieved by sub-leasing, which can be seen on Figure 2.8 which shows the trends for all distributions with 20 bids. The highest gains are achieved for the distribution with the most clusters, while the distribution with the fewest clusters brings the lowest gain. It is interesting to note that uniform geographical distribution of bids results in an achieved gain which is between the highest and the lowest. This can be seen in more detail on Figure 2.9, which shows the average profit increase percentage for all distribution and all considered numbers of bids.

2.4 Concluding Remarks on Spectrum Sub-leasing

This chapter has shown a brief overview of the three-tier spectrum sharing frameworks, what motivated them and when and where there are used. It pointed out some of its shortcomings, and addressed the one concerning the relatively large areas in which spectrum licenses were awarded, as well as long time periods the licenses were valid for. With the ongoing spectrum scarcity, operators could be pressed to obtain the licenses before the network deployment is planned in the entire licensed area, at least in the immediate future. This chapter has introduced a spectrum sub-leasing marketplace as a proposed solution to the given problem. It presented potential models for sub-license allocation by the SAS, and finally showed the results of those models presented as potential profits for the license holder. The detailed analysis showed that the geographical distribution of spectrum users can affect the optimization of the utilized spectrum sub-licensing models, i.e., in cases where the users were distributed in more regional clusters the providers would be able to increase their profits by up to 35%, whereas in regional distributions with fewer clusters this profit increase could only be 17%. Additional analysis showed that, assuming a high density of spectrum users, with limited flexibility of spectrum renters by allowing their required regions of rented licenses to be reduced by up to 20%, they could significantly increase the probability of obtaining a sub-license, while the primary license holders could increase their profits by an average of 25%.

Chapter 3

DARPA SC2 – CIL and SCATTER PHY

The increasing number of wireless communication protocols, with no apparent convergence on a global scale, introduces the problem of coordinating spectrum access for multiple heterogeneous networks in the same spectrum bands. This problem has motivated DARPA to organize the Spectrum Collaboration Challenge (SC2), a 3 year long competition for teams across the world, with the final goal of achieving various methods for heterogeneous networks to not only coexist in the same spectrum, but also to cooperate, collaborate and help each other increase their spectrum usage efficiency and achieved data throughputs. Team SCATTER has participated in the competition since its began in 2016. The team name was motivated partly because of the intention of team members to opportunistically use different bands, "scattered" across the entire spectrum, and also by the geographical distribution of its members, who were "scattered" around the world, with members from imec, Ghent University and University of Antwerp, as well as a part of the team at WINLAB, Rutgers University, New Jersey. This chapter will describe the physical layer designed and implemented by team SCATTER, and showcase its features and benefits. It will also evaluate the benefits of the network interaction language used by teams to coordinate spectrum access, developed during the challenge by DARPA and all participating teams. The described work was published in [23] and [24].

While the DARPA Spectrum Collaboration Challenge system design and implementation was a team effort, and the conceptual design of the system is a collaborative effort that presents the results of numerous discussions between all the team members, the particular personal contributions of the thesis author can be emphasized in the components of FPGA based filtering to block out-of-band transmissions, establishing

timed commands, providing support in UHD framework abilities which helped pinpoint the feasible options that can be used in the final system as well as the design of the collaboration channel strategy. This strategy was subsequently used as the basis for the development of the collaboration implementation and evaluation on the ORBIT Testbed. In terms of the system implementation, the author's specific contributions can be summarized as follows:

- Integration of srsLTE and C++ based UHD codebase to enable full UHD functionality while utilizing the srsLTE functionality in the physical layer.
- Integration of required RFNoC FPGA blocks into the codebase.
- Testing and evaluation and instrumentation of the system, performed on ORBIT Wireless Testbed.
- Integrating the basic ZeroMQ functionality to enable communication with other components of the network.
- Various specific physical layer functionality implemented in UHD to support the requirements of the overlaying system design, such as bursty traffic, scheduled transmissions, etc.
- Overall system performance evaluation on COLOSSEUM testbed.

3.1 SCATTER PHY

The high-level architecture of the SCATTER PHY is depicted in Figure 3.1. The SCATTER PHY is implemented using Universal Software Radio Peripheral (USRP) Hardware Driver (UHD) software Application Programming Interface (API) [25] and runs on top of Ettus USRP X family of software-defined radio (SDR) devices including NI's RIO platforms [26, 27] and communicates with it through the UHD driver and its APIs [28]. As can be seen in the figure, the individual PHY modules are connected to the ZeroMQ (Data/Control) module, also known as 0MQ, which interconnects the SCATTER PHY with the MAC layer through the ZeroMQ bus [29]. This module

manages the exchange of control and statistics messages between the SCATTER PHY and MAC layer.

Communication with the SCATTER PHY is entirely implemented through a well-defined interface designed with Google's Protocol Buffers (protobuf) [30] for data serialization coupled with the ZeroMQ messaging library [29] for distributed exchange of control, statistics and data messages. Implementing the ZeroMQ push-pull pattern allows local or remote MAC layer's real-time configuration of several parameters and reading of several pieces of information/statistics provided by the SCATTER PHY. Based on the ZeroMQ logic, PHY and MAC layers are able to exchange control and data messages following a non-blocking communication paradigm. The SCATTER PHY was designed to be completely decoupled and independent of the MAC layer module, not posing any constraints on hardware, software and/or programming language adopted by it.

The SCATTER PHY contains the following set of main features:

- Bursty transmissions: with discontinuous transmissions, it is possible to improve the use of an available spectrum band and to coordinate its usage with other networks/radios in an opportunistic/intelligent/collaborative way.
- Dual-Concurrent PHYs: having two physical interfaces simultaneously transmitting and receiving at independent frequencies enables Multi-Concurrent-Frequency Time-Division Multiple Access (McF-TDMA) scheme to be implemented by the MAC layer. This ability for concurrent allocations allows for a smarter spectrum utilization as vacant disjoint frequency chunks can be concurrently used.
- FPGA-based filtered transmissions: filtering the transmitted signal effectively minimizes out-of-band emissions (OOBE), which allows for better spectrum utilization, by enabling radios to have their transmissions closer to each other in the frequency domain, reducing spectrum wastage.
- Out-of-Band Full-Duplex operation: both PHYs operate completely independently, meaning that Tx and Rx modules are able to transmit and receive at different channels, set different gains and use different PHY BWs.

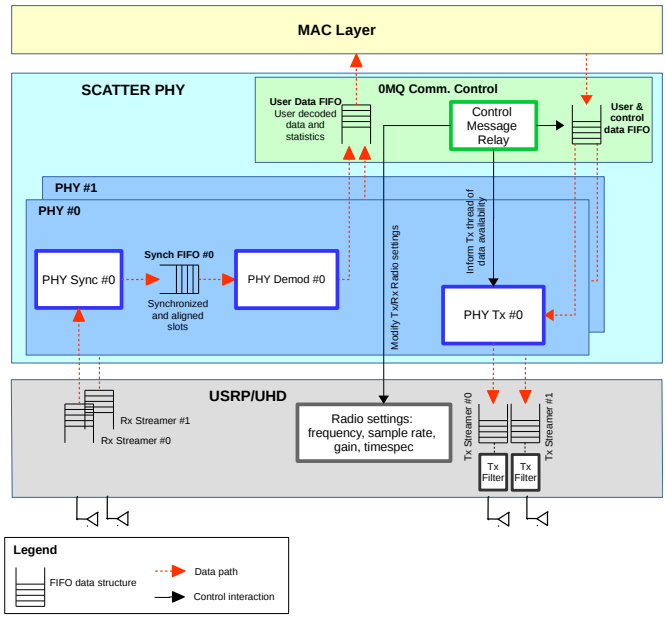

Figure (3.1) High-level architecture of SCATTER PHY. The grey block contains the implementation considering the USRP device itself, including the developed FPGA blocks. The light blue block showcases the software portion that handles and processes the data coming from the radio, as well as the control signals coming from higher layers (MAC).

- Timed-commands: this feature allows the configuration of the exact time in the future to (i) start a transmission and (ii) change Tx/Rx frequencies/gains. This allows the MAC layer to implement a TDMA scheme.

Figure 3.1 illustrates the different software/hardware layers composing the SCATTER PHY and the threads within each one of them. Red dashed arrows indicate data paths while black arrows indicate control/information interaction between threads.

SCATTER PHY is a discontinuous transmission-based PHY, which transmits data bursts in small transport units called subframes. A subframe is a container through which user and control data is exchanged in the network. The SCATTER PHY is built upon the srsLTE library [31], and therefore absorbs and evolves on top of the existing LTE features. srsLTE is a free and open-source LTE software library developed by Software Radio Systems (SRS) [31]. We adopt Orthogonal Frequency-Division Multiplexing (OFDM) as the SCATTER PHY waveform. OFDM is a mature technology, widely implemented in a variety of products due to its several advantages, such as robustness to severe multipath fading, low implementation complexity, easy integration with MIMO,

facilitated channel estimation, etc. [32]. User and control data are mapped into subcarriers over 14 OFDM symbols spanning 1 ms. The control data signal carries the used modulation and coding scheme (MCS) for the current transmission and the number of subsequent subframes with user data modulated with that specific MCS.

The control data signal is used at each one of the PHY receivers to automatically detect the number of allocated RBs, the location of the allocated RBs in the resource grid, and the MCS used to transmit data of a specific user. By embedding the control information into the transmitted signal, the Medium Access Control (MAC) layer does not need to know in advance the number of subframes and MCS in a given channel occupancy time (COT). Upon correct user data decoding, each PHY informs the MAC layer of the number of received bytes, the corresponding MCS, and decoding statistics (e.g. RSSI, SINR, number of detection/decoding errors, etc.).

SCATTER PHY employs discontinuous (i.e., bursty) transmission of subframes. A subframe is the basic transmission unit of the SCATTER PHY and each one is 1 ms long. SCATTER PHY works with two types of subframes, namely, synchronization and data-only subframes. A synchronization subframe carries the synchronization signal, reference signal, control, and user data. A data-only subframe carries the reference signal and user data. The synchronization signal is a 72 symbols long sequence that is generated using Zadof-Chu sequences [33]. The control signal is based on maximum length sequences (M-sequences) [34], where two M-sequences, each of length 31, are used to transmit the information necessary to decode the user data. The reference signal is used to estimate the channel and then equalize the received signal so that channel interference to the desired signal is minimized.

The SCATTER PHY allows bursty transmissions with variable COT, i.e., the number of subframes to be transmitted in sequence without any gap (i.e., idle time) between them is variable. The number of subframes in a COT is derived based on MCS, PHY BW and data length (i.e., number of bits to be transmitted) parameters sent by the MAC layer in the control message. The minimum COT is equal to 1 ms and is equivalent to the synchronization subframe. Variable COT enables the support of different traffic loads and channel occupancy. Every subframe can carry a pre-defined number of

bits, which is based on the MCS and current PHY BW.

Each one of the PHY modules is split into three sub-modules, namely, PHY Tx, PHY Rx Synchronization and PHY Rx Demodulation where each one of them runs on an exclusive, standalone thread. The reason for having a multi-threaded PHY implementation is that it allows independent critical and/or time-consuming tasks to be executed simultaneously (i.e., concurrently), which improves computing performance and efficiency. Allied with multi-core enabled Central Processing Units (CPUs), the multi-threaded PHYs naturally support full-duplex communications mode, i.e., each one of the two PHYs can simultaneously transmit and receive at different frequencies, which consequently results in higher throughput. The PHY Tx threads (#0 and #1) are responsible for modulation and transmission of data (i.e., user and control data). The PHY Rx Synchronization threads are responsible for the detection of the synchronization (Synch) signal, decoding of the control data, carrier frequency offset (CFO) estimation/correction and subframe time-alignment tasks. Detection of the Synch signal is carried out through a two-stage detection algorithm, which at the first stage correlates the received signal with a locally-stored version of the synchronization subframe with no data and control signals. If the peak-to-side-lobe ratio (PSR) is greater than a constant threshold, then, the second stage applies a Cell-Average Constant False Alarm Rate (CA-CFAR) algorithm to the OFDM symbol carrying the Synch signal [33]. The two-stage approach employed by the SCATTER PHY improves the Synch signal detection when compared to a detection approach that only uses the PSR of the correlation calculated at the first stage.

The CFO estimation task is split into coarse and fine estimations/corrections, where the coarse estimation is based on the Synch signal and the fine estimation is based on the Cyclic Prefix (CP) portion of the OFDM symbols [35]. The integer part of the frequency offset (i.e., integer multiples of the subcarrier spacing) is estimated and corrected by the coarse CFO algorithm, which is based on the maximization of the correlation of the received synchronization signal with several locally generated frequency offset versions of it. On the other hand, the fractional frequency offset (i.e., offset values less than one half of the subcarrier spacing) is estimated and corrected by the fine CFO algorithm,

which is based on the phase difference of the correlation between the CP and the last part of the OFDM symbol (i.e., the portion used to create the CP).

The PHY Rx Demodulation threads take care of user data demodulation, i.e., OFDM demodulation (FFT processing and CP removal), channel estimation/equalization, resource demapping, symbol demodulation, de-scrambling, de-interleaving/de-rate matching, turbo decoding, de-segmentation and, CRC checking. Each PHY receives data and control messages from the ZeroMQ Data/Control module. Decoded user data and statics related to the PHY operation (Rx/Tx statistics) are sent directly to the upper layers through the 0MQ bus. Regarding numerology, the SCATTER PHY uses subcarrier spacing of 15 kHz, a CP of 5.2 μs and supports 1.26, 2.7, 4.5, and 9 MHz bandwidths. Each PHY channel bandwidth (BW) can be set through the command line at start up, and changed in real-time through Tx control messages.

SCATTER PHY frame structure is depicted in Figure 3.2. As can be seen, synchronization and control data signals are added only to the very first subframe of a COT. With this frame structure, synchronization (i.e., detection of synchronization signal, time-alignment and CFO estimation/correction) and control data decoding only happen for subframe #0, i.e., the synchronization subframe.

The Payload Data Unit (PDU) adopted by SCATTER PHY is a Transport Block (TB). Therefore, a TB is the payload coming from the MAC layer and given to the individual PHYs to be encoded and transmitted over the air. One TB consists of a number of bits that can be accommodated within a 1 ms long subframe given the selected PHY BW and MCS. Therefore, given the PHY BW and the desired MCS, the MAC layer can find the number of bits that can be handled by a 1 ms long subframe. Table 3.1 presents the coding rate for each one of the 32 defined MCS values.

The communication between the SCATTER PHY and the MAC layer is carried out through the exchange of four pre-defined messages. The first two, namely, Tx and Rx Control messages, are used to manage subframe transmission and reception respectively. The parameters carried by these two messages can be configured and sent to the individual PHYs by the MAC layer before the transmission of every subframe, hence allowing runtime configuration. The other two messages, namely, Tx and Rx

Table (3.1) MCS values and their Respective Code Rate.

		Code Rate			
MCS	**Modulation**	**1.26 MHz**	**2.7 MHz**	**4.5 MHz**	**9 MHz**
0	QPSK	0.0857	0.0900	0.0940	0.0940
1	QPSK	0.1143	0.1200	0.1220	0.1220
2	QPSK	0.1357	0.1467	0.1480	0.1480
3	QPSK	0.1714	0.1933	0.1920	0.1920
4	QPSK	0.2107	0.2400	0.2400	0.2440
5	QPSK	0.2571	0.2933	0.2960	0.2920
6	QPSK	0.3071	0.3400	0.3440	0.3440
7	QPSK	0.3571	0.4000	0.4160	0.4130
8	QPSK	0.4071	0.4667	0.4640	0.4690
9	QPSK	0.4714	0.5200	0.5360	0.5330
10	16QAM	0.2357	0.2600	0.2680	0.2665
11	16QAM	0.2571	0.2933	0.2920	0.2905
12	16QAM	0.3000	0.3267	0.3280	0.3305
13	16QAM	0.3357	0.3667	0.3800	0.3825
14	16QAM	0.3857	0.4267	0.4290	0.4298
15	16QAM	0.4286	0.4667	0.4770	0.4718
16	16QAM	0.4429	0.5000	0.5090	0.5078
17	64QAM	0.2952	0.3333	0.3393	0.3385
18	64QAM	0.3238	0.3600	0.3553	0.3625
19	64QAM	0.3524	0.3911	0.4033	0.4087
20	64QAM	0.3905	0.4411	0.4353	0.4407
21	64QAM	0.4286	0.4678	0.4727	0.4727
22	64QAM	0.4571	0.5122	0.5047	0.5100
23	64QAM	0.4952	0.5478	0.5570	0.5642
24	64QAM	0.5333	0.5833	0.5970	0.6042
25	64QAM	0.5714	0.6189	0.6210	0.6308
26	64QAM	0.5905	0.6633	0.6690	0.6770
27	64QAM	0.6190	0.6900	0.7010	0.7010
28	64QAM	0.6571	0.8056	0.8067	0.8178
29	64QAM	0.6952	0.8322	0.8280	0.8552
30	64QAM	0.7333	0.8617	0.8493	0.8925
31	64QAM	0.7714	0.8883	0.8707	0.9240

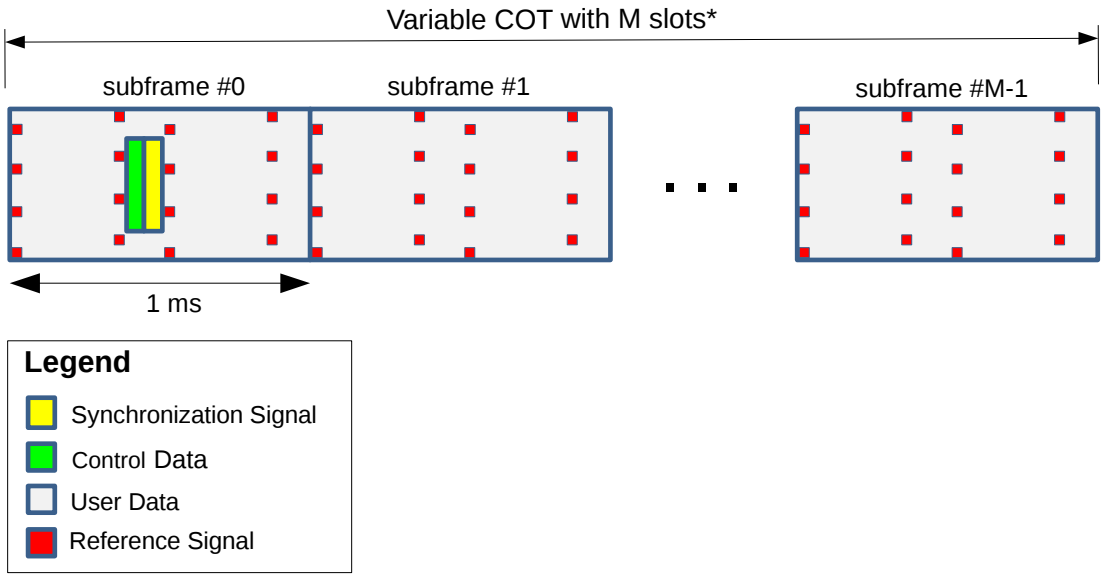

Figure (3.2) SCATTER PHY frame structure. The number of subframes depends on the channel occupancy time (COT). The first subframe contains synchronization and control data. All subframes contain the payload and reference signals, used for channel estimation and equalization.

statistics messages, are used to provide real-time feedback from each PHY to the MCA layer, yielding vital information necessary for such layer to take actions.

Tx control messages carry the user data (i.e., TB) to be transmitted and Tx parameters related to that transmission, namely, PHY ID, MCS, data length, Tx gain, Tx channel, Tx PHY BW, and transmission timestamp. The transmission timestamp parameter enables time-scheduled transmissions, which allows the MAC layer to implement a Multi-Frequency (MF) Time Division Multiple Access (TDMA) medium access scheme. **Rx control** messages are used to configure Rx channel, Rx gain, and Rx PHY BW of a specific PHY, which addressed through the PHY ID parameter. The PHY ID parameter is used to specify to each one of the two PHYs a control message is meant to.

The other two messages, namely, Rx and Tx statistics, are used to inform the MAC layer of the Rx and Tx processing results of each PHY, respectively. **Rx statistics** messages carry the PHY ID, received data and reception statistics related to the received data such as Channel Quality Indicator (CQI), Received Signal Strength Indication

(RSSI), decoded MCS, subframe error counter, decoding time, etc. **Tx statistics** messages inform upper layers of transmission statistics like coding time, the total number of transmitted subframes of each specific PHY. Table 3.2 summarizes all the real-time configurable parameters and statistics offered by the SCATTER PHY.

3.1.1 COT-based Filtering

OFDM-based waveforms are not suited for spectral coexistence due to their poor spectral localization [36]. This problem is caused by the rectangular pulse-shape used in OFDM, which leads to a sync-pulse property in the frequency domain with a very low second lobe attenuation of -13 dB [37]. Therefore, in order to guarantee a better spectral localization, i.e., lower out-of-band (OOB) emissions, and maintain the complex-domain

[1] Depends on the USRP daughterboard installed [28].

[2] Bandwidths: 1.26, 2.7, 4.5, and 9 MHz respectively.

Table (3.2) Real-time configurable parameters and statistics.

Message	Parameter	Type	Unit	Range
Tx control	PHY ID	uint32	-	0-1
	MCS	uint8	-	0-28
	Tx gain	uint32	dB	depends on HW [1]
	Tx channel	uint32	-	≥ 0
	Tx PHY BW	uint8	MHz	0-3 [2]
	Transmission timestamp	uint64	s	≥ 0
	User Data length	uint32	-	> 0
	User data	uchar[]	-	uchar range
Rx control	PHY ID	uint32	-	0-1
	Rx channel	uint32	-	≥ 0
	Rx gain	uint32	dB	depends on HW [1]
	Rx PHY BW	uint8	MHz	0-3 [2]
Rx statistics	PHY ID	uint32	-	0-1
	CQI	uint8	-	0-15
	RSSI	float	dBW	float range
	Noise	float	dBW	float range
	Decoded MCS	uint8	-	0-28
	Subframe error counter	uint32	-	≥ 0
	Decoding time	uint32	ms	≥ 0
	Data length	uint32	-	≥ 0
	Received data	uchar[]	-	uchar range
Tx statistics	PHY ID	uint32	-	0-1
	Coding time	uint32	ms	≥ 0
	Number of transmitted subframes	uint32	-	≥ 0

orthogonality of the OFDM symbols is to apply some sort of filtering to the time domain subframes. The filtering process is applied to each COT of each PHY independently. The subframes comprising a COT are generated at the SW level and then filtered at HW level, by a 128 order FPGA-based FIR filter. The COT-based filtering improves closer coexistence with other radios (either belonging to our team or others), allowing radio transmissions to be closer in frequency. The filter used in the SCATTER PHY is designed and explained in [38].

This co-design SW/HW is used so that fast-processing high order filters can be implemented adding up very low latency to the transmission chain and still allowing the flexibility of the software-defined PHYs. The filter's coefficients applied to the COT are automatically selected according to the configured Tx PHY BW (i.e., the coefficients are selected in real-time based on the Tx PHY BW field in the Tx control message) as it needs to have its cut-off frequency changed to exactly filter the desired signal's bandwidth.

3.1.2 Benefits of the SCATTER PHY

The first advantage of the proposed PHY is that the COT-based filtering makes the SCATTER PHY more spectral efficient as the OOB emissions are reduced. OOB emissions interfere with nearby radios (i.e., radios with nearby channels), decreasing the quality of their received signals, which consequently impacts the throughput experienced by those radios. The reduced OOB emissions make SCATTER PHY ideal for coexistence with other radios, allowing it to operate closer to them in the frequency domain and consequently, reducing spectrum wastage while increasing the spectral efficiency. The second advantage offered by SCATTER PHY is the possibility to configure in real-time all PHY parameters through the control messages.

3.2 Experiment Results

In this section, we present some experimental results in order to demonstrate the effectiveness and usability of the SCATTER PHY. All the experiments presented here were

Figure (3.3) Comparison of McF-TDMA feature. (a) FPGA-based filters disabled. (b) 128 order FPGA-based FIR filter enabled results in significantly lower out-of-band transmissions and thus reduced interference in adjacent bands.

carried out with the framework running on servers with Intel Xeon E5-2650 v4 CPUs (@2.2 GHz,30 M cache, 9.60 GT/s QPI, Turbo, HT, 12 Cores/24 Threads, 105 Watts) with 128 GB of RAM memory connected to x310 USRPs with 10 Gigabit Ethernet links, and equipped with CBX-120 RF daughterboards [39].

Figure 3.3 shows two spectrograms collected during 40 ms over a 31.25 MHz bandwidth with both PHYs set to operate concurrently, at each instant, at two of six 4.5 MHz channels. This experiment intends to show that the SCATTER PHY is able to generate two concurrent and independent transmit channels per node. In this experiment, each PHY transmits a random number of subframes at randomly selected channels. The channel number and the number of transmitted subframes, i.e., COT, are randomly selected between the ranges 0-5 and 1-3, respectively. Here, a gap of 1 ms between consecutive transmission is used. As can be noticed in both figures, (a) and (b), SCATTER PHY is able to independently transmit at two distinct channels with a different number of subframes. Figure 3.3 (a) shows the case where no filtering is enabled. As can be noticed, OOBE might cause interference to adjacent channels and consequently decrease other radios' throughput. Figure 3.3 (b) shows the case when the FPGA-based FIR filters are enabled. As can be seen, when filtering is enabled, the OOBE is mitigated and consequently, interference to and from adjacent channels is also mitigated. Additionally, another consequence of the filtering use is that the channel spacing could be made smaller and consequently decrease the spectrum wastage.

Figure 3.4 depicts throughput measurements taken with the SCATTER PHY for several PHY BW, MCS, and Duty Cycle (DC) values working in full-duplex mode (i.e., the two independent PHYs are simultaneously transmitting and receiving). We employ the full-duplex mode in order to check if it impacts somehow the measured throughout once in full-duplex mode the SCATTER PHY is being fully utilized. The measurement is taken for a COT of 20 ms with a gap of 1 ms between subsequent transmissions. This means a DC of 95.24 %. The throughput is averaged over 10 measurement intervals of 10 seconds each. During one measurement interval, the number of received bits is counted and then divided by the interval to calculate the throughput measured during that interval. The SNR on the link was set to 30 dB so that the packet reception rate for all MCS values was equal to 1. In the figure, for comparison reasons, we also add the theoretical maximum throughput achieved by the *Streaming* mode, where there is no gap between subsequent transmissions, i.e., a DC equal to 100 %. The theoretical maximum throughput is calculated dividing the TB size in bits for each MCS by 1 ms.

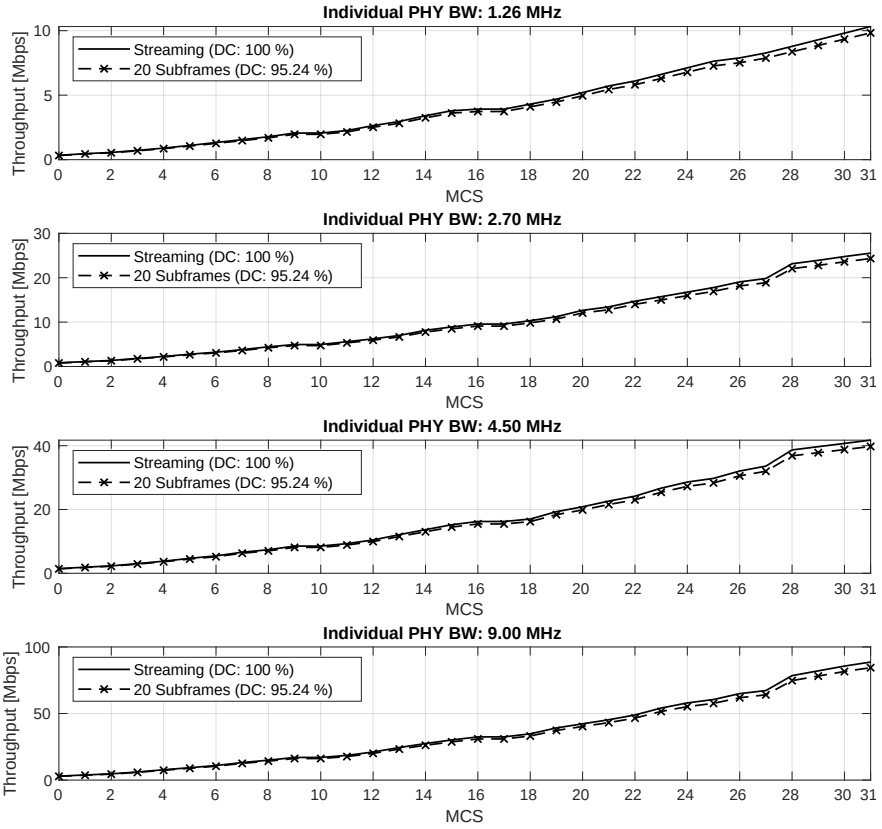

Figure (3.4) Throughput for different supported bandwidths, MCS values and duty cycles (DC) in full-duplex mode. The *Streaming* mode indicates the theoretical maximum throughput (duty cycle of 100%).

As expected, the measured SCATTER PHY's throughput is close to the theoretical maximum throughput for all MCS and PHY BW values, yielding more than 84 Mbps for a BW of 9 MHz and MCS 31. Additionally, the operation in full-duplex mode has no visible impact on the achieved throughput. This is due to the powerful server, with 12 cores, used to run the SCATTER PHY.

Figure 3.5 compares the mean absolute percentage error (MAPE) of the fine CP-based CFO estimation method against estimation based on the synchronization signal. The results were obtained by connecting the Tx port to the Rx port of the same USRP so that frequency offset caused by the HW is minimal or nonexistent and adding the desired frequency offset plus AWGN at the SW level just before the subframes are transmitted. As can be seen, the fine CFO estimation algorithm, which is based on the CP portion of the OFDM symbols, outperforms the synch-based CFO algorithm even when only two consecutive CPs are averaged. Additionally, we see that the performance of

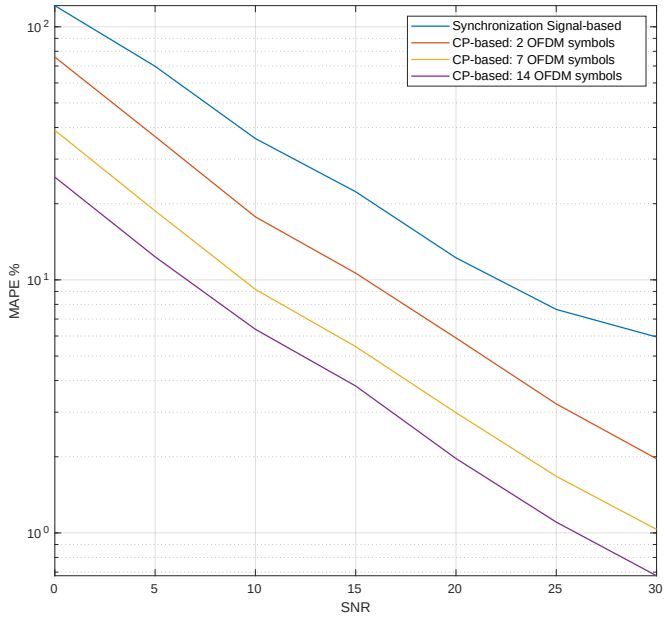

Figure (3.5) CFO estimation error improvements achieved by averaging specified numbers of cyclic prefix (CP) portions of consecutive OFDM symbols.

the CP-based estimation improves as the number of averaged CPs increases, however, the downside of averaging more CPs is an increase in the processing time. The CP-based CFO estimation method employed by the SCATTER PHY is an improvement over the CFO estimation implemented by the srsLTE. The SCATTER PHY implementation makes use of single-instruction multiple data (SIMD) instructions in order to decrease the CFO estimation and correction computational time allowing it to use more number of averaged CPs. Additionally, SCATTER PHY implements a more accurate and fine-grained complex exponential signal for generating the correction signal, making it possible to generate signals with frequencies that are closer to the estimated ones.

Figure 3.6 presents the comparison of the packet reception rate (PRR) for each one of the signals carried by a SCATTER PHY subframe, namely, synchronization, control and data signals over several SNR and PHY BW values. The SNR is calculated based on the power of a 1 ms subframe, therefore, before adding noise to the transmitted subframe, the subframe power is calculated and then the necessary noisy power to achieve the desired SNR is calculated. The MCS used for modulating the user data is set to 0, which is the most robust coding scheme, allowing SCATTER PHY to decode data in low SNR scenarios. The purpose of this experiment is to identify the lowest possible SNR

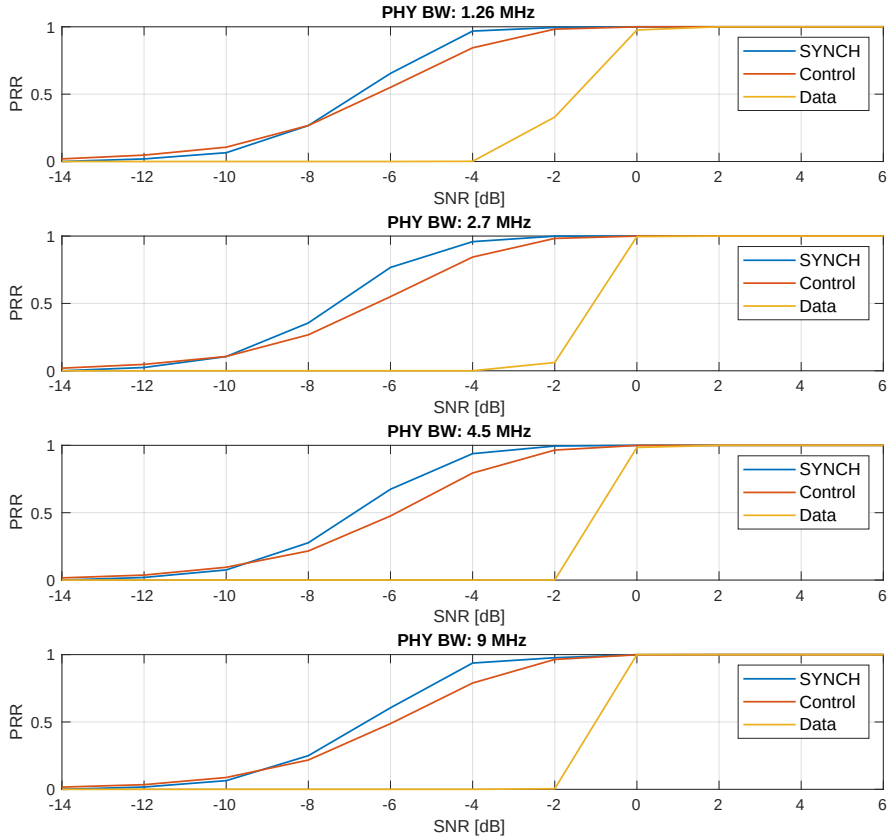

Figure (3.6) Packet reception rate comparison for supported signal types, identifying the lowest SNRs at which data can be correctly decoded at MCS 0.

at which SCATTER PHY can still correctly decode the user data. This experiment is run by adding a channel emulator between the Tx and Rx sides of a single PHY instance. At the Tx side, the generated subframes, instead of being sent to the USRP HW are sent to an abstraction layer that emulates the HW and adds AWGN noise to the transmitted signal, next abstraction layer transfers the noisy signal to the receiving side of the PHY. The PRR is averaged over 10^4 trials, where at each trial, the Tx side of the PHY sends a single synchronization subframe. As can be seen, synchronization and control signals have a better PRR performance than that of the data decoding, however, SCATTER PHY's PRR performance is limited by the ability of correctly decoding the data section of a subframe. Additionally, we see that the Data PRR is better for the 1.26 MHz case, which is due to the fact that compared to the other PHY BW values, MCS 0 for the 1.26 MHz case, carries more redundancy bits, as shown by Table 3.1, making it more robust against noise. The other two signals, synchronization, and control present similar PRR

curves for all PHY BW values. Moreover, it is noticeable that the data PRR is equal to 1 for SNR values greater than or equal to 0 dB.

3.3 Spectrum Access Coordination with CIL

Spectrum sharing and coordination has been an attractive research topic for a long time. A copious amount of work has been done in this field, with various ideas and technologies emerging to tackle the issue of spectrum scarcity from the point of view of coexistence with other wireless networks and protocols.

In [40], several different categories of heterogeneous spectrum sharing are mentioned – database mediated, TV white space, actively managed by a Spectrum Access System, LTE for heterogeneous networks, Bell Labs white cell architecture, Ericsson architecture. Most of these categories assume spatial reuse and centralized coordination and licensing.

A distributed coordination scheme was proposed in [41]. Authors first describe a rendez-vous protocol to form localized groups of users which would exchange information mutually, and then this information can be propagated to other groups.

A system optimization of LTE and WiFi to enable spectrum usage coordination has been studied in [42].

Authors have proposed a distributed spectrum coordination architecture via the Internet in [43].

Finally, [44] proposes a spectrum etiquette protocol called "common spectrum coordination channel" (CSCC), which includes announcing radio and service parameters to other networks.

Multiple other papers have been published, which affirms the attractiveness of the subject. Many of the proposed ideas are similar to the CIL. However, no practical, physical implementation and usage of such system was ever noted by the authors. This work presented in this thesis bases its evaluation on a protocol that has been used throughout the course of the SC2.

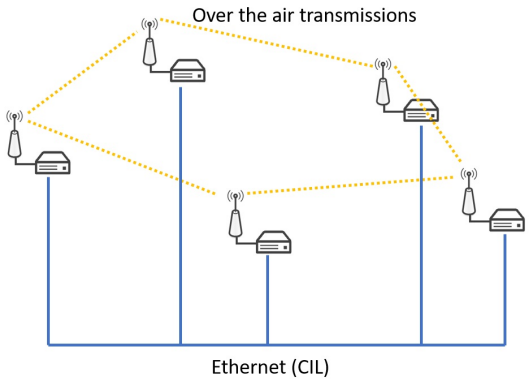

Figure (3.7) CIL design schematic. Yellow dashed lines present over-the-air transmitted signals between the nodes; blue solid lines represent Ethernet wires.

3.4 CIRN Interaction Language (CIL) Design

Collaboration between independent heterogeneous wireless networks is one of the many innovative ideas and contributions of the DARPA Spectrum Collaboration Challenge. The main goal of the collaboration is to improve the spectrum usage efficiency and coexistence by communicating intentions for future spectrum usage, informing other networks of the spectral needs and requirements and exchanging other types of information and requests. CIL's design, itself being a collaboration protocol, is also the result of the collaboration between all teams participating in the DARPA SC2.

CIL has evolved throughout the challenge, directed by input and feedback from all teams and DARPA itself, and has grown into a complex protocol that involves exchanges of various types of information and requests, and uses well-defined message structures. CIL was implemented as a publish-subscribe messaging system, where any CIRN can publish a message and all other networks are able to immediately receive it. There were multiple types of CIL messages, some used more than others. However, analyzing all types of messages and all layers of CIL functionality is beyond the scope of this work. This thesis will focus on evaluating the impact of one of the most widely used functions of the CIL – announcing the spectrum usage intentions by each CIRN. The thesis will show an attempt to measure the impact of this function on the achieved ensemble throughput of all CIRNs.

As shown on Figure 3.7, CIL is designed to handle the communication between

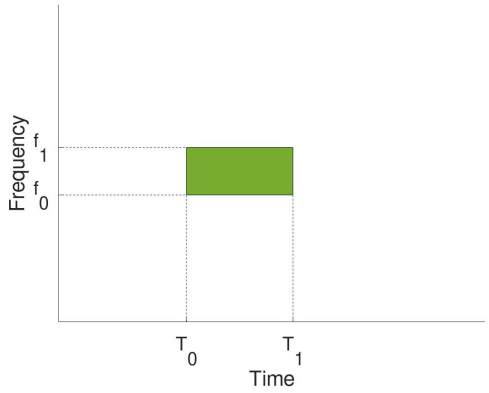

Figure (3.8) A spectrum voxel. In simple terms, it is defined by the exact time when a signal will be transmitted, and the precise frequency band it will use.

CIRNs independent of their over-the-air transmissions. Internal communication and data transfer between nodes within each CIRN was wireless, but all CIL related message exchange, as a protocol for direct communication with other CIRNs, was done over the wire. This ensures no interference in the spectrum caused by the CIL itself, as well as guaranteed delivery of CIL messages to all other CIRNs.

CIRNs are required to announce their spectrum usage by periodically sending the spectrum voxel specification to the CIL. A spectrum voxel, shown on Figure 3.8, can include varying degrees of information about the CIRNs intended spectrum occupancy. The required specifications are the exact time and frequency bands that the CIRN is planning to use for its transmission. Periodicity can be defined either by giving the exact start and end times of each transmission, or by specifying the duty cycle. The exact geographical location (latitude and longitude) of each transmitter can be given, along with the intended transmission power using any absolute units. This provides the opportunity for spatial reuse of the spectrum when nodes are geographically separated. [45]

3.5 CIL Evaluation Framework

The SC2 was run on Colosseum, a radio signal emulation testbed designed and built by DARPA. Colosseum contains a large number of nodes with USRP devices, with the

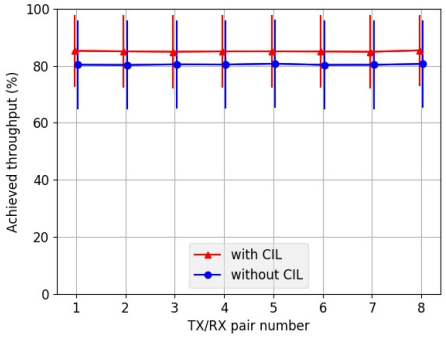

(a) Total throughput for 8 TX/RX pairs in 32 channels

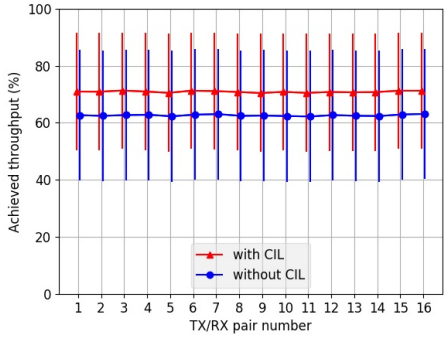

(b) Total throughput for 16 TX/RX pairs in 32 channels

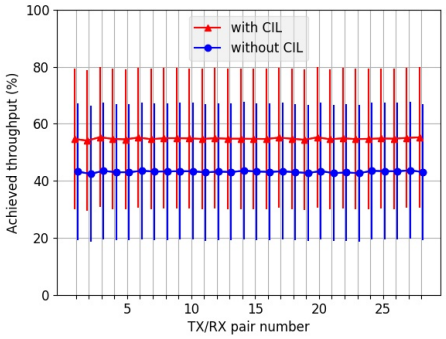

(c) Total throughput for 28 TX/RX pairs in 32 channels

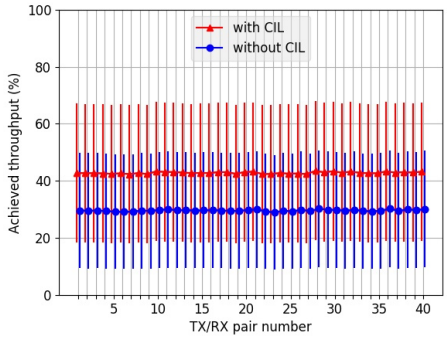

(d) Total throughput for 40 TX/RX pairs in 32 channels

Figure (3.9) Achieved combined throughput for 8, 16, 28 and 40 TX/RX pairs with 32 available channels. Throughput expressed as percentage of maximal theoretically achievable throughput by each TX/RX pair assuming no interference and 100% successful packet receiving and decoding.

entire wireless framework being emulated. The environment enables competition organizers to emulate any desired topological or spectral configuration, thus enabling them to run different scenarios and test teams' performance in various circumstances. Testbed users (the teams in case of SC2) can build their own LXC containers, which then run on Colosseum nodes and have direct access to the connected USRP devices. System clocks of all nodes are synchronized using NTP. CIL related communication is implemented and runs via Ethernet, and is thus independent of the wireless communication.

CIL was successfully used throughout the DARPA SC2. However, its performance and impacts upon the decisions made were never evaluated, as an environment to run similar competition scenarios without the usage of CIL was not available. The described environment in the Colosseum testbed did not include the possibility to disable CIL

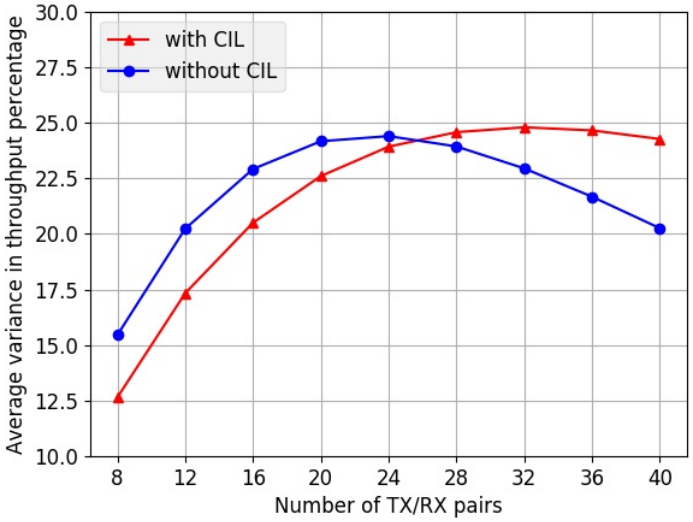

Figure (3.10) Average variance in throughput for all experiments, shown by the number of pairs. A higher number of pairs increases the variance

and run experiments with other teams' containers without it. Hence, it is unclear how significant the impact of CIL usage on the entire performance may be. Indeed, a proper evaluation would require comparison of two equal competition scenarios, which implies equal geographical distribution of nodes, exact same sets of participating teams, and equivalent traffic loads and requirements for each CIRN and their nodes. Without the proper environment provided by DARPA, it is not trivial or straightforward to design a framework that would support the entire CIL functionality and provide all resources and conditions that SC2 CIRNs were developed for and would expect to encounter in an official SC2 scenario.

Thus, the presented work focused on evaluating only the spectrum usage announcement functionality of the CIL. First, a basic evaluation system was built to emulate the simplified functionality of the CIL. The system consists of multiple USRP X310 devices, communicating over the air. Transmit/receive pairs of nodes were implemented with GNU Radio and used instead of complex wireless communication networks. The CIL was emulated by feeding the information about selected frequency bands by each TX/RX pair for future time slots to all transmitters. The evaluation was performed by running multiple throughput tests with varying number of nodes, with the simulation of CIL presence turned on and off. Each transmitter would repeatedly send a packet

with a pre-determined payload, and the receiver would attempt to decode the received message and check whether the value is correct. The throughput was measured by the number of successfully decoded messages for each TX/RX pair. A throughput of 100% would indicate that all transmitted packets were successfully decoded; similarly, a throughput of 10% would indicate that only one tenth of all packets were successfully decoded. It can be seen (shown in section 3.6) that with two transmitters using the same channel, the interference is such that their respective receivers can occasionally successfully decode a packet, but are unsuccessful for more than 99% of the transmitted packets. The results also show and analyze the variance in the achieved throughput over multiple runs of the experiment depending on the number of TX/RX pairs present in the experiment, and how this variance is affected by CIL.

The evaluation system was implemented under the assumption that spectrum usage announcement messages are simplified. The spectrum voxels in the evaluation do not include any geographical location data, so there is no spatial spectrum reuse, as the node distance can be easily calculated and all nodes with negligible interference can simply be ignored, so for the purpose of evaluating the impact of CIL on data throughput, including location information would be of little value. To simplify the experiments, time and frequency bands were observed in discrete non-overlapping slots, as any loosening of these constraints could significantly increase the complexity of spectrum usage decision making and the entire evaluation of spectrum voxel announcement impacts could be severely affected by the potentially non-deterministic nature of the AI subsystems of the CIRNs.

3.6 Evaluation Results

The evaluation system was designed as a set of transmitter/receiver pairs, where all pairs select a channel in each time slot and transmit data. Each run of the experiment was performed with and without the simulated CIL. During a run without the simulated CIL, each pair would select a channel and attempt to transfer data. If any other pair happened to randomly select the same channel, the two transmitters would interfere

with each other and receivers would be unable to properly decode the received packets, causing significant drops in throughput for each transmitter/receiver pair. A larger number of pairs (relative to the number of available channels) results in the increased probability of occurrence of such interference. The testing environment had no external interference caused by any transmissions other than the transmitters used in the experiment itself. A CIL simulated run would include announcements of the selected channels for future time slots to each transmitter/receiver pair. With the available information on channel occupancy in the future, each pair can now make an informed decision on whether to transmit data in the following time slot, or back off. The evaluation system assumes that for each transmitter/receiver pair, when usage of their selected channel for an upcoming time slot has also been announced by N other pairs, they will transmit with the probability of $\frac{1}{N+1}$. There are many possible, less naive techniques with varying degrees of complexity to further increase the combined throughput by making an informed decision on using a different channel, one which wasn't found in the announced spectrum voxels. However, the discussion and evaluation of a multitude of such techniques is beyond the scope of this thesis.

Figure 3.9 shows the average throughput for all time slots per TX/RX pair in a scenario with 8, 16, 28 and 40 different TX/RX pairs and 32 available frequency channels, for 500 experiment runs with 60 time slots, with each time slot being of 1 second duration. The achieved throughput is expressed as percentage of maximal theoretical throughput, which is achieved when data transmission in each time slot is performed with no interference and every single transmitted packet gets successfully received and decoded. The throughput for each experiment is shown on the y axis, and each plot has the y axis range of the entire 100% to emphasize the dependence of the benefits of CIL on the number of existing transmitters (i.e. on the spectrum shortage). With 8 pairs in 32 channels, spectrum voxel announcements and the naive back-off algorithm used by transmitter result in an increase of overall achieved throughput of about 5%, while with 40 existing transmitters this increase is above 12%.

Additionally, it is interesting to observe the effect that all parameters have on the

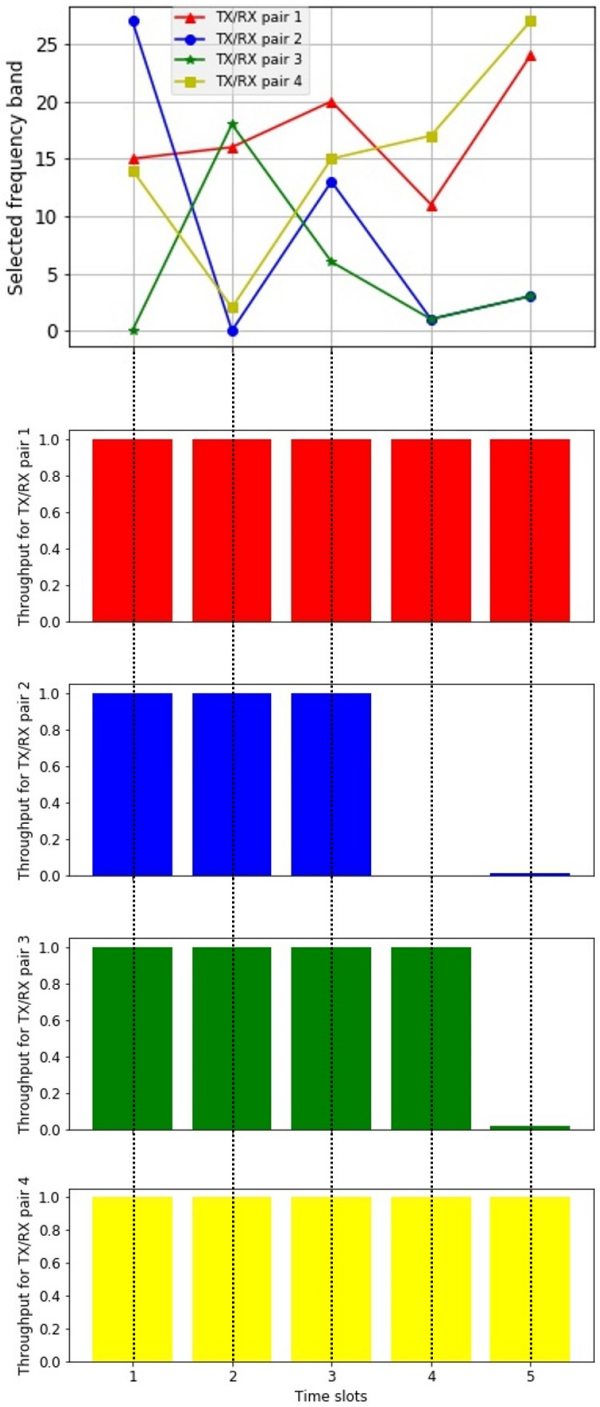

Figure (3.11) A schedule and throughput per pair of 4 TX/RX pairs, showing an example of channel selection collision and the resulting throughput.

variance in throughput through multiple runs of the experiment. With a higher number of TX/RX pairs it is notable that the variance also increases, which can easily be explained by the fact that a lower number of nodes that need spectrum access result in

a more orderly and predictable environment. The more detailed values of throughput variance per the number of pairs in the environment is shown on Figure 3.10. These values can be explained by the observation that a lower demand for spectrum results in a less variable environment. Presence of CIL further decreases the volatility of the environment in this case. However, with a higher number of TX/RX pairs, and thus proportionally higher demand for spectrum, even though CIL helps in increasing the ensemble throughput of all pairs, it also increases the volatility of the environment, indicated by the variance being higher when CIL is present. The explanation for this phenomenon can be found in the fact that the highest variance in an environment without CIL is seen for 24 pairs seeking spectrum access in 32 available spectrum channels, while with CIL the number of pairs with peak variance is exactly 32 – the same as the number of channels. Once the spectrum utilization reaches a saturation point, any further attempts at spectrum access will more often than not result in failure, which again increases the predictability of the system.

Figure 3.11 shows the schedule (selected channels) and achieved throughput of 4 TX/RX pairs over 5 consecutive time slots in a single run. It can be seen from the schedule plot (the first of five plots) that pairs 2 and 3 end up selecting the same channel in the last two time slots, namely slots 4 and 5. The throughput bar plots show that the pair 2 (blue) has no throughput in time slot 4, with some minor throughput in time slot 5, while pair 3 (green) has high throughput in time slot 4 with similarly minor throughput in time slot 5. This could indicate the possibility that in slot 4, the blue pair has randomly chosen to back off, while the green pair has decided to transmit data. In the last time slot, both nodes have made the decision to transmit data, but due to the interference they posed to each other they have both been able to decode only a negligible number of packets, which resulted in non-zero but negligible throughput for both pairs.

Overall, the performed experiments show that announcing the intended spectrum voxels used for transmission can lead to an increase in combined throughput of all nodes and systems of between 4 and 12%, depending on the spectrum demand and availability. These gains were achieved with the naive random back-off algorithm described in section

3.5.

3.7 SC2 Aftermath and Outcomes

The DARPA Spectrum Collaboration Challenge ended in October 2019, at the Mobile World Congress in Los Angeles, CA. It concluded by leaving the wireless research community with a powerful wireless testbed, Colosseum, as well as new research ideas and topics, mainly related to improving wireless spectrum usage efficiency by using machine learning and AI. This chapter explained the main ideas behind the challenge, showcased the design and implementation of the physical layer of team SCATTER, and finally addressed the benefits a network interaction protocol may have on the combined ensemble throughput in a given wireless spectrum, as opposed to scenarios when networks are not able to interact with each other in any way.

Chapter 4

Experimental Spectrum Access Coordination With CIL and SCM

As the previous chapter already mentioned, DARPA's CIL was one of the first practical implementations of a spectrum access coordination protocol that was used by real communication systems. This represents an important milestone, as such protocols have recently started gaining traction at an accelerated pace in the wireless communications research community. Specifically, the bulk of the effort was focused on coexistence of LTE and WiFi in particular, with several proposed solutions for LTE in unlicensed spectrum.

It is important to note that the idea of coordinated spectrum access is not recent. In addition to the work presented in the previous chapter, several protocols of spectrum coordination in various forms were suggested over the past decades. [46] proposed a spectrum marketplace in 2005, with a regional spectrum access broker named DIMSUMnet coordinating spectrum access. In this case the spectrum access resolution is centralized. Similar ideas are still evolving today, such as the one described in Chapter 2. In the same year, [47] introduced another protocol named DSAP, which was envisioned to complement DIMSUMnet, a broker for larger regions, and serve as a spectrum broker for smaller, densely populated areas.

[48] presents an opportunistic protocol for spectrum access coordination between independent networks and wireless protocols. The protocol includes a central Cognitive Radio (CR) terminal which assigns spectrum in order to establish "fairness" based on data flows. [49] introduces a simple message exchange protocol, in many ways similar to CIL. The protocol was named Common Spectrum Coordination Channel, and operated in a separate, narrow frequency band to allow networks to exchange simple message types

and announce their spectrum usage. The evaluation includes an ns-2 based simulation. [50] shows a semantics based algorithm which uses FFT analysis and energy detection with semantic reasoning to determine available frequency bands for transmission. The paper additionally distinguishes itself with a practical implementation of their algorithm, using OpenAirInterface (OAI) to deploy their LTE communication nodes. Another semantic approach is shown in [51]. [52] serves as an example of a decentralized protocol. The authors introduced this protocol with 5G in mind. No monetary transactions are assumed, but the protocol uses a virtual spectrum access currency which serves as a bargaining chip for spectrum access negotiation, which can happen according to a given set of rules.

A large number of publications in the past decades address this issue. Their approaches vary greatly, and can range from centralized to decentralized, may or may not include cooperation, and use improvements and modifications of existing protocols in different communication layers to reach the common goal of improving spectrum utilization for everyone. With this in mind, CIL can be seen as an evolution of a set of these protocols, designed as an interaction language between heterogeneous networks, enabling them to negotiate but also cooperate to find the optimal ways to utilize the available spectrum bands.

CIL's implementation was targeted for the specific needs of teams participating in the DARPA Spectrum Collaboration Challenge, and as such used a relatively simple method to describe spectrum usage. Spectrum access was described using spectrum voxels, which in simple terms mainly serve to define the specific bands a transmitter intends to use, transmission times and geographical location. The described voxels served their purpose during the competition, but their simplicity leaves a lot of room for improvement. For example, some protocols could coexist in the same bands with better spectrum utilization than others, but voxels don't allow the spectrum usage description to be this specific.

Spectrum Consumption Model (SCM), an IEEE standard 1900.5.2-2017 described in [53], includes spectrum masks as significantly more detailed descriptions of each transmitter's characteristic. The standard also defines the specific algorithms that can

be used for compatibility calculations between different transmitters and/or receivers. Networks can use this to determine whether their intended transmissions could cause harmful interference to other networks' nodes. DARPA's CIL, now open sourced and publicly available, was used and modified to utilize SCM messages (SCMs) instead of Spectrum Voxel messages as part of the NSF funded Spectrum Innovation Initiative program, a joint project between Rutgers University, Columbia, NYU, U. Arizona, UT Austin, Oregon State, Princeton, and U. Wisconsin-Madison. The described implementation showcases successful exchange and utilization of SCMs via CIL to coordinate spectrum access in the ORBIT Wireless Testbed. As opposed to Spectrum Voxels from the DARPA SC2, SCM allows for a significantly more detailed explanation and characterization of spectrum occupancy. For example, rather than just guessing and using trial and error to assess whether the out-of-band emissions from other networks' transmitters will interfere with its own, a network can use the provided algorithms and calculation methods to determine the exact interference other transmissions will present to its precise spectrum mask, and fit its transmissions as close as possible in frequency to other networks to minimize unusable "holes" in spectrum occupancy.

In addition to the greater potential to improve spectrum utilization with the detailed spectrum masks and usage characterization, SCM also represents a more standardized, well documented solution, facilitating spectrum access coordination for new networks. The advantages that a network interaction language such as CIL, along with SCMs, offers for spectrum coordination seem intuitively clear. However, similarly to the aforementioned issues with Spectrum Voxel usage in CIL, it is unclear how big of an improvement this could be. These advantages have never been measured or quantified. This thesis aims to not only provide some measurements of network performance, expressed as ensemble throughput of all wireless networks present in a given collision domain, using transceiver topologies, traffic requirements and AI based spectrum occupancy prediction methods that resemble real world scenarios, but it also attempts to provide an extensible framework that can be used to facilitate similar measurements in the future.

With a significant amount of research efforts aimed at solving and improving co-existence of heterogeneous networks in the same spectrum bands, it is important to note that the bulk of the work was theoretical, occasionally backed with simulations (often in MATLAB or a network simulation framework such as ns-3), but that very few scientific papers mention any practical implementations of their research. This chapter presents work that aims to fill that gap, and show an implementation of a coordination protocol, based on CIL as a prime example of an already successfully used protocol for spectrum access coordination during the DARPA Spectrum Collaboration Challenge, used to exchange SCM messages which are now part of an IEEE standard. The relevant publication about the presented work can be found in [54].

4.1 Background: Spectrum Consumption Models

Spectrum consumption models (SCMs) provide an information model that can capture the boundaries of the use of spectrum by RF systems so that their compatibility (i.e. non-interference) can be arbitrated by efficient and standardized computational methods [53, 55, 56]. The information captured in SCMs allows for efficient determination of aggregate interference levels and of aggregate compatibility interference between many devices.

The IEEE 1900.5.2 standard for modeling spectrum consumption specifies 11 constructs for an SCM:

1. Reference power: This value provides a reference power level for the emission of a transmitter or for the allowed interference in a receiver. It is used as the reference power value several other SCM constructs (i.e. spectrum mask, underlay mask, and power map).

2. Spectrum mask: Data structure that defines the relative spectral power density of emissions by frequency.

3. Underlay mask: Data structure that defines the relative spectral power density of allowed interference by frequency.

4. Power map: Data structure that defines a relative power flux density per solid angle.

5. Propagation map: Data structure that defines a path loss model per solid angle.

6. Intermodulation mask: Data structure that defines how co-located signals generate intermodulation products in a transmitter or receiver.

7. Platform name: A name or list of names of platforms that are attributed to a particular site (i.e., ship, airplane, etc.). They are useful in identifying when multiple systems are co-located.

8. Schedule: Construct that specifies the time in which the model applies (start time, end time). Periodic activity can also be defined.

9. Location: The location where an RF device may be used. Several types of locations and trajectory/orbit descriptions are supported.

10. Minimum power spectral flux density: A power spectral flux density that when used as part of a transmitter model, implies the geographical extent in which receivers in the system are protected.

11. Policy or protocol: A named protocol or policy with parameters that define behaviors supported by a device or systems that allow different systems to be co-located and to coexist in the same spectrum.

These constructs can be used to build different types of SCMs that follow an aggregation hierarchy as shown in Figure 4.1. It is worth noting that depending on the type of model and its purpose, not all constructs are required. Figure 4.1 shows the relationships between different types of SCMs as defined in the IEEE 1900.5.2 standard [55]. A transmitter model captures the extent of RF emissions of an active radio device, including but not limited to: spectral emission mask, propagation map, antenna radiation pattern, possible locations of the device and times of operation. A receiver model conveys what is harmful interference to an RF device, providing a limit to the aggregate interference that transmitter devices can cause to a receiver in the temporal, spatial and

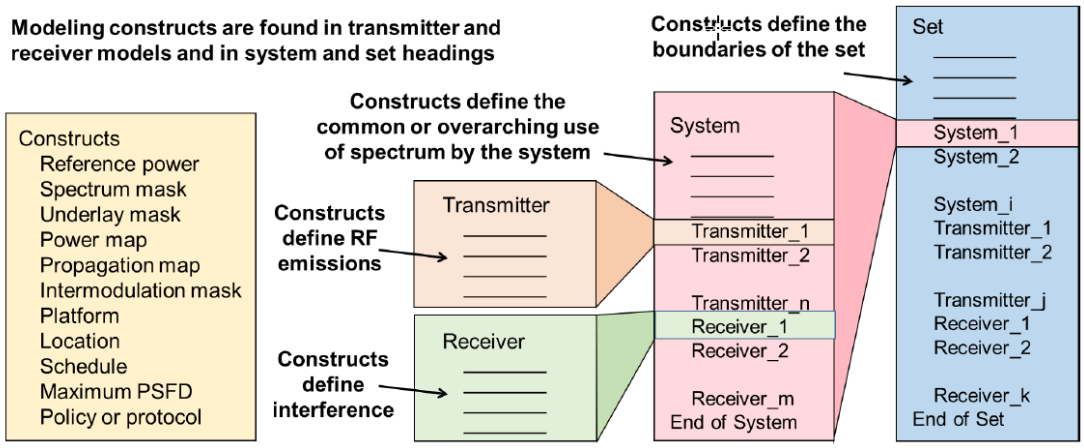

Figure (4.1) Spectrum Consumption Model types

spectrum dimensions. System models are a collection of transmitter and receiver models that collectively capture the spectrum use of an RF system. An SCM set is a collection of system, transmitter and receiver SCMs. SCM sets can be used to structure lists that describe spectrum that is available for use (Spectrum authorization sets), identify constraints to spectrum use (Spectrum constraint sets), and to list the spectrum being consumed (used) by a group of systems and devices (Collective consumption set)

In addition to a definition of the constructs for SCMs, the IEEE 1900.5.2 standard specifies a method for computing the compatibility of spectrum use between devices and/or systems that have expressed the boundaries of their spectrum use via SCM [53, 55]. Depending on the locations of the devices for which compatibility is to be assessed and if there is overlap in their spectrum use operations (in time and in frequency), the information conveyed by their SCMs related to transmitter spectrum masks, receiver underlay mask, reference power, power map and propagation map, among other constructs, determine the details of a link budget computation. In the case of a single transmitter-receiver pair, if the RF signal power from the transmitter is determined to be below the interference limit value specified in the receiver model, the pair is determined to be compatible (i.e. they can share spectrum). This computation can extended to multiple transmitter and receiver pairs making use of the information in the SCMs to compute aggregate interference values. Other spectrum use characteristics such as intermodulation and frequency hopping behavior can also be taken into account. More

details are provided in [55].

4.2 Architecture for Dynamic Spectrum Management

Our experimentation on DSA mechanisms is part of our ongoing research in the development of a spectrum management architecture composed of four distinct functional planes: monitoring and measurement plane, wireless data plane, spectrum control plane and cloud based spectrum service plane as shown in Fig. 4.2. The radio devices and network on the data plane are represented by wireless domain (WD) controller in the control plane. A WD controller can represent one or more wireless networks operating in a single administrative domain. Spectrum consumption information from individual wireless networks is aggregated at the domain controller level and the spectrum control plane is heavily based on DARPA's Collaborative Interaction Language (CIL), allowing for the exchange of SCMs messages as mentioned earlier between the domains as well as to control radio nodes in the data plane. In addition to peering of SCM exchange enabled by the control plane, a cloud based spectrum service layer is introduced to accommodate hierarchical control with benefits of centralized optimization involving complex AI/ML algorithms. The cloud service layer provides spectrum management, monitoring and

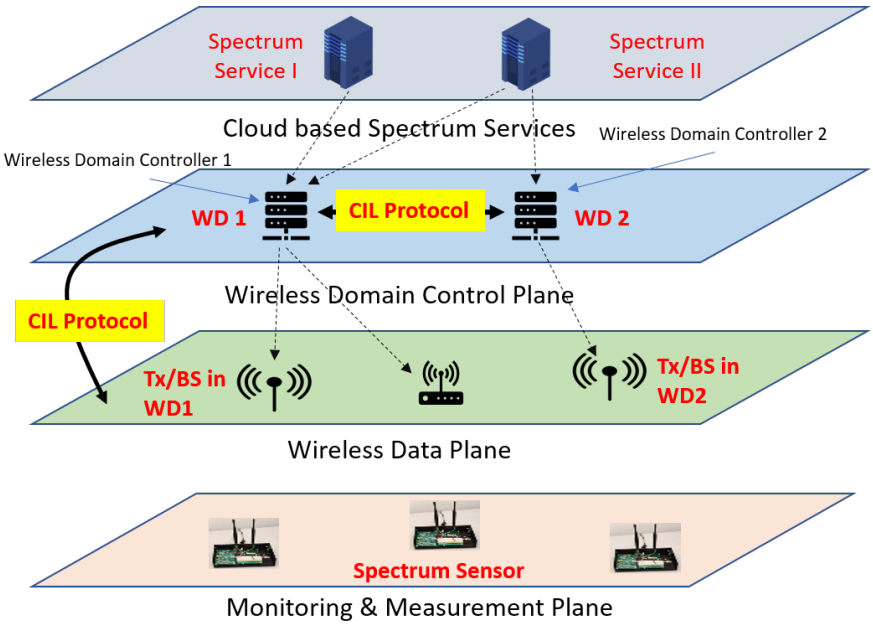

Figure (4.2) Overview of decentralized spectrum management architecture

marketplace capabilities to which WD's can subscribe. The data, control and cloud service planes are further supported by an independent spectrum monitoring infrastructure intended to provide accountability for actual spectrum use. The monitoring plane collects and aggregates sensor data which is then passed up to a spectrum analytic service in cloud layer for further processing. The analytic/monitoring application in the cloud disseminate this information to WD controllers using information centric PUB/SUB techniques.

The CIL protocol between the domain controllers and radio nodes supports several types of messages as listed below. More details on protocol exchange can be found in Section 4.4.

1. ***Register ()*** : Generated by WD to register with collaboration server/system
2. ***Inform ()*** : Informs newly joined peer about existing peers
3. ***Notify ()*** : Notifies existing peers about the new joined peer
4. ***SCM Request ()*** : Message to request SCMs from peers
5. ***SCM ()*** : Message to send SCM to the requester
6. ***CT Report ()*** : Sends compatibility test report to peers
7. ***Calibrate Radios ()*** : Message to calibrate SDRs with respective gain, frequency, modulation, etc.
8. ***Leave ()*** : Generated by WD to exit the system

4.3 SCM Generation and Processing

The interactions in our Wireless Domain Control plane will rely on the exchange of SCM messages that characterize the transmitters and receivers of each WD. The required constructs for the transmitter models are: reference power, spectrum mask, power map, propagation map, schedule and location. For the receivers, the spectrum mask is not required but an underlay mask is. For our experiments, we used a set of USRP devices

as transmitters and another set as receivers. The schedule and location of each transmitter and receiver device was well known. For simplicity, all transmission used BPSK modulation with a channel bandwidth of 1 MHz. Additional details on the characterization of the devices and the processing of their SCMs are mentioned in the following subsections.

4.3.1 Transmitter characterization

In order to characterize and build the SCM for a transmitter (Tx) USRP, we first obtain the power spectral density (PSD) plot that will help us build the spectrum mask for the device. This is performed by setting Tx gain to 10dB, resolution bandwidth to 100KHz or lower and amplitude 500mV. The received samples are captured at a distant receiver (Rx) setup to operate at the same central frequency as the Tx and the distance between Tx and Rx is reported. Next, we change the amplitude values at the Tx to understand how well it translates to radiated power. The separation distance between the devices should always be at least 1 meter. Finally, we repeat the same experiment again but with a different receiver and report the distances between Tx and Rx. Capturing the Tx's radiated power at receivers located at different distances provides details to elaborate the propagation map construct needed in the SCMs.

4.3.2 Receiver characterization

To characterize the receiver USRP, determining the shape of its underlay mask is key. For this purpose, we fix the Tx amplitude to 500m, Rx and Tx gain to 10dB and the center frequency at receiver to 2 GHz. From the Tx, transmit a BPSK modulated signal while varying the central frequency and capture the PSD image and SNR value at each frequency value at the receiver. The central frequency of the transmission is varied in steps of 200KHz from 1997MHz to 2003 MHz. Next we gather data to determine the allowable interference on the receiver considering Fig. 4.3 as an example topology. From Tx1, we transmit a BPSK modulated signal with a center frequency of 2 GHz and gain 10dB. At the receiver, we make sure that we can demodulate Tx1's BPSK signal and capture the signal at the receiver while Tx1 is ON and Tx2 is OFF. Continuing with the

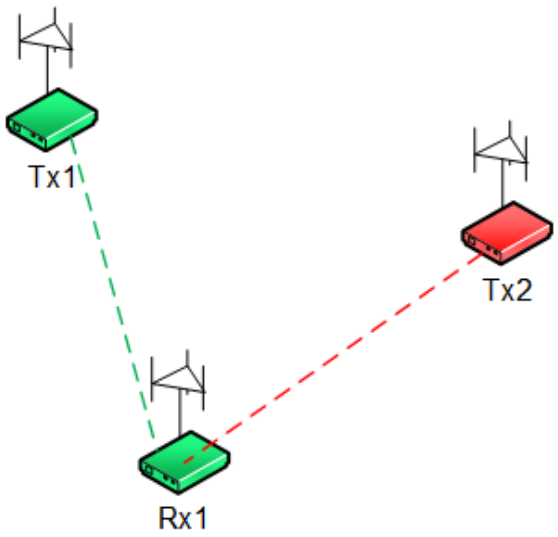

Figure (4.3) Receiver Characterization

same setup as before, while Tx1 is still transmitting, Tx2 is turned on and transmits a QPSK modulated signal (interfering signal) with a center frequency of 2 GHz. The initial Tx2 gain should be 10dB, and later we change Tx2's gain until we get to a value which we will call X where Rx1 will stop demodulating Tx1 transmissions and if Tx2's gain is later lowered by 0.5 dB (i.e. Tx2's gain becomes X-0.5dB), Rx1 will be able to demodulate some of Tx1 transmissions. We report the value of X and generate a signal capture on Rx1 when Tx2 is operating with gain X and Tx1 is ON. Finally we also determine the available interference on the receiver considering a displaced frequency of 1999.6MHz. The above described procedure is repeated using this frequency and the value of X and signal is captured.

4.3.3 Compatibility computations

An overview of the structure of the transmitter spectrum mask and receiver underlay mask obtained after our characterization process is shown in Fig. 4.4.

Using the data contained in the SCMs, calculations are performed to determine a bound that specifies by how much power can the spectrum mask of a transmitter be adjusted to be at the threshold of compatibility (i.e. non-interference) with the underlay mask of a receiver. This computation is known as the power margin computation, and its result can be used to evaluate how spectrum reuse and/or sharing opportunities can be

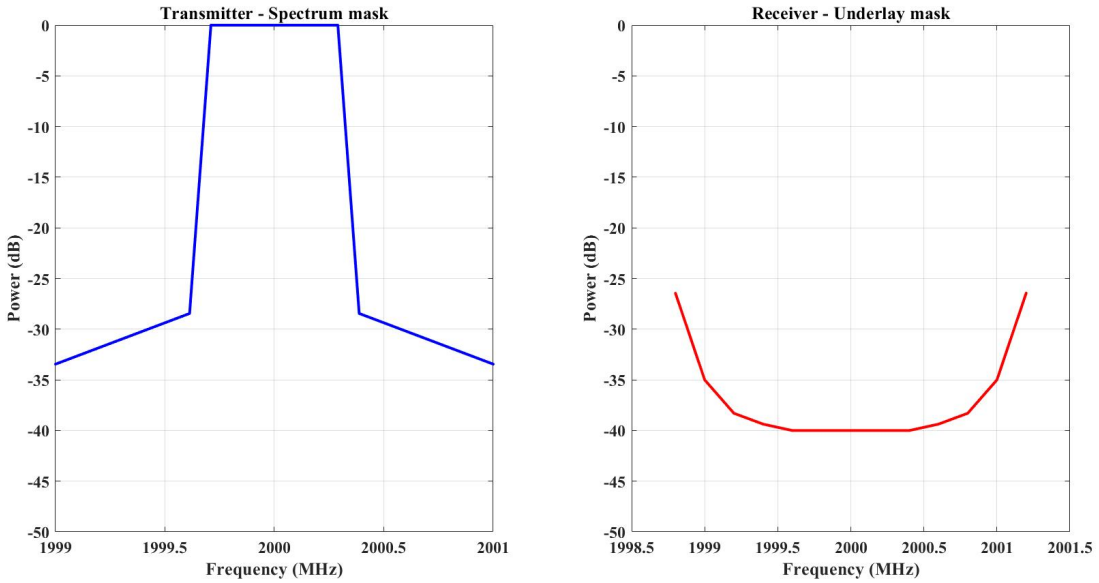

Figure (4.4) Transmitter spectrum mask and receiver underlay mask

leveraged. Two types of power margin computations are specified in the IEEE 1900.5.2 standard: (a) maximum power density and (b) total power. The maximum power density method determines if the maximum power spectral density of a transmitter's spectrum mask (after propagation and antenna gains) exceeds any threshold level of the receiver's underlay mask. The total power method uses the underlay mask to determine the total power from the transmitter that enters the receiver and check if it is less than some threshold that would otherwise be characterized as unacceptable interference. We used the total power method in our experiments.

4.4 Experiment Design and Implementation

Our experimentation framework for dynamic spectrum access interactions was built on top of the COSMOS Wireless testbed. As described earlier, the framework aims to enable wireless networks to exchange messages in order to coordinate and synchronize their spectrum access, as well as facilitate the computations necessary to determine available spectrum bands and avoid collision with other networks present in the same spectrum.

The network interaction language was built on top of DARPA's implementation of

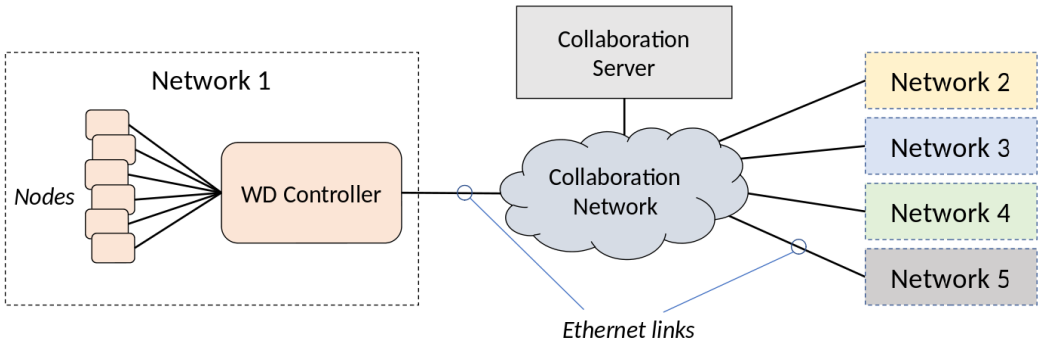

Figure (4.5) Network Interaction Language design schematic. The domain controller of a network can exchange messages with other networks' domain controllers.

the CIRN Interaction Language (CIL). CIL was originally designed as a PUB-SUB message queuing system, with several types of messages being broadcast, and all networks able to receive them. The new interaction language developed for our experiments has introduced several adaptations and improvements. Figure 4.5 shows the design schematic of the language. Similarly to CIL, each network has a designated node, a Wireless Domain controller (WD), which can use the interaction language to communicate with other networks' WDs. To maintain the high performance and efficiency of the message queuing service, Google's Protocol Buffers are still used to convert all supported messages into binary blobs. However, the message types were adapted to accommodate newer spectrum usage description standards. The networks are now able to exchange Spectrum Consumption Model (SCM) messages which provide a standardized and detailed description of spectrum usage [53].

Another advantage gained by using SCM based messages, in addition to the level of detail the messages provide, is that the standard on which they are based, also includes the algorithms for computations necessary to determine spectral compatibility between different transmitters or receivers. The framework heavily utilizes these computations. Compatibility computations are included as a component of the interaction language. First, with each new network joining the framework, the interaction server performs the compatibility computations, and only sends its SCM information to the networks in the same interference domain. Additionally, the networks themselves can invoke

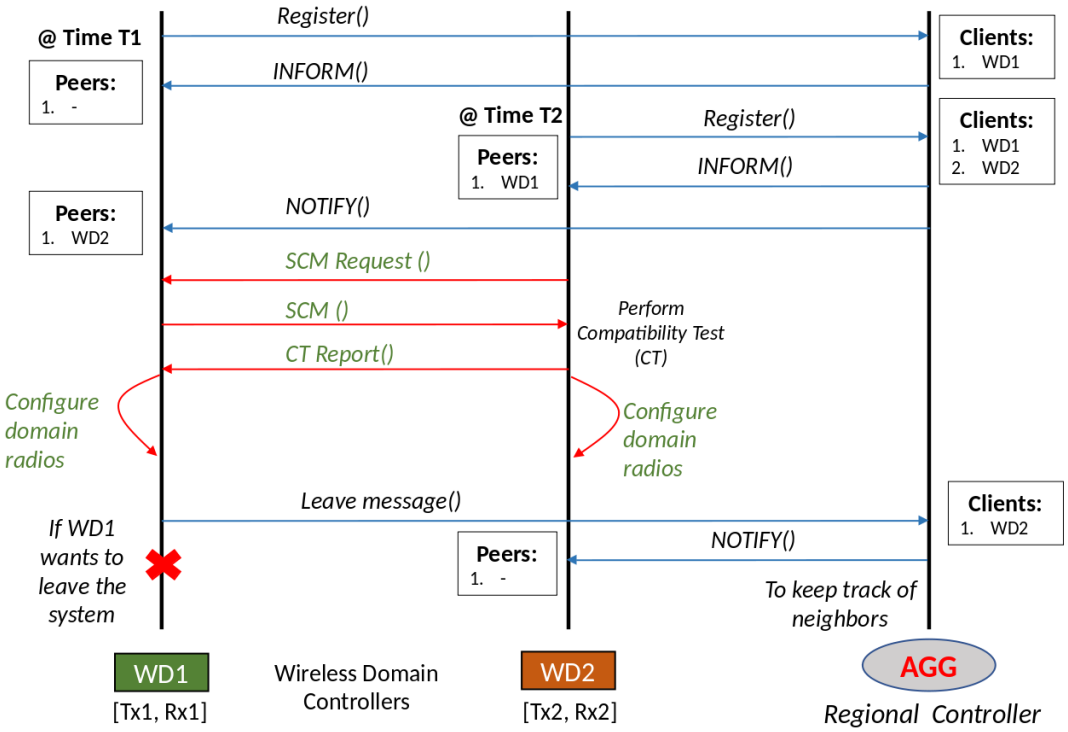

Figure (4.6) Network interaction language protocol details, showing the message types exchanged when a wireless network joins, interacts with other networks, and leaves the domain.

compatibility computations for any pair of nodes, and use the obtained information to determine the optimal bands for their transmissions.

The details of the protocol used by the interaction language are shown in Figure 4.6.

4.4.1 DSA experimentation in COSMOS

This subsection describes an experiment designed to demonstrate the capabilities of the developed framework. In the experiment, there are three wireless networks that have transmitters and receivers near each other's local area. The goal is for the networks to dynamically configure their wireless transmission characteristics, in this case, their central frequency of operation so that there is no harmful interference between the devices of different networks. Each network contains only a single transmitter and a single receiver. Transmitters and receivers are implemented using USRP X310 and B210 devices, while on the controller plane WD controllers operate on separate nodes, with CIL message exchanges conducted via Ethernet links between the WD controller

Parameter	Value
Central Freq	2.0 GHz
No. of Channels	3
Bandwidth	1 MHz
Modulation	BPSK
Bitrate	0.5M
Gnuradio	v3.7
USRPs	x310s and b210s

Table (4.1) Experimentation Settings

nodes. Additionally, all transmitters will use the same spectrum mask and will attempt transmission using a 1 MHz wide channel. The decision-making logic (whether to transmit and which frequency and power to use) of the network is implemented as part of the WD node, and GNU Radio scripts were utilized to implement the physical layer functionality, including packetization and USRP over-the-air transmission.

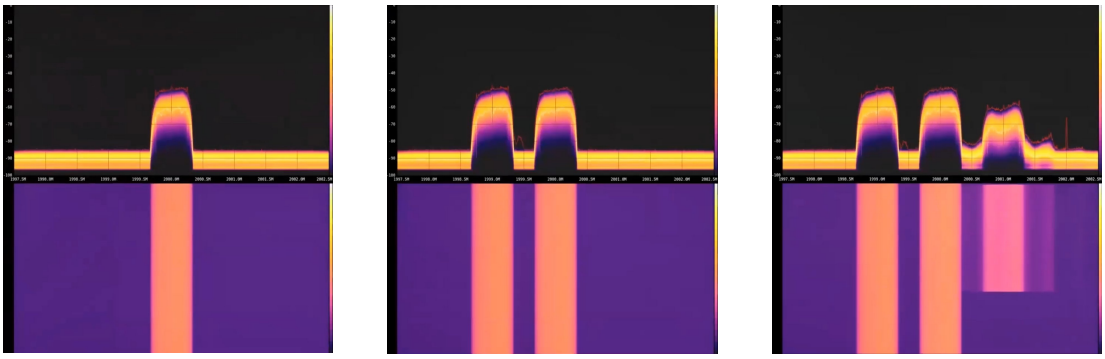

(a) One TX (b) Two TX (c) Three TX

Figure (4.7) Fosphor visualization showing the spectrum after the first, second and third network join the domain. To optimize spectrum usage, the second network has selected a band near the one used by the first network.

The experiment showcases a simple scenario. At the beginning, there are no networks present in the spectrum. At time t_1, Network 1 joins, and since the entire spectrum is available it is free to select any channel. A 1 MHz channel centered around 2 GHz is selected. Figure 4.7a shows a GNU Radio Fosphor visualization of the spectrum with only a single transmitter. After that, at time t_2 Network 2 joins the domain. Its default frequency setting is 2 GHz, the same central frequency used by Network 1. The interaction language server (Regional Topology Manager) determines that the networks are in the same wireless collision domain, and notifies Network 1 that a peer

is now present. Networks 1 and 2 establish a link for interaction, and immediately exchange a pair of SCM messages. After performing the compatibility checks, the results of which are shown on Figure 4.8a, Network 2 determines that it can not use the intended frequency, and finds an alternative optimal center frequency – 1999 MHz. Similarly, at time t_3, Network 3 joins, with the assistance of the Regional Topology Manager establishes connections with Networks 1 and 2, and after exchanging a set of SCM messages and performing its spectrum compatibility calculations determines that the optimal center frequency to use is 2001 MHz. Figure 4.7c shows all three transmitters at three orthogonal central frequencies. Figure 4.8c shows a summary of the SCM information that the controller node at Wireless Domain 3 (WD3) had to process to carry out the compatibility tests necessary to locate a central frequency where its Transmitter (Tx3) could operate without causing interference to receivers already present in the environment and where its receiver (Rx3) could also operate without being interfered by existing transmitters.

Figure 4.9 shows the scheme of the experiment – the networks with their WDs, node IDs of the USRPs utilized, and the established interaction links. Figure 4.10 shows the physical topology of the USRP nodes used within the COSMOS Sandbox, as well as the location of the Sensor node (used for GNU Radio Fosphor visualization). It also provides the location of the columns in the room, which are built out of concrete and metal and may affect the wireless signal propagation within the room.

4.5 Summary

This chapter describes the design and capabilities of the framework, intended to run on COSMOS Testbed. The framework utilizes the CIRN Interaction Language (CIL) developed by DARPA for the Spectrum Collaboration Challenge to enable interaction between independent wireless networks, adapted to use Spectrum Consumption Model (SCM) messages in order to provide peer networks with detailed information about the spectrum masks, transmission frequencies and other parameters, and provides a platform to quickly and efficiently perform compatibility calculations to determine usable

spectrum bands and avoid collision with other networks. Finally, to showcase the capabilities of the developed framework, the paper shows an experiment with 3 networks, joining the wireless communication environment at different times and exchanging SCM messages to autonomously determine optimal frequencies for their communication channels.

Future work will focus on developing and implementing advanced algorithms that are able to dynamically optimize the configuration of the networks. In particular, the focus will be on developing traditional and Machine Learning-based algorithms with different computational complexities and different performances in terms of spectrum utilization, power consumption, data throughput, reconfiguration delay, and transient congestion due to reconfigurations. Notice that, in practice, network reconfiguration is not instantaneous and may lead to transient congestion. Depending on the duration and magnitude of the congestion, data packets may be severely delayed or even lost, which may degrade the performance of delay-sensitive applications. Both the transient and steady-state effects of algorithms that reconfigure the Network will be taken into account. Two examples of algorithms are: (i) a centralized global reconfiguration algorithm in which a central server/node receives information from every associated Network and solves an optimization problem to compute a near-optimal resource allocation; and (ii) a distributed local reconfiguration algorithm in which a local disturbance triggers neighboring Networks to start negotiating aiming to agree on a new improved resource allocation. In the case of decentralized algorithms, it is important to evaluate the time that the algorithm takes to converge. In general, the framework described in this chapter can be leveraged to evaluate and compare these and other algorithms with the state-of-the-art in the literature.

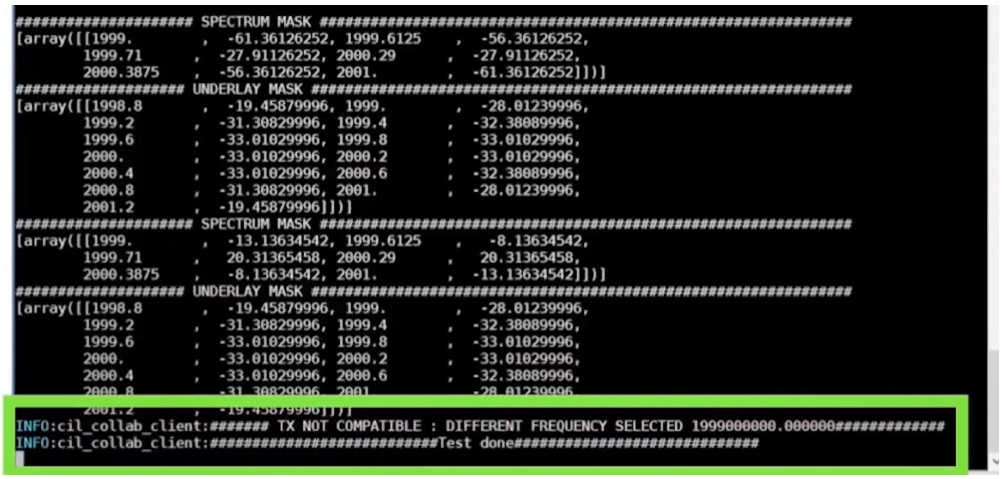

(a) Compatibility test result, reporting that the newly added transmitter would not be compatible with existing receivers with its intended parameters.

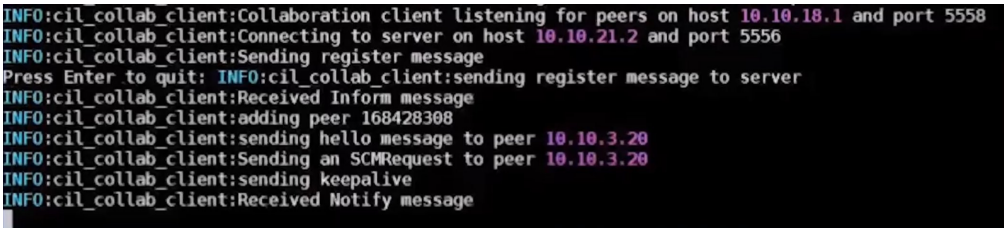

(b) Exchange of 'Inform' and 'Register' messages between domain controllers.

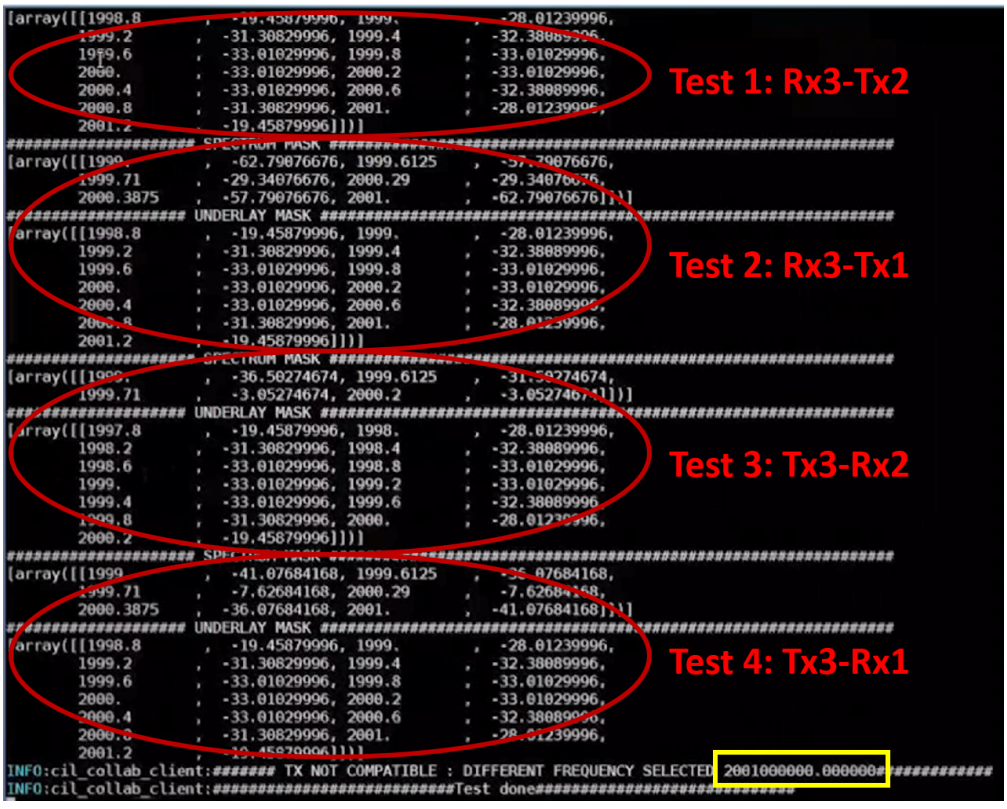

(c) List of compatibility tests between WD3's devices (Tx3 and Rx3) and existing Tx/Rx devices

Figure (4.8) Screenshots of specific message types supported by the system.

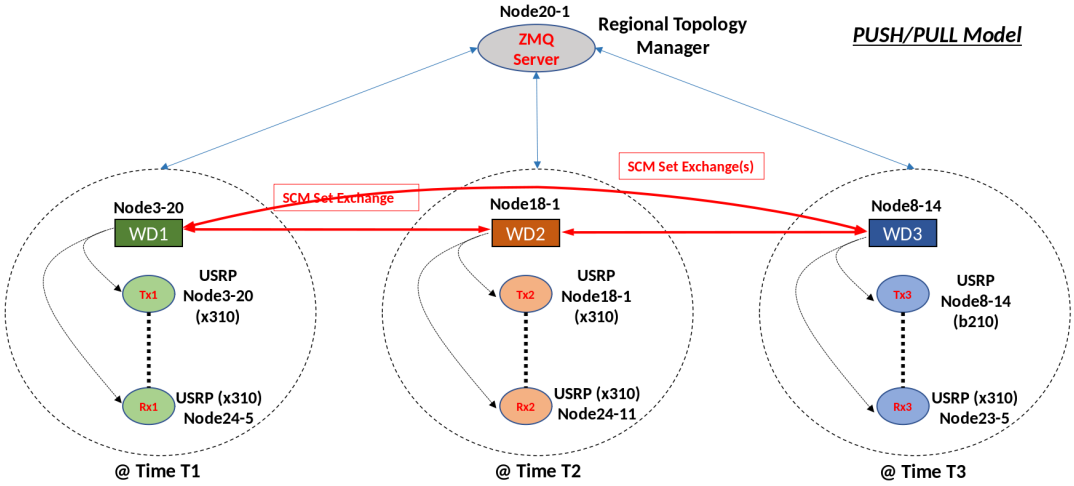

Figure (4.9) An experiment scheme, showing the timeline of networks joining and leaving the domain, as well as the directions of message exchange.

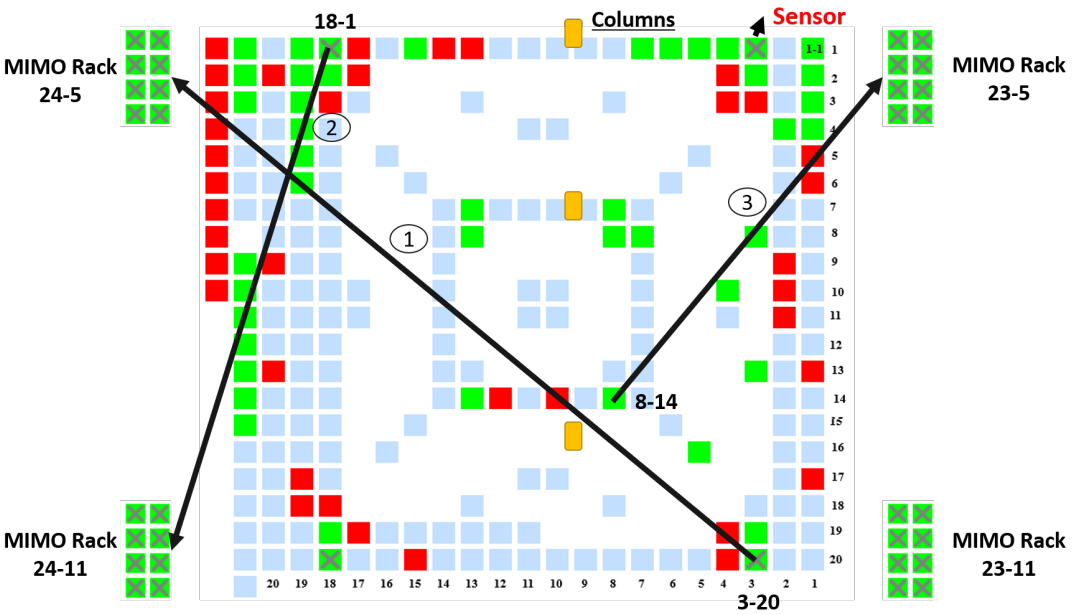

Figure (4.10) Topology of transmitters and receivers in the COSMOS Sandbox (ORBIT Grid). Yellow boxes represent columns made of concrete and metal in the room.

Chapter 5

Proactive Spectrum Usage – Predicting Short-Term Spectrum Availability

Previous chapters have shown various methods to address the issue of spectrum shortage and increase the spectrum usage efficiency. To recap, these methods include a policy driven approach where a spectrum marketplace is designed which enables short-term and small area auctions for spectrum access, which could benefit both the spectrum license holder over a larger area with no plans to use all portions of the area, and the users with localized spectrum access needs. Afterwards, the thesis describes the design of a highly efficient physical layer for Software Defined Radio, implemented for use during the DARPA Spectrum Collaboration Challenge for team SCATTER. The same chapter also showcases CIRN Interaction Language (CIL), a language design for inter-network communication where networks would collaboratively determine and coordinate their spectrum usage to improve efficiency. The following chapter shows an adaptation of CIL for usage of Spectrum Consumption Model (SCM) messages, which are standardized to contain more detailed and precise information for each transmitter and specify methods to compute spectrum usage compatibility between different transmitter/receiver pairs.

All of the previously shown spectrum access methods are reactive. Reactive approaches rely on continuous monitoring, whether by directly scanning the spectrum for activity or collecting current transmission data from other sensors, or from transmitters themselves. However, reaction imposes a delay. While many widely methods focused on reducing the reaction delay, it would be beneficial to develop a predictive method which would help avoid the delay entirely. Newer spectrum usage efficiency improvement methods which have gained more focus and motivated significant effort

recently include Machine Learning and AI driven approaches to predict spectrum availability in the short-term future, which could range from a few milliseconds up to several seconds. This method was also an important aspect of the DARPA Spectrum Collaboration Challenge, where teams built their neural networks and similar architectures, and collected training data during experiments in order to predict spectrum availability in the SC2 matches. These approaches included transmitter recognition, which could be based either on utilized modulation schemes and protocols, or on spectral behavior, which would then affect the prediction of the corresponding transmitter's future emissions. While several different implementations were used throughout the competition, it is hard to obtain details and specific measurements due to the closed-source nature of the implemented radio communication systems.

This chapter describes an effort to implement some of the available prediction methods to forecast spectrum availability, and compares their effectiveness in specific cases.

5.1 Testing Framework

The framework was designed similarly to the ones described in Sections 3.5 and 4. Certain modifications and improvements were made to achieve a more detailed and accurate evaluation, with a larger number of variables related to the system design characteristics now taken into account.

Once again, ORBIT Wireless Testbed was used to implement the framework. Using ORBIT with communications implemented as real wireless signal transmissions and receptions, via SDRs, in itself distinguishes this thesis from a significant amount of previous work, performed largely using theoretical calculations, synthetically generated data or in emulated environments.

In terms of the implementation of the physical layer, certain considerations were taken into account. Using complex protocols such as WiFi or LTE could provide an advantage with a better representation of real world systems. However, it would also significantly increase implementation complexity in several ways. Using off-the-shelf hardware does not provide the flexibility that SDRs do, i.e., commercial LTE and WiFi

devices are not designed to operate in the same bands, which practically eliminates the possibility of using any devices other than SDRs, where the user is allowed to specify the exact operating frequency. However, after stress testing several open-source software implementations of LTE and WiFI for SDR, in addition to the lack of complete and stable WiFi implementations, the USRP based implementations were shown to have significant overheads during their establishment of a connection with an external radio device (USRP X310) which made scheduling data transmission with dynamic amounts of data in time domain infeasible.

Ultimately, one of the main focuses of the presented work was to quantify the benefits that wireless communications systems would derive from using a spectrum access coordination protocol. Additionally, it was important to collect real-world measurements that would subsequently be used for short-term spectrum usage predictions. These benefits, as well as the collected data for spectrum usage predictions, would largely depend on the presence of transmitted signals and the communications system's adaptability to the environment, i.e., spectrum availability, rather than the exact physical layer protocol. Therefore, instead of complex systems such as WiFi or LTE, simpler single modulation based systems were used for testing.

Another important consideration in the testing framework was the amount of data traffic that particular transmitters will create. In other words, predicting spectrum availability would be highly dependent on having realistic network traffic, in regards to the time periods in which data bursts would occur, and the size (duration) of the bursts. In [57], the authors describe collecting data burst throughput traces from various users over several weeks, and generating random distributions of data bursts sizes and 'off' periods duration (amount of time spent waiting for the next data burst) that resembles the traffic generated in a real-world network. The testing framework for network interaction language evaluation and availability prediction uses the distributions provided by the paper to generate its own traffic. The random distribution means and variances are scaled to accommodate the testing scenarios where data can only be collected in shorter time periods, given the storage limitations.

One of the important aspects of the testing framework are the internal decision-making algorithms that networks can use to determine the spectrum usage parameters. This is analogous to precisely defining and accurately describing a spectrum voxel in section 3.4. In particular, several aspects of spectrum occupancy were considered. One of these aspects is time, which could be taken continuously or discreetly. It was determined that discrete time could significantly simplify the occupancy prediction process, and with a higher time resolution the accuracy of spectrum availability would not be significantly affected. Another important aspect is spectrum occupancy itself. If a band is occupied, i.e. a transmission is present, the energy of the transmission can matter to a varying degree. In terms of spectrum occupancy, it only matters whether choosing to transmit would interfere with existing transmissions, so the energy of the existing transmissions can be taken as simple logical values. Additionally, should the underlying protocol and/or modulation scheme of the transmission also be taken into consideration? This is a significantly more complex consideration, as any additional information could help improve the predictions. In terms of predicted values, if each spectrum band is shown as a time-series of values representing the status of that band at a given moment, each modulation could simply be represented with a different number. No occupancy could then be represented with the number "0". The final consideration was the frequency range observed, where, similarly to time, it was determined that frequency should also be discretely divided.

5.2 Prediction Methods

Since spectrum occupancy can be represented by a time series, as previously explained, the spectrum occupancy prediction problem can for this purpose be considered equivalent to time series prediction. Time series prediction is a well-researched topic, with a wide range of work and publications presenting effective and accurate solutions. Recently, in [58], the author applied and compared several time series prediction methods to predict internet traffic. It was shown that Long Short-Term Memory (LSTM) based prediction, using the method described in [], produced comparable or better results than statistical methods, such as the one using ARIMA models. This inspired the effort

presented in this chapter, where a similar method was used to attempt short-term spectrum occupancy prediction. However, compared to the effort presented in this thesis, the mentioned work focused on long-term prediction, where trends such as traffic increases during certain parts of the day can be noticed. In short-term spectrum occupancy prediction, no such trends are present, and randomness of the traffic significantly affects the ability to accurately predict the time-series. As a proof of concept, similar to the results shown in the mentioned publications, figure 5.1 shows LSTM based predictions of common periodic signals, which were all highly accurate and consistent.

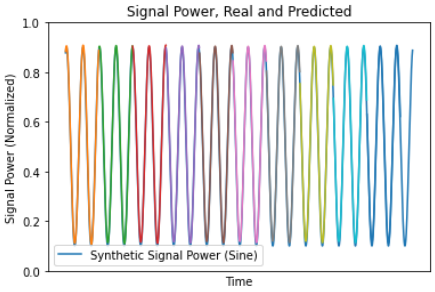

(a) Short-term predictions for a sine signal, very high prediction accuracy

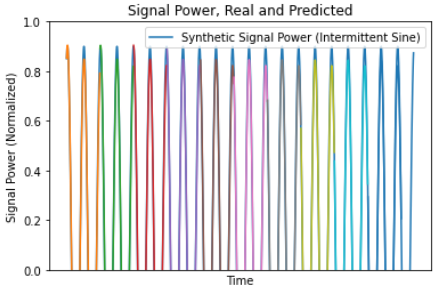

(b) Short-term predictions for an intermittent sine signal, high prediction accuracy

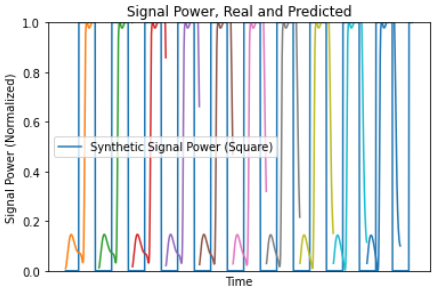

(c) Short-term predictions for a square signal, period longer than the prediction duration which affects prediction accuracy

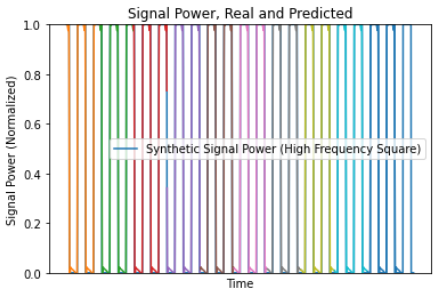

(d) Short-term predictions for a square signal with higher frequency, high prediction accuracy

Figure (5.1) A showcase of LSTM predictions of values for several common periodic signal types.

Using the same methods with random data, obtained by sensing the spectrum with transmissions initiated following the data generation using the results shown in [57], the short-term prediction outcome is significantly worse. An example is shown in figure 5.2, where the predicted short-term usage has no correlation with actual data, and the resulting root mean-square error (RMSE) of 0.929, compared to the RMSE's of periodic

signals of 0.017, 0.061, 0.516 and 0.068.

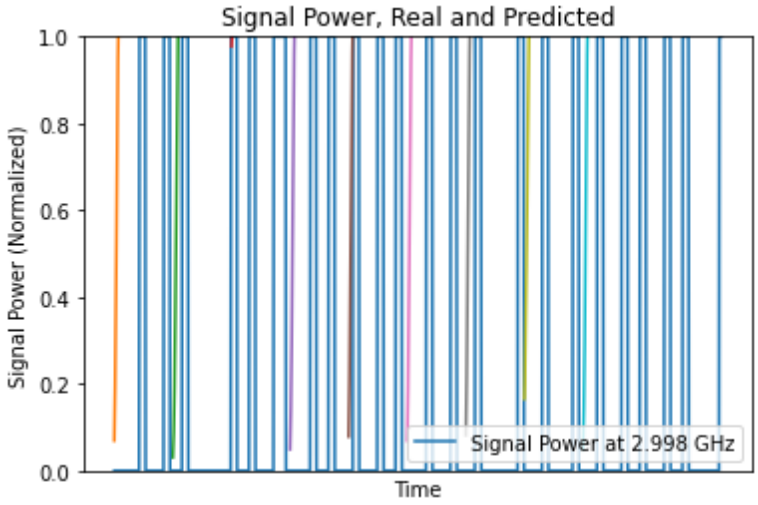

Figure (5.2) LSTM-based spectrum occupancy prediction with real spectrum data collected using USRP X310 devices on COSMOS Testbed, showing low prediction accuracy due to the randomness of the data.

Indeed, LSTM-based prediction methods rely on existing patterns and periodicity in the values they are predicting. It is also important to note that the duration of the period of signal patterns, relative to the predicted duration, also affects the prediction accuracy, as seen when predicting square signals of different frequencies, shown in Figures 5.1c and 5.1d.

5.3 Performance

To better understand the performance of the LSTM-based predictive spectrum access method in different cases, two example scenarios are shown in Figure 5.3. The first scenario with random bursts of traffic is shown in the left column. In the real world, many protocols include short periodic bursts, such as beacon frames, which is the second scenario shown in the right column. As expected, prediction accuracy is very high for a periodic signal, and low for the random one. The first row of the figure shows the real signal's spectrum occupancy in blue, with short predictions shown in various other colors. The second row of the figure shows the achieved cumulative throughput throughout the experiment. The simulation assumes that the transmitter which relies on prediction of spectrum availability to access the spectrum, will utilize the channel in any

time slot that is predicted to be available. High prediction accuracy increases channel utilization and maximizes the achieved throughput (right column). However, lower prediction accuracy not only reduces the achieved throughput, but also results in "missed opportunities" – time slots in which the channel was available but was not utilized due to incorrect prediction, and can cause interference which reduces the effective throughput of all transmitters in time slots where channel was wrongly predicted to be available. The third row shows the achieved cumulative throughput for a transmitter which uses CIL and accurately receives spectrum utilization information from all other transmitters.

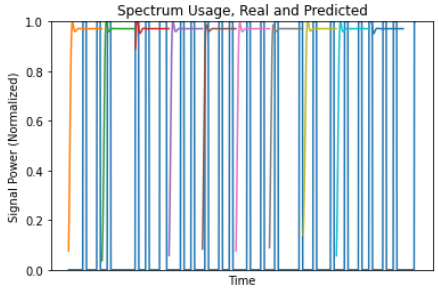

(a) Predictions for a random signal

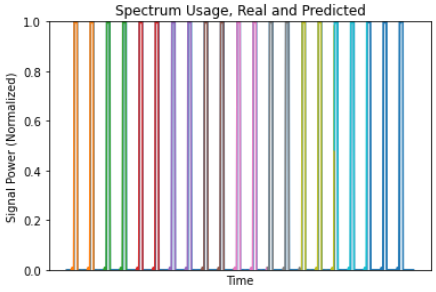

(b) Predictions for a periodic signal

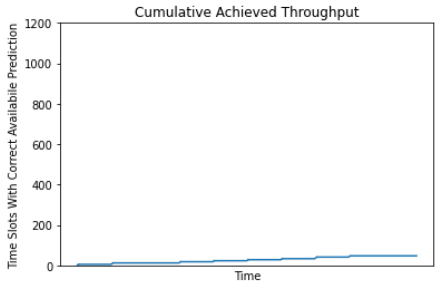

(c) Cumulative achieved throughput with predictions for a random signal

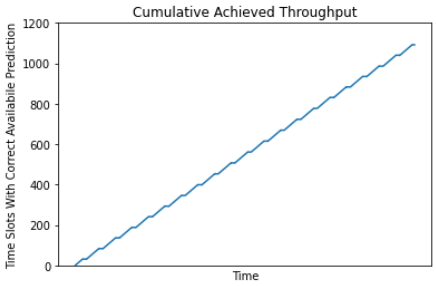

(d) Cumulative achieved throughput with predictions for a periodic signal

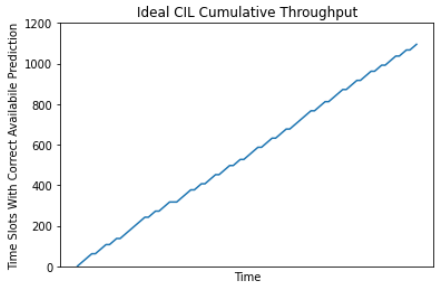

(e) Ideal throughput assuming full spectrum utilization with a random signal

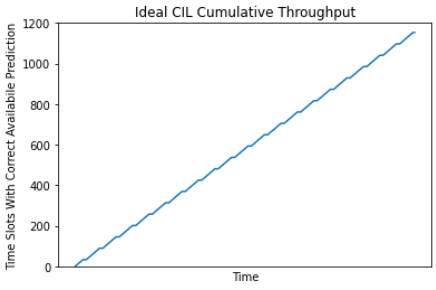

(f) Ideal throughput assuming full spectrum utilization with a periodic signal

Figure (5.3) Spectrum availability prediction performance for random and periodic signals.

In summary, one of the main benefits of spectrum availability prediction as a proactive spectrum access mechanism is independence from other spectrum users. In other words, no compatibility with particular protocols or specific features are required from anyone else. The shortcoming of such an approach is the fact that in many cases, traffic can be modeled with random bursts, and in such scenarios predictive methods are ineffective. On the other hand, when supported by all spectrum users, a coordination protocol will almost always result in better spectrum utilization, but it requires standardization and design of the protocol, which includes the existence of a reliable dedicated channel for spectrum access coordination, as well as the messages exchanged between all spectrum users.

With many legacy protocols still being used, it is not feasible to always expect all participants to support a new protocol. However, the DARPA Spectrum Collaboration Challenge showed that a hybrid of the two described methods could have most of the advantages of both approaches, with very few of the shortcomings. Spectrum users could rely on the coordination protocol, while spectrum is still being actively monitored, and in cases where incumbents or simply legacy protocol users attempt transmissions, other networks could collaboratively conclude that the spectrum is being utilized by a non-supported party, and either use predictive methods to determine the next available time slot of the channel, or simply use the spectrum reactively after the incumbent user, or the user that does not support the collaboration protocol, finishes its transmissions.

Chapter 6

Conclusions

With the ever-growing number of wireless communication devices, the number of which is now increasing exponentially with the emergence of 5G networks and IoT, this thesis focuses on the increasingly important issue of spectrum shortage. In order to address the concern of the lack of available spectrum, the thesis takes several different approaches, assesses their benefits, compares their value and provides insight into the ways they can be combined to take advantage of their benefits while circumventing their shortcomings. The spectrum usage efficiency improvement methods include a policy-driven approach (a 3-tier spectrum renting marketplace), design of an adaptive physical layer for a wireless communication system, collaboration and spectrum access coordination between wireless networks in the same collision domain and, finally, spectrum occupancy prediction, which benefits by being a proactive approach which does not rely on other users to support and be compatible with the utilized spectrum access coordination protocol.

6.1 Policy-driven Spectrum Utilization Improvements

The 3-tier spectrum sharing framework is a very beneficial way of sharing the spectrum because it ensures interference protection to higher tier users and at the same time enables lower tier users to access the spectrum opportunistically. The shortcoming of such system is that spectrum licenses are obtained for relatively long periods and are assigned for large areas. This can result for sub-optimal utilization of the spectrum by the license holder until the network is deployed in the entire county. To mitigate the licensing costs, the license holder can sub-lease the unused portions of licensed areas with the help of the Spectrum Access Server. This thesis proposes several models of efficient spectrum license sub-leasing. The area not used by the primary license holder is opened

for sub-leasing, and interested parties submit bids for the license sub-lease in defined portions of the entire area. In the basic model, the bid scenario is represented as a graph, and the efficient method of license sub-lease assignment are then found as the maximum weighted independent set of the graph. The presented work proposes two improvements to this base model, by allowing bid adjustments to accommodate for more users. If bidders allow limited reduction of their defined regions to avoid interference with other operators, or coexistence with other operators in smaller portions of their regions, they can improve their chances of getting a sub-leasing license assigned. The work also considers the effect of geographical distribution of the bidders on the achieved efficiency of channel allocation. In particular, two different realistic geographical distributions are tested - a uniform distribution, where the density of bidders is similar across the entire county, and clustered distribution, where most bidders are clustered around given points. Clustered distribution is considered with numbers of clusters ranging from 3 to 8. The implementation of the basic method and the improvements is described and tested, and the achieved gains are shown to be between 15 and 35% in financial compensation to the Primary Access License holder, depending on the geographical distribution of the bidders and the shapes and sizes of their defined regions. In summary, the thesis proposes a solution to the attractive problem of spectrum access license sub-leasing, and showcases potential gains achieved by using that solution.

6.2 Wireless Network Collaboration and Spectrum Access Coordination

Homogeneous systems, such as LTE, can rely on a slot scheduler as a fully centralized spectrum access coordination mechanism. However, with the increasing number of novel wireless communication methods and protocols, with many of the older legacy protocols still actively used, heterogeneous systems are becoming more and more ubiquitous. This makes centralized spectrum access mechanisms becomes significantly harder to implement. To address this issue and attempt a design of an effective distributed spectrum access coordination protocol, DARPA has established the Spectrum Collaboration Challenge (SC2) as a competition between multiple teams that build intelligent

and collaborative wireless communication networks - CIRNs (Collaborative Intelligent Radio Networks), with the goal to further motivate researchers to work on improved and novel ways to increase spectrum usage efficiency. To enable the collaboration between teams' networks, DARPA has introduced the CIRN Interaction Language – a language used for communication between different networks in a heterogeneous system to announce their spectrum usage intentions. This thesis describes the design of CIL and ideas behind it, affirms the fact that CIL's benefits, while clearly improving the achieved throughput of all networks, were never measured and quantified. Further, the thesis proposes a simulation framework that would help evaluate the CIL's gains, describes its implementation and showcases the results of the performed evaluation experiments. The naive back-off algorithm used in the experiments results in the increase of the achieved combined throughput of all networks between 4 and 12%, depending on the number of transmitters and available channels, i.e. the spectrum demand and availability. The variance in achieved throughput for each transmitter was also shown to depend on both the number of transmitters in the spectrum and the existence of CIL. In particular, a higher number of transmitters, as well as the presence of CIL spectrum usage announcement messages, both result in increase in throughput variance, meaning that CIL not only potentially increases the ensemble throughput, but also reduces the uncertainty of expected number of successful transmissions.

6.3 Adaptive Physical Layer Design

Additionally, the thesis describes the design, features, and benefits of the SCATTER PHY, the physical layer utilized for DARPA SC2 by team SCATTER. The main features of the design include efficient interaction with higher layers of the system, ability to transmit/receive two different types of packets which contain synchronization signals, control data, and user data. The PHY layer communication operates with various modulation and coding schemes, several different channel bandwidths, and automatically improves the CFO correction by utilizing the OFDM CP in addition to the synchronization subframe detected in the signal. SCATTER PHY minimizes OOBE by utilizing custom-built FPGA-based FIR filters on the transmitter side, configured for

each channel bandwidth at run time. Additionally, two independent physical interfaces are available, which effectively enable usage of full-duplex and theoretically double the possible throughput achieved with one interface. The described benefits were verified experimentally. The verification shows that using full-duplex makes no impact on the achieved throughput in each interface. Spectrometer measurement results showcase the capabilities of the FPGA-based FIR filter, resulting in negligible OOBE, which in turn increases channel utilization and facilitates communication in neighboring channels by other radios. Using CP-based CFO estimation and correction is demonstrated to outperform the sync-signal based correction in all cases, at the expense of higher processing time. Finally, the packet reception rate for all types of transmitted signals is presented, showing that the SCATTER PHY can successfully decode user data at SNR values as low as 0 dB.

6.4 Spectrum Occupancy Prediction

Traditionally, cognitive radio actively scans the spectrum and attempts transmission only in channels which are no longer occupied. This is a reactive approach, which naturally imposes a delay between the end of the previous transmission, and the start of the new one. Additionally, reactive approaches struggle in scenarios when traffic is characterized by random short bursts of large amounts of packets [59], where the bursts themselves experience interference in the beginning until cognitive radio reacts and stops transmission, and after every burst stops its own transmission it will cause a new delay until cognitive radio reacts. Predictive approaches can be used to avoid these delays. The thesis assesses the accuracy of predicting network traffic with machine learning and Long Short-Term Memory (LSTM) neural networks, presented in recent scientific efforts. The prediction accuracy is highly dependent on the signal characteristics, since prediction relies on signals patterns and periodicity. It is shown that while periodic signals can easily be predicted with very high accuracy, which can be used in long-term traffic prediction where certain traffic patterns can be observed – such as higher traffic during certain times in the day, as well as short-term predictions with highly regular signals such as brief periodic beacons, the scenarios where traffic is random

significantly affect the effectiveness of this method. Collaboration can be utilized as a mechanism capable of resolving the issue of traffic randomness. However, collaboration requires other networks to be compatible and participate. As a counterexample, simple IoT transmitters may not have this capability, but will have regular easily predictive periodic transmissions. With a wide variety of wireless network and transmitter types and their characteristics, predictive and collaborative approach can be used in tandem to produce best spectrum utilization.

6.5 Future Work

In the future, each of these approaches can be addressed independently with further improvements. Indeed, there is an abundance of active effort focusing on all of these approaches, with a strong emphasis of Machine Learning and Artificial Intelligence based architectures, which have recently gained increasing attention of the research community. ML and AI based approaches have also shown to be effective in detection and mitigation of malicious network usage. However, this thesis shows that taking into account complementary methods can be beneficial in developing new spectrum access mechanisms and improving existing ones. The work will therefore focus not only on individual methods and designs but on how to achieve improvements by expanding them with different architectures, both complementary and orthogonal, with the goal of building more complex but complete spectrum access systems, capable of adapting to adverse conditions, proactively utilizing the available spectrum and efficiently coordinating access with other networks. Additionally, a more detailed analysis of various degrees of spectrum access coordination capabilities will help in determining the effectiveness of this method and aid in shaping the design of similar approaches in the future.

Bibliography

[1] Li Xu and Song Junde. "A kind of new frequency planning method". In: *Fifth Asia-Pacific Conference on ... and Fourth Optoelectronics and Communications Conference on Communications,* vol. 1. 1999, 491–494 vol.1.

[2] G. Staple and K. Werbach. "The end of spectrum scarcity [spectrum allocation and utilization]". In: *IEEE Spectrum* 41.3 (2004), pp. 48–52.

[3] J.M. Peha. "Approaches to spectrum sharing". In: *IEEE Communications Magazine* 43.2 (2005), pp. 10–12.

[4] Artur Pereira, Helder Fontes, and Atilio Gameiro. "Simulation Platform for Opportunistic Radio Systems". In: *2006 IEEE 17th International Symposium on Personal, Indoor and Mobile Radio Communications*. 2006, pp. 1–5.

[5] 30 F.C.C. Rec. 3959. *Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550-3650 MHz Band*. visited on 11/30/2021. 2015. URL: `https://www.federalregister.gov/documents/2017/06/12/2017-12117/amendment-of-the-commissions-rules-with-regard-to-commercial-operations-in-the-3550-3650-mhz-band`.

[6] Preston Marshall. *Three-Tier Shared Spectrum, Shared Infrastructure, and a Path to 5G*. Cambridge University Press, 2017.

[7] Y. Yi et al. "Spectrum Leasing to Multiple Cooperating Secondary Cellular Networks". In: *2011 IEEE International Conference on Communications (ICC)*. 2011, pp. 1–5.

[8] ZeroMQ. ØMQ - The Guide. 2021. URL: `https://zguide.zeromq.org/` (visited on 11/30/2021).

[9] srsLTE. *Open source SDR LTE software suite from Software Radio Systems (SRS)*. 2021. URL: `https://www.srslte.com/` (visited on 11/30/2021).

[10] Dragoslav Stojadinovic and Milind Buddhikot. "Design of a Secondary Market for Fractional Spectrum Sub-leasing in Three-Tier Spectrum Sharing". In: *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2019, pp. 1–8.

[11] C. W. Kim, J. Ryoo, and M. M. Buddhikot. "Design and implementation of an end-to-end architecture for 3.5 GHz shared spectrum". In: *2015 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2015, pp. 23–34.

[12] C. Liu and R. A. Berry. "Tiered spectrum sharing and price competition". In: *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. 2016, pp. 586–591.

[13] G. Saha, A. A. Abouzeid, and M. Matinmikko. "Competitive online algorithm for leasing wireless channels in 3-Tier Sharing Framework". In: *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. 2016, pp. 792–799.

[14] G. Saha, A. A. Abouzeid, and M. Matinmikko-Blue. "Online Algorithm for Leasing Wireless Channels in a Three-Tier Spectrum Sharing Framework". In: *IEEE/ACM Transactions on Networking* 26.6 (2018), pp. 2623–2636.

[15] X. Wang et al. "Adaptive video streaming over whitespace: SVC for 3-Tiered spectrum sharing". In: *2015 IEEE Conference on Computer Communications (INFOCOM)*. 2015, pp. 28–36.

[16] S. Basnet et al. "Opportunistic Access to PAL Channel for Multi-RAT GAA Transmission in Spectrum Access System". In: *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*. 2017, pp. 1–5.

[17] Marko Palola et al. "The first end-to-end live trial of CBRS with carrier aggregation using 3.5 GHz LTE equipment". In: *2017 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2017, pp. 1–2.

[18] H. Wang et al. “Design of Contour Based Protection Zones for Sublicensing in Spectrum Access Systems”. In: *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*. 2017, pp. 1–5.

[19] Ming-Shing Yu, Lin Yu Tseng, and Shoe-Jane Chang. “Sequential and parallel algorithms for the maximum-weight independent set problem on permutation graphs”. In: *Information Processing Letters* 46.1 (1993), pp. 7–11.

[20] Stefano Basagni. “Finding a Maximal Weighted Independent Set in Wireless Networks”. In: *Telecommunication Systems* 18.1 (Sept. 2001), pp. 155–168.

[21] Gabriel Valiente. “A New Simple Algorithm for the Maximum-Weight Independent Set Problem on Circle Graphs”. In: *Algorithms and Computation*. Ed. by Toshihide Ibaraki, Naoki Katoh, and Hirotaka Ono. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 129–137.

[22] Ekkehard Köhler and Lalla Mouatadid. “A linear time algorithm to compute a maximum weighted independent set on cocomparability graphs”. In: *Information Processing Letters* 116.6 (2016), pp. 391–395.

[23] Felipe A. P. de Figueiredo et al. “SCATTER PHY: An Open Source Physical Layer for the DARPA Spectrum Collaboration Challenge”. In: *Electronics* 8.11 (2019). URL: https://www.mdpi.com/2079-9292/8/11/1343.

[24] Spilios Giannoulis et al. “Dynamic and Collaborative Spectrum Sharing: The SCATTER Approach”. In: *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2019, pp. 1–6.

[25] Ettus Research. *USRP Hardware Driver*. 2021. URL: http://files.ettus.com/manual/ (visited on 11/30/2021).

[26] Ettus Research. *USRP X Series*. 2021. URL: https://www.ettus.com/product/category/USRP-X-Series (visited on 11/30/2021).

[27] National Instruments. *Overview of the NI USRP RIO Software-Defined Radio*. 2021. URL: http://www.ni.com/white-paper/52119/en/ (visited on 11/30/2021).

[28] Ettus Research. *Products*. 2021. URL: `https://www.ettus.com/product` (visited on 11/30/2021).

[29] P. Hintjens. *ZeroMQ: Messaging for Many Applications*. Oreilly and Associate Series. O'Reilly Media, Incorporated, 2013.

[30] Google. *protobuf - Protocol Buffers - Google's data interchange format*. 2021. URL: `http://code.google.com/p/protobuf/` (visited on 11/30/2021).

[31] Ismael Gomez-Miguelez et al. "srsLTE: An Open-Source Platform for LTE Evolution and Experimentation". In: *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization*. WiNTECH '16. New York City, New York: Association for Computing Machinery, 2016, pp. 25–32.

[32] S. Kanchi et al. "Overview of LTE-A technology". In: *2013 IEEE Global High Tech Congress on Electronics*. 2013, pp. 195–200.

[33] Felipe Pereira de Figueiredo et al. "Multi-Stage Based Cross-Correlation Peak Detection for LTE Random Access Preambles". In: *Revista Telecomunicacoes* 15 (Oct. 2013).

[34] D. V. Sarwate and M. B. Pursley. "Crosscorrelation properties of pseudorandom and related sequences". In: *Proceedings of the IEEE* 68.5 (1980), pp. 593–619.

[35] F. Wang and Y. Zhu. "An efficient CFO estimation algorithm for the downlink of 3GPP-LTE". In: *2011 International Conference on Wireless Communications and Signal Processing (WCSP)*. 2011, pp. 1–6.

[36] Q. Bodinier, F. Bader, and J. Palicot. "On Spectral Coexistence of CP-OFDM and FB-MC Waveforms in 5G Networks". In: *IEEE Access* 5 (2017), pp. 13883–13900.

[37] F. -L. Luo and C. Zhang. "Major 5G Waveform Candidates: Overview and Comparison". In: *Signal Processing for 5G: Algorithms and Implementations*. 2016, pp. 169–188.

[38] F. A. P. De Figueiredo et al. "A Spectrum Sharing Framework for Intelligent Next Generation Wireless Networks". In: *IEEE Access* 6 (2018), pp. 60704–60735.

[39] Ettus Research. *CBX 1200-6000 MHz Rx/Tx (120 MHz, X Series only*. 2021. URL: https://www.ettus.com/product/details/CBX120 (visited on 11/30/2021).

[40] Wireless Innovation Forum. *Dynamic Spectrum Sharing Annual Report – 2014*. visited on 11/30/2021. 2014. URL: https://www.wirelessinnovation.org/assets/work_products/Reports/winnf-14-p-0001-v1.0%20dynamic%20spectrum%20sharing%20annual%20report%202014.pdf.

[41] Jun Zhao, Haitao Zheng, and Guang-Hua Yang. "Distributed coordination in dynamic spectrum allocation networks". In: *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005.* 2005, pp. 259–268.

[42] Shweta S. Sagari. "Coexistence of LTE and WiFi Heterogeneous Networks via Inter Network Coordination". In: *Proceedings of the 2014 Workshop on PhD Forum*. PhD forum '14. Bretton Woods, New Hampshire, USA: Association for Computing Machinery, 2014, pp. 1–2.

[43] P. Karimi et al. "SMAP: A Scalable and Distributed Architecture for Dynamic Spectrum Management". In: *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2018, pp. 1–10.

[44] D. Raychaudhuri and Xiangpeng Jing. "A spectrum etiquette protocol for efficient coordination of radio devices in unlicensed bands". In: *14th IEEE Proceedings on Personal, Indoor and Mobile Radio Communications, 2003. PIMRC 2003.* Vol. 1. 2003, 172–176 Vol.1.

[45] D. Stojadinovic et al. "SC2 CIL: Evaluating the Spectrum Voxel Announcement Benefits". In: *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. 2019, pp. 1–6.

[46] M. M. Buddhikot et al. "DIMSUMnet: new directions in wireless networking using coordinated dynamic spectrum". In: *Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks*. 2005, pp. 78–85.

[47] V. Brik et al. “DSAP: a protocol for coordinated spectrum access”. In: *First IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005.* 2005, pp. 611–614.

[48] Hang Qin and Yanrong Cui. “Spectrum Coordination Protocol for reconfiguration management in cognitive radio network”. In: *2009 IEEE Youth Conference on Information, Computing and Telecommunication*. 2009, pp. 407–410.

[49] Xiangpeng Jing and Dipankar Raychaudhuri. “Spectrum Co-Existence of IEEE 802.11b and 802.16a Networks Using Reactive and Proactive Etiquette Policies”. In: *Mob. Netw. Appl.* 11.4 (Aug. 2006), pp. 539–554.

[50] Milorad Tosic et al. “Semantic coordination protocol for LTE and Wi-Fi coexistence”. In: *2016 European Conference on Networks and Communications (EuCNC)*. 2016, pp. 69–73.

[51] Bhausaheb Shinde and V Vijayabaskar. “LTE and Wi-Fi Coexistence Using New Semantic Co-ordination Protocol”. In: (Jan. 2020), pp. 3959–3971.

[52] Bikramjit Singh et al. “Coordination Protocol for Inter-Operator Spectrum Sharing in Co-Primary 5G Small Cell Networks”. In: *IEEE Communications Magazine* 53 (May 2015).

[53] C. E. C. Bastidas et al. “IEEE 1900.5.2: Standard Method for Modeling Spectrum Consumption: Introduction and Use Cases”. In: *IEEE Communications Standards Magazine* 2.4 (2018), pp. 49–55.

[54] Dragoslav Stojadinovic et al. “A Spectrum Consumption Model-Based Framework for DSA Experimentation on the COSMOS Testbed”. In: *Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & CHaracterization*. WiNTECH’21. New Orleans, LA, USA: Association for Computing Machinery, 2022, pp. 77–84.

[55] IEEE. “IEEE 1900.5.2-2017 - Standard for Method for Modeling Spectrum Consumption”. In: (2017).

[56] John A. Stine and Carlos E Caicedo Bastidas. “Enabling spectrum sharing via spectrum consumption models”. In: *IEEE Journal on Selected Areas in Communications* 33.4 (2015), pp. 725–735.

[57] A. Varet and N. Larrieu. “How to Generate Realistic Network Traffic?” In: *2014 IEEE 38th Annual Computer Software and Applications Conference*. 2014, pp. 299–304.

[58] Tejashri Kuber. “Learning and prediction using radio and non-radio attributes in wireless systems”. Rutgers University, 2021.

[59] Md Jalil Piran et al. “Multimedia communication over cognitive radio networks from QoS/QoE perspective: A comprehensive survey”. In: *Journal of Network and Computer Applications* 172 (2020), p. 102759.