Reactive Spectrum Sharing with Radio Dynamic Zones

Aarushi Sarbhai*, Frost Mitchell*, Sneha Kasera*, Aditya Bhaskara*, Jacobus Van der Merwe*, Neal Patwari†
*Kahlert School of Computing, University of Utah, Salt Lake City, USA
†McKelvey School of Engineering, Washington University in St. Louis, St. Louis, USA
*aarushis,*fmitch,*kasera,*bhaskara,*kobus@cs.utah.edu,†npatwari@wustl.edu

***Abstract*—Radio Dynamic Zones (RDZs) are emerging as a means to enable dynamic spectrum sharing. Passive services like remote satellite sensing, radio astronomy, and earth sciences are vital candidates to share spectrum with RDZs. RDZs must protect sensitive receivers outside the zone from undesirable interference from secondary spectrum use inside the zone. We develop a scalable, novel reactive framework to minimize interference at the sensitive receivers while maximizing spectrum utilization within the zone. We utilize interference supervision at the sensitive receiver site to manage allocation decisions. We present a complete, viable, and deployable spectrum management solution and evaluate its operation both in over-the-air experiments using the POWDER wireless testbed and by simulating a real spectrum-sharing scenario with a sensitive receiver and varying sizes of RDZs at long distances. By incorporating location information, propagation characteristics, and an exponentially weighted moving average of the number of RDZ users sharing the band we achieve lower interference periods at the sensitive receiver and high spectrum utilization in the RDZ.**

***Index Terms*—Radio Dynamic Zones, Sensitive Passive User, Zone Management System**

I. Introduction

There is a growing effort to reclaim some of the statically allocated sub-6 GHz radio frequency spectrum that is significantly underutilized [1], [2]. The reclaimed spectrum is being opened to secondary use through dynamic spectrum access [3]. Radio Dynamic Zones (RDZs) enable dynamic spectrum sharing with consumer broadband, special transmitters such as directed energy systems and high-power microwave transmitters [4], and other experimental radio systems [5]. RDZs are envisioned as a space for a diverse set of services and applications, ranging from long-term spectrum usage (spanning years) to opportunistic use over shorter duration (minutes to months). Users in the RDZ must share spectrum with incumbents outside the zone, especially highly *sensitive* receivers used in passive services like remote satellite sensing, radio astronomy, and earth sciences. These services have frequency bands allocated for exclusive access. However, there are substantial periods when these bands are not fully utilized, offering an opportunity for secondary use. RDZs must ensure that the sensitive receivers outside the zones are protected from radio interference from secondary spectrum use inside the zone. Certain transmitters in the RDZ may be able to share the spectrum without causing harmful interference to the sensitive receivers. Identifying the optimal combination of transmitters that can simultaneously operate to maximize spectrum utilization within the zone and ensure the sensitive receiver is protected is NP-hard [6]. Furthermore, determining the precise solution also requires the exact transmission behavior of the RDZ users and an accurate model of the dynamically changing environment.

This research is supported in part by the NSF grants #2232463, #2232464, and #1827940, and the PAWR Project Office grant #10046930.

Existing solutions rely on computationally expensive scheduling and conservative interference estimates, which are not suitable for a highly dynamic RDZ. Furthermore, the sensitive receiver can report its operational parameters seconds before it starts operating and allows only a brief acceptable interference period. An adaptive approach is essential to react quickly to interference and strategically select transmitters in the RDZ to maximize spectrum reuse.

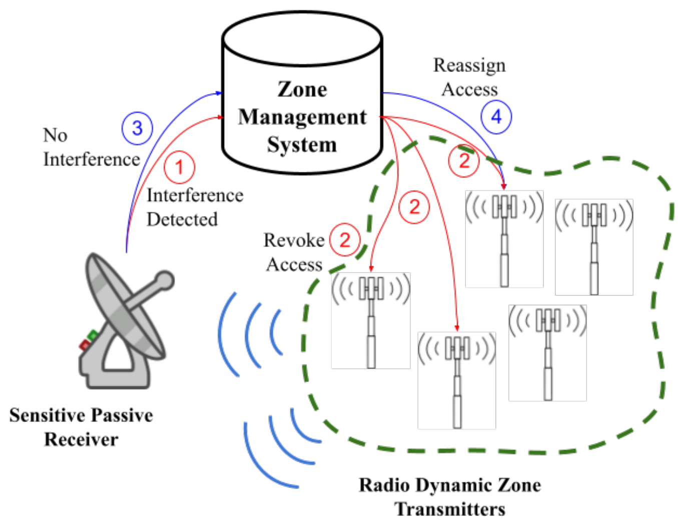

Fig. 1. Zone Management System (ZMS) facilitating coordination between sensitive receivers and interfering RDZ transmitters.

In this paper, we develop a scalable, novel reactive framework for spectrum sharing between RDZ transmitters and a sensitive receiver. Figure 1 shows the conceptual abstraction of our framework to coordinate spectrum access between a sensitive receiver outside the RDZ and transmitters inside an RDZ. We utilize interference supervision by the sensitive user to manage spectrum allocation decisions in the RDZ. The sensitive user employs one or more monitors to identify interference using energy detector-based sensing [7], [8]. When the sensitive user experiences interference in a frequency band, it informs an entity called the *Zone Management System* (ZMS) that regulates spectrum access to RDZ transmitters. The ZMS revokes spectrum access from a set of transmitters in the RDZ using that band to mitigate interference. While interference persists, the sensitive user continues to notify the ZMS, which further reduces the number of RDZ transmitters that have access to the shared band. Once interference is not detected at the sensitive receiver, the ZMS reassigns access to a subset of RDZ transmitters and continues to add more RDZ transmitters until interference is detected. The key performance goal is to minimize the duration of interference at the sensitive receiver and maximize the shared spectrum use in the RDZ. A related performance goal is to minimize interruptions to spectrum availability in the RDZ and also avoid frequent interference periods at the sensitive receiver. The challenge lies in swiftly deciding which RDZ transmitters to revoke or grant access to achieve these goals. Note that we are only concerned with harmful interference outside the zone and do not control interference among RDZ users. Our approach is agnostic to individual signals and only focuses on aggregate impact at the receiver.

Our framework generates an *impact list* of potential interferers using the following approaches: (i) random selection when transmitter characteristics are unknown, (ii) sorting transmitters based on the distance to the sensitive receiver when locations are known, and (iii) sorting transmitters based on the expected signal strength from RDZ transmitters at the sensitive receiver when propagation characteristics are known using a Digital Spectrum Twin (DST) [9] that uses propagation modeling with terrain and elevation maps. A diverse RDZ is expected to have varying spectrum usage patterns over time. Therefore, our framework incorporates computing an exponentially weighted moving average (EMWA) of the number of RDZ transmitters operating simultaneously without causing interference. EWMA represents the overall spectrum use trend in the RDZ and exponentially decays hysteresis as it becomes less relevant. We use weighted variance to capture recent changes in spectrum use. When reacting to interference, the framework generates an impact list where the number of transmitters is calculated using the EWMA estimate.

We implement the ZMS using OpenSAS [10] to provide a complete and deployable open-source spectrum management solution for RDZs. We adapt the OpenSAS messaging protocol and add the ZMS framework modules. We emulate the RDZ and sensitive users with the ZMS-OpenSAS on the POWDER wireless testbed [11]. The experiment shows successful spectrum sharing between test RDZ transmitters and the emulated sensitive receiver, maintaining interference below 0.1%. These results validate the proof-of-concept for our reactive control mechanism in the ZMS prototype.

To comprehensively evaluate our approach, we simulate a real spectrum-sharing scenario with a sensitive receiver and varying sizes of RDZs at long distances. The evaluation focuses on the trade-off between spectrum reuse in the RDZ and interference control at the sensitive receiver. Our proposed approach successfully maintains interference periods below the 0.1% threshold. Notably, utilizing EWMA for interference reaction results in significantly lower interference periods when incorporating location information of RDZ transmitters and propagation estimates from the DST. Our approach demonstrates a substantial advantage, retaining over 80% of RDZ transmitters at scale, compared to if the EWMA estimate is not utilized. Moreover, our approach improves the stability of spectrum access for both RDZ transmitters and the sensitive receiver, with a 50% reduction in interference response time and fewer interruptions on average for RDZ users.

In summary, the contributions in this paper are as follows. We describe a mutually beneficial spectrum-sharing scenario between the sensitive receiver and RDZ transmitters. We propose and implement a framework to coordinate spectrum sharing in this scenario. We examine various scenarios to examine the trade-off between minimizing prolonged interference at the sensitive receiver site and maximizing spectrum reuse in the RDZ.

II. Related Work

This section provides an overview of existing spectrum-sharing ecosystems and their limitations that we address in Section IV.

A. Spectrum Access System

The Citizens Broadband Radio Service (CBRS) [12] facilitates secondary users in sharing the 3.55 - 3.7 GHz band with incumbents, including US Department of Defense (DoD) radar systems. CBRS operates with three tiers of spectrum access, managed by a centralized entity called the Spectrum Access System (SAS). Incumbents are guaranteed interference protection, followed by Priority Access License holders and General Authorized Access users. CBRS has gained widespread adoption among cellular providers. However, it has limitations in terms of its spectrum-sharing approach, imposing constraints on transmitters (e.g., power level, height) and requiring professional installation, resulting in extended installation times. The computational cost for spectrum access from the SAS is high, involving the collection of information about all secondary users and computing aggregate interference for various combinations of lower-tier users. CBRS reserves several hours each night for allocation decisions [13]. Large protection zones (150-400 km) during incumbent operations lead to low spectrum utilization, as fewer secondary users are permitted. The SAS also waits for 2 hours of incumbent radar inactivity before notifying secondary users of channel availability. CBRS users are expected to operate over extended periods (months to years). Therefore, it is acceptable to have long times before new users can operate (hours to days).

B. Automated Frequency Coordinator

Automated Frequency Coordinator (AFC) is a simple and fast solution that coordinates the 6 GHz shared spectrum for standard-power unlicensed devices [14]. It establishes exclusion zones around the incumbent without aggregate interference computations, ensuring a rapid response time in seconds for secondary users. Real-time operations involve swift database access to verify protection zones. However, static and conservative exclusion zones result in low spectrum utilization for secondary users.

III. Problem Settings

In this section, we describe our spectrum-sharing scenario between a sensitive receiver (passive spectrum user) and active transmitters in the RDZ. The following outlines the anticipated behavior within the RDZ. Subsequently, we describe the spectrum requirements, sensing patterns, and interference constraints for one type of sensitive receiver, a remote terrestrial satellite sensing station.

A. RDZ Transmitters

RDZs offer diverse opportunities to users based on the region and applications within and around the zone. In remote areas, there's potential for broadband expansion, while federal closed-loop connectivity is feasible in other regions. Due to the isolation provided by building structures, dense spectrum use areas may facilitate more opportunistic reuse, particularly in indoor and industrial applications. We expect RDZs to play a vital role in current and future-generation cellular networks, offering opportunities for critical infrastructure technologies as well. Special transmitters, such as directed energy systems and high-power microwave transmitters [4], along with experimental radio technology [5], can undergo testing in these zones. An RDZ can cover a non-contiguous space and frequency range.

Currently, two schools of thought exist for potential RDZs.

1) *The "Wild West" RDZ:* This scenario entails minimal regulation, allowing RDZ users not to disclose operational information due to security or privacy concerns. The only enforced condition is to prevent harmful interference outside the zone.
2) *The Managed RDZ:* This scenario adopts a more conventional spectrum-sharing approach, requiring information about RDZ transmitters (e.g., location and operational parameters).

Given these scenarios, the RDZ can have users with new or unknown specifications. Consistent spectrum access is a universal need for all potential RDZ users.

B. Satellite Remote Sensing Stations

A station engaged in satellite passive remote sensing of the Earth and its atmosphere is an important representative of sensitive passive spectrum users. Challenges faced by satellite sensing stations when sharing spectrum with other networks are outlined in [15]. Fixed satellite sensing stations, designed for observing faint directional signals with high-gain antennas, are highly susceptible to interference. While mitigation techniques are applied in rare circumstances, this results in sensitivity and data loss [16]. Sensitive users express concern about potential secondary use over distances ranging from tens to hundreds of kilometers [17], [18].

These sensitive receivers require periodic access to the spectrum at specific frequencies, depending on the type of observations, with applications ranging from scientific to commercial operations. The time to compute the next frequency range to scan and the start of operations varies from tens of seconds to a few minutes. We consider the operations in the 2.64-2.70 GHz band allocated to satellite sensing as a potential range of interest for sharing. This band is scanned in 10 MHz chunks during a 30-minute sensing session. The required sensitivity for receivers is -146 dBm, above which signals are considered interference. While some interference is permissible, it should not exceed 0.1% of the observation time within a 24-hour period [19].

IV. Spectrum Management Approach

This section introduces the Zone Management System (ZMS), designed to govern operations within a Radio Dynamic Zone (RDZ) with the objective of minimizing interference outside the zone to facilitate spectrum-sharing with sensitive receivers. The ZMS makes allocation decisions based on the Spectrum Management Framework that ensures a quick, best-effort response to interference detected at the sensitive receiver. By leveraging this framework, the ZMS enables real-time allocation decisions that efficiently accommodate varying numbers of RDZ users. The framework avoids complex aggregate interference calculations, eliminating a common bottleneck. This allows users with diverse parameters to seamlessly enter and exit the RDZ as needed. Subsequently, we elaborate on how our framework effectively manages scenarios with multiple sensitive receivers. Finally, we underscore how our framework maintains the interference threshold for sensitive receivers while overseeing a large RDZ.

A. Zone Management System

The ZMS is the central entity that enforces the RDZ boundary constraints and regulates spectrum use within the zone. It is a trusted party that maintains the operational information of the RDZ users and the external systems interacting with the RDZ through a uniform interface.

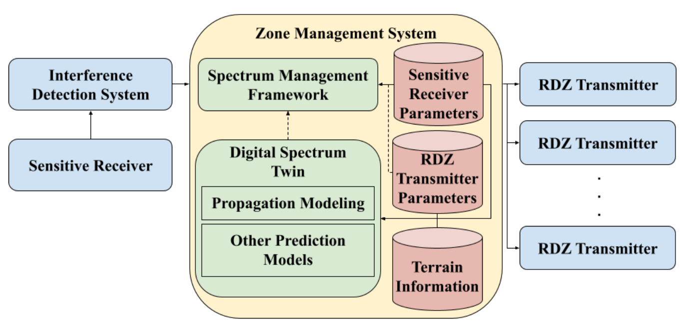

Fig. 2. ZMS modules and interfaces.

Figure 2 shows the modular structure of the ZMS. The main decision-making engine is the spectrum management framework, with various supporting components enabling the ZMS to regulate RDZ user access. The modular design of the framework allows for easy switching between different management approaches. The essential information for the ZMS is the interference indication from the sensitive receiver. Additional information about the sensitive receiver, RDZ transmitters, and the environment enhances the accuracy of spectrum access decisions. We assume the RDZ users adhere to operational instructions and are not malicious. The ZMS can access information about RDZ transmitters depending on the RDZ type (as described in Section III-A). The final ZMS component is a Digital Spectrum Twin (DST) [9] along with the necessary databases. The DST predicts propagation characteristics and harmful interference outside the zone using propagation modeling and information about the environment and users.

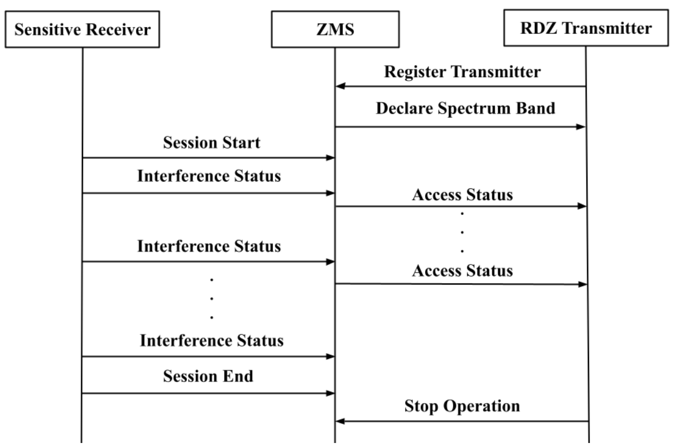

Fig. 3. Spectrum access control communication between ZMS, sensitive receiver, and RDZ transmitter.

Algorithm 1 Poll sensitive receiver and framework calls

if $active$ **then**
  if $interferenceDetected$ **then**
    RevokeAccess($grants, impactList, c_{est}$)
  else
    $c_{est} \leftarrow$ *UpdateEstimate*(n_{grants})
    ReassignAccess($grants, impactList$)
  end if
end if

Algorithm 2 RDZ Heartbeat

for all $rdzTx$ in $grants$ **do**
  if $rdzTx$ in $impactList$ **then**
    UpdateRDZStatus($pause$)
  else
    UpdateRDZStatus($granted$)
  end if
end for

Figure 3 depicts the messages between the ZMS and the two spectrum users. Transmitters register their spectrum use with the ZMS and declare operational parameters in a Managed RDZ (Section III-A). The ZMS receives updates from the sensitive receiver at the beginning and at the end of a sensing session and when interference is detected (Algorithm 1). Based on this information, the ZMS instructs the RDZ transmitters to either resume or pause operations (Algorithm 2). When instructed to pause spectrum use, the RDZ transmitter must halt all transmissions until the ZMS grants permission to resume activity. Finally, the RDZ transmitter notifies the ZMS when it permanently ceases operations.

B. Spectrum Management Framework

The spectrum management framework iteratively identifies RDZ transmitters eligible to share the spectrum with the sensitive user. The framework minimizes real-time computation by avoiding complex planning, a priori knowledge of usage patterns, and the need for extensive propagation modeling.

1) Reactive RDZ access control: In response to observed interference outside the zone, the framework ranks RDZ transmitters based on their interference potential in an *impact list*. The ZMS instructs transmitters on the impact list to pause spectrum use in the shared band. While interference persists, the ZMS updates the impact list and reduces spectrum use in the RDZ further as a best-effort approach to minimize prolonged interference at the sensitive receiver. Once interference stops, access is selectively restored to RDZ transmitters, beginning with those most recently paused. This adaptive approach accommodates changes in spectrum use and environmental conditions, improving spectrum reuse in the zone. Over time, the framework refines estimates of RDZ transmitters likely to cause harmful interference. Complete spectrum access is restored for all RDZ transmitters when the sensitive receiver vacates the band.

Algorithm 3 Update Estimate (n_{grants})

$\mu_i \leftarrow (1-\alpha) * \mu_{i-1} + \alpha * n_{grants}$
$\sigma_i \leftarrow (1-\beta) * \sigma_{i-1} + \beta * |n_{grants} - \mu_i|$
$c_{est} \leftarrow \mu_i - \sigma_i$
return c_{est}

Algorithm 4 Revoke Access ($grants, impactList, c_{est}$)

if $n_{grant} > c_{est}$ **then**
  $n_I \leftarrow n_{grant} - c_{est}$
else
  $I_\Delta \leftarrow I_i / I_T$
  $n_I \leftarrow n_{grant} * I_\Delta$
end if
append(impactList, SortByRDZType($grants, n_I$))

The *Update Estimate* function (Algorithm 3) is called when the sensitive receiver is active and no interference is detected. It records the number of users successfully sharing the band with the sensitive receiver without causing harmful interference. The number of transmitters with active grants in the shared band at iteration i is denoted by n_{grant}. The algorithm updates the estimated number of users (c_{est}) using an exponentially weighted moving average (EWMA) of both the mean (μ_i) and the variance (σ_i) of n_{grant}. The estimate

c_{est} is computed as the lower bound of the expected weighted average, i.e., $\mu_i - \sigma_i$, consistently converging toward the largest set of transmitters that can safely operate with the sensitive receiver. To avoid repeated interruptions to service for all spectrum users, the algorithm uses c_{est} to determine the impact list size.

In the *Revoke Access* operation (Algorithm 4), the framework estimates the impact list size n_I. If the current number of users with spectrum access (n_{grant}) exceeds the estimate c_{est}, n_I is set to the estimated number of users that can simultaneously operate without causing harmful interference ($n_{grant} - c_{est}$). Otherwise, additional revoke operations are necessary due to continued interference. We use the interference duration threshold (I_T) defined by the sensitive user to determine the impact list size. I_T is the acceptable interference period over the total operation duration. The impact list size is increased proportional to $I_\triangle$, the ratio of the observed interference duration I_i to the interference threshold I_T.

Algorithm 5 Reassign Access ($grants, impactList$)

```
1: I_△ ← I_i/I_T
2: n_I = n_grant * (1 − I_Delta)
3: reverse(impactList, n_I)
4: append(grants, impactList)
```

In the absence of interference, the *Reassign Access* function (Algorithm 5) reinstates spectrum access to transmitters, beginning with those least recently revoked. The number of transmitters that regain access is inversely proportional to $I_\triangle$.

Through this iterative process, the framework can adapt to changes in spectrum utilization in real time while maintaining interference below the desired threshold.

2) Creating/Managing Impact List: Our reactive spectrum-sharing solution is designed to operate in both RDZ scenarios III-A. The accuracy of predicting the impact list varies with available information. Three approaches are proposed to create/manage the impact list.

- The *Randomized* approach is used with no prior RDZ transmitters information, picking transmitters randomly. This approach is effective when all RDZ transmitters have similar operations, mitigating variations through randomization. Transmitters would have a similar impact, especially at large distances from the sensitive user.
- The *Distance-based* approach is used when the ZMS has access to location information, utilizing proximity to the sensitive receiver to rank potential interferers. Additional parameters, such as transmission power, are considered if available.
- The *DST-based* approach uses received signal strength estimates at the sensitive receiver from each transmitter. The DST utilizes terrain maps and environmental information to calculate a point-to-point received signal strength estimate at the sensitive receiver. This approach requires more information management by the ZMS than in the previous cases. The impact list's accuracy depends on the accuracy of the DST.

C. Multiple Sensitive Receivers

Here, we consider multiple sensitive receivers operating concurrently with similar or distinct performance requirements. Such scenarios are expected at large sites where multiple sensitive receivers operate in tandem or independently, for example, at various radio astronomy sites [20]. The ZMS manages each receiver's interference reports independently. If multiple receivers encounter interference simultaneously, their impact lists are computed based on prior estimates, including any overlapping users. If interference persists, spectrum access for the remaining RDZ transmitters is revoked in accordance with the algorithm described above. This can prolong the interference period experienced by the sensitive receiver until the estimated impact list adjusts. Spectrum access is not restored until all interference periods are below the desired threshold. Consequently, the maximum number of transmitters that can simultaneously operate is bound by the spectrum reuse when only one sensitive receiver operates outside the RDZ.

D. Design Outcomes

Finally, we highlight how the framework achieves two pivotal design objectives. First, maintaining interference below the acceptable threshold, and second, the ability to manage a large number of RDZ users.

1) Interference Guarantee: The acceptable interference period plays a crucial role in the real-time interaction between the sensitive receiver and RDZ transmitter. Our framework continuously adjusts the number of RDZ transmitters sharing the band with the sensitive receiver based on the remaining interference capacity. This ensures that the interference period does not cross this threshold. However, implicit in this guarantee is the assumption that the sensitive receiver can tolerate the time it takes for the system to receive the first interference notification and send access control notifications to RDZ transmitters. Thereafter, no further reassignments occur if the interference period approaches the threshold.

2) Scalability: Unlike existing spectrum-sharing solutions that rely on conservative propagation models to compute aggregate interference for all combinations of secondary users [12], our approach takes advantage of a sorted list to efficiently handle a large number of RDZ transmitters. This step is the main bottleneck for scaling the operations of existing systems in real-time. In our proposed framework, the impact list is recomputed only when a new RDZ transmitter is added, or the location is changed. This near real-time method allows for scalability at the cost of tolerating a small interference period at the protected spectrum user.

V. ZMS Implementation Using OpenSAS

We implement the ZMS by extending OpenSAS [21] APIs and spectrum allocation algorithms to interface with spectrum users in our scenario and validate our proposed spectrum management framework, shown in Figure 4. OpenSAS exposes the bare-bones SAS (Section II-A) software stack. Operations of the sensitive receiver and RDZ transmitters are emulated on USRP B210 Software-Defined Radios (SDRs) in the POWDER wireless testbed [11].

A. Modified OpenSAS Operation

Modifications to the OpenSAS messaging protocol to communicate with the spectrum users are described next, along with the supporting events triggered by incoming messages.

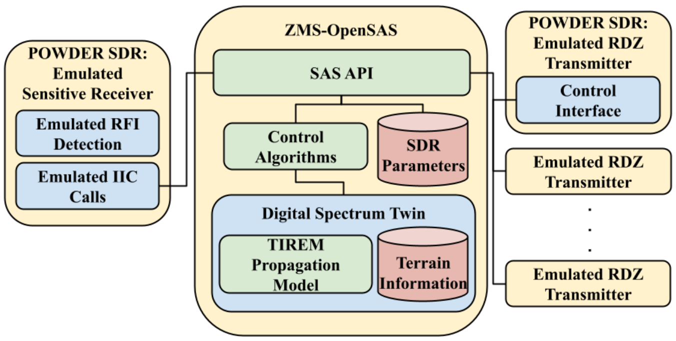

Fig. 4. Software components of experiment with the ZMS-OpenSAS.

1) RDZ Transmitter Interface: We leverage OpenSAS APIs [22] to interface with RDZ transmitters. First, the RDZ transmitters send a Registration Request to ZMS-OpenSAS. We extend the *CBSD category* to identify RDZ transmitters. OpenSAS is modified to accept messages with missing transmitter parameters. A successful registration is acknowledged by a Registration Response from ZMS-OpenSAS. Subsequently, RDZ transmitters send a Grant Request with their desired spectrum range. ZMS-OpenSAS updates spectrum use and sends the Grant Response to the RDZ transmitter with an updated *heartbeat interval* parameter. The RDZ transmitter can begin spectrum use and send regular Heartbeat Requests to ZMS-OpenSAS. The *grant status* in the Heartbeat Response from ZMS-OpenSAS is modified to indicate whether to pause or resume operations. When a transmitter ceases operation, it informs ZMS-OpenSAS via a Deregistration Request. ZMS-OpenSAS removes it from active grants and responds with a Deregistration Response.

2) Sensitive Receiver Interface: We provide two methods for a sensitive user to communicate with ZMS-OpenSAS. First, similar to the RDZ transmitter interface, the Registration Request uses *CBSD category* to identify the sensitive receiver. *Grant status* in Heartbeat Request denotes session status and interference flags. Optionally, the Received Power Measurement Report [22] is used to declare the interference power at the sensitive receiver. Second, OpenSAS polls the Incumbent Informing Capability (IIC) [23] service for incumbent occupancy information via the Incumbent Information call. We add an interference indicator to this call. ZMS-OpenSAS polls the IIC service at regular intervals.

3) ZMS Algorithm: We bypass the OpenSAS allocation algorithm and aggregate interference computation. Instead, we incorporate the ZMS framework and, if available, get the expected signal strength from an RDZ transmitter at the sensitive receiver from the DST. Upon receiving an interference flag, the ZMS framework calculates the probable impact list of RDZ transmitters. ZMS-OpenSAS then marks these to pause operations in the band.

4) Digital Spectrum Twin: We use the TIREM engine for path loss [24] as the DST in our experiment region. We augment the model with a digital surface map from the State of Utah LiDAR survey [25]. The map includes buildings, trees, elevation, and terrain data. TIREM uses information about the sensitive receiver's and RDZ user's locations and other transmission parameters to make propagation predictions.

B. POWDER experiment

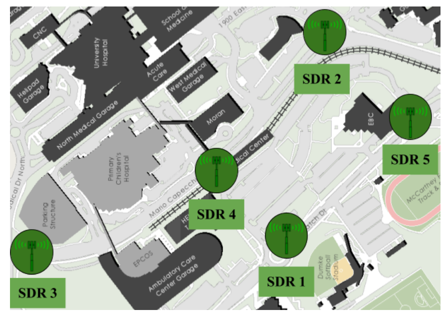

Fig. 5. SDRs in POWDER emulating spectrum sharing scenario.

We emulate our scenario on the POWDER wireless testbed [11] shown in Figure 5. One SDR emulates the sensitive receiver's operation, and the other SDRs emulate test RDZ transmitter behavior. Distances between SDRs range from around 200 m to 800 m. We operate in the 3510-3520 MHz range under test experimental license. Heartbeat intervals are set to 10 seconds for all spectrum users. Given the experiment's small scale, throttling at the ZMS is not a concern. Each test RDZ transmitter has variable spectrum occupancy independent of others, modeled as a simple ON/OFF signal transmission [26]. We transmit a narrowband continuous wave sine function with a maximum transmission gain of 80 dB.

The sensitive receiver uses heartbeat messages to share session and interference information. The maximum receiver gain is set to 30 dB for the entire experiment. First, the sensitive receiver SDR monitors the environment without any RDZ transmitters operating to detect the radio noise floor. The average noise floor observed was around -70 dBm. The interference threshold is set to +5 dB from the noise floor. The session duration and spectrum occupancy adhere to the pattern described in Section III-B, with each scan being 30 minutes and an occupancy rate of approximately 50% over time.

C. Performance Evaluation

We evaluate the trade-off between interference at the sensitive receiver and spectrum reuse in the zone by comparing the three approaches described in Section IV-B – random, distance- and DST-based. For comparison, we emulate the sensitive receiver with one of the three SDRs (SDR 1, 2, 3 in Figures 5 and 6) while the remaining four SDRs simulate test RDZ operations, with each experiment lasting 12 hours.

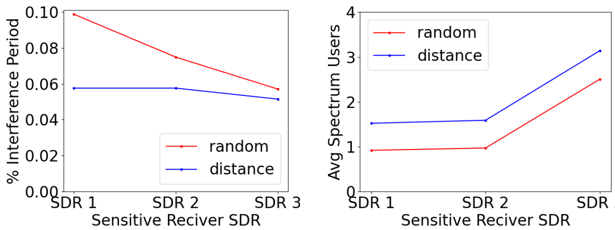

(a) Interference period (b) Spectrum reuse in RDZ.

Fig. 6. Trade-off between interference and spectrum reuse.

Figure 6 shows that the system successfully keeps interference below the 0.1% threshold. Sorting users based on distance leads to reduced interference periods and consistently higher spectrum reuse among the test RDZ transmitters. The DST-based approach resulted in an equivalent performance to the distance-based approach and is therefore omitted.

The ZMS prototype and evaluation of the three proposed methods with emulated users in a real environment serve as a proof-of-concept of the feasibility of our reactive approach, even though it does not cover all possible physical scenarios due to transmit power and other restrictions. The following section covers a diverse and large-scale simulated scenario.

VI. Large Scale Simulation

This section details the modeled spectrum-sharing environment and operations of the two spectrum users – the RDZ transmitters and the sensitive user in a large region. We evaluate the ZMS approaches to find the optimal configuration.

A. *Propagation Model and Region Map*

We detect harmful interference periods at the sensitive receiver using the TIREM propagation model [24]. This model incorporates terrain and elevation maps to predict received signal strength, providing a realistic environmental simulation. TIREM requires various parameters, such as condensation, humidity, surface refractive index, transmitter and receiver antenna specifications, frequency, location, height, gain, and polarization. We model an area spanning approximately 40 km x 30 km, utilizing data from the LiDAR survey [25] with a resolution of 20 meters per point. The elevation map encompasses diverse terrains, including cities, mountains, plateaus, canyons, and basins, as illustrated in Figure 7. We select median values for atmospheric and climate conditions, including refraction index, conductivity, and humidity for the region.

B. *Satellite Remote Sensing Station*

As outlined in Section III, we focus on satellite sensing operations in the 2.64 - 2.70 GHz range. The sensing station scans the 60 MHz range in 10 MHz sections. The interference threshold is set at -146 dBm, and each scanning session lasts 30 minutes, with general occupancy time around 50% over time. Initiating a new session and the band of operation is determined probabilistically. For the performance evaluation in Section VI-E, the sensitive receiver site is situated 35 km

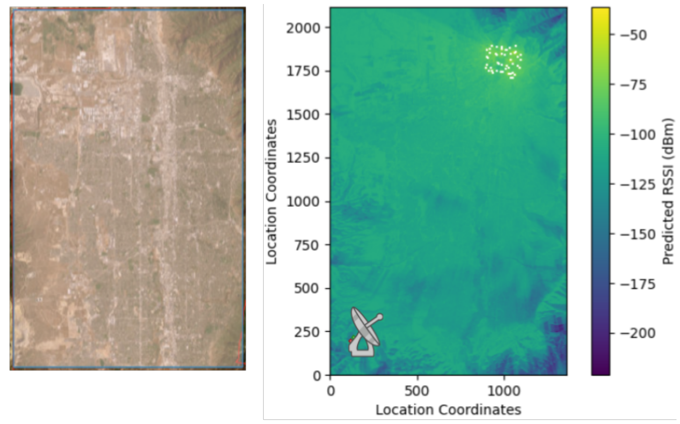

Fig. 7. LiDAR imagery (left) and modeled map with predicted signal strength (right) showing RDZ transmitters and sensitive receiver locations.

from the nearest RDZ transmitter, with an elevation 550 m higher than the RDZ. Additionally, a second sensitive receiver is considered 30 km from the nearest RDZ transmitter and 2 km from the first site in Section VI-F.

C. *RDZ Transmitters*

We simulate the RDZ in a metropolitan region of the map, and transmitters are placed at random locations. Figure 7 shows the expected aggregate power levels across the region for one sample distribution of the transmitters in the map's top right corner. The radio parameters and locations are modeled after the distribution described in [27] that provides expected secondary spectrum usage in CBRS. RDZ transmitter parameters are detailed in Table I. Two sizes of the RDZ are defined, each with varying numbers of transmitters. Each transmitter is assigned one of six 10 MHz channels, and their spectrum occupancy is modeled as a simple ON/OFF signal [26]. The peak case for RDZ use is considered, where all transmitters operate simultaneously at full power.

TABLE I
RDZ Transmitter Parameters

RDZ Size (km x km)	RDZ Transmitters per channel	
4 x 4	{8, 13, 19, 27}	
8 x 8	{32, 53, 75, 107}	
Transmitters	**EIRP (dBm)**	**Antenna height (m)**
45%	26	3
45%	26	6 to 18
10%	40 to 47	6 to 30

D. *ZMS Operation*

In our simulation, we model the ZMS interaction with both the sensitive receiver and RDZ transmitter. To achieve fine-grained control at the considered scale, the heartbeat interval is set to 1 minute, effectively staggering the requests. The EWMA weights α and β are set to 0.125 and 0.25, respectively.

For the DST case, the propagation modeling method employed is the TIREM propagation model [24]. Since the

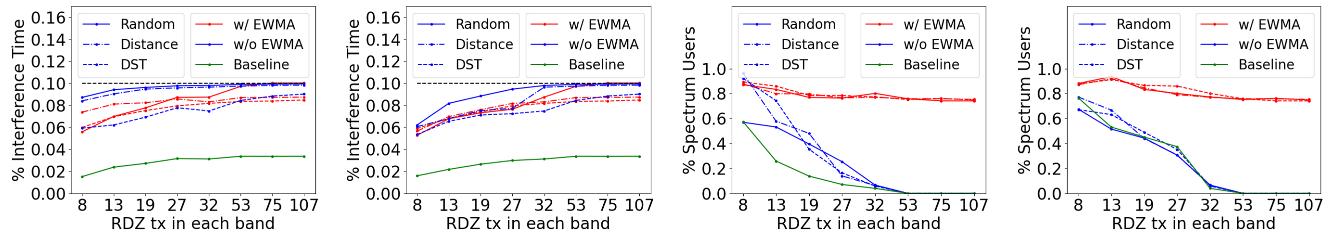

(a) Interference w/ peak RDZ spectrum use (b) Interference w/ dynamic RDZ spectrum use (c) Spectrum utilization w/ peak RDZ spectrum use (d) Spectrum utilization w/ dynamic RDZ spectrum use

Fig. 8. Comparison of interference period at the sensitive receiver and spectrum reuse in the RDZ

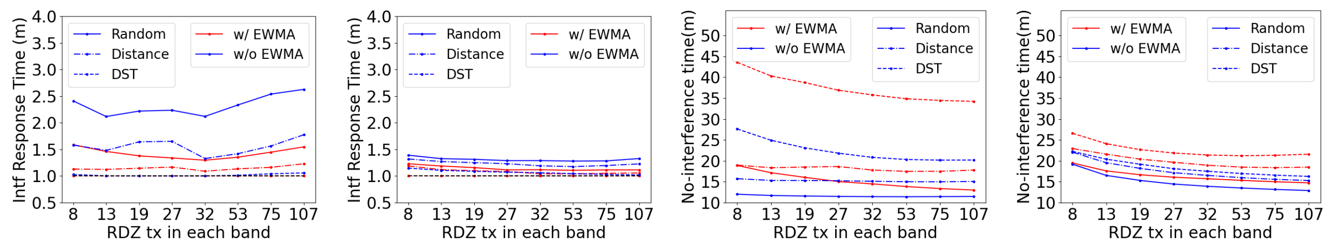

(a) Interference response time w/ peak RDZ spectrum use (b) Interference response time w/ dynamic RDZ spectrum use (c) Interference-free periods w/ peak RDZ spectrum use (d) Interference-free periods w/ dynamic RDZ spectrum use

Fig. 9. Interference response times and continuous interference-free periods observed at the sensitive receiver

simulation uses TIREM as the ground truth, this case depicts the scenario where we obtain perfect information through the DST, which is not feasible with existing models. However, this scenario serves as an optimal benchmark when comparing the performance of the framework.

E. Performance Evaluation

We evaluate the performance of our proposed framework for the three approaches to identify the impact list – at random, by distance, and using a DST. We assess the impact of the EWMA averaging mechanism on performance by comparing it with a reactive mechanism that doesn't use the EWMA estimate c_{est}. Each scenario operates for three simulated days with an event granularity of 1 minute, and the presented results are averages across 100 simulations for each case.

First we focus on the central trade-off crucial for both spectrum users – the interference duration at the sensitive receiver and the number of RDZ transmitters concurrently operating with the sensitive receiver. Next, we investigate the impact of our approach on continuous spectrum availability for both types of spectrum users.

1) Interference period: The acceptable interference period for a sensitive satellite sensing receiver is 0.1% of the time the sensitive receiver operates in 24 hours (Section III-A). This threshold is represented by the dashed black line in Figures 8(a) and 8(b). We compute the interference period as the ratio of the time the sensitive receiver detects interference to the total time it operates in the frequency range over 24 hours. Figures 8(a) and 8(b) illustrate how the average interference period changes with the number of transmitters in each band for peak and dynamic spectrum use by the RDZ transmitters, respectively.

Given that our spectrum-sharing system relies on real-time feedback, there will always be a non-zero interference period at the sensitive receiver. At a minimum, this period includes the time for the sensitive receiver to detect and report interference to the ZMS. To establish the minimum bound for the interference period, we consider a baseline scenario where the ZMS revokes access from all transmitters in the RDZ upon receiving an interference notification, shown in Figures 8(a) and 8(b).

Our results demonstrate successful interference mitigation as the interference threshold is consistently maintained across all cases while the interference period increases with the size of the RDZ, as expected. In scenarios involving dynamic RDZ use, a lower interference period is observed due to the reduced cumulative power received at the sensitive receiver compared to peak RDZ use. Randomly selecting users consistently results in a higher interference period. The interference period is lower when utilizing EWMA, as the framework reacts more quickly to interference. The DST-based approach shows comparable performance to EWMA, with the optimal case being the approach that utilizes the DST. The distance-based approach has a comparable outcome with EWMA at scale, but without EWMA, the interference period is close to the randomized case.

2) Spectrum utilization: We quantify spectrum utilization the number of RDZ transmitters that share the band with the sensitive receiver. Figures 8(c) and 8(d) depict the fraction of users with access to the spectrum when the sensitive receiver operates in the same band, considering peak and dynamic spectrum use in the RDZ, respectively. Similar to the previous section, we include the baseline case, representing the period

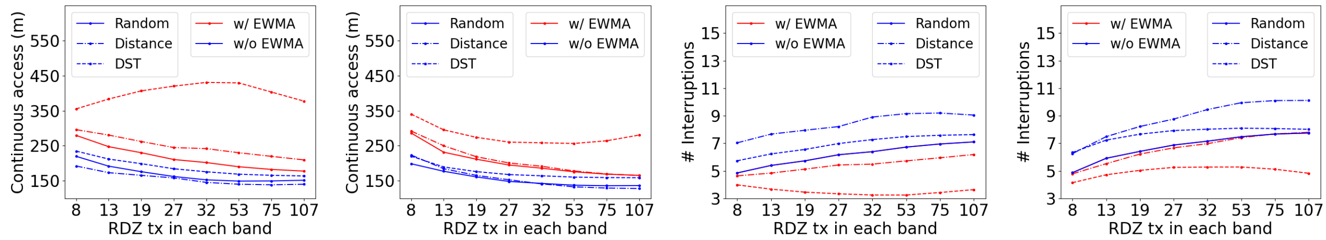

(a) Spectrum access time per user w/ peak RDZ use
(b) Spectrum access time per user w/ dynamic RDZ use
(c) Number of interruptions per user w/ peak RDZ use
(d) Number of interruptions per user w/ dynamic RDZ use

Fig. 10. Continuous spectrum access time and interruptions per user in the RDZ

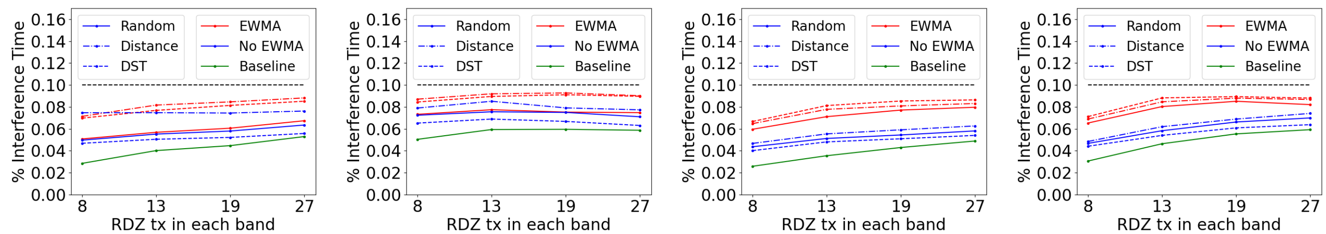

(a) Interference at sensitive user 1 w/ peak RDZ spectrum use
(b) Interference at sensitive user 2 w/ peak RDZ spectrum use
(c) Interference at sensitive user 1 w/ dynamic RDZ spectrum use
(d) Interference at sensitive user 2 w/ dynamic RDZ spectrum use

Fig. 11. Interference times at both sensitive receivers.

of time when RDZ transmitters are operating before receiving a notification to revoke spectrum access from the ZMS.

For both operational models in Figures 8(c) and 8(d), our approach demonstrates a significant improvement in spectrum utilization, consistently above 80% of spectrum users, compared to the scenario where the EWMA estimate c_{est} is not utilized. Without EWMA, the ZMS struggles to adapt to variations in spectrum use in the RDZ, resulting in utilization similar to the baseline case.

In our approach, each time an interference event occurs, the ZMS swiftly reacts by revoking access from transmitters based on the recent average. This allows for more efficient utilization of the acceptable interference period, enabling attempts to reassign transmitters. In contrast, the interference response is slower without EWMA, and the utilization precipitously drops as the interference period continues to increase. Once the interference stops, the low magnitude of the remaining acceptable interference period results in little to no room to increase utilization. The slower recovery from the low utilization in the RDZ occurs due to the persistence of this period. This cycle repeats for each sensing session of the sensitive receiver.

3) Interruptions at sensitive receiver: Figures 9(a) and 9(b) illustrate the average continuous interference time at the sensitive receiver, representing the time it takes for the ZMS to respond effectively to an interference notification. The dashed black line represents the 1-minute heartbeat interval, achievable when the ZMS leverages the combination of rapid response with EWMA and accurate predictions with the DST under ideal network conditions. On average, EWMA-enabled approaches respond twice as fast as those without EWMA for peak RDZ usage. A similar trend is observed in dynamic RDZ use. However, the difference is less than a minute.

Figures 9(c) and 9(d) show the average continuous interference-free periods at the sensitive receiver. The pattern observed here complements the previous result. As expected, the DST-based approach exhibits longer interference-free periods of up to four times, followed by the approaches using distance and at random. This shows the accuracy of the impact list when using EWMA. With dynamic use, shown in Figure 9(d), we note a smaller improvement when using EWMA of 15 minutes on average. However, this is still significant as that would correspond to half of a sensing session. Considering these results along with average interference periods (Figures 8(a) and 8(b)), we can conclude that utilizing EWMA also reduces the number of interference events.

4) Spectrum availability in the RDZ: Figures 10(a) and 10(b) show the spectrum access stability in the RDZ by comparing the continuous access duration per transmitter. Utilizing EWMA maximizes the mean duration but introduces a higher standard deviation, indicating varied access durations among RDZ users. Figures 10(c) and 10(d) reveal the number of interruptions per transmitter. In both peak and dynamic use, utilizing EWMA reduces interruptions on average. The randomized approach performs well when EWMA is not used as each user had fewer interruptions, but as we see from Figures 8(c) and 8(d), more users have to be removed. When using EWMA, sorting by DST results in the fewest interruptions, followed by distance. The randomized case performs similarly to that without using EWMA.

F. Performance with Multiple Sensitive Receivers

Figures 11(a) to 11(d) demonstrate that our framework maintains the interference period below the 0.1% threshold for both sensitive receivers, under peak and dynamic RDZ spectrum use. The spectrum utilization in the RDZ shown in Figures 12(a) and 12(b) consistently exceeds 80% for spectrum users. These results align with those from the single sensitive user case in Section VI-E, demonstrating the applicability of our solution to multiple sensitive receiver scenarios.

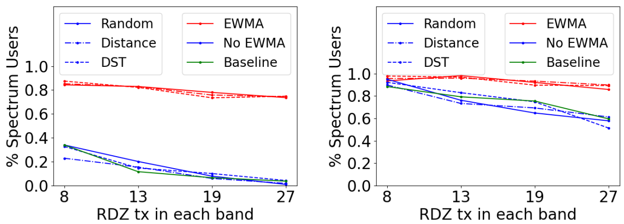

(a) Spectrum utilization w/ peak RDZ spectrum use.

(b) Spectrum utilization w/ dynamic RDZ spectrum use.

Fig. 12. Spectrum reuse with two sensitive receivers.

VII. Conclusion

We developed a novel scalable framework to share the spectrum between a sensitive passive receiver and transmitters in an RDZ. Our framework reacts to interference observed at the sensitive receiver to manage allocation decisions in the RDZ. We provided an open-source spectrum management solution using OpenSAS and examined its operation by emulating the spectrum users in the POWDER wireless testbed. Furthermore, we evaluated our approach by simulating a real spectrum-sharing scenario with a sensitive user and varying sizes of RDZs over long distances. We examined the trade-off between spectrum utilization in the RDZ and interference control at the sensitive receiver. We found that using the location information of the RDZ transmitters or propagation characteristics along with an exponentially weighted moving average to estimate the number of RDZ users capable of sharing the band results in lower interference periods at the sensitive receiver and high mean spectrum utilization in the RDZ.

References

[1] A. Ghosh, S. Kasera, and J. Van Der Merwe, "Spectrum usage analysis and prediction using long short-term memory networks," in *Proceedings of the 24th International Conference on Distributed Computing and Networking*, ser. ICDCN '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 270–279. [Online]. Available: https://doi.org/10.1145/3571306.3571412

[2] Y. Chen and H.-S. Oh, "A survey of measurement-based spectrum occupancy modeling for cognitive radios," *IEEE Communication Surveys & Tutorials*, vol. 18, no. 1, pp. 848–859, 2016.

[3] C. B. Papadias, T. Ratnarajah, and D. T. Slock, *Spectrum sharing: the next frontier in wireless networks*. John Wiley & Sons, 2020.

[4] T. Kidd, "National radio quiet and dynamic zones," in *Information Technology Magazine*. The Department of the Navy, 2018. [Online]. Available: https://www.doncio.navy.mil/CHIPS/ArticleDetails.aspx?ID=10299

[5] M. Zheleva, C. R. Anderson, M. Aksoy, J. T. Johnson, H. Affinnih, and C. G. DePree, "Radio dynamic zones: Motivations, challenges, and opportunities to catalyze spectrum coexistence," *IEEE Communications Magazine*, pp. 1–7, 2023.

[6] C. Peng, H. Zheng, and B. Y. Zhao, "Utilization and fairness in spectrum assignment for opportunistic spectrum access," *Mobile Networks and Applications*, vol. 11, pp. 555–576, 2006.

[7] N. R. A. Observatory. (2023, May) National radio dynamic zone (nrdz). [Online]. Available: https://info.nrao.edu/do/spectrum-management/national-radio-dynamic-zone-nrdz

[8] S. Tschimben, A. Aradhya, G. Weihe, M. Lofquist, A. Pollak, W. Farah, D. DeBoer, and K. Gifford, "Testbed for radio astronomy interference characterization and spectrum sharing research," in *2023 IEEE Aerospace Conference*, 2023, pp. 1–16.

[9] G. D. Durgin, M. A. Varner, N. Patwari, S. K. Kasera, and J. Van der Merwe, "Digital spectrum twinning for next-generation spectrum management and metering," pp. 1–6, 2022.

[10] A. P. DaSilva and M. McDonald. (2023, May) Cci xg testbed pioneers cbrs research. [Online]. Available: https://cyberinitiative.org/xg-testbed/cci-xg-testbed-pioneers-end-to-end-o-ran-research–.html

[11] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas *et al.*, "Powder: Platform for open wireless data-driven experimental research," in *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization*, 2020, pp. 17–24.

[12] M. H. Dortch, "Amendment of the commission's rules with regard to commercial operations in the 3550-3650 mhz band," Federal Communications Commission, https://docs.fcc.gov/public/attachments/FCC-16-55A1.pdf, Order on Reconsideration and Second Report and Order FCC 16-55, May 2016.

[13] G. Cloud. (2023, May) Spectrum access system. [Online]. Available: https://cloud.google.com/spectrum-access-system/docs/

[14] F. C. Commission, "Oet announces conditional approval for 6 ghz band afc systems," https://www.fcc.gov/document/oet-announces-conditional-approval-6-ghz-band-afc-systems, Public Notice ET Docket No. 21-352, November 2022.

[15] M. Höyhtyä, A. Mämmelä, X. Chen, A. Hulkkonen, J. Janhunen, J.-C. Dunat, and J. Gardey, "Database-assisted spectrum sharing in satellite communications: A survey," *IEEE Access*, vol. 5, pp. 25 322–25 341, 2017.

[16] T. E. Gergely, "Spectrum access for the passive services: the past and the future," *Proceedings of the IEEE*, vol. 102, no. 3, pp. 393–398, 2014.

[17] M. Höyhtyä, "Sharing fss satellite c band with secondary small cells and d2d communications," in *2015 IEEE International Conference on Communication Workshop (ICCW)*, 2015, pp. 1606–1611.

[18] I. T. Union, "Studies on compatibility of broadband wireless access (bwa) systems and fixed-satellite service (fss) networks in the 3400–4200 mhz band," https://www.itu.int/pub/R-REP-S.2199, Report ITU-R S.2199, November 2010.

[19] ——, "Performance and interference criteria for satellite passive remote sensing," https://www.itu.int/rec/R-REC-RS.2017/en, Report ITU-R RS.2017-0, August 2012.

[20] NRAO. (2023) National radio astronomy observatory. [Online]. Available: https://public.nrao.edu/

[21] S. Kikamaze, "Design, deployment and performance of an open source spectrum access system," Ph.D. dissertation, Virginia Tech, 2018.

[22] C. W. Standards. (2023, May) Success case studies in the citizens broadband radio service (cbrs) band. [Online]. Available: https://cbrs.wirelessinnovation.org/cbrs-success-stories

[23] M. DiFrancisco, E. Drocella, P. Ransom, and C. Cooper, "Incumbent informing capability (iic) for time-based spectrum sharing," in *The 48th Research Conference on Communication, Information and Internet Policy*, december 2020.

[24] M. A. Varner, T. A. Rodriguez, and G. D. Durgin, "Rf coverage mapping of bistatic radio links using the terrain integrated rough earth model (tirem)," in *2022 16th European Conference on Antennas and Propagation (EuCAP)*, 2022, pp. 1–5.

[25] OpenTopography. (2015, Jan.) State of utah acquired lidar data - wasatch front. [Online]. Available: https://doi.org/10.5069/G9TH8JNQ

[26] L. Csurgai-Horváth and J. Bitó, "Primary and secondary user activity models for cognitive wireless network," in *Proceedings of the 11th International Conference on Telecommunications*. IEEE, 2011, pp. 301–306.

[27] W. Gao and A. Sahoo, "Performance impact of coexistence groups in a gaa-gaa coexistence scheme in the cbrs band," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 1, pp. 184–196, 2021.