

Scalable Dynamic Spectrum Access with IEEE 1900.5.2 Spectrum Consumption Models

Prasad Netalkar, Carlos E. Caicedo Bastidas, Igor Kadota,
Gil Zussman, Ivan Seskar, and Dipankar Raychaudhuri

Abstract—Dynamic Spectrum Access (DSA) is a key mechanism for meeting the ever-increasing demand for emerging wireless services. DSA involves managing and assigning available spectrum resources in a way that minimizes interference and allows RF coexistence between heterogeneous devices and systems. Such co-existence mechanisms, if they are to succeed when heterogeneous RF devices managed by different entities need to operate in a given area and frequency band (licensed and/or unlicensed), require a common mechanism for expressing the boundaries of spectrum use of each device so that spectrum use deconfliction methods can be built and verified. Spectrum Consumption Models (SCMs) – defined in the IEEE 1900.5.2 standard – offer a mechanism for RF devices to: (i) declare the characteristics of their intended spectrum use and their interference protection needs; and (ii) determine compatibility (non-interference) with existing devices. In this paper, we propose a novel SCM-based Spectrum Deconfliction (SD) algorithm that dynamically configures RF operational parameters (e.g., center frequency and transmission power) of a target transmitter-receiver pair aiming to minimize interference with existing devices/systems. We also propose sequential and distributed DSA methods that use the SD algorithm for assigning spectrum in large-scale networks. To evaluate the performance of our methods in terms of computation time, spectrum assignment efficiency, and overhead, we use two custom-made simulation platforms. Finally, to experimentally demonstrate the feasibility of our methods, we build a proof-of-concept implementation in the NSF PAWR COSMOS wireless testbed. The results reveal the advantages of using SCMs and their capabilities to conduct spectrum assignments in dynamic and congested communication environments.

Index Terms—DSA, SCM, spectrum management, wireless experimentation, research testbed.

I. INTRODUCTION

THE rapid growth in wireless technologies, policies, and user requirements is driving the need for more granular dynamic spectrum access in time, frequency, and geographic location [3]–[6]. Spectrum is a scarce resource and this has

motivated many researchers to develop new spectrum sharing technologies and procedures that aim to achieve high data rates and spectrum use efficiency. Significant effort has also been made to develop co-existence techniques to address interference and the coordinated operation of heterogeneous devices in bands with different operational requirements (licensed / unlicensed) [7]–[10]. For 5G and future wireless systems that support spectrum aggregation and multi-band radio capabilities, dynamic spectrum sharing will be a key enabler to support high user traffic demands and service requirements [11]–[14].

Given the increasing demand for spectrum from commercial services/networks along with the spectrum use and protection demands of passive users (e.g., radio astronomy), unlicensed users (e.g. WiFi 6E, 5G NR-U) and non-commercial active users (e.g., weather radar), coexistence and Radio Frequency (RF) interference management are key challenges in the implementation and design of future spectrum sharing mechanisms [15], [16]. Various schemes for managing spectrum use and co-existence have been proposed in the literature. Collaborative spectrum sharing schemes that leverage a dedicated coordination channel for the exchange of spectrum information were proposed in [8], [17]. A well-known centralized scheme for spectrum management is demonstrated by the Spectrum Access System (SAS) developed for the 3.5 GHz CBRS band [18], [19]. The benefits of such centralized architectures are significant, but they suffer from scalability and single-point of failure issues. Moreover, centralized spectrum management can lead to excessive concentration of market power and insufficient local autonomy needed to drive spectrum use innovation.

Motivated by DARPA’s efforts on the Spectrum Collaboration Challenge (SC2) and also taking inspiration from the decentralized protocols developed for the Internet, this work presents novel approaches to perform large-scale spectrum deconfliction operations for Dynamic Spectrum Access (DSA) environments based on the *Spectrum Consumption Models* (SCMs) defined in the IEEE 1900.5.2 standard [20]. An SCM describes the characteristics and boundaries of spectrum usage of an RF device and it facilitates the autonomous and dynamic selection of spectrum resources across a large number of devices/networks that are sharing spectrum to establish non-interfering wireless communication operations. SCMs simplify spectrum use coordination when compared to other DSA methods that solely rely on sensing to avoid interference between RF devices [21].

This work builds upon our previous research on DSA methods using SCMs [1], [2], further proposing a new DSA proto-

This work was supported in part by NSF grants CNS-2148128, CNS-1827923, OAC-2029295, AST-2232455, AST-2232456, AST-2232459, EEC-2133516, SES-2332054. Partial and preliminary versions of this work appeared in Proc. ACM WiNTECH 21 [1] and Proc. IEEE WCNC 23 [2]. (Corresponding author: P. Netalkar.)

P. Netalkar, I. Seskar, and D. Raychaudhuri are with WINLAB, Rutgers University, North Brunswick, NJ, USA (e-mail: {pnetalka, seskar, ray}@winlab.rutgers.edu).

C. E. C. Bastidas is with the School of Information Studies, Syracuse University, Syracuse, NY, USA. (e-mail: ccaicedo@syr.edu).

I. Kadota is with the Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL, USA (e-mail: kadota@northwestern.edu).

G. Zussman is with the Department of Electrical Engineering, Columbia University, New York, NY, USA (e-mail: gil.zussman@columbia.edu).

col for deconflicting large-scale spectrum use in a distributed manner and effectively highlighting the performance differences between various SCM-based spectrum access schemes. Our focus is on spectrum use deconfliction for fixed location RF devices, and mobility-related issues are left for future research. The main contributions of this paper can be summarized as follows:

- We propose a novel SCM-based Spectrum Deconfliction (SD) algorithm which dynamically configures RF operational parameters (e.g., center frequency and transmission power) of a target transmitter-receiver pair aiming to minimize local and aggregate interference with existing devices/systems.
- Based on the SD algorithm, we develop sequential and distributed DSA methods to assign spectrum to a set of devices within dense large-scale networks.
- We build two custom-made simulation platforms to evaluate the performance of the proposed SCM-based DSA methods in scenarios with a large number of RF devices.
- Finally, to demonstrate the feasibility of our DSA methods, we built a proof-of-concept implementation in the NSF PAWR COSMOS wireless testbed [22]–[24].

The rest of the paper is organized as follows: Section II provides a brief description of related work on dynamic spectrum access and sharing mechanisms. Section III and IV provide an introduction to SCMs and an overview of the compatibility computation process, respectively. Section V describes the spectrum management architecture underlying our framework and experimentation. Section VI describes our proposed deconfliction algorithm and spectrum access methods based on the use of SCMs. Section VII explains the simulation platform created to assess SCM-based DSA algorithms and presents the performance results for the proposed algorithms. Section VIII briefly discusses our experimentation framework to evaluate DSA with SCMs on the COSMOS wireless testbed. Section IX offers directions for future experimentation, and Section X concludes the paper.

II. RELATED WORK

Numerous methods for dynamic spectrum access (DSA) have been proposed and evaluated in the past, including DSA methods for unlicensed spectrum [25]–[27]. The authors in [28], extensively discuss 5G spectrum sharing techniques that utilize DSA, categorizing them based on their architecture (centralized or distributed), spectrum assignment behavior (cooperative or non-cooperative) and method (dynamic exclusive, open access, and hierarchical). The survey covers various spectrum sharing techniques that draw on concepts from game theory, information theory, stochastic modeling, and database-assisted algorithms. Yet, challenges remain in terms of practical implementation, standardization, privacy, interference management, and system architectural design.

In [9], the authors proposed SMAP, a distributed spectrum management architecture that utilizes a policy-driven approach and an aggregated radio map for exchanging spectrum information among peers to assign spectrum. However, the authors did not provide specific details regarding the syntax and structure of the radio maps, and their evaluation was limited to

a small topology where nodes were only allowed to adjust their frequencies. In a similar vein, authors in [29], [30] presented a Radio Environment Map (REM) based spectrum access architecture to protect primary incumbents and share the available spectrum. These REMs are constructed from sensor measurements and are used to determine channel availability and estimate interference levels at specific locations of interest. However, these architectures rely on centralized databases, and the methods for determining compatibility among incumbents and secondary users are not standardized. All of these reported methods often depend on the tools used to perform spectrum use/occupancy analysis, and their results may be subject to debate by each of the parties involved [20]. The research described in [31] proposes an opportunistic protocol for coordinating spectrum access between independent and heterogeneous wireless networks. This protocol involves a central Cognitive Radio (CR) terminal which assigns spectrum aiming to establish fairness based on data flows. In the context of 5G, the work in [32] presents a survey of how 5G networks can benefit from DSA and spectrum sharing techniques.

With the rise in importance of AI/ML methods in the area of Telecommunications [13], [33], in recent years, numerous data-driven AI/ML-based spectrum access solutions have been proposed, e.g., [34]–[39]. In [34], the authors propose a Q-learning based DSA scheme for cognitive industrial IoT networks, enabling spectrum utilization in various scenarios, including orthogonal multiple access (OMA) and non-orthogonal multiple access (NOMA), to ensure guaranteed throughput and reduced interference. Similarly, the authors in [39] present a Q-learning-based framework to address the challenges of opportunistic coexistence between unlicensed LTE and WiFi networks. The authors in [37], present a DSA scheme based on deep reinforcement learning (RL), leveraging inter-node collaboration to enhance spectrum utilization and access efficiency. The work in [38] introduces an innovative multi-agent spectrum access strategy that integrates federated learning (FL) and multi-agent deep reinforcement learning (MADRL) to improve both learning efficiency and privacy. While these AI-driven solutions are insightful, they often face limitations due to the potential mismatch between training mechanisms and real-world environments, high computational complexity, and scalability to large, heterogeneous networks with numerous devices and time-varying conditions.

Despite the extensive research efforts dedicated to enabling the coexistence of diverse networks in the same spectrum bands, it should be noted that much of the work has been theoretical and often only supported by simulations. There are limited scientific papers that evaluate their findings through practical implementations. In our work, we use a combination of simulation and hardware-based implementations to evaluate the performance, scalability, and practical feasibility of our proposed SCM-based deconfliction algorithm and spectrum access methods for coordinating spectrum use in dynamic and dense communication environments. Our approach is easily generalizable and applicable to DSA operations between devices in licensed and unlicensed bands with homogeneous and heterogeneous spectrum use and operational characteristics.

III. SPECTRUM CONSUMPTION MODELS (SCMs)

SCMs provide an information model that can capture the boundaries of the use of spectrum by RF devices and systems so that their compatibility (i.e., non-interference) can be arbitrated by efficient and standardized computational methods [20], [40]. The information captured in SCMs allows for the efficient determination of aggregate interference levels and aggregate compatibility between many devices. It also allows for accounting the impact on spectrum use of RF parameters and phenomena that can be impacted by direction or orientation such as antenna radiation patterns, antenna placement, and propagation conditions. SCMs use a set of 11 data elements, referred to as constructs, which are described below. Unless otherwise stated, the constructs must be used in both transmitter and receiver models.

- 1) **Reference Power:** Value that provides a reference power level for the emission of a transmitter or for the allowed interference in a receiver.
- 2) **Spectrum mask:** Defines the relative spectral power density of emissions by frequency. This construct is mandatory for Transmitter models only (see Fig. 6(a)).
- 3) **Underlay mask:** Defines the relative spectral power density of allowed interference by frequency. This construct is mandatory for Receiver models only (see Fig. 6(b)).
- 4) **Power map:** Defines a relative power flux density per solid angle. It conveys the dispersion of electromagnetic energy from a transmitter's antenna or the concentration of energy at a receiver antenna.
- 5) **Propagation map:** Defines a rate of RF signal power attenuation per solid angle (direction).
- 6) **Schedule:** Specifies the time in which the model applies (start time, end time). Periodic activity can also be defined.
- 7) **Location:** Specifies where an RF device may be used. Locations can be specified as points, areas, volumes or path trajectories. Areas and volumes can have a transmitter density value associated to them. Paths can be used to provide a description of the 3D movement of a device.
- 8) **Policy or Protocol:** (optional) Provides a means to specify a policy or protocol that systems will use for spectrum access and co-existence (e.g. a listen before talk protocol, CSMA/CD, etc.).
- 9) **Intermodulation mask:** (optional) Defines how co-located signals generate intermodulation products in a transmitter or receiver.
- 10) **Platform name:** (optional) 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.
- 11) **Minimum power spectral flux density:** (optional) When used as part of a transmitter model, this value helps determine the geographical extent in which receivers in the system are protected (e.g. coverage area in broadcasting).

Several of the SCM constructs can have a confidence value associated with them. This provides flexibility in the expression of the boundaries of spectrum use by a device so that they don't have to be stated in terms of the worst case scenario.




Fig. 1: Spectrum Consumption Model (SCM) types [1].

The SCM constructs can be used to build different types of SCMs that follow an aggregation hierarchy as shown in Fig. 1 as defined in the IEEE 1900.5.2 standard [40]. 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 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 also be used to structure lists that describe the spectrum that is available for use (spectrum authorization sets), identify constraints to spectrum use (spectrum constraint sets), and list the spectrum being consumed (used) by a group of systems and devices (collective consumption set).

IV. COMPATIBILITY COMPUTATIONS WITH SCMS

The IEEE 1900.5.2 standard not only defines the constructs for SCMs, but also specifies a method for determining the compatibility of spectrum use (i.e. non-interfering) between devices or systems that have expressed their spectrum use boundaries through SCMs [20], [40]. We refer to this method as the Compatibility Test (CT) computation. Two or more RF devices/systems are compatible if their transmitter models do not violate the interference boundaries of any of the receiver models. The CT computation takes into account the locations of the devices being assessed for compatibility, as well as any overlaps in their spectrum utilization in terms of time and frequency. In general, the information conveyed by SCMs determine the details of a link budget computation that is used to assess if a particular transmitter will interfere with a receiver.

A. Compatibility of a single transmitter-receiver pair

In its most basic form, the CT computation is designed to assess whether a transmitter model is compatible with a receiver model. The CT process begins by checking if the SCMs overlap in both time and frequency. If no overlap is found, the devices are deemed to be compatible. However, if there is an overlap, the evaluation continues. The IEEE 1900.5.2 standard describes how to compute the power spectral flux density (PSFD) from the transmitter at the location of the receiver and the corresponding maximum allowed interference power for the receiver. If the transmitter's power at the receiver's

location does not exceed the receiver's maximum allowed interference power, the devices are declared compatible. The difference between the maximum allowed interference power and the transmitter's power at the receiver's location is referred to as the power margin. This power margin can be used to determine the extent to which the transmitter's power could be increased, if necessary (e.g., for coverage expansion), while still maintaining compatibility with the receiver. On the other hand, the power margin can be used to calculate the amount of attenuation required at the transmitter to achieve compatibility if the devices are initially determined to be incompatible.

B. Compatibility for multiple devices

In situations where multiple transmitters and receivers may interfere with each other, the CT can be enhanced so that it is based on the computation of the *aggregate interference* caused by the transmitters under consideration at a particular receiver. Aggregate compatibility is achieved when the aggregate interference at every receiver under consideration is below each receiver's maximum allowed interference power. The IEEE 1900.5.2 standard [40] provides guidelines and a method for computing aggregate compatibility using SCMs for scenarios involving multiple transmitters and receivers. When the locations of transmitters and receivers are fixed, the CT for a single transmitter-receiver (Tx/Rx) pair can be extended to cases with multiple transmitters and receivers relatively easily. However, when there is mobility involved, finding the most constraining configuration for a specific receiver becomes necessary before evaluating compatibility. This configuration should be a feasible positioning and configuration of all devices in the scenario that maximizes the aggregate interference on that receiver [41]. In wireless or Wi-Fi networks that use unlicensed spectrum where multiple devices follow a behavior or policy (e.g. listen before talk) to determine the single device that can access and use a spectrum resource (channel) at a specific moment in time, the determination of the contributions to interference from the active transmitter device can be simplified and will likely follow a spatially dependent probability function. The computation of such function can be done via external tools that leverage the information of the SCMs of the unlicensed devices such as their locations (or density of devices in an area) and transmitter characteristics.

V. SPECTRUM MANAGEMENT ARCHITECTURE AND PROTOCOL

In this section, we provide a description of the architecture and protocol developed for the coordination of spectrum use between multiple independent networks by utilizing SCMs. Fig. 2 shows the layout of the architecture we used for our analysis. We consider that several wireless networks want to operate around the same geographical area and will collaborate to deconflict their spectrum use so that none of the devices within each network interferes with those of other networks. Since each network can be managed by different entities and use different wireless transmission protocols, we refer to them as Wireless Domains (WDs). Spectrum use coordination




Fig. 2: Network architecture overview [1].

between WDs makes use of a network/wireless domain interaction language built on top of DARPA's implementation of CIL (Collaborative Intelligent Radio Networks Interaction Language).

Each network has a designated node, a wireless domain controller, which uses our version of CIL to communicate with other network's WD controller. The controller helps carry out the necessary computations to identify spectrum resources that are available and prevent any interference with other networks operating in the same spectrum. Originally, CIL was designed with a PUB-SUB message queuing system, and in our work, we introduced several adaptations and improvements to support SCMs as the main component of the messages in our WD interaction/coordination language. For high performance and efficiency of the messaging queuing service, Google's Protocol Buffers (protobuf) are used to convert all supported messages into binary blobs. The message types were adapted to support the exchange of SCMs and compatibility reports between the WDs. An example of the interactions between WDs using our CIL variant is shown in Fig. 3. The regional aggregator (RA) helps maintain the topology between multiple WDs and aggregates SCM information.

Some of the messages used for spectrum coordination and the exchange of SCMs are listed below.

- `Register()`: Generated by a WD controller to register with the regional aggregator
- `Inform()`: Regional aggregator informs newly joined WD about existing WDs in the network
- `Notify()`: Regional aggregator notifies existing WDs about the newly joined peer WD
- `SCM_request()`: Message to request SCMs from the peer WDs
- `SCM_response()`: Reply message to send SCM to a requesting peer
- `CT_report()`: Send CT report to peer WDs
- `Calibrate_radios()`: Message to change the configuration parameters of Software-Defined Radios (SDRs) to a specific gain, frequency, modulation, etc.
- `Leave()`: Generated by WD when exiting the system

A. SCM exchange overhead

Whenever a new device joins an operational area, we need to perform CT computations to determine its operational parameters (e.g. frequency, power, etc) so that it will not cause interference to existing receivers, for the case of transmitter or, in the case of a receiver, its operational parameters should allow it to operate reliably when taking into account the aggregate interference caused by other transmitters. As previously discussed, this computation requires the exchange of SCMs either centrally or in a distributed manner (discussed

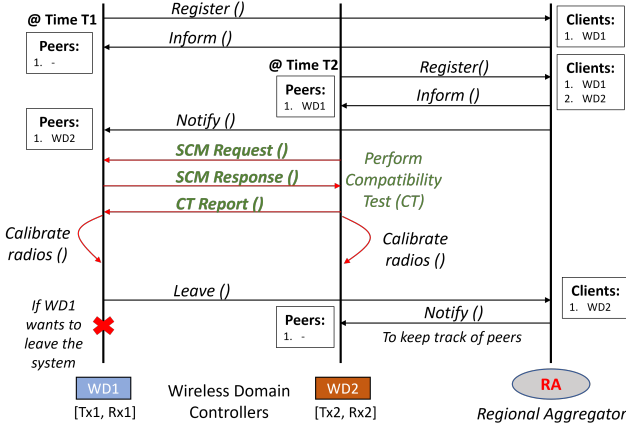


Fig. 3: Network interaction language timing diagram [1].

in Section VI), depending on the chosen spectrum assignment method. However, the overhead from the SCM exchange process is generally minimal. For instance, a new wireless domain composed of one transmitter and one receiver joining a dense operational area only needs to exchange two SCMs (Tx and Rx) with the controllers of its peer WDs. Assuming 20 peer WDs and each SCM being 5KB in size, in a centralized setup where all SCMs are necessary for the CT computation, the overhead due to SCM exchange to allocate spectrum for Tx and Rx in the WD is $20 \text{ peers} * 2 \text{ SCMs (1 for Tx and 1 for Rx)} * 5\text{KB} = 200 \text{ KB}$. This information overhead amount is small when compared to the capacity of the control channel. More details on the CT computation overhead can be found in Section VI-C5.

B. Security mechanisms

The coordination of spectrum use using SCMs described so far is susceptible to attacks from malicious users that issue inaccurate SCMs to hoard more spectrum than they need or to prevent others from using certain spectrum resources, among many other different types of attacks. SCMs and their exchange between peers could be protected via the use of X.509 certificates or similar methods. However, a detailed discussion on the security mechanisms to secure the exchange of SCMs and their content is beyond the scope of this work.

VI. DYNAMIC SPECTRUM ACCESS WITH SCMS

In this section, we describe a novel SCM-based spectrum deconfliction algorithm and a set of spectrum access methods that leverage the proposed algorithm to deconflict spectrum use in networks with a large number of RF devices. The algorithm leverages CIL and SCM to dynamically determine the transmission parameters of the devices, particularly their central frequency and transmission power levels, in order to achieve *aggregate compatibility*. We further propose two spectrum access methods that use the SD algorithm, namely, (1) **LCS** - Logically Centralized Sequential spectrum access and (2) **LND** - Local Neighborhood based Distributed spectrum access. For sequential spectrum access, we compute the aggregate interference from *all* transmitters at each receiver and monitor it using a centralized registry, whereas

TABLE I: Sequential versus distributed DSA methods.

Features	Sequential (LCS)	Distributed (LND)
Spectrum Assignment	One-by-one	Parallel
Required SCMs	All existing devices	Only assigned peers
Compatibility tests	All devices	Peer devices
Algorithm backoff	Based on priority	Timers and priority
Complexity (in steps)	# of Tx/Rx link-pairs	Chromatic number (lower bound)

for distributed spectrum access, aggregate compatibility is only ensured between the systems that are part of the same neighborhood. Table I provides a brief comparison between the sequential and distributed spectrum access methods as proposed in this work. It is important to highlight that we will focus the description of our spectrum access methods and their performance evaluation (in section VII) on static devices. Our methods can be extended to scenarios where RF devices are mobile but the details of how location information in the SCMs should be refreshed and how aggregate interference values should be updated are left for future work.

Our methods for the deconfliction of spectrum use for networks having a large number of RF devices are based on the targeted use of SCM CT computations between new and existing devices. For simplicity, we assume each network is only composed of a single transmitter-receiver pair that has a designated wireless domain (WD) controller, which interacts with other WD controllers for spectrum access. We propose a novel **SCM-based Spectrum Deconfliction (SD)** mechanism, as described in Algorithm 1, with key terms defined in Table II.

The SD algorithm running at the corresponding WD controller requests the SCMs from all existing (i.e. for sequential) or neighboring (i.e. for distributed) link pairs and initializes the following parameters: compatibility score ($ct_score = 0$), total power margin ($total_{PM} = 0$) and max power margin ($max_{PM} = 0$). The score is increased every time a new Tx or Rx under evaluation is determined to be compatible with a pre-existing RF counterpart (Tx vs. Rx or Rx vs. Tx) and is used to verify the compatibility of the new Tx/Rx link pair with the entire system. In contrast, the power margins are used to assess if small adjustments in the power level of the new transmitter (Tx_n) can make it compatible with the existing receivers.

TABLE II: Key Parameters of Algorithm 1 [2].

Parameter	Description
P_{ij}	Interference power at j^{th} RX from i^{th} Tx
R^*x_{interf}	Interference at each Rx from all existing Txs
$R^*x_{interf_n}$	Interference at each Rx from Tx_n
$total_{TP}$	Total power at Rx_n from all existing Txs
$power_{MT}$	Power margin threshold
max_{PM}	Max power margin from all Rxs
$curr_{PM}$	Current power margin
$P_{allowRx_n}$	Allowable power at Rx_n

A. Spectrum Deconfliction (SD) algorithm

As the WD controller performs CTs between the new receiver Rx_n and existing transmitters ($Tx_1 \dots Tx_{n-1}$), if there is compatibility, the WD controller computes and updates

Algorithm 1: Spectrum Deconfliction (SD) [2] involving aggregate interference, frequency, and power adjustments at WD_n .

```

1 if LCS then
2   Input:  $WD_n$  collects SCMs from all existing devices
3   Initialize:  $R^*x_{interf} = []$ 
4 else if LND then
5   Input:  $WD_n$  collects SCMs from assigned peer devices
6 Output: Deconflicted  $Tx_n$  and  $Rx_n$  pair
7 Initialize:  $power_{MT}$ 
8 Initialize:  $ct\_score = 0$ ,  $total_{TP} = 0$  and  $max_{PM} = 0$ 
9 for  $Rx_n$  perform CT between  $Tx_1..Tx_i..Tx_{n-1}$  do
10    $total_{TP} = total_{TP} + P_{in}$ 
11   if compatible then
12      $ct\_score = ct\_score + 1$ 
13 if  $ct\_score == n - 1$  and  $total_{TP} \leq P_{allowRx_n}$  then
14   Initialize:  $ct\_score = 0$  and  $R^*x_{interf\_n} = []$ 
15    $R^*x_{interf\_n}[n] = total_{TP}$ 
16   for  $Tx_n$  perform CT between  $Rx_1..Rx_j..Rx_{n-1}$  do
17      $R^*x_{interf\_n}[j] = P_{nj}$ 
18      $curr_{PM} = R^*x_{interf}[j] + P_{nj} - P_{allowRx_j}$ 
19      $max_{PM} = \max(max_{PM}, curr_{PM})$ 
20     if compatible and  $curr_{PM} \leq 0$  then
21        $ct\_score = ct\_score + 1$ 
22 if  $ct\_score == n - 1$  then
23    $R^*x_{interf} = R^*x_{interf} + R^*x_{interf\_n}$ 
24 else if  $max_{PM} \leq power_{MT}$  then
25   Adjust  $Tx_n$  reference power based on  $max_{PM}$ 
26   Update  $R^*x_{interf}$  based on  $max_{PM}$  and  $R^*x_{interf\_n}$ 
27   Verify if the new link is reachable or not
28   if link not reachable then
29     Move  $Rx_n$  &  $Tx_n$  frequency and recompute CTs
30 else
31   Move  $Rx_n$  &  $Tx_n$  frequency and recompute CTs
32 else
33   if  $total_{TP} > P_{allowRx_n}$  and  $ct\_score == n - 1$  then
34     Aggregate interference detected
35     Move  $Rx_n$  &  $Tx_n$  frequency and recompute CTs
36 Update  $Rx_n$  &  $Tx_n$  SCM and setup the link

```

the total aggregate interference value ($total_{TP}$) at Rx_n . With the total aggregate interference value, an aggregate evaluation of compatibility is performed. If Rx_n is not compatible with the existing transmitters or if the $total_{TP}$ is greater than the allowable interference power at Rx_n (identified as $P_{allowRx_n}$), the Tx and Rx pair are moved to a different frequency/channel and the CTs start over (more details on the frequency assignment scheme will be discussed later). On the other hand, if Rx_n is compatible with the existing transmitters, the WD controller verifies whether the corresponding new transmitter Tx_n is compatible with the existing receivers ($Rx_1..Rx_{n-1}$). To that end, it computes the interference caused by Tx_n at the existing receivers. If Tx_n is compatible with the receivers already operating in the operational area, the WD controller sets up the new link. Otherwise, the WD controller tries to achieve compatibility by decreasing Tx_n 's power by no more than the value of the power margin threshold ($power_{MT}$). In case the power adjustment fails to maintain reachability (i.e. effective coverage) between Tx_n and Rx_n , the WD controller chooses a different frequency channel for the Tx and Rx pair and the compatibility tests start over. At the end of the algorithm, once aggregate compatibility is achieved, the WD

Algorithm 2: LCS - Sequential spectrum access

```

1 Input: Sequence IDs or priorities of WDs (set  $S$ )
2 Output: Deconflicted Tx/Rx pairs in a system
3 for each  $WD_i$  over the sequence set  $S$  do
4   Obtain SCMs from all existing RF devices
5   Run Spectrum Deconfliction algorithm

```

controller sets up the link and updates the SCMs of Tx_n and Rx_n based on the frequency and power values found to achieve compatibility so that they are ready for future CTs. Note that in the case of distributed spectrum access, aggregate compatibility is only ensured among the devices that are part of the same neighborhood, and there is no centralized variable (such as Rx^*_{interf} and $Rx^*_{interf_n}$) as in the case of sequential spectrum access to update the contribution of the existing and new transmitters to the aggregate interference seen by each receiver.

Regarding frequency assignment: Upon entering the operational area, an RF device attempts operation at a default center frequency f_c and, if compatibility is not achieved, the device moves to new frequencies in Δf increments until a compatible frequency for the operation of the new Tx/Rx pair is found. The objective here is to achieve high spectrum efficiency by minimizing the number of different frequency channels used. Although, other approaches to determining f_c and Δf could be used to speed up frequency assignment operations, they are left for future research.

B. LCS - Sequential spectrum access

In this spectrum access method, WDs allocate spectrum resources sequentially as described in Algorithm 2. The assumption here is that existing devices are given high priority, new devices are added in sequence and their spectrum use request is deconflicted reaching a device configuration that is compatible with the previously existing systems. By considering all the SCMs from all the existing devices, sequential spectrum access ensures global compatibility. At first, the links/WDs are prioritized either based on their sequence IDs or a certain policy, and the deconfliction is performed using Algorithm 1. The interference caused by new and existing transmitters at each of the receivers is computed as each CT is performed and stored centrally (Rx^*_{interf}) at the regional aggregator. This value is then used by the SD algorithm to detect the presence of aggregate interference and to compute the ideal max power margin (line 18 & 19 of Algorithm 1) required to limit the power of the new transmitter.

C. LND - Distributed spectrum access

The previously described spectrum access method, LCS, assigns channels sequentially based on CT results whenever a new pair of RF devices joins the system. This process introduces delays, as each device must wait for prior assignments, creating a bottleneck in large networks. As the number of competing devices grows, the frequency of compatibility tests increases, resulting in substantial overhead and latency, rendering LCS inefficient for large-scale networks.

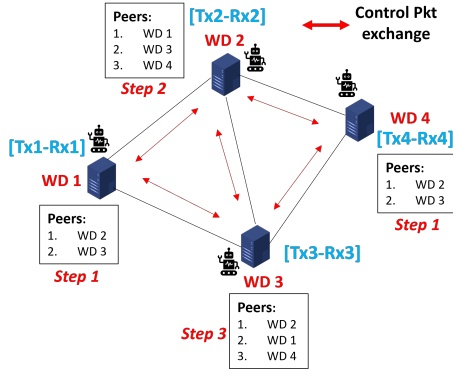


Fig. 4: Illustration of the step-by-step operation of the proposed distributed DSA method (LND).

To this end, we propose LND, a fully distributed spectrum access method that still uses Algorithm 1 (SD) to perform dynamic frequency and power assignments while effectively addressing the challenges posed by LCS (see Table I). In LND, the coordination between multiple independent WDs is still performed using CIL but with additional messaging types required to support distributed operation. In LND, the WD controller (i.e., the node) only maintains local information, unlike LCS, which uses a regional aggregator to record the aggregate interference at each receiver. Additionally, in LND each node computes an interference graph during the bootstrapping phase from the exchange of beacons with its neighbors [42]. To ensure global compatibility, the selection of peers/neighbors is critical. Once the node status and peer list are established, each WD independently assigns spectrum, but only if its assignment does not cause interference to any member of its peer group. In our implementation, the peer group is determined by the proximity (distance based) of each transmitter or receiver to other devices. Developing a more adaptive peer selection mechanism that considers interference awareness requires further study. Looking at the scenario in Fig. 4 as an example, it takes three time steps for LND to resolve conflicts. During the first time step, WD 1 and WD 4, which do not interfere with each other, make their spectrum assignments. At time steps 2 and 3, WD 2 and WD 3, respectively, perform their assignments since they would cause interference if they tried to assign spectrum at the same time as WD 1 and WD 4. The specifics of LND's operations are outlined below.

1) *Node status (or state)*: This determines the current status of the node, hence allowing the node to either wait or perform spectrum assignment. A WD can only be in one of the following states at a given time, and they include:

- *Unassigned*: Spectrum not assigned to the WD
- *Pre-assignment*: WD wants to perform deconfliction
- *Pause*: Temporally pause pre-assignment (if any)
- *Running*: WD currently running SD algorithm
- *Assigned*: WD has completed spectrum assignment

2) *Message exchange types*: In addition to the messages we added to CIL to support LCS, the list of additional coordination messages required to perform distributed spectrum deconfliction in LND include:

- *Node_status_request()*: Message to obtain the current node status from a peer node [NSR]
- *Node_status_response()*: Response message to an NSR. The reply provides the current status of a WD. If the WD is in the assigned state, an SCM is also included in the reply [NSRp]
- *ACK()*: Message to acknowledge peers after the reception of an NSRp, endorsing that the node will be entering the running state next [ACK]
- *Assigned()*: Message to inform peers on the completion of running the SD algorithm [A]
- *Reset()*: Message to inform peer WDs to perform re-assignment if there are any local changes within the domain of the issuer of the message [RESET]

3) *Timer types*: We also introduce the following list of timers to prevent any contention between the nodes and also to avoid infinite waiting in a particular state.

- *NSR Timer*: Timer for NSRp and NSR to be received from another peer to prevent any conflict and infinite waiting in a pre-assignment state
- *ACK Timer*: Timer for an ACK to be received, required to prevent infinite waiting in the pause state if a peer WD fails to perform spectrum assignment
- *Restart Timer*: Timer for the reception of the *Assigned()* message from a peer node performing spectrum assignment to prevent infinite waiting in the pause state after receiving an ACK from that node.

4) *Protocol details*: We describe in detail the state transitions of LND. The assumption here is: after a node performs an assignment, it never changes provided there is no change in spectrum demand by the node and no *Reset()* message has been received.

Algorithm 3: Pre-assignment state transition

```

1 Enter PRE-ASSIGNMENT state
   • If WD needs to perform spectrum assignment
Send NSR and initialize NSRp timer
while Timer expired or all NSRp received do
  if NSR received from a peer node then
    Based on the conflict resolution procedure determine
    priority
    if Low priority then
      Enter PAUSE state, delete timer
Evaluate all the received NSRps
if any NSRp == RUNNING then
  Enter PAUSE state, delete timer
else
  Send ACK to the peer WDs
  Enter RUNNING state

```

a) **Pre-assignment**: The WD enters the pre-assignment state if it needs to determine its spectrum assignment. At first, it sends an NSR to its peers and starts the NSR timer. Next, it waits to receive NSRp (for the corresponding NSR) from all its peers. If any of the peers is in a running state, the WD enters the pause state immediately. Moreover, if within the NSR timer window, another NSR from a peer WD is received, based on the conflict resolution procedures described later, the WD may enter the pause state to avoid further conflict

Algorithm 4: Pause state transition

```

1 Enter PAUSE state
    • If WD in UNASSIGNED and received an NSR from a peer
    • If WD in PRE-ASSIGNMENT, and received an NSR from
      a high priority peer
    • If WD in PRE-ASSIGNMENT, and received NSRp ==
      RUNNING from a peer

Initialize ACK timer
while ACK timer expired or ACK received do
|   Wait

if ACK timer expired then
|   Enter PRE-ASSIGNMENT or UNASSIGNED state
else if ACK received then
|   Initialize Restart timer
|   while Restart timer expired or Assigned received do
|   |   Wait
|   Enter PRE-ASSIGNMENT or UNASSIGNED state

```

in performing spectrum assignment. Finally, if all the NSRps have been received within the timer window and the WD is not in a pause state, it enters the running state to perform spectrum assignment (see Algorithm 3).

b) Pause: The WD enters a pause state when one of the following events occurs (see Algorithm 4):

- WD in an unassigned state and received an NSR from a competing peer
- WD in pre-assignment state and within the NSR timer window, it received an NSR from a higher priority node (priority decided based on conflict resolution procedure)
- WD in the pre-assignment state but one of its peers entered the running state

In the pause state, the WD initializes the ACK timer and waits for an acknowledgment from the peer (that is in the running state) performing a spectrum assignment. If the ACK message is received, it initializes the restart timer and waits for an `Assigned()` message. If the assigned message is received or any of the timers (ACK/Restart) expires, the WD enters the pre-assignment state only if it requires an assignment of spectrum, otherwise it stays unassigned.

c) Running: A WD enters the running state, if it was in the pre-assignment state before and all the NSRps have been received thus allowing it to proceed with its spectrum assignment computation. The WD executes the SD algorithm and finally issues an `Assigned()` message to its peer WDs when it completes the computation (see Algorithm 5).

Algorithm 5: Running state transition

```

1 Enter RUNNING state
   • If WD in PRE-ASSIGNMENT and received NSRp  $\neq$  RUNNING
Run Spectrum deconfliction (Algorithm 1)
if Complete then
    | Send assigned message to the peer WDs
    | Enter ASSIGNED state

```

d) Example illustration: Considering Fig. 5, let's assume WD 1 is in a pre-assignment state intending to determine its spectrum assignment. It requests the node status from its peers WD 2 and WD 3. On receiving this request message,

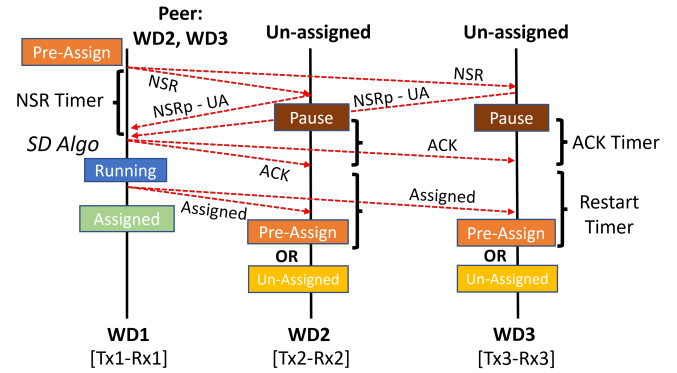


Fig. 5: Message exchanges of the LND protocol.

the peers, currently in an unassigned state will reply with NSRp conveying their current node status (but will not send any SCMs) to WD 1. Next, WD 1 will acknowledge peers of it entering the running state and performing a spectrum assignment determination. Since in this particular example, none of the peers of WD 1 are in a running state and no SCMs have been received, WD 1 will setup the link with default transmission parameters. Also, various timers as described before are initialized to prevent infinite waiting in the current state. Furthermore, if any of the peers are in an assigned state, WD 1 will use its SCM to perform spectrum deconfliction using Algorithm 1. Finally, when complete, WD 1 issues an assigned message to its peers.

e) Conflict resolution: We introduce a procedure to avoid any conflicts between peers performing spectrum assignments at the same time by determining which of the interacting nodes should have precedence. In this work, we select the WD with the lowest ID value (min ID), to resolve any conflict among the peers. Evaluation of the other conflict resolution strategies is left for future work.

5) *Complexity*: LND performs spectrum assignments in steps. In each step, the protocol selects a subset of unassigned WDs that do not interfere with one another (i.e., that are not peers) and allows this subset to perform spectrum assignments simultaneously. Spectrum assignment terminates when all WDs are assigned. A key observation is that WDs can only perform assignment in the same step if they are not peers. From this observation, we draw a connection between LND and vertex coloring which leads to the conclusion that the *minimum number of steps to achieve spectrum assignment de-confliction with LND is given by the chromatic index $\chi(\mathcal{G}) \in \mathbb{N}$ of the underlying peer interference graph \mathcal{G} .*

In vertex coloring, colors are assigned to vertices such that peer vertices have different colors. The chromatic index $\chi(\mathcal{G})$ is the *smallest number of colors* needed in the vertex coloring of a graph \mathcal{G} . Given a vertex coloring of an interference graph, it is always possible to construct an associated sequence of spectrum assignment steps that follows the rules of LND. Specifically, given a vertex coloring with colors $1, 2, \dots, C$, LND can select the subset of WDs with color 1 during step 1, then color 2 during step 2, and so on, until all WDs are assigned in the last step C . Similarly, notice that given an instance of LND, by assigning different colors to the subsets of WDs selected in different steps, it is always possible to




Fig. 6: (a) Tx spectrum mask; and (b) Rx underlay mask [2].

construct a vertex coloring. Naturally, the lower the number of colors in a vertex coloring, the lower the number of steps in the corresponding LND based spectrum assignment. Hence, it follows that the minimum number of steps to run LND in a given interference graph is equal to its chromatic index, i.e., $C \geq \chi(\mathcal{G})$. According to Brooks' theorem [43], for any connected undirected peer interference graph \mathcal{G} with maximum degree $\Delta(\mathcal{G})$, the chromatic index is such that $\chi(\mathcal{G}) \leq \Delta(\mathcal{G}) + 1$. It is easy to see that the maximum number of steps C to run LND is equal to the number of vertices in the graph \mathcal{G} .

VII. PERFORMANCE EVALUATION

To evaluate the performance of our proposed SD algorithm and spectrum access methods, we use a Python-based simulator that assigns Tx/Rx pairs to fixed or random locations in a given operational area and then uses Algorithm 1 to deconflict spectrum use. We start our analysis with a fixed topology scenario and then we progress towards more complex random scenarios that include a significant number of radio frequency (RF) devices. For sequential spectrum access, the simulation uses a discrete event-based framework in which Tx/Rx pairs join the system one after the other. For distributed spectrum access, the simulation uses an agent-based modeling framework, called MESA [44], in which all Tx/Rx pairs join the system simultaneously (at the beginning of the simulation) and are represented as independent agents making autonomous spectrum assignment decisions. To perform CTs, we leverage Octave code from the Spectrum Consumption Model Builder and Analysis Tool (SCMBAT) [45], which interfaces with our Python simulator using *oct2py* [46].

The parameters for our simulations are summarized in Table III. The SCMs used in the simulator are similar to the SCMs from our experimental work as described in [1], except that, we are limiting the transmit power such that the coverage radius for each Tx is around 100 meters. The structure of the transmitter spectrum mask and the receiver underlay mask used is illustrated in Fig. 6. For simplicity, all transmissions used BPSK modulation with a channel bandwidth of 1 MHz. Omnidirectional antennas were used at the transmitter and in the receiver. It is worth mentioning that the operations of our proposed SD algorithm can be generalized to any spectrum use deconfliction scenario as long as the appropriate SCMs are used.

TABLE III: Simulation Parameters [C = indicates parameters used only in the evaluation of complex scenarios in Sec. VII.B].

Parameter	Value
Operational Area	0.8 Km * 0.8 Km [C]
Min Tx-Tx separation	10 m [C]
Tx-Rx separation	<i>Uniform</i> (10m, 100m) [C]
Frequency shift (Δf)	1 MHz
Power margin threshold	3 dB
Number of trials	100 [C]
Python version	3.8
Noise floor	-114 dBm
Modulation	BPSK
Bitrate	0.5M
Machine configuration	Intel i7-4790 (3.60 GHz) Cores: 8, Memory: 15 GB




Fig. 7: Fixed topology setup.

For the fixed topology scenarios, the position of the transmitter and receiver pairs is as illustrated in Fig. 7. For the complex random topology scenarios, the position of the transmitters is chosen uniformly at random in a 0.5 square mile area, subject to the constraint of a minimum distance of 10 meters between any two transmitters. Moreover, the separation between a transmitter and its associated receiver is chosen uniformly at a random interval between 10 and 100 meters. In our simulations, we assume (i) an ideal network/mechanism for exchanging SCMs; and (ii) each WD is either using spectrum obtained from a previous SD operation or is seeking a new spectrum assignment.

A. Fixed topology scenario

In this subsection, we evaluate the spectrum usage of sequential and distributed spectrum access methods by considering a fixed topology composed of five WDs, each having a single Tx-Rx link-pair as shown in Fig. 7.

Table IV provides various performance metrics such as global compatibility error (CE), system throughput, and step count, for our proposed DSA methods. The throughput per link C_i (in bits per second) is computed based on the Shannon capacity using the following equations.

$$NP = kTB \quad (1)$$

$$SINR_i = \frac{Power_{signal_i}}{NP + \sum_{n=0, n \neq i}^{N-1} Interf_{ni}} \quad (2)$$

$$C_i = B * \log_2(1 + SINR_i) \quad (3)$$

TABLE IV: Performance evaluation for fixed topology.

Spectrum access methods	Global CE	Throughput (Mbps)	Step count
Sequential (all devices)	0%	36.88	5
Distributed (p=100m)	80%	33.87	2
Distributed (p=200m)	40%	36.98	4
Distributed (p=100m, 0.1 dB Adj)	40%	33.85	2
Distributed (p=200m, 0.1 dB Adj)	0%	36.98	4

where NP is the total noise power at the receiver (in watts), $k = 1.38 \times 10^{-23}$ is Boltzmann's constant, $T = 290$ is the temperature (in Kelvin), B is the channel bandwidth (in hertz), $SINR_i$ is the Signal to Interference plus Noise Ratio of link i , $Power_{signal_i}$ is the average signal power of link i (in watts), $Interf_{ni}$ is the interference at the Rx of link i from Tx n (in watts), and N is the total number of links.

We observe the CE to be 0 % for sequential spectrum access since all devices are considered for deconfliction, but it requires a higher number of step counts to complete deconfliction. The step count improves with distributed spectrum access, as only peer devices are considered and the devices belonging to different peer groups (neighborhoods) will be performing deconfliction in parallel, but with a high penalty on the global CE.

The CE can be improved further by reducing the Tx power by a very small adjustment value of 0.1dB on top of the power margin (PM) value for the devices performing power adjustments. The main reason behind doing this is to have a small separation margin of error added to the power margin to account for any potential interference which gets missed out by only considering devices in a peer-to-peer basis, unlike in sequential spectrum access where a centralized variable is used to monitor all the interference contributions at each receiver.

In our simulation results, the distributed spectrum access procedure having peer group (p) defined by distance $p=200m$ and 0.1 dB power adjustment outperforms other distributed schemes, as the compatibility error is found to be 0% and the throughput is close to that of the sequential spectrum access algorithm. An interesting observation is that the throughput is slightly higher for the distributed approach ($p=200m$) compared to the sequential approach, primarily due to transmission power adjustments. In the sequential approach, all devices are considered for CT computation, limiting the transmission power of newly added devices to prevent interference. In the distributed approach, not all devices are considered, and the power is not optimally limited, resulting in high throughput but also a high CE of up to 40%. Based on the sensitivity of a receiver and on how much interference it can tolerate, an appropriate peer group (p value) should be selected to achieve low CE and efficient computational and spectrum use performance.

B. Complex topology scenario

We will compare sequential and distributed spectrum access methods by considering a complex topology also using the

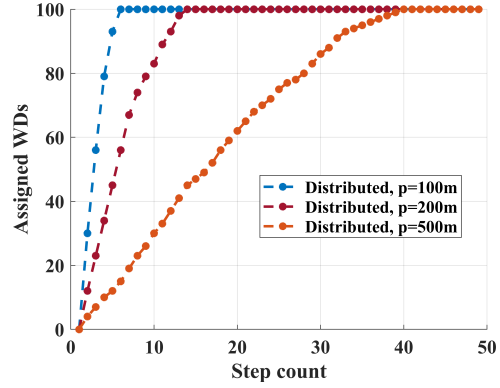
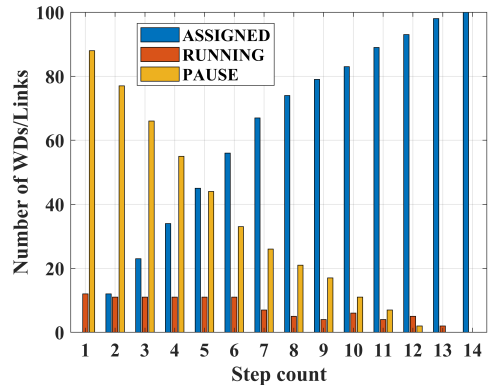


Fig. 8: Convergence of distributed spectrum access for different peer groups (100 links/WDs & 1 trial).

Fig. 9: Histogram depicting WDs in respective states at each time step (100 links/WDs, single trial, $p=200m$).

same simulation parameters listed in Table III. The evaluation is based on convergence, average step count/computation time, algorithm efficiency, and channel usage. Based on the results provided in section VII-A for a simple topology, a 0.1 dB power adjustment is applied for the analysis of the distributed spectrum access method.

1) *Convergence*: We evaluate the convergence of the LND distributed protocol. From Fig. 8, for a network composed of 100 pre-existing links and a single trial run, we observe that the convergence highly depends on the selection of the p value. Scenarios where WDs have a large number of peers (i.e. the value of p is high), require a higher number of steps to perform assignment as more devices are considered for CT evaluations. As an example, a system of WDs with $p=100m$ requires 5 time steps compared to 13 steps required to perform deconfliction with $p=200m$. Furthermore, Fig. 9 depicts the operation of LND at each time step for 100 WDs and $p=200m$. Since all the WDs have the intent to perform spectrum assignment at the same time, we observe a large number of them unable to perform due to conflicts with peer nodes and hence they enter the *pause* state. This situation improves with time, and the WDs enter the *running* state, and once assigned, the WD changes its state to *assigned*. On average, around 8 WDs perform assignments simultaneously at every time step, demonstrating the parallelism achieved by our protocol. Overall, the step count for LND to achieve spectrum deconfliction is observed to be similar to the chromatic

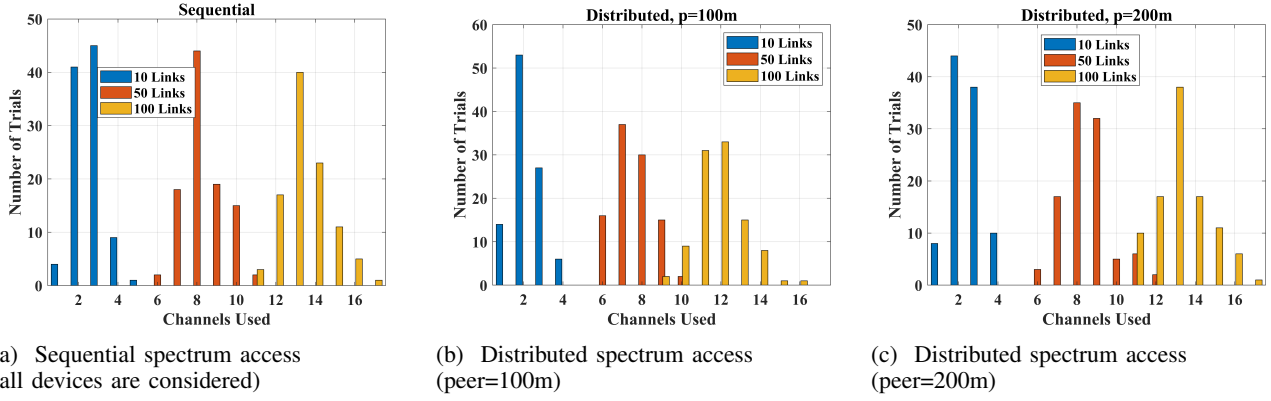


Fig. 10: Channel allocation of sequential and distributed spectrum access schemes for networks of different sizes deployed in a 0.5 sq mile area (averaged over 100 trials).

number of the peer interference graph when evaluated against standard graph coloring algorithms.

2) *Spectrum usage*: We evaluate the spectrum usage of the sequential and distributed spectrum access methods based on the number of channels used. Considering sequential spectrum access, as shown in Fig. 10a, for 100 trials, the *maximum* number of channels required to configure 10 links is 5 and for 100 links, it is 17. The *mode* of the number of channels required for deconflicting 10, 50, and 100 links are 3, 8, and 13, respectively. With distributed spectrum access, the channel usage depends on the selection of the peer group distance parameter p (see Fig. 10b and Fig. 10c). With $p=100m$, the *mode* for the number of channels required to deconflict 10, 50, and 100 links are 2, 7, and 12, respectively. However, the global compatibility error (as shown in Table V) is highest for $p=100m$, implying we need to select a larger p to reduce it. We also observe that the number of channels used for distributed spectrum access with $p \geq 200m$ is similar to sequential spectrum access (where all devices are taken into account for CT evaluations). This suggests that a peer group distance of 200m or more, which is twice the considered transmitter coverage distance, is appropriate for selecting the set of peers that can effectively capture the total aggregate interference from other devices that may influence channel selection.

For a better understanding of the performance behavior for sequential spectrum access, we generated a modified version of LCS that only adjusts the frequency of operation of a device to deconflict spectrum use without using power adjustments as a variable to achieve deconfliction. We call this the **baseline** access method, which resembles sensing-based spectrum assignment [47], [48], where devices first sense a channel before use and select a different channel if the first one is occupied.

In Fig. 11, we compare the channel usage of the sequential and baseline procedures considering a network of 100 links. When looking at the *mode/max* values, a higher number of channels is required for baseline when compared to our proposed scheme since the baseline algorithm just operates on assigning frequencies without adapting Tx power levels. Furthermore, sequential spectrum access without power margin threshold (PMT) when compared against the baseline algorithm achieves better spectral efficiency since it will achieve a fewer number of aggregate interference situations [2].

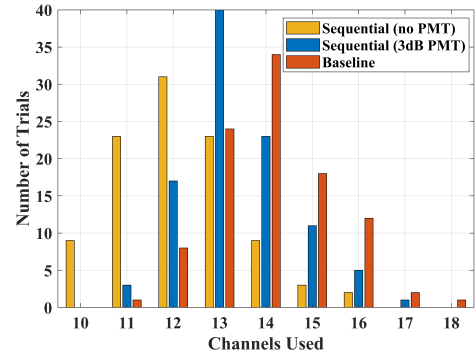


Fig. 11: Number of channels used - Baseline (no power adj.) vs. Sequential for 100 links & 100 trials.

3) *Algorithm performance*: We measure several performance metrics such as achieved global compatibility error, average system throughput, average channel usage, and average computation time to validate and compare the efficiency of our proposed algorithms. These metrics are averaged over 30 trials/runs for networks of different sizes and presented in Table V. Next, we briefly describe our results on the compatibility error and system throughput.

a) **Global compatibility error**: Considering distributed spectrum access, we observe that inappropriate selection of peer group distance leads to high compatibility errors and we compute this error with respect to the exhaustive (sequential) case where the evaluation of compatibility is done against all existing devices not just with those covered by the peer group distance. We see that the compatibility error decreases with the increase in p since more devices will be considered for CT computations, thus allowing us to more accurately capture the impact of aggregate interference at each receiver. Also, a high compatibility error is observed for $p=100m$, and the error drops significantly beyond that value. With $p=100m$, the CTs are considering only interference from devices that are within the typical coverage area of each Tx. At greater values of p , devices that are outside the coverage area that can generate interference to a particular receiver (i.e. hidden nodes) are also taken into account. For sequential spectrum access (and also for baseline), the compatibility error is 0% since all devices are considered for CTs. These results indicate that based on the sensitivity of the new receiver, an appropriate interference

TABLE V: Global compatibility error (CE%), system throughput (TPT), average channel usage (CH), and average computation time per link (in seconds) for different spectrum access methods, averaged over 30 random trials in networks of different sizes.

Methods	20 Links				50 Links				100 Links			
Performance Metrics	CE (%)	TPT (Mbps)	CH (Avg.)	Time (Sec)	CE (%)	TPT (Mbps)	CH (Avg.)	Time (Sec)	CE (%)	TPT (Mbps)	CH (Avg.)	Time (Sec)
Distributed, $p=100m$	12.16	91.93	3.9	0.01	17.86	187.87	7.66	0.02	22.9	344.06	12.26	0.04
Distributed, $p=200m$	0	96.23	4.43	0.04	0	196.38	8.3	0.08	0	372.46	13.4	0.17
Distributed, $p=500m$	0	92.96	4.46	0.15	0	198.81	8.33	0.47	0	363.25	13.56	1.20
Sequential (no PMT)	0	88.97	4.4	0.35	0	182.16	7.9	1.15	0	318.71	12.6	3.03
Sequential	0	92.76	4.26	0.37	0	204.64	8.6	1.26	0	362.76	13.66	3.25
Baseline (only freq.)	0	94.77	4.23	0.38	0	209.43	8.73	1.30	0	384.03	14.1	3.51




Fig. 12: Number of steps required by sequential and distributed spectrum access methods to deconflict spectrum use.

graph based on the peer group distance value p should be determined to reduce the number of CT computations and minimize the probability of error.

b) System throughput: We compute the Shannon capacity per link considering a 1 MHz channel bandwidth using the equations as described before (Eq. 2 to Eq. 3) and then aggregated across all the links. The throughput is directly proportional to the channel usage and we observe a high throughput for baseline as it operates only in the frequency domain with high Tx power. Considering sequential spectrum access without PMT, for 100 links the throughput is found to be 318.71 Mbps. It's worth noting that if a CT indicates that the Tx power of a transmitter, if reduced by an amount that does not exceed the value indicated by the PMT, will lead to compatibility with all previous receivers, the adjustment is applied, otherwise, the transmitter will need to move to a different frequency/channel. Without using a PMT value (no PMT), the Tx power of the new transmitter is reduced to an ideal setting where it achieves compatibility with existing receivers and can maintain a link with its intended receiver. Consequently, the Tx power for the device in this case is relatively low and the throughput achieved is lower when compared to other schemes where a 3dB PMT value is considered. Moreover, distributed spectrum access achieves a throughput performance comparable to that of the sequential and baseline access methods.

4) *Step count and computation time:* We evaluate the step count (# of time-steps) and the time required to deconflict spectrum use, i.e., the time to determine compatible parameters for sequential and distributed spectrum access in an operational area. The computation time does not include: (i) pauses




Fig. 13: Average time per link required by sequential and distributed methods to deconflict spectrum use (in seconds).

between the appearance of new link pair instances; (ii) time to turn on the RF devices, which depends on the radios being used; (iii) SCM transmission times, which usually take a few milliseconds.

Fig. 12 and Fig. 13 provide the step count and computation time respectively, averaged over 100 trials for networks with different numbers of Tx/Rx pairs. We observe that distributed spectrum access requires a lower number of steps as compared to sequential spectrum access since multiple WDs will be performing spectrum assignments simultaneously among their non-interfering peers. We also observe that to configure a scenario composed of 20 pre-existing links, sequential requires (on average) 0.37 seconds as compared to 3.25 seconds required for a network of 100 pre-existing link pairs. Distributed on the other hand requires lower computation time to perform deconfliction, a gain corresponding to 19x is observed with 100 link-pairs and $p=200m$ as compared to sequential. The increase in computation time with the increase in the total number of Tx/Rx pairs is expected, and it is mainly due to the larger number of CTs that need to be performed. The computation time can be significantly reduced by using more powerful CPUs and selecting appropriate peer groups for the distributed access method.

VIII. DSA EXPERIMENTATION ON COSMOS

To demonstrate the feasibility of our SCM-based DSA methods, we developed a proof-of-concept experiment in the ORBIT grid [22] which is part of the NSF PAWR COSMOS [23], [24] wireless testbed. In our experiment, we have three co-located wireless networks, each containing one




Fig. 14: Map of the ORBIT grid [22] that highlights (using black arrows) the position of the three transmitter-receiver pairs of SDRs utilized during the experiment [1].

transmitter-receiver pair, positioned as illustrated in Fig. 14. Transmitters and receivers are implemented using NI USRP X310 and B210 SDRs. All transmitters use the same spectrum mask and attempt transmissions with a 1 MHz wide channel centered around 2000 MHz. WD controllers use CIL message exchanges via an out of band Ethernet interface to relay SCMs and CT reports. We start the experiment with no networks occupying the spectrum. At time T_1 , WD 1 joins the experiment and selects the center frequency of 2000 MHz. After that, at time T_2 , WD 2 joins the experiment. The Regional Aggregator determines that the networks are in close proximity to one another, and notifies WD 1 that a peer is now present. WD 1 and 2 establish a peering relationship and immediately exchange their SCMs and perform compatibility computations. WD 2 determines that it can't use its initial center frequency of 2000 MHz, and selects an alternate center frequency of 1999 MHz. Similarly, at time T_3 , WD 3 joins the experiment and, with the assistance of the Regional Aggregator, WD 3 exchanges SCMs with WD 1 and 2, and performs compatibility tests after which it determines that the center frequency to use is 2001 MHz. Fig. 15 shows the fosphor visualization of the frequency selection of the three networks that joined the experiment at times T_1 , T_2 , and T_3 . With this experiment, we show that each WD was able to autonomously select a different channel in order to deconflict the spectrum use and avoid harmful interference. For more information, please refer to [1] and/or the video in [49].




Fig. 15: Fosphor visualization showing the spectrum occupancy as the 1st, 2nd, and 3rd networks start to operate [1].

IX. FUTURE WORK

SCMs provide great flexibility and detail for expressing the characteristics and boundaries of spectrum use of RF systems, this facilitates the design of many DSA mechanisms and for the automation of spectrum management procedures. In particular, they can be used to design new spectrum access mechanisms in licensed and unlicensed spectrum that avoid static RF operational limits providing instead a platform for interested parties (e.g. telecom operators) to negotiate limits under the supervision of a regulator (see [50]). In future work we will analyze how SCMs can be used to model WiFi operations in the 6 GHz band to evaluate aggregate interference effects in outdoor and indoor scenarios and determine if regulations for these operations can have more flexible operational limits.

SCMs can capture directional effects/parameters that impact RF operation and interference, we have shown some of these capabilities in [51]. In our future work, we will also explore networks that incorporate mmWave nodes with phased array antennas and develop improved SCM-based DSA algorithms that consider antenna directionality. We will evaluate their performance in terms of spectrum utilization, power usage, data transfer rate, and convergence time, using both a simulator and RF systems the COSMOS testbed [23], [24]. Spectrum use deconfliction algorithms involving mmWave/directional cases, heterogeneous systems, and mobility will also be addressed in the future.

Finally, we will also conduct research work on building AI/ML based spectrum use deconfliction methods that are trained with data from SCM-based spectrum use cases. The training data will be generated from experiments to be carried out in the COSMOS testbed and synthetic scenarios generated with our simulators. We expect, that the declarative nature of SCMs will provide some tractability on the way spectrum use deconfliction is achieved specially in large-scale scenarios with devices that have heterogenous spectrum needs and operational characteristics.

X. CONCLUSION

In this work, we introduced and evaluated new collaborative DSA methods for spectrum sharing that leveraged SCMs, which provide a standardized way for networks to declare their intended spectrum use and/or their interference protection needs. We modified the CIL interaction language developed by DARPA to enable the exchange of SCM messages among wireless networks. Then, we developed a Spectrum Deconfliction (SD) algorithm that dynamically configures frequency and power for wireless links aiming to minimize aggregate interference, thereby resolving spectrum use conflicts. Further, we developed sequential and distributed DSA methods based on the SD algorithm that assigns spectrum in large-scale networks. To evaluate the performance of our DSA methods, we developed a custom-made simulation platform and built a proof-of-concept implementation in the NSF PAWR COSMOS wireless testbed.

The simulation results demonstrate that in the context of sequential and distributed spectrum access, the selection of

appropriate power margin thresholds plays a crucial role in achieving a reduction in the overall utilization of the spectrum. Importantly, the results also demonstrate that the distributed DSA method can achieve up to a 7x improvement in spectrum deconfliction speed when compared to the sequential approach in large scale scenarios. Finally, the computational performance of our methods is reasonable and can be improved based on the computational hardware configuration, which should facilitate their incorporation into real-world platforms.

REFERENCES

- [1] D. Stojadinovic, P. Netalkar *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*, 2022.
- [2] P. Netalkar, A. Zahabee *et al.*, “Large-scale dynamic spectrum access with IEEE 1900.5. 2 spectrum consumption models,” in *2023 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2023.
- [3] O. Ileri and N. B. Mandayam, “Dynamic spectrum access models: toward an engineering perspective in the spectrum debate,” *IEEE Communications Magazine*, vol. 46, no. 1, 2008.
- [4] “National Spectrum Strategy,” online: https://www.ntia.gov/sites/default/files/publications/national_spectrum_strategy_final.pdf, 2023.
- [5] “National Spectrum Strategy Implementation Plan,” online: <https://www.ntia.gov/sites/default/files/publications/national-spectrum-strategy-implementation-plan.pdf>, 2024.
- [6] W. Lehr, R. Berry *et al.*, “Automating spectrum sharing from the ground up,” in *Proc. of Research Conference on Communications, Information, and Internet Policy (TPRC)*, 2024.
- [7] Z. Guan and T. Melodia, “CU-LTE: Spectrally-efficient and fair coexistence between LTE and Wi-Fi in unlicensed bands,” in *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*. IEEE, 2016.
- [8] S. Sagari, S. Baysting *et al.*, “Coordinated dynamic spectrum management of LTE-U and Wi-Fi networks,” in *2015 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 2015.
- [9] P. Karimi, W. Lehr *et al.*, “SMAP: A scalable and distributed architecture for dynamic spectrum management,” in *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 2018.
- [10] G. Naik, J.-M. Park, J. Ashdown, and W. Lehr, “Next generation Wi-Fi and 5G NR-U in the 6 GHz bands: Opportunities and challenges,” *IEEE Access*, vol. 8, pp. 153 027–153 056, 2020.
- [11] “Making 5G NR a reality, in white paper,” <https://www.qualcomm.com/media/documents/files/whitepaper-making-5g-nr-a-reality.pdf>, 2016.
- [12] “Unlocking spectrum innovations for 5G Advanced and 6G to support future wireless growth,” online: <https://www.qualcomm.com/news/onq/2023/05/unlocking-spectrum-innovations-for-5g-advanced-and-6g-to-support-future-wireless-growth>, 2023.
- [13] “Co-creating a cyber-physical world,” <https://www.ericsson.com/4a32a8/assets/local/reports-papers/white-papers/2024/co-creating-a-cyber-physical-world.pdf>, 2024.
- [14] “Spectrum access mechanisms,” online: https://nextgalliance.org/white_papers/spectrum-access-mechanisms/, 2024.
- [15] C. Walker, S. Di Pippo, M. Aubé *et al.*, “Dark & quiet skies II,” 2021.
- [16] N. R. Council *et al.*, *Spectrum management for science in the 21st century*. National Academies Press, 2010.
- [17] D. Stojadinovic, F. A. De Figueiredo *et al.*, “SC2 CIL: Evaluating the spectrum voxel announcement benefits,” in *IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2019.
- [18] C. W. Kim, J. Ryoo *et al.*, “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.
- [19] M. M. Sohul, M. Yao *et al.*, “Spectrum access system for the citizen broadband radio service,” *IEEE Communications Magazine*, 2015.
- [20] J. A. Stine and C. E. C. Bastidas, “Enabling spectrum sharing via spectrum consumption models,” *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 4, 2015.
- [21] B. Benmammam, A. Amraoui, and F. Krief, “A survey on dynamic spectrum access techniques in cognitive radio networks,” *International Journal of Communication Networks and Information Security (IJCNIS)*, vol. 5, no. 2, 2013.
- [22] D. Raychaudhuri, I. Seskar *et al.*, “Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols,” in *Proc. of IEEE WCNC*, 2005.
- [23] D. Raychaudhuri, I. Seskar, G. Zussman *et al.*, “Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless,” in *Proc. of ACM MobiCom*, 2020.
- [24] “Cloud enhanced open software defined mobile wireless testbed for city-scale deployment (COSMOS),” <https://cosmos-lab.org/>, 2021.
- [25] X. Jing and D. Raychaudhuri, “Spectrum co-existence of IEEE 802.11b and 802.16a networks using reactive and proactive etiquette policies,” *Mobile Networks and Applications*, vol. 11, no. 4, 2006.
- [26] M. Parvini, A. H. Zarif, A. Nouruzi, N. Mokari, M. R. Javan, B. Abbasi, A. Ghasemi, and H. Yanikomeroğlu, “A comprehensive survey of spectrum sharing schemes from a standardization and implementation perspective,” *arXiv preprint arXiv:2203.11125*, 2022.
- [27] Z. Zhou, Y. Jia, F. Chen, K. Tsang, G. Liu, and Z. Han, “Unlicensed spectrum sharing: From coexistence to convergence,” *IEEE Wireless Communications*, vol. 24, no. 5, pp. 94–101, 2017.
- [28] W. S. H. M. W. Ahmad, N. A. M. Radzi *et al.*, “5G technology: Towards dynamic spectrum sharing using cognitive radio networks,” *IEEE Access*, 2020.
- [29] F. Paisana, Z. Khan *et al.*, “Exploring radio environment map architectures for spectrum sharing in the radar bands,” in *2016 23rd International Conference on Telecommunications (ICT)*. IEEE, 2016.
- [30] H. B. Yilmaz, T. Tugcu *et al.*, “Radio environment map as enabler for practical cognitive radio networks,” *IEEE Communications Magazine*, 2013.
- [31] H. Qin and Y. Cui, “Spectrum coordination protocol for reconfiguration management in cognitive radio network,” in *Proc. of IEEE YCICT*, 2009.
- [32] L. Zhang, M. Xiao, G. Wu, M. Alam, Y.-C. Liang, and S. Li, “A survey of advanced techniques for spectrum sharing in 5g networks,” *IEEE Wireless Communications*, vol. 24, no. 5, pp. 44–51, 2017.
- [33] “Ai-native and sustainable communication,” <https://bit.ly/4htcYIC>, 2025.
- [34] X. Liu, C. Sun, W. Yu, and M. Zhou, “Reinforcement-learning-based dynamic spectrum access for software-defined cognitive industrial internet of things,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 6, 2021.
- [35] S. Zhang, K.-Y. Lam, B. Shen, L. Wang, and F. Li, “Dynamic spectrum access for internet-of-things with hierarchical federated deep reinforcement learning,” *Ad Hoc Networks*, vol. 149, 2023.
- [36] S. Zhang, Z. Ni, L. Kuang, C. Jiang, and X. Zhao, “Traffic priority-aware multi-user distributed dynamic spectrum access: A multi-agent deep rl approach,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 6, 2023.
- [37] F. Li, B. Shen, J. Guo, K.-Y. Lam, G. Wei, and L. Wang, “Dynamic spectrum access for internet-of-things based on federated deep reinforcement learning,” *IEEE Transactions on Vehicular Technology*, vol. 71, no. 7, 2022.
- [38] Y. Song, H.-H. Chang, and L. Liu, “Federated dynamic spectrum access through multi-agent deep reinforcement learning,” in *GLOBECOM 2022-2022 IEEE Global Communications Conference*. IEEE, 2022, pp. 3466–3471.
- [39] N. Rastegardoost and B. Jabbari, “A machine learning algorithm for unlicensed lte and wifi spectrum sharing,” in *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*. IEEE, 2018, pp. 1–6.
- [40] IEEE, “IEEE 1900.5.2-2017 - IEEE standard for method for modeling spectrum consumption,” 2017.
- [41] D. Altner, S. Schmitz *et al.*, “Determining aggregate compatibility using dynamic spectrum consumption models via global optimization,” in *Optimization Online*, 2016.
- [42] Y. Cai, K. A. Hua, and A. Phillips, “Leveraging 1-hop neighborhood knowledge for efficient flooding in wireless ad hoc networks,” in *PCCC 2005. 24th IEEE International Performance, Computing, and Communications Conference*, 2005. IEEE, 2005.
- [43] R. L. Brooks, “On colouring the nodes of a network,” *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 37, no. 2, p. 194–197, 1941.
- [44] “MESA: Agent-based modeling in python 3+,” <https://mesa.readthedocs.io/en/stable/>, 2023.
- [45] “Spectrum consumption model builder and analysis tool (SCMBAT),” <https://github.com/cacicedo/SCMBAT>, 2021.
- [46] S. Silvester *et al.*, “Oct2py,” URL <https://pypi.org/project/oct2py>.

- [47] X. Kang, Y.-C. Liang, H. K. Garg, and L. Zhang, "Sensing-based spectrum sharing in cognitive radio networks," *IEEE Transactions on Vehicular Technology*, vol. 58, no. 8, 2009.
- [48] N. Muchandi and R. Khanai, "Cognitive radio spectrum sensing: A survey," in *2016 International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT)*. IEEE, 2016.
- [49] "IRIS cooperative dynamic spectrum access demo," <https://www.youtube.com/watch?v=7Ra1n5cV8sw&t=800s>, 2021.
- [50] W. Lehr, R. Berry, I. Kadota, C. Caicedo Bastidas, K. Mu, Z. Xie, and I. Tamim, "Automating spectrum sharing from the ground up: A bottom-up approach and research agenda," *Telecommunications Policy*, 2025.
- [51] I. Tamim, C. Caicedo Bastidas, I. Kadota, and G. Zussman, "Cooperative dynamic spectrum access for large-scale networks using directional antennas," in *WiOpt'25 - RAWNET (Workshop on Resource Allocation and Cooperation in Wireless Networks)*, 2025.



Prasad Netalkar received the Ph.D. degree in electrical and computer engineering from Rutgers University, The State University of New Jersey, under the guidance of Prof. Dipankar Raychaudhuri in May 2023. Prior to this, he earned his Master's degree in electrical engineering from NYU Tandon School of Engineering and a Bachelor of Engineering degree in electronics and communication engineering from PES Institute of Technology, India. Currently, he is working as a Senior Modem Technologies Engineer at Qualcomm, San Diego. His research

aims to advance next-generation network performance by leveraging cross-layer optimization techniques. Specifically, his work leverages edge cloud infrastructure and key concepts from the "named-object" information-centric network (ICN) architecture to enhance data transfer during mobility, enable efficient task offloading, and support collision-free dynamic spectrum access. Throughout his academic journey, he collaborated with industry mentors from Qualcomm, Samsung, and InterDigital, where he gained valuable insights into various aspects of network designs.



Carlos E. Caicedo Bastidas (Senior Member, IEEE) is an Associate Professor and Director of the Center for Emerging Network Technologies (CENT) at the School of Information Studies at Syracuse University, NY, USA. He has a Ph.D. in Information Science from the University of Pittsburgh and M.Sc. degrees in Electrical Engineering from the University of Texas at Austin and from the Universidad de los Andes, Colombia. He is also the vice-chair of the IEEE DySPAN-SC (Dynamic Spectrum Access and Networks Standardization Committee) 1900.5

Working group on Policy Language and Policy Architectures for Managing Cognitive Radio for Dynamic Spectrum Access Applications. Dr. Caicedo's research interests and areas of expertise are in the areas of Dynamic Spectrum Access, Radio Frequency spectrum management, Internet of Things (IoT), Network and Information System design, and agent-based modeling.



Igor Kadota (Senior Member, IEEE) received the B.S. degree in electronic engineering and the S.M. degree in telecommunications from the Aeronautics Institute of Technology (ITA), Brazil, in 2010 and 2013, the S.M. degree and the Ph.D. degree in communication networks from the Massachusetts Institute of Technology (MIT) in 2016 and 2020. He was a Postdoctoral Research Scientist with Columbia University from 2020 to 2023. He is an Assistant Professor with the Department of Electrical and Computer Engineering, Northwestern University.

His research is on modeling, analysis, optimization, and implementation of emerging communication networks, with the emphasis on wireless networks and time-sensitive traffic. He was a recipient of several research, teaching, and mentoring awards, including the Best Paper Award at IEEE INFOCOM 2018, the Best Paper Award Finalist at ACM MobiHoc 2019, the Best Student Paper Award at WiOpt 2024 and at WiOpt 2025, the MIT School of Engineering Graduate Student Extraordinary Teaching and Mentoring Award of 2020, and the 2019–2020 Thomas G. Stockham Jr. Fellowship.



Gil Zussman (Fellow, IEEE) received the Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. He has been with Columbia University since 2007, where he is the Kenneth Brayer Professor of Electrical Engineering and Chair of the Electrical Engineering Department. His research interests are in the area of networking, and in particular in the areas of wireless, mobile, and resilient networks. He is a co-recipient of 7 paper awards including the ACM SIGMETRICS'06 Best Paper Award, the 2011 IEEE Communications Society Award for Advances in Communication, and the ACM CoNEXT'16 Best Paper Award. He received the Fulbright Fellowship, the DTRA Young Investigator Award, two Marie Curie International Fellowships, and the NSF CAREER Award.



Ivan Seskar (Senior Member, IEEE) is currently the Chief Technologist of the WINLAB, Rutgers University, for experimental systems and prototyping projects. He is also the Program Director of COSMOS Project responsible for the New York City NSF PAWR Deployment, the PI of the NSF GENI Wireless Project, which resulted in campus deployments of LTE/WiMAX base stations at several U.S. universities, and the PI for the NSF CloudLab deployment at Rutgers. He has been the Co-PI and the Project Manager for all three phases of the NSF-supported ORBIT mid-scale testbed project at WINLAB, successfully leading technology development and operations since the testbed was released as a community resource, in 2005, for which the team received the 2008 NSF Alexander Schwarzkopf Prize for Technological Innovation. He is a Co-Chair of the IEEE Future Networks Testbed Working Group, a member of ACM, and the Co-Founder and the CTO of Upside Wireless Inc.



Dipankar Raychaudhuri (Life Fellow, IEEE) is Distinguished Professor, Electrical & Computer Engineering and Director-Emeritus, WINLAB (Wireless Information Network Lab) at Rutgers University. As WINLAB's Director from 2001–2024, he was responsible for an internationally recognized industry-university research center specializing in wireless technology. He has served as Principal Investigator for several large multi-institutional U.S. National Science Foundation grants including the "ORBIT" wireless testbed and the "MobilityFirst"

future internet architecture (FIA) project, and the "COSMOS" project under the NSF Platforms for Advanced Wireless Research (PAWR) program. Dr. Raychaudhuri has previously held corporate R&D positions including: Chief Scientist, Iospan Wireless (2000–01), Assistant General Manager & Department Head, NEC Laboratories (1993–99) and Head, Broadband Communications, Sarnoff Corp (1990–92). He obtained the B.Tech (Hons) from IIT Kharagpur in 1976 and the M.S. and Ph.D. degrees from SUNY, Stony Brook in 1978, 79. He is a Fellow of the IEEE and the recipient of several professional awards including the Rutgers School of Engineering Faculty of the Year Award (2017), IEEE Donald J. Fink Award (2014), Indian Institute of Technology - Kharagpur, Distinguished Alumni Award (2012), and the Schwarzkopf Prize for Technological Innovation (2008).