Using Incumbent Channel Occupancy Prediction to Minimize Secondary License Grant Revocations

Divya Ramanujachari

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

Master of Science

in

Computer Engineering

Yaling Yang, Chair

Gang Wang

Lingjia Liu

November 16th, 2018

Blacksburg, Virginia

Keywords: Dynamic Spectrum Access, Spectrum Access System, Spectrum Occupancy

Prediction, Spectrum Grant Revocation

Using Incumbent Channel Occupancy Prediction to Minimize Secondary License Grant Revocations

Divya Ramanujachari

ABSTRACT

With commercial deployment of the Citizens Band Radio Service commencing in the last quarter of 2018, efforts are in progress to improve the efficiency of the Spectrum Access System (SAS) functions. An area of concern as identified in recent field trials is the time-bound evacuation of unlicensed secondary users from a frequency band by the SAS on the arrival of an incumbent user. In this thesis, we propose a way to optimize the evacuation process by reducing the number of secondary spectrum grant revocations to be performed. The proposed work leverages knowledge of incumbent user spectrum occupancy pattern obtained from historical spectrum usage data. Using an example model trained on 48 hours of an incumbent user transmission information, we demonstrate prediction of future incumbent user spectrum occupancy for the next 15 hours with 94.4% accuracy. The SAS uses this information to set the time validity of the secondary spectrum grants appropriately. In comparison to a case where spectrum grants are issued with no prior knowledge, the number of revocations declines by 87.5% with a 7.6% reduction in channel utilization. Further, the proposed technique provides a way for the SAS to plan ahead and prepare a backup channel to which secondary users can be redirected which can reduce the evacuation time significantly.

Using Incumbent Channel Occupancy Prediction to Minimize Secondary License Grant Revocations

Divya Ramanujachari

GENERAL AUDIENCE ABSTRACT

Studies on spectrum occupancy show that, in certain bands, licensed incumbent users use the spectrum only for some time or only within certain geographical limits. The dynamic spectrum access paradigm proposes to reclaim the underutilized spectrum by allowing unlicensed secondary users to access the spectrum opportunistically in the absence of the licensed users. In the United States, the Federal Communications Commission (FCC) has identified 150 MHz of spectrum space from 3550-3700 MHz to implement a dynamic spectrum sharing service called the Citizens Broadband Radio Service (CBRS). The guiding principle of this service is to maximize secondary user channel utilization while ensuring minimal incumbent user disruption. In this study, we propose that these conflicting requirements can be best balanced in the Spectrum Access System (SAS) by programming it to set the time validity of the secondary license grants by taking into consideration the incumbent spectrum occupancy pattern. In order to enable the SAS to learn incumbent spectrum occupancy in a privacy-preserving manner, we propose the use of a deep learning model, specifically the long-short term memory (LSTM). This model can be trained by federal agencies on historical incumbent spectrum occupancy information and then shared with the SAS in a secure manner to obtain prediction information about possible incumbent activity. Then, using the incumbent spectrum occupancy information from the LSTM model, the SAS could issue license grants that would expire before expected arrival time of incumbent user, thus minimizing the number of revocations on incumbent arrival. The scheme was validated using simulations that demonstrated the effectiveness of this approach in minimizing revocation complexity.

Acknowledgments

First, I would like to thank my advisor, Dr. Yaling Yang, whose support enabled me to resume my graduate studies. Her advice and guidance has been invaluable in shaping this thesis. I would also like to thank my committee members, Dr. Wang and Dr. Liu. Finally, I would like to thank my lab mates at CNSR and Yousi Lin for sharing their insights on various aspects of my research.

Dedication

This thesis is dedicated to my teachers, family, and friends.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The increase in the number of networked wireless devices and demand for data-intensive content to support new technologies has led to an increased in spectrum requirements of operators. According to a recent mobile data traffic growth report by Ericsson, traffic per smartphone will increase from 7.1 GB per month in 2017 to 48 GB by 2023 in North America [1]. The total mobile data traffic will double every year in the same time. This situation, along with the static spectrum allocation approach that is being currently followed, has sparked concerns about an approaching ‘spectrum crunch’ – a shortage of frequency spectrum for future allocation.

In this context, the idea of spectrum sharing has emerged as a viable approach for increasing the available spectrum effectively. This approach involves opening up bands that are currently underutilized by their licensed users [2]. New subscribers will access the reclaimed spectrum in an opportunistic manner. This paradigm, called dynamic spectrum access (DSA), is the focus of this study. The Federal Communications Commission (FCC) manages spectrum allocations in the United States of America while the National Telecommunications and Information Administration (NTIA) enacts regulations on spectrum usage. These agencies have identified the 3.5 GHz Band, which encompasses the frequency range from 3550 to 3700 MHz, for shared wireless broadband use. Shared spectrum access will be provided via a service termed the Citizens Broadband Radio Service (CBRS) [3]. Currently, various stakeholders have developed hardware and software solutions for supporting CBRS operation. Once the FCC certifies these solutions, commercial deployment will begin.

Next, the DSA paradigm is explained in detail and the key aspects of the CBRS are outlined. This is followed by a discussion on open research questions. Our research objective is defined and an overview of the proposed approach is provided. This is followed by a summary of our original research contributions and an outline of the remainder of this report.

1.1 Background

The DSA paradigm, as originally proposed, envisioned a scenario in which unlicensed secondary users (SUs) used cognitive radios to access licensed spectrum bands in the absence of the licensed primary users (PUs). The SUs would sense the radio environment, arrive at a decision about the presence/absence of PU, and reconfigure themselves to transmit over an available frequency. However, owing to technological challenges and administrative concerns regarding a completely distributed architecture, the CBRS proposed by the FCC differs in some key aspects.

The CBRS is a three-tiered shared access model managed by a centralized frequency coordinator [3, 4]. The three tiers of service are listed below.

- Incumbent Access: This type of access is the privilege of the subscribers who hold the original license to transmit over a band or the incumbent users (IUs) . IUs include both existing federal and non-federal licensed users. IU transmissions have the highest priority and are protected against interference from lower tier transmissions.
- Priority Access (PA): PA users hold short-term licenses for transmission within certain geographical limits. They are guaranteed spectrum on-demand when the IU is absent but the frequency of operation is allocated dynamically. They are protected against interference from unlicensed users in the lowest tier.
- General Authorized Access (GAA): Unlicensed users constitute the lowest tier. GAA users access the spectrum opportunistically and are not guaranteed any interference protection.

The secondary users in the CBRS, i.e., the PA users and GAA users, are collectively referred to as the Citizens Broadband Radio Service Devices (CBSDs). The frequency coordinator in the CBRS is termed the Spectrum Access System (SAS). The SAS is the heart of the CBRS and acts as the intermediary that tracks spectrum occupancy and performs managerial functions like coordinating frequency and power allocations for secondary users. Its operation is highly autonomous and crucial for maintaining a stable and fair spectrum environment. The SAS operation will be supported by a network of sensors near fixed IU installations with Environment Sensing Capability (ESC). This subsystem will monitor the radio environment and notify the SAS of impending IU transmissions.

1.2 Motivation

There are several challenges evident in the CBRS proposal outlined above. The foremost among them is IU security and privacy. Since federal agencies constitute the IUs on several of the underutilized channels, the privacy and security of the CBRS operation must be considered. Further, there is a strict time bound imposed by the FCC for the CBSDs to stop transmitting and clear the channel for IU transmission. The original FCC guideline states that the SAS must either suspend or relocate CBSDs within 60 seconds of receiving a notification about an IU's arrival. This was later extended to 300 seconds as field studies showed that the previous limit could not be met even under ideal conditions [5].

On the other hand, the CBSDs must be guaranteed reasonable spectrum utilization for the spectrum sharing system to be fair and commercially viable. This precludes taking a conservative approach toward CBSD spectrum allocation. The FCC proposal addresses this aspect by setting threshold levels for the minimum amount of spectrum that must be assigned to CBSDs in the absence of the IU. This differs for the PA and GAA tiers. Further, the FCC mandates that PA users must be assigned an alternate frequency for transmission on arrival of an IU.

The above mentioned aspects reveal the need for efficient and streamlined SAS algorithms. These must be designed with a view toward balancing the trade-off between protecting IU's privacy and upholding the FCC guidelines and maximizing spectrum utilization. A recent field trial of the 3.5 GHz showed that while the extended CBSD evacuation time limit is met, the major delay comes from reassigning the CBSD to an alternate frequency [6]. This aspect is of utmost important and should be addressed to minimize the inconvenience to CBSDs and to optimize the required time for clearing the spectrum.

1.3 Proposed Work

In this thesis, we explore the possibility of the SAS learning the IU's spectrum usage pattern and using this knowledge to improve its efficiency. Specifically, we explore the use of IU occupancy prediction in setting the time validity of spectrum grants issued by the SAS to CBSDs. This is done with a view toward minimizing revocation overhead in case of IU arrival. The effect of this on channel utilization by CBSDs is also investigated.

The first part of the study deals with predicting IU spectrum occupancy. Since one of the main concerns in CBRS is the privacy of IU activity information, it is unlikely that federal agencies would share details about their impending transmissions. Further, FCC regulations forbid the SAS from storing information about past IU transmissions. The possibility of federal agencies maintaining a separate database and having a secure querying between commercial SAS and federal SAS has also been rejected. The FCC envisions a situation wherein arrival of the incumbent is only detected and conveyed by the network of sensors with ESC. In this thesis, we propose the use of a deep learning model trained on historical IU spectrum usage information to predict when an IU is likely to arrive. This model can be trained by the IUs or authorized ESCs. Thus, this approach avoids storage of IU occupancy data as-is by the SAS while providing the benefits of improved SAS function efficiency.

The second part of the study investigates the impact of knowledge of IU arrival on spectrum grant revocation complexity. We propose that the SAS issue time-limited spectrum grants to CBSDs such that the number of revocations to be issued when an IU arrives is minimized. There are several advantages to this approach. This will eliminate the need to broadcast a message about the IU's arrival, thus preventing privacy leakage from SAS's notification messages and ensuring fast evacuation. Since IU's arrival is computed ahead of time, CBSDs that require longer transmission times can request extra license on a different band in advance. The proposed approach can also help provide a seamless transition of CBSDs to an empty channel, thereby reducing the time taken to clear the channel and minimizing interruptions to CBSD transmission.

1.4 Research Contribution

With regard to SAS algorithms, the aspects discussed in the FCC proposal serve only as guidelines, and the exact implementation is left to the discretion of the stakeholders [4]. Since the SAS is being developed by competing stakeholders and commercial deployments are pending FCC approval, the SAS operation and algorithms as used in the field tests are not available in the public domain. The details discussed in this thesis derive from proceedings at the FCC workshop on SAS architecture and operation and technical specifications proposed by the Wireless Innovation Forum (WInn Forum) [4, 7].

To the best of our knowledge, there is no work on improve SAS functions based on IU spectrum usage prediction. Previous studies have analyzed improving the efficiency of the SU cognitive functions based on knowledge about PU activity. However, all these studies consider a decentralized architecture in which SUs monitor the spectrum independently by performing individual sensing. The prediction algorithms are customized for implementation on low memory and low power devices, which is not true of SAS. Moreover, SUs cannot track incumbent usage over a long time/geographical region as is possible for a SAS and hence their knowledge will be limited.

There is also no work on techniques to minimize revocation overhead from the SAS perspective. As mentioned earlier, the FCC does not provide any details on how the SAS is expected to perform spectrum management. The protocol specification published by the WInnForum [7] is used as a guide with reference to the spectrum grant issued by the SAS for CBSDs. Therefore, the current work is in direct response to published findings from field trials and gaps in existing literature.

1.5 Report Outline

A description of the architecture and operation of the CBRS is provided in Chapter 2. Aspects that are key to our research work, namely, the mandates regarding the privacy of IU data and the procedure followed in issuing and revoking spectrum grants to CBSDs, are emphasized. The proposed work is introduced in Chapter 3 along with a high-level discussion on what we intended to achieve. The first aspect of our proposed work, IU spectrum occupancy prediction, is detailed in Chapter 4. A description of the design of the LSTM model and its implementation for a real-world spectrum occupancy dataset is provided. An investigation on the impact of this knowledge of IU spectrum occupancy on SAS functions, specifically, minimization of the number of CBSD license grant revocations on the arrival of an IU, is undertaken in Chapter 5. Results from a simulation of a CBRS are presented and analyzed. The conclusion and possible future directions for research are presented in Chapter 6.

Chapter 2

Citizens Broadband Radio Service

The DSA paradigm has been adapted for implementation in North America by the FCC. The main objective of this effort is to reclaim spectrum that was previously set aside nationwide for federal and non-federal purposes but is in use only in specific locations. This spectrum would be made available to commercial operators for both short-term licensed access and unlicensed access. The 3.5 GHz Band has been identified as being suitable for this purpose and the CBRS has been proposed to investigate the feasibility of spectrum sharing [3]. In this chapter, the architecture and operation of the CBRS is described in detail. The current status of the commercial deployment of the CBRS is also discussed.

2.1 CBRS Architecture

The FCC has recommended that the CBRS be implemented as a three-tier service providing varying levels of access, namely, incumbent access, priority access (PA), and general authorized access (GAA).

Incumbent access is provided to the current federal and non-federal licensed users in the 3.5 GHz band. These include radiolocation service and aeronautical radio-navigation service radars operated by the Department of Defense (DoD) and receive-only operations of Fixed Satellite Service (FSS) earth stations. The DoD radars comprise both fixed and mobile systems. The FSS incumbents constitute the non-federal users. In any case, IU transmissions have first priority and these users are protected from interference from users in other tiers.

In the absence of IU transmission, the spectrum is shared among the PA and GAA users. The former includes entities like hospitals, utilities, and first responders, all of whom have critical spectrum needs. These users are protected against interference from the GAA users to a certain extent. The FCC has mandated that PA users be guaranteed spectrum in the absence of an IU. Further, even on arrival of an IU, PA users must be provided with an alternate frequency for transmission. The GAA users access the spectrum opportunistically and are not protected from interference from any type of users. In order to encourage adoption, a portion of spectrum is reserved for the GAA users subject to incumbent usage.

The centralized frequency coordinator in the CBRS spectrum sharing architecture is termed as the Spectrum Access System (SAS). The SAS maintains the radio environment status in a database and ensures appropriate interference protection for eligible users. It also ensures that spectrum use is maximized by coordinating dynamic frequency assignment for the CBSDs. The SAS monitors the spectrum environment at a given geographical location and time and computes the maximum permissible power levels for CBSDs. It ensures that, once assigned a power, the CBSDs operate within specified parameters and do not cause interference. SAS functions will be highly automated in order to use the spectrum to full potential. The SAS also performs administrative functions such as the authentication and logging of CBSDs. From the above discussion it is evident that SAS functioning is critical to maintain a stable and fair spectral environment.

Finally, information about federal incumbent transmissions will be relayed in real time by a network of infrastructure-based and device-based sensors deployed near federal radar facilities with environment sensing capability (ESC). This network will be administered by a commercial entity and neither the DoD nor the NTIA will oversee operations on a day-to-day basis. The overall architecture is shown in Figure 2.1.

2.2 CBRS Operation

Next, the CBRS operation is described with regard to the interaction between the SAS and the CBSDs. This is in accordance with the protocol document issued by the WInnForum

Figure 2.1: CBRS Architecture

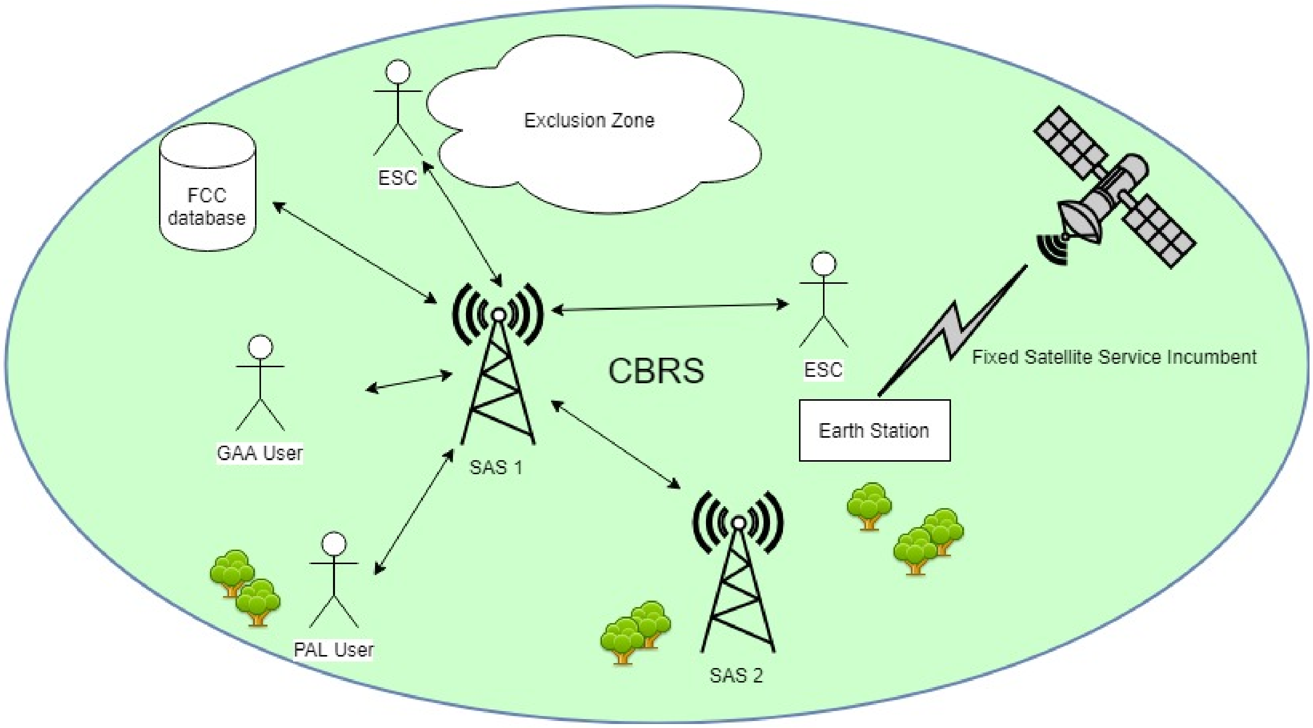

[7]. This protocol has also been followed in implementing field tests [6, 5]. The SAS procedures, as described in this document, are shown in Figure 2.2.

First, a CBSD locates the nearest SAS and registers by providing location and operational parameter details. Then, it requests spectrum from the SAS by submitting a spectrum grant request. The SAS executes internal functions to find an appropriate frequency and transmit power. It responds to the CBSD request with the granted frequency, bandwidth, and duration that the grant is valid for. Following this, the CBSD regularly checks in with the SAS by using a periodic heartbeat request. The heartbeat response sent in reply enables the SAS to monitor the CBSD and issue instructions. If a heartbeat request is not

Figure 2.2: SAS Functions

sent, connectivity to SAS is assumed to be lost, and the grant is automatically invalidated. If the CBSD completes transmission before the end of grant time, it sends a relinquish request to the SAS asking for the spectrum to be released back into the pool. In case the SAS needs to revoke/reassign a spectrum grant, it sends a grant revocation notice over the heartbeat response when the CBSD checks in next. This might or might not have information on an alternate spectrum channel that can be used by the CBSD. In all of these scenarios, the SAS can only respond during the next heartbeat response and is not capable of initiating communication with the CBSD directly.

The relevant SAS-CBSD operational messages and their fields are listed in table 2.1.

Table 2.1: SAS-CBSD Operational Messages

#	Message Type	Fields
1	CBSD Grant Request	CBSD ID, Protection level (PAL/GAA), Peak transmit power
2	CBSD Grant Response	CBSD ID, Grant ID, Success (Request approved/disapproved), Grant time, Lifetime of the grant, CBSD polling interval (heartbeat duration), Measurement report configuration, Operation parameters, Error Message
3	Heartbeat Request	CBSD ID, Grant ID, Operation state (Granted/Transmission), Operation parameters
4	Heartbeat Response	CBSD ID, Grant ID, Success (Approved/disapproved), Grant suspension (True/false), Grant termination (True/false), Operation parameters (a new set of operating parameters), CBSD polling interval, Measurement report configuration, Error Message

2.3 SAS Field Trials

Several field trials have been conducted to verify the optimality of the proposed architecture and the feasibility of upholding the proposed operational guidelines [6, 5, 8]. One such field trial [5] was instrumental in determining that an evacuation period of 60 seconds was difficult to meet even under ideal conditions, resulting in the evacuation time being changed to 300 seconds. A recent field trial by the same group [6] measured the time taken to evacuate CBSDs from a frequency band when an incumbent arrives. The procedure followed in assigning the evicted CBSD to an alternative channel was also timed. The former was found to take approximately 90 seconds, while the total time taken in the case of executing the latter took a maximum of four minutes. They found that an effective technique to reduce the evacuation time would be to keep track of alternate channels to which a CBSD could be reassigned. According to the study, this would reduce the evacuation and reassignment time to near two minutes, a reduction of about 70%.

2.4 Summary

In this chapter, the CBRS architecture and operation were described in detail. The conflicting requirements of having to protect IU from interference while enabling SAS to reuse maximum amount of spectrum is evident from this background information. The implications of such trade-offs, as seen in SAS functions, will be illustrated in the next chapter.

Chapter 3

Proposed Work

In this chapter, the challenges in the CBRS system are described. This is followed by our research statement and a detailed explanation of the proposed approach. The two focal points of our proposed approach, namely, the IU spectrum occupancy prediction using deep learning and CBSD spectrum grant revocation are elaborated. Finally, the advantages and limitations of the proposed approach are discussed.

3.1 Challenges in CBRS

As mentioned in the previous chapter, the FCC only provides mandates but does not restrict the operators in terms of the implementation approach or algorithms to be used. In this section, we discuss three key aspects of the CBRS system: IU operational privacy, IU interference protection, and CBSD channel utilization.

The FCC recommends taking a phased approach to IU operation protection since the operation of federal agencies is a matter of national security and must be dealt with cautiously. In the first phase, large geographical areas around the federal facilities, termed exclusion zones, will be exempt from spectrum sharing. Once the ESC network is up and running, the second phase will be executed. In this, the exclusion zones will be opened up for spectrum sharing as protected zones. Spectrum availability within these zones will be conveyed automatically by the ESC network to the SAS.

In terms of the information storage, the FCC instructs that the SAS must log CBRS user activity. Non-federal incumbent locations and other transmission parameter information can be stored to facilitate SAS operation. However, with respect to federal radar installations, the SAS can only store exclusion zone information and not log any information about federal IU spectrum use. In the second phase, once the ESC network is deployed and reports to the SAS about IU transmission, the SAS will not be allowed to store or monitor federal user transmission. It will only receive a notification about the IU transmission and cease CBSD operation. The exact data to be reported by the ESC network is yet to be established by the DoD/NTIA. Since information security is essential to the CBRS operation, the SAS databases must be secured cryptographically. Transmissions from SAS to CBSDs must be encrypted.

With regard to the IU interference protection, the FCC mandates that all CBSDs must be cleared from a channel with incoming incumbent transmission in a timely manner.

Earlier the maximum time to evacuate the channel was set to 60 seconds; following a field test showing the infeasibility of this timeout period, this has been now extended to 300 seconds.

From the above described mandates, it is clear that the success of the CBRS initiative is dependent on ensuring that the IUs are not affected by the spectrum sharing, both in terms of interference and privacy and security. On the other hand, keeping in mind the original objective behind introducing dynamic spectrum access and to attract CBSD stakeholders to invest in this setup, it is essential that the underutilized spectrum be reused to the maximum. According to the FCC proposal, in terms of the CBSD spectrum availability, in the absence of IUs, the PA users can be allocated 70 MHz at the most, while a minimum of 80 MHz should be set aside for the GAA users.

These conflicting requirements, specifically, the IU interference protection and maximization of CBSD spectrum reuse, must be balanced by the SAS. SAS functions must ensure that the spectrum availability guarantee is met in the absence of IUs. However, on arrival of the IUs, the SAS must clear the channel quickly and in a manner that does not compromise the IU's operational privacy.

3.2 Research Objective

In this work, we address the trade-off between providing IU interference and privacy protection and ensuring maximum spectrum reuse. We propose that the SAS extract knowledge about the IU spectrum occupancy patterns from historical data. This data can be used to set the time validity of the CBSD spectrum grants in a manner that reduces the revocation complexity. Lesser number of revocations would enable quicker evacuation of CBSDs on the arrival of an IU. Since the SAS cannot directly monitor the IU transmissions or store them, we propose a privacy-preserving IU spectrum occupancy prediction method using deep learning. The design of the prediction method and the impact of using IU occupancy prediction to set spectrum grant validity are described next.

3.3 IU Spectrum Occupancy Prediction using Deep Learning

The concept of using IU spectrum occupancy information to enhance the efficiency of DSA functions is well-studied and several models have been proposed for predicting the IU spectrum occupancy pattern [9, 10, 11, 12, 13, 14]. However, most of these works predate the CBRS proposal and consider a distributed DSA scenario. Therefore, the objective of these works is different in scope and mostly relate to cognitive cycle functions, namely, sensing the spectrum, decision-making about IU presence, managing shared access to the spectrum, and providing spectrum mobility. The associated metrics dealt with minimizing the spectrum sensing time, energy expended, switching probability, etc. A comprehensive survey of the prediction methods employed for this purpose is provided in [15]. The categorization proposed by the authors is shown in Figure 3.1.

Figure 3.1: Categorization of Spectrum Occupancy Prediction Methods

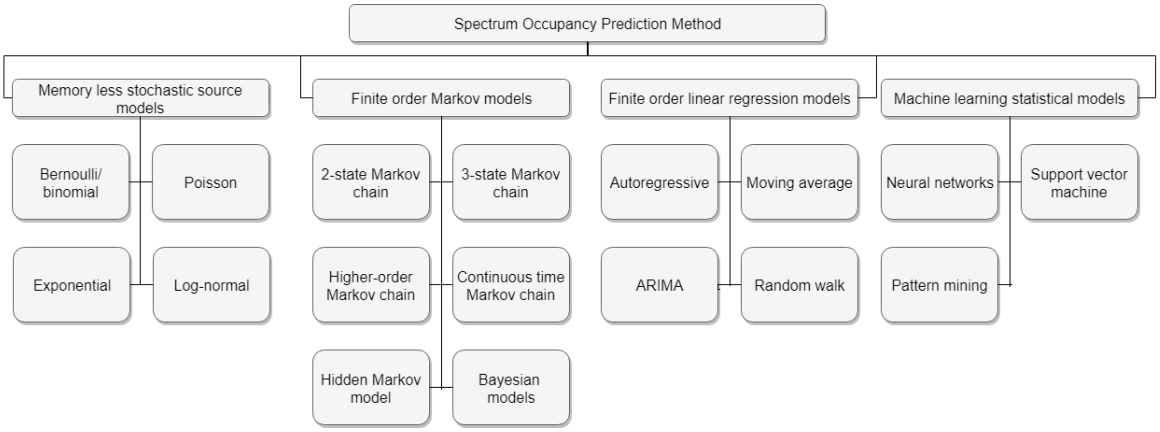

Deep learning models, especially the long short term memory (LSTM) variant which is suitable for time-series prediction, have been proposed for IU spectrum occupancy prediction [16, 17, 18, 19, 20, 21]. As in the case of the above mentioned studies, these too are also from the perspective of a secondary user. The objectives cover a wide range; from improving cognitive function efficiency to detection of primary user emulation attack. However, there are several drawbacks to these works even when considering only the prediction function. Many of these works use simulated data as IU data leading to inherent bias. Further, most of these works attempt to only predict the next time slot.

In this study, we attempt to use LSTM to predict IU spectrum occupancy for several time slots in the future. We envision that this model will be trained either by the IUs themselves or by the ESC network nodes that are authorized to monitor IU signals. The trained model is shared with the SAS, which uses this to obtain future incumbent arrival times.

3.4 CBSD Spectrum Grant Revocation

As described earlier, the SAS is responsible for CBSD spectrum management. Once a CBSD raises a spectrum request, the SAS considers the present spectrum environment and communicates a frequency for transmission as well as the time till which this spectrum grant is valid. In case an IU requires the spectrum, the SAS issues a revocation request to all the CBSDs that are transmitting at that time. This evacuation must be completed under 300 seconds. In this study, we focus on this particular SAS function of issuing a time-valid spectrum grant and revoking it. Specifically, we use the IU spectrum occupancy prediction information to set the time validity of the spectrum grant such that the number of revocations to be issued in case of IU arrival is minimized. To the best of our knowledge, there is no study on improving the efficiency of SAS functions by using IU spectrum occupancy pattern knowledge. Further, there is also no work on manipulating the time validity of the spectrum grant to minimize revocation on the arrival of an IU.

3.5 Analysis of Proposed Work

The proposed work is advantageous from several aspects.

- IU privacy protection: In the proposed IU spectrum occupancy prediction technique, the IU data is not stored as-is. Further, till date, there is no known work on ways to reverse engineer deep learning models to obtain the input dataset. The state-of-the-art in this domain is limited to being able to detect whether a particular record was one of the training inputs given access to the model; this is called a membership attack [22]. However, we do not consider this attack since we assume the SAS to be a secure and trusted party. Finally, a lesser of revocations supports the possibility of notifying the CBSDs individually instead of broadcasting a message about the IU's arrival. This prevents potential privacy leakage from SAS notifications.

- IU interference protection: Minimizing the revocation complexity would enable faster evacuation times and reduce the possibility of interfering with the IU. Further, a knowledge of IU arrival times would help the SAS plan ahead and maintain a list of unoccupied channels in reserve. Once the IU arrival is notified, the SAS could simply redirect the CBSDs to switch to a different frequency, thus minimizing the evacuation and spectrum reassignment procedure.
- CBSD channel utilization: As mentioned above, knowledge of the IU spectrum occupancy pattern would help the SAS assign alternate channels for CBSD transmission, thus reducing the number of CBSD transmission drops. Since the spectrum grants clearly indicate a time validity, CBSDs that require longer transmission times can request extra license on a different band in advance. This would enable seamless handoff and uninterrupted transmission from the CBSD perspective.

The limitations of the proposed work mainly derive from having to monitor the IU transmissions in advance to gather enough information to train the LSTM model. However, since extensive spectrum monitoring has already been undertaken in 3.5 GHz Band to determine its suitability for hosting CBRS, this information could be used. In some cases, there might be no discernable pattern to the IU transmission and for the model to learn and predictions might be entirely wrong. However, as it is shown in Chapter 4, even a model with low accuracy is preferable to providing random assignments.

3.6 Summary

This chapter provided a detailed discussion about the open challenges in the CBRS operation. The research problem was elaborated and the proposed work was described in terms of its merits in addressing the above mentioned challenges. In the subsequent chapters, the two aspects of the proposed work are described along with the investigations on implementing these.

Chapter 4

Privacy-preserving Spectrum Occupancy Prediction

In this chapter, the working of the LSTM model is described in detail. Details about the implementation such as the input dataset and model parameters are provided. The output from the model is shown and the results are analyzed.

4.1 Deep Learning Models for Time-series Prediction

As mentioned in the previous chapter, the state-of-the-art among the previous studies that use deep learning models for incumbent spectrum occupancy prediction use a model called the long-short term memory (LSTM). LSTMs are a type of recurrent neural network

(RNN). RNNs are neural networks equipped with memory elements designed to process sequential inputs. This makes them suitable for modelling applications with temporal dependencies such as in speech recognition, time-series prediction, natural language processing, and gesture recognition. However, when the time dependency extends over long sequence lengths, RNNs suffer from the vanishing gradient problem in which information from previous states decay geometrically over time. LSTMs were proposed to overcome this issue by selectively retaining past observations.

The overall LSTM network architecture follows the same pattern as a typical RNN except that all the neurons are replaced by LSTM cells as shown in 4.1. However, unlike RNNs in which only one memory state is carried forward, two types of memory are carried forward from the previous stage in LSTMs, specifically, long term memory (LTM) and short term memory (STM). Each LSTM cell comprises of four gates, namely, the forget, learn, remember, and use gates.

Figure 4.1: LSTM Cell

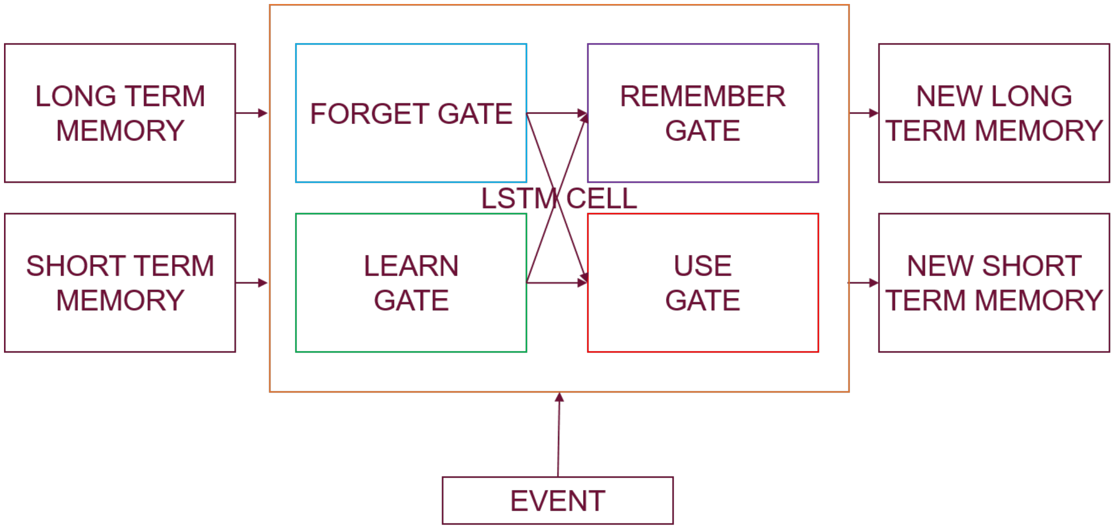

- Learn gate: The learn gate combines information from the STM and current input to generate new learning and is shown in 4.2.

Figure 4.2: Learn Gate

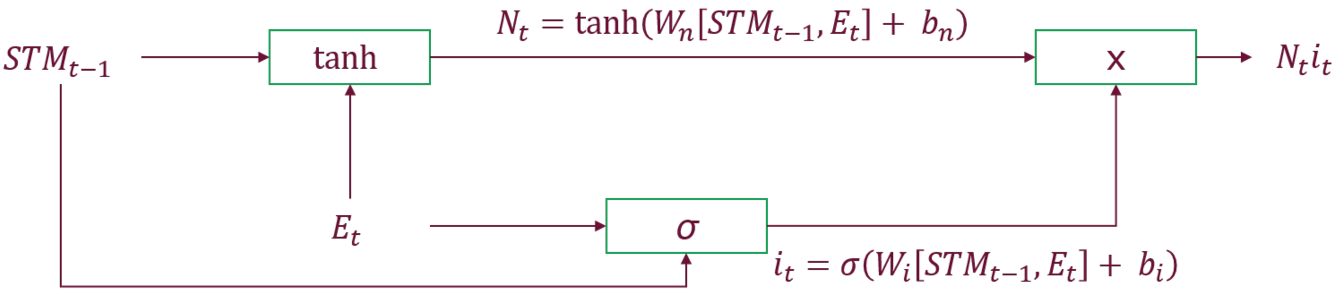

- Forget gate: The forget gate is setup as in 4.3 and involves modifying the LTM data with the objective of removing all unnecessary historical information.

Figure 4.3: Forget Gate

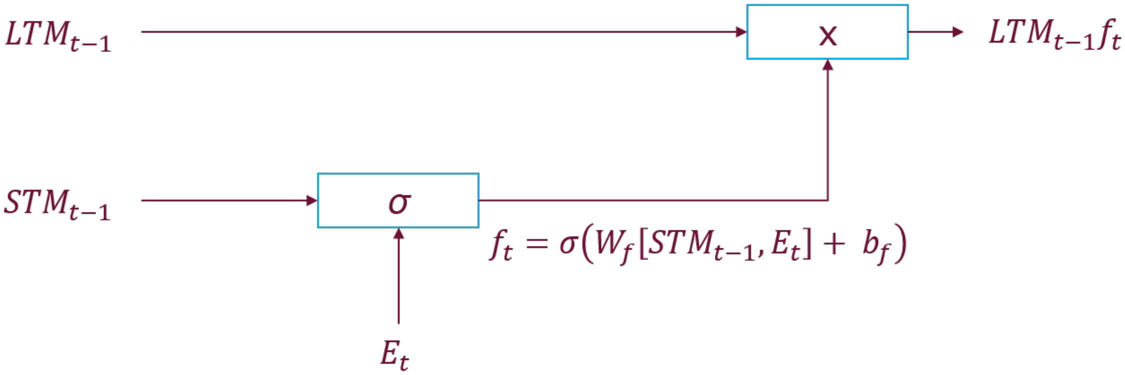

- Remember gate: The remember gate generates the updated LTM by combining the output of the forget gate (useful historical information) with the output of the learn gate (learning from STM and current input). The gate construction is depicted in 4.4.

Figure 4.4: Remember Gate

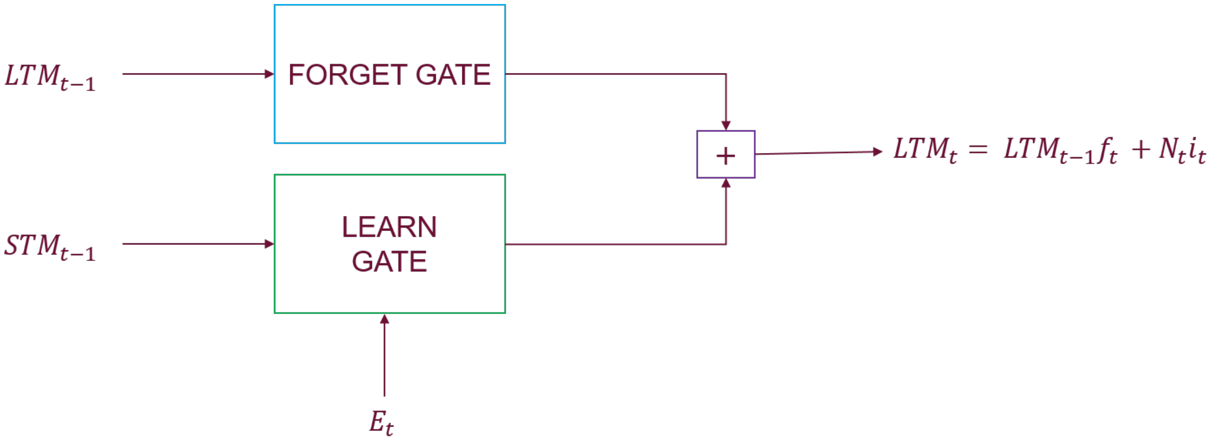

- Use gate: The use gate generates the updated STM which is also the current output by combining the output of the forget gate (useful historical information) with a selective learning from the STM and current input. It is shown in 4.5

Figure 4.5: Use Gate

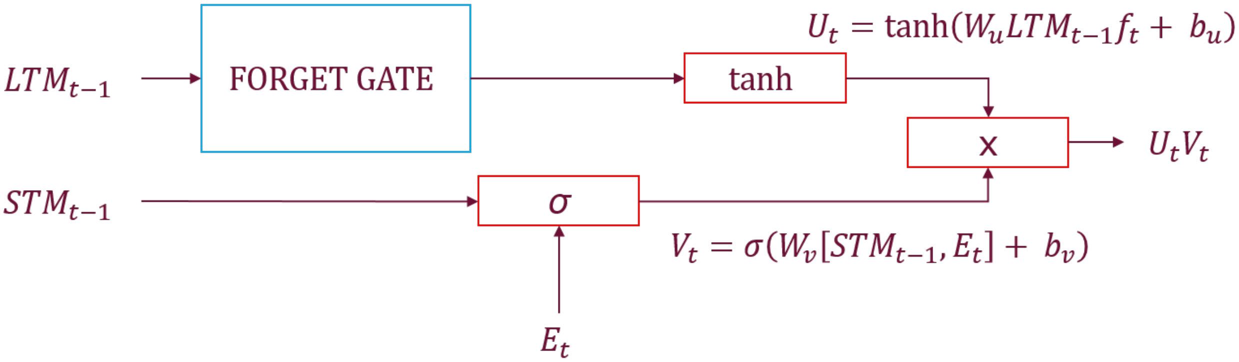

In order to implement prediction over several future time slots, the LSTM network is configured as a sequence-to-sequence model. This setup comprises an encoder-decoder framework to implement prediction over many future time slots. In this, the functions of extracting insights from historical information (encoder) and using it to provide future predictions (decoder) are segregated. The encoder observes the complete input sequence till a certain end point and updates its internal state; this provides a context for the next sequence to be predicted by the decoder. The decoder uses both this context information as well as immediate previous input from the previous step to predict next value for that output window.

4.2 Implementation

The most important aspect in a deep learning model is the data on which it is trained. In this case, the input data would ideally represent measurements of the radio spectrum environment over a period of time. While examining the feasibility of implementing spectrum reuse, several spectrum occupancy measurement campaigns were undertaken across different frequency bands and geographical locations. One such campaign was conducted by the RWTH Aachen University at three different locations in Aachen, Germany, and Maastricht, Netherlands [23]. The measurements from this study are available in the public domain. The readings are in the form of power spectrum density (PSD) measurements at different times of the day. The measurement campaign collected data over several days. We use the PSD values for one of the frequency bands in this dataset as the input to our LSTM model.

The design of the LSTM model used to predict IU spectrum occupancy is described next. The input was divided into 3 batches of sequence length 60. The model was designed with 2 hidden layers with 40 nodes in each layer. The 'adam' optimizer was used to tune the model with the 'root mean square' value as the metric to be minimized. The learning rate was set to 0.001. The mini-batch size was set to 256 and the model was trained for 200 epochs.

4.3 Results and Discussion

The input to the LSTM module is shown below. This represents the PSD values captured over three weekdays and aggregated into values over five minutes. PSD values over two days were used to train the LSTM model. This is shown in Figure 4.6.

Figure 4.6: Training Input

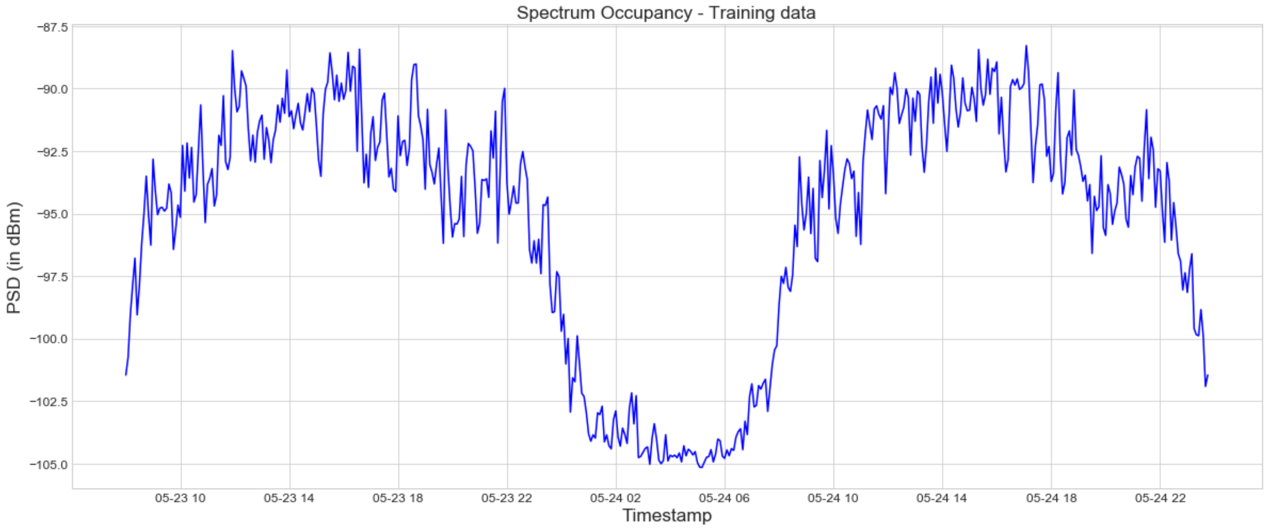

The accuracy of the trained model was tested by making predictions over the third day. The test dataset is shown in Figure 4.7. The test dataset was divided into three sessions

– forenoon, afternoon, and evening. Predictions were made separately for each of these segments.

Figure 4.7: Test Input

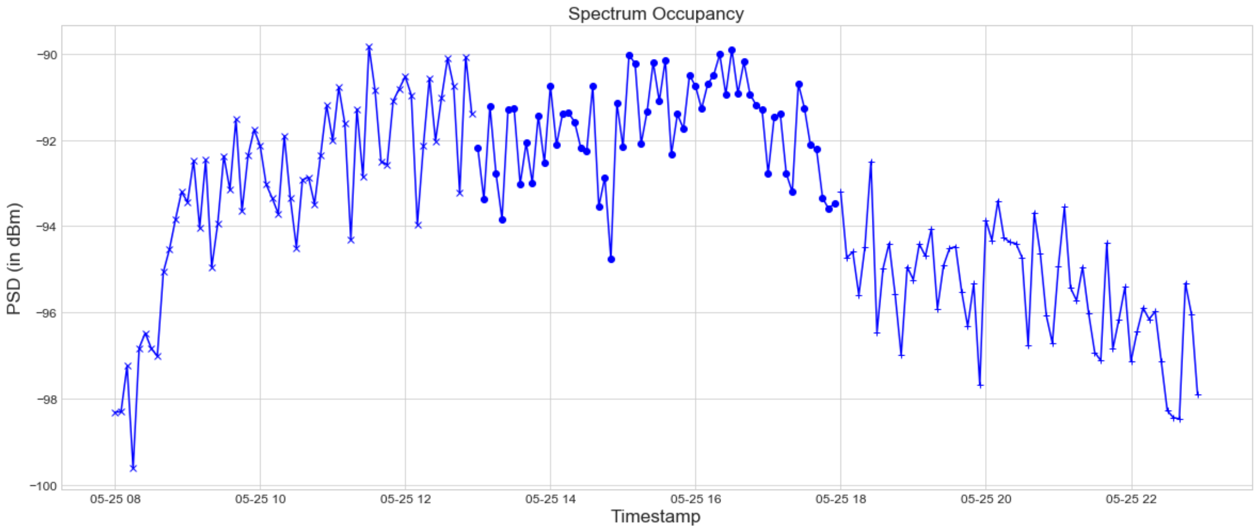

The predictions from the LSTM model are shown below along with the real future values for three different segments. In total, 180 future predictions were made. The continuous PSD measurements were converted to binary states (ON/OFF) using a threshold of -94 dBm. The results showed that the prediction misclassified 10 points over the entire sequence, yielding an accuracy of 94.4%.

There are some limitations to the model proposed here. We consider the variation in PSD across three days for only one frequency. Since the observations were recorded in Europe, the IUs in the band are not the same and their characteristics may vary. In a practical scenario, the prediction model parameters will need to be recalibrated for each frequency and geographical zone. Moreover, the model does not consider the variations

Figure 4.8: LSTM Model Prediction for Batch 1

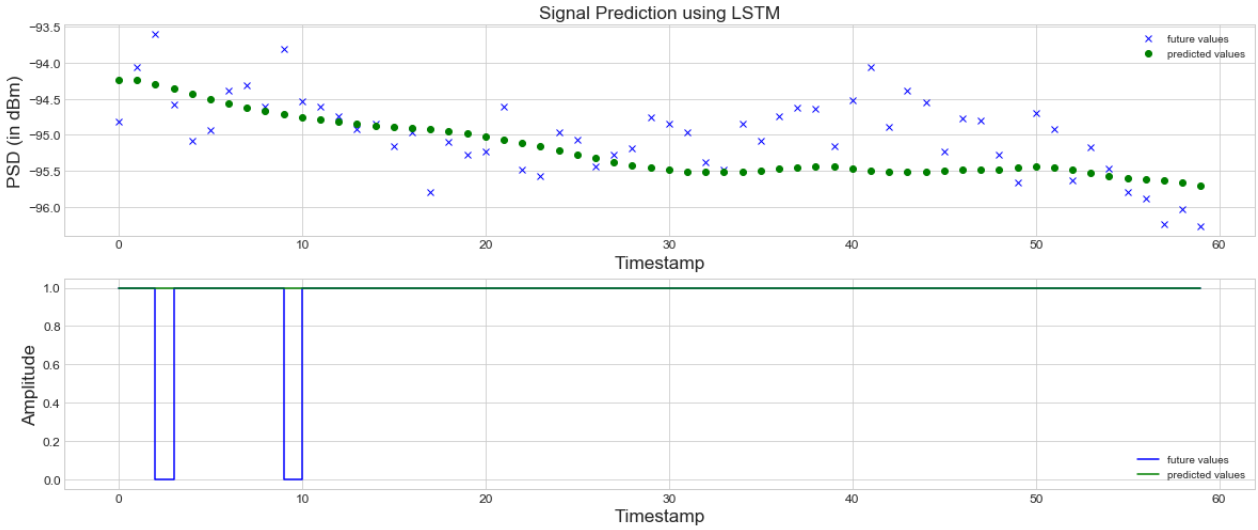

Figure 4.9: LSTM Model Prediction for Batch 2

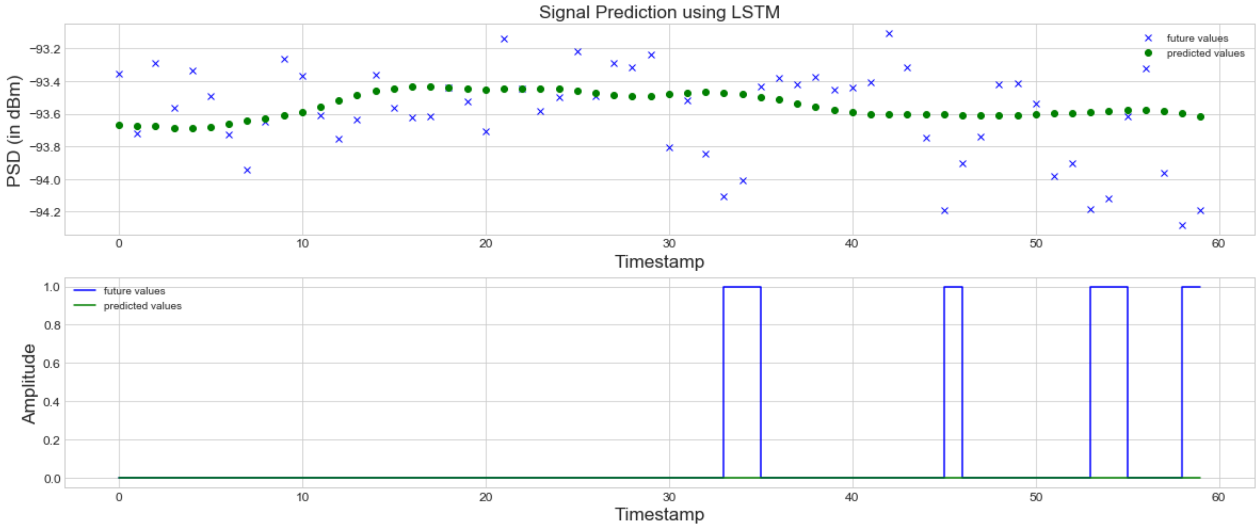

generated by special events that might result in out-of-ordinary IU data traffic levels.

Figure 4.10: LSTM Model Prediction for Batch 3

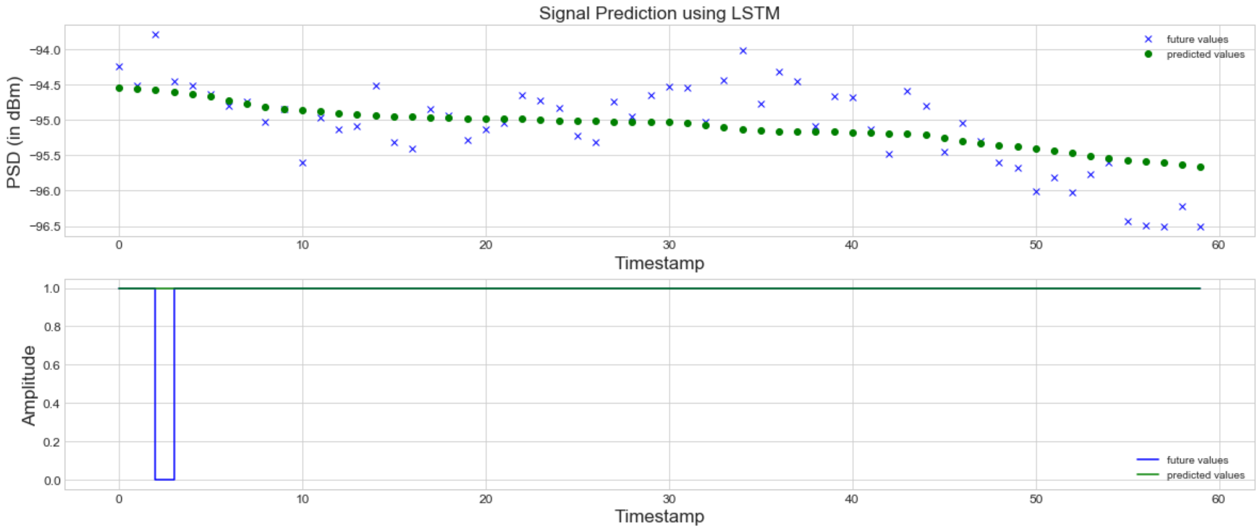

4.4 Summary

This chapter described the design of the LSTM model and the results obtained by training the model on 48 hours of IU transmission information. The model successfully predicted future spectrum occupancy for the next 15 hours with 94.4% accuracy. The next chapter describes how this prediction information is used in a CBRS simulation to set the spectrum grant time validity. The impact on the number of revocations to be performed is analyzed.

Chapter 5

License Grant Revocation Based on Spectrum Occupancy Prediction

In this chapter, the simulation performed to evaluate the impact of knowledge about IU arrival on SAS functions is described in detail. The quantized spectrum occupancy measurement dataset used in Chapter 4 is used as input to the simulation. The prediction information from the LSTM model is used to set the spectrum grant time validity and the complexity of revocation is simulated for this case. In order to provide a baseline for comparison, knowledge about IU arrival times is simulated with differing levels of accuracy and revocation complexity for these are computed as well. The impact on CBSD channel utilization is studied as well.

5.1 Simulation

The simulation was coded in Python 3 using SimPy, a process-based discrete-event simulation framework. Using this framework, operations like IU transmission and CBSD transmission can be modelled as processes. Processes are suspended by yielding to an event and once the same event is triggered, the suspended process resumes. Processes are put to sleep for a given time by using the timeout event and this is used for modeling actual transmission events. The framework also provides support for shared resources. This is used to simulate a channel that has a limited capacity and is accessed by processes in an exclusive or non-exclusive basis. In this simulation, we set the maximum channel capacity to be 10. The IU is assumed to always occupy full capacity.

5.2 Simulation Setup

The classes used in the simulation are shown below.

These classes are used to implement the following sequence of simulation events.

- On the arrival of an IU, a timeout of one unit is provided as a buffer to clear the channel. In this time, the following actions are initiated. The channel log is checked for presence of CBSDs. If there are any present, an interrupt is raised and their spectrum grant is revoked. The CBSD IDs are cleared from the channel log.
- Once the revocation is complete and the channel is cleared, the IU is appended to

Table 5.1: Simulation Classes

#	Class	Fields
1	Channel	users_log, revocation_log
2	License	start_time, end_time
3	PrimaryUser	ID, channel, events: busy_over (triggered on successful transmission completion to notify SAS that channel is free), finished (triggered on successful simulation completion)
4	SecondaryUser	ID, channel, license, events: finished (triggered on successful transmission completion and also after forced termination)

the front of the queue. A timeout is used to simulate the IU transmission. Once this is complete, a busy_over event is raised and the IU ID is removed from the log. This serves as a notification to the SAS that the IU is leaving the channel.

- The busy_over event also triggers the CBSDs to poll the SAS for spectrum grants. These are served in a first-come-first-served basis till the channel capacity permits. The start time for the issued grants are taken as the current time and the end time is set to 1 time unit before the next anticipated IU arrival time.
- The CBSDs that are issued a spectrum grant are queued in the channel user log. Once their transmission is complete they trigger a finished event to notify the SAS that they are relinquishing their grant. This event is also triggered when the CBSD operation is interrupted either owing to end of spectrum grant or arrival of IU.

As mentioned before, IU arrival time and transmission duration is derived from the real-world dataset. On the other hand, since the CBRS is not yet active, there are no real-time measurements of CBSD traffic. Therefore, we use a statistical model in order to simulate CBSD traffic. Specifically, we use the approach followed in /cite MosbahHSA17, and assume that the CBSDs follow the Erlang-loss queuing system. In accordance with this model, the CBSDs are assumed to arrive with an aggregate Poisson arrival process of rate λ. The CBSD transmission times are modelled as being exponentially distributed with a service rate ρ. The ratio is called load factor and is set to 1 in this simulation. Further, it is assumed that there is no queuing of CBSD requests; if the channel capacity is full, no spectrum grant is issued by the SAS. The spectrum grant/response communication time is assumed to be negligible when compared to the CBSD transmission time and is not accounted for.

5.3 Results and Discussion

The input dataset is shown in Figure 5.1.

Figure 5.2 shows the channel occupancy scenario for the case predicted by the LSTM model. Figure 5.3 shows a simulated channel occupancy scenario over the dataset period when there is no prior information about the IU arrival. The SAS attempts to guess this based on prior arrival times and transmit durations and assigns spectrum grants accordingly. However, this scenario results in poor spectrum utilization when the arrival time

Figure 5.1: Simulation Input

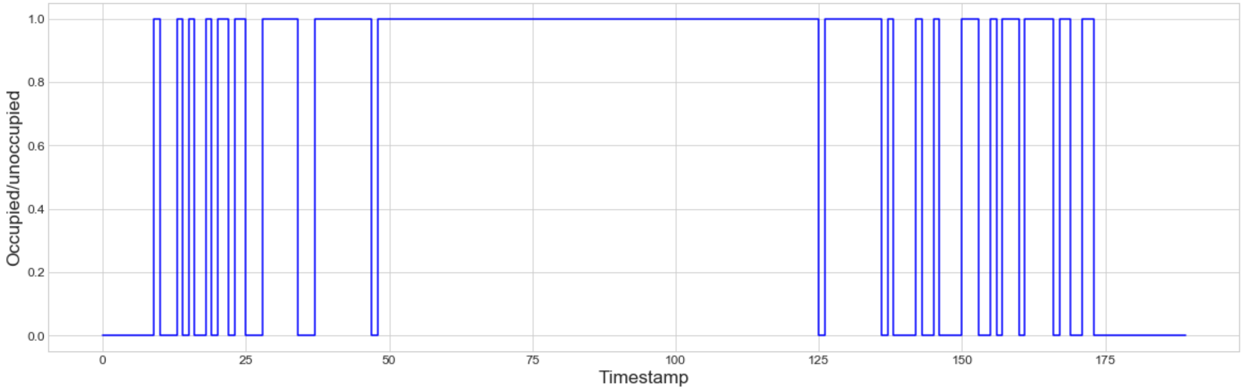

is predicted incorrectly. Further, since the CBSD spectrum grant license validity is not computed keeping in mind the IU occupancy, the complexity associated with revoking the spectrum grants when an IU actually arrives is large. This entails issuing revocation responses to several nodes and might affect the evacuation time adversely. In comparison, Figure 5.4 shows the case in which the SAS has complete knowledge of IU spectrum occupancy. This allows it to generate spectrum grants that end just before the IU is expected. This minimizes the number of revocations to be performed when an IU actually arrives, enabling evacuation to be performed in a timely manner.

In order to analyze the variation in revocation complexity with accuracy of LSTM prediction, IU predictions with varying levels of accuracy were generated. The total number of revocations across differing levels of IU knowledge are plotted in Figure 5.5. A steady decline is observed in the revocation complexity as accuracy of IU prediction increases. The average number of revocations per IU arrival is also shown in the same graph.

Additionally, the spectrum occupancy level was monitored from the CBSD perspective to

Figure 5.2: Channel Occupancy using IU Occupancy Prediction from LSTM Model

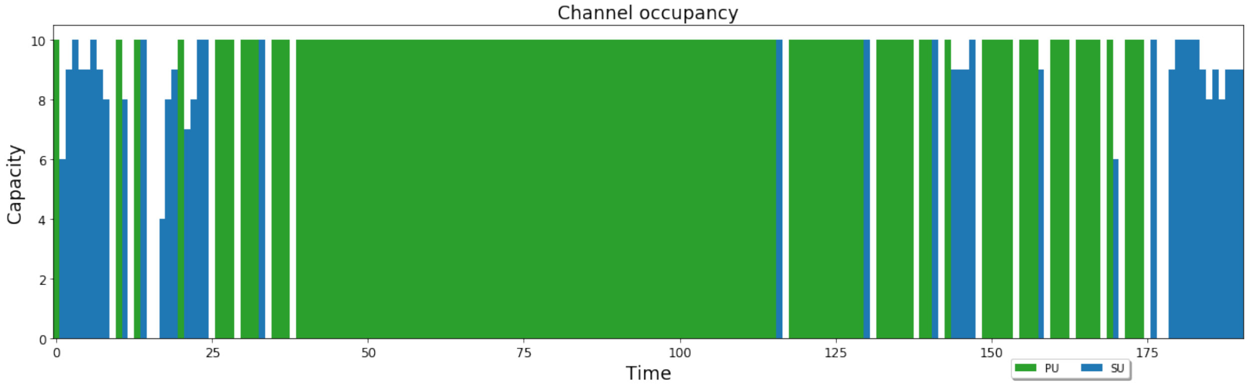

Figure 5.3: Channel Occupancy without Prediction

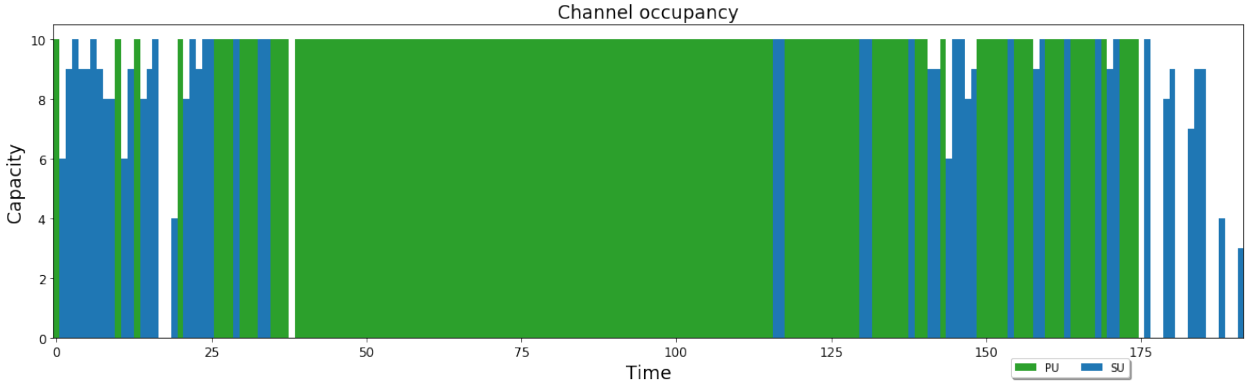

Figure 5.4: Channel Occupancy with Complete Knowledge of IU Activity

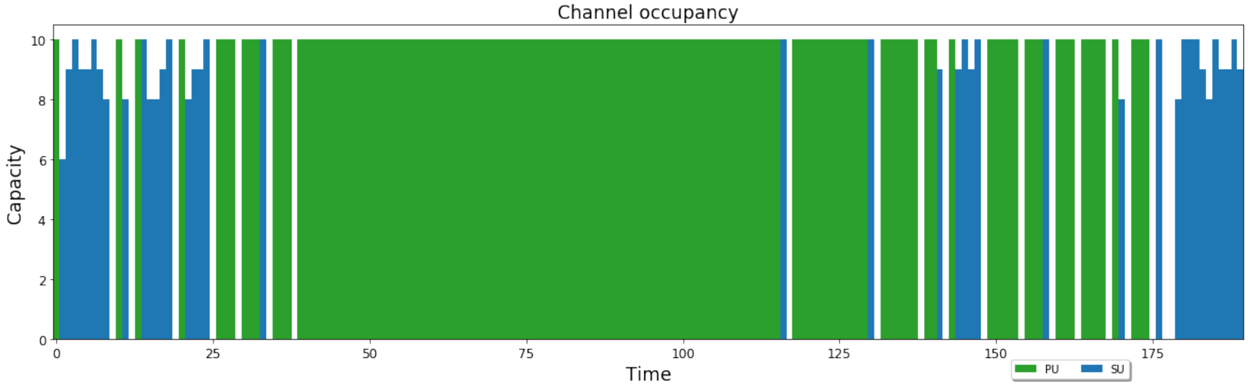

Figure 5.5: Number of Revocations Across Differing Levels of IU Occupancy Knowledge

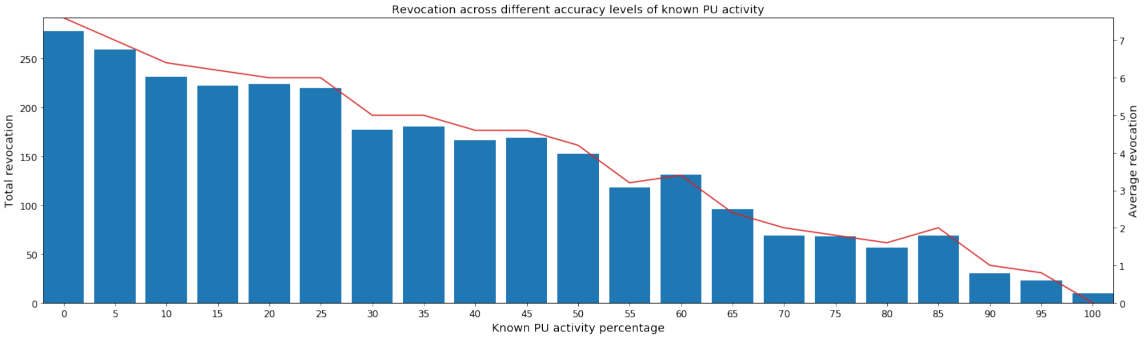

study the impact of refraining from issuing spectrum grants near the expected time of IU arrival.

Figure 5.6: CBSD Utilization Across Differing Levels of IU Occupancy Knowledge

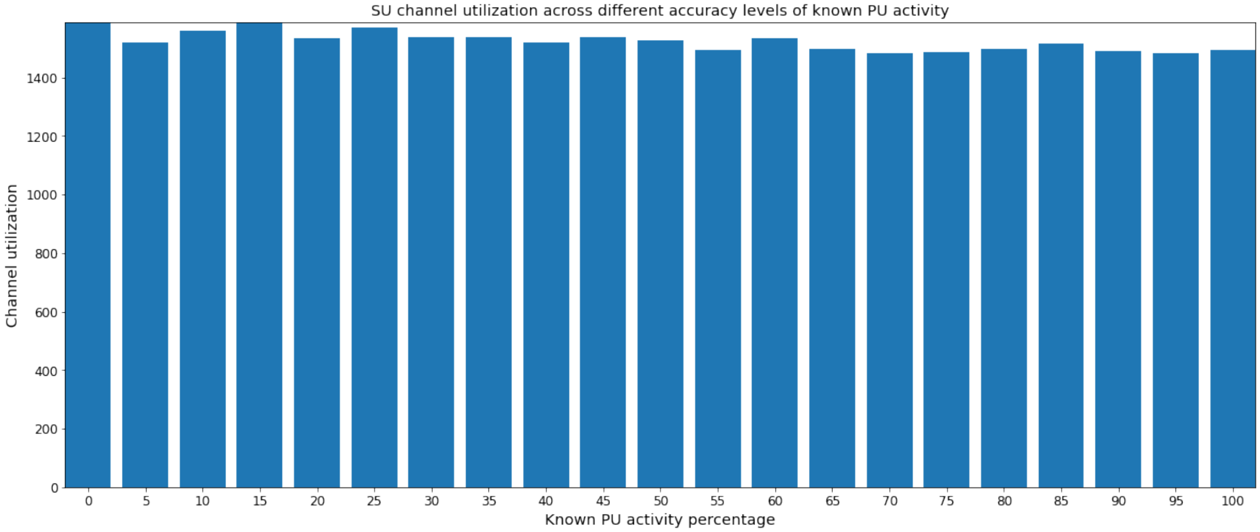

While there is some decrease in the level of CBSD channel utilization, it is not as drastic as expected. This is because in the case of incomplete information, the SAS was avoiding issuing spectrum grants near expected IU arrival time, resulting in utilization loss. In case of prediction accuracy of 90% and more, the loss in utilization is offset by the decrease in revocation complexity. This is the revocation complexity-utility tradeoff.

5.4 Summary

The simulation results showed that knowledge of IU spectrum occupancy pattern could be used to good effect in reducing the number of revocations to be performed on arrival of an IU while having a minimal impact in terms of decrease in utility. Specifically, by incorporating the prediction information with 94% accuracy, the number of revocations decreased by 87.5% when compared to a case in which spectrum grants were issued with random time validity. The corresponding reduction in channel utilization was 7.6%. A summary of the thesis is provided in the next chapter and possible directions for future work are discussed.

Chapter 6

Conclusion and Future Work

6.1 Conclusion

With commercial deployment of the Citizens Band Radio Service commencing in the last quarter of 2018, efforts are in progress to improve the efficiency of the Spectrum Access System (SAS) functions. An area of concern as identified in recent field trials is the time-bound evacuation of unlicensed secondary users from a frequency band on the arrival of an incumbent user (IU). In this thesis, we propose a way to optimize the evacuation process by reducing the number of secondary user spectrum grant revocations to be performed. The proposed work leverages knowledge of incumbent user spectrum occupancy pattern obtained from historical spectrum usage data. Using an example model trained on 48 hours of an incumbent user transmission information, we demonstrate prediction of future incumbent user spectrum occupancy for the next 15 hours with 94% accuracy. The SAS uses this information to set the time validity of the secondary spectrum grants appropriately. In comparison to a case where spectrum grants are issued with no prior knowledge of IU spectrum occupancy, the number of revocations declines by 87.5% with a 7.6% reduction in channel utilization. Further, the proposed technique provides a way for the SAS to plan ahead and prepare a backup channel to which secondary users can be redirected; this is estimated to reduce the evacuation time by 70%.

6.2 Future Work

This work is only a preliminary analysis of the techniques by which IU spectrum occupancy can be predicted and SAS functions can be enhanced by using this information. In terms of spectrum occupancy prediction, in this work, we only consider time-series prediction. However, there is a possibility of correlation between occupancy in adjacent frequencies and in frequencies within the same geographical area. Therefore, a model that takes into consideration these factors could potentially yield more accurate results on the whole. This could involve different architectures like a convolutional neural network or a combination of these. Further, since the ultimate objective is for the SAS to be autonomous in operation, reinforcement learning can be used to enable the SAS to learn for itself the arrival of an IU and seamlessly integrate this learning with other functions.

In terms of SAS functions, this work considers only the feasibility of reducing the number of spectrum grant revocations. However, this idea of using IU spectrum occupancy information to enhance SAS algorithms can also be extended to other functionalities. For example, the SAS could use the spectrum occupancy pattern to compute the expected system capacity and vary the probability of issuing the spectrum grants appropriately. Also, as discussed previously, this approach could play an important role in preplanning the reassignment of CBSDs to alternate bands and optimize handoff to a different SAS.

Bibliography

[1] Mobile data traffic growth outlook – ericsson. [Online]. Available: https://www.ericsson.com/en/mobility-report/reports/november-2017/mobile-data-traffic-growth-outlook

[2] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey," *Computer Networks*, vol. 50, no. 13, pp. 2127–2159, Sep. 2006. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1389128606001009

[3] F. GN Docket No. 12-354, Further Notice of Proposed Rulemaking, "In the Matter of Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550-3650 MHz Band," Tech. Rep., 2014.

[4] R. GN Docket No. 12-354, Order, and F. Second Further Notice of Proposed Rulemaking, "Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550-3650MHz Band," Tech. Rep., 2015.

[5] M. Palola, V. Hartikainen, M. Makelainen, T. Kippola, P. Aho, K. Lahetkangas, L. Tudose, A. Kivinen, S. Joshi, and J. Hallio, "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)*. Piscataway, NJ: IEEE, Mar. 2017, pp. 1–2. [Online]. Available: http://ieeexplore.ieee.org/document/7920763/

[6] M. Palola, M. Hoyhtya, P. Aho, M. Mustonen, T. Kippola, M. Heikkila, S. Yrjola, V. Hartikainen, L. Tudose, A. Kivinen, R. Ekman, J. Hallio, J. Paavola, M. Makelainen, and T. Hanninen, "Field trial of the 3.5 GHz citizens broadband radio service governed by a spectrum access system (SAS)," in *2017 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. Piscataway, NJ: IEEE, Mar. 2017, pp. 1–9. [Online]. Available: http://ieeexplore.ieee.org/document/7920760/

[7] H. D. WINNF-15-H-0023, "Interim SAS to CBSD Protocol Technical Report-A," Tech. Rep., 2017.

[8] Google division touts milestones in cbrs 3.5 ghz sas tests, only deployments left. [Online]. Available: https://www.rcrwireless.com/20170224/policy/Google+division-touts-milestones-in-cbrs-3-5-ghz-sas-tests-only-deployments-left-tag2

[9] X. Xing, T. Jing, W. Cheng, Y. Huo, and X. Cheng, "Spectrum prediction in cognitive radio networks," *IEEE Wireless Communications*, vol. 20, no. 2, pp. 90–96, Apr. 2013. [Online]. Available: http://ieeexplore.ieee.org/document/6507399/

[10] A. V. Kordali and P. G. Cottis, “A Reinforcement-Learning Based Cognitive Scheme for Opportunistic Spectrum Access,” *Wireless Personal Communications*, vol. 86, no. 2, pp. 751–769, Jan. 2016. [Online]. Available: http://link.springer.com/10.1007/s11277-015-2955-4

[11] J. Lee and H.-K. Park, “Channel prediction-based channel allocation scheme for multichannel cognitive radio networks,” *Journal of Communications and Networks*, vol. 16, no. 2, pp. 209–216, Apr. 2014. [Online]. Available: http://ieeexplore.ieee.org/document/6812085/

[12] M. Hawa, D. I. Abu-Al-Nadi, O. M. K. Alsmadi, and I. F. Jafar, “On Using Spectrum History to Manage Opportunistic Access in Cognitive Radio Networks,” *IEEE Access*, vol. 4, pp. 5293–5308, 2016. [Online]. Available: http://ieeexplore.ieee.org/document/7539227/

[13] X. Tan, H. Zhang, Q. Chen, and J. Hu, “Opportunistic channel selection based on time series prediction in cognitive radio networks,” *Transactions on Emerging Telecommunications Technologies*, vol. 25, no. 11, pp. 1126–1136, Nov. 2014. [Online]. Available: http://doi.wiley.com/10.1002/ett.2664

[14] Guangxiang Yuan, R. C. Grammenos, Yang Yang, and Wenbo Wang, “Performance Analysis of Selective Opportunistic Spectrum Access With Traffic Prediction,” *IEEE Transactions on Vehicular Technology*, vol. 59, no. 4, pp. 1949–1959, May 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5361432/

[15] H. Eltom, S. Kandeepan, R. J. Evans, Y. C. Liang, and B. Ristic, “Statistical spectrum occupancy prediction for dynamic spectrum access: a classification,” *EURASIP Journal on Wireless Communications and Networking*, vol. 2018, no. 1, p. 29, Feb 2018. [Online]. Available: https://doi.org/10.1186/s13638-017-1019-8

[16] S. Rajendran, W. Meert, D. Giustiniano, V. Lenders, and S. Pollin, “Deep Learning Models for Wireless Signal Classification With Distributed Low-Cost Spectrum Sensors,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 4, no. 3, pp. 433–445, Sep. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8357902/

[17] “Deep Recurrent Neural Network for Multiple Time Slot Frequency Spectrum Predictions of Cognitive Radio,” *KSII Transactions on Internet and Information Systems*, vol. 11, no. 6, Jun. 2017. [Online]. Available: http://itiis.org/digital-library/manuscript/1722

[18] H. Villapol and J. Luis, “Spectrum Analysis and Prediction Using Long Short Term Memory Neural Networks and Cognitive Radios,” p. 62.

[19] L. Yu, Q. Wang, Y. Guo, and P. Li, “Spectrum availability prediction in cognitive aerospace communications: A deep learning perspective,” in *2017 Cognitive Communications for Aerospace Applications Workshop (CCAA)*. Cleveland, OH, USA: IEEE, Jun. 2017, pp. 1–4. [Online]. Available: http://ieeexplore.ieee.org/document/8001877/

[20] L. Yu, J. Chen, and G. Ding, "Spectrum prediction via long short term memory," in *2017 3rd IEEE International Conference on Computer and Communications (ICCC)*. Chengdu: IEEE, Dec. 2017, pp. 643–647. [Online]. Available: http://ieeexplore.ieee.org/document/8322623/

[21] Q. Dong, Y. Chen, X. Li, and K. Zeng, "Explore Recurrent Neural Network for PUE Attack Detection in Practical CRN Models," *arXiv:1805.00428 [cs]*, May 2018, arXiv: 1805.00428. [Online]. Available: http://arxiv.org/abs/1805.00428

[22] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE, 2017, pp. 3–18.

[23] M. Wellens and P. Mahønen, "Lessons learned from an extensive spectrum occupancy measurement campaign and a stochastic duty cycle model." IEEE, 5 2009.