Building Radio Dynamic Zones with OPENZMS

David Johnson*, Dustin Maas*, Serhat Tadik†, Alex Orange*, Leigh Stoller*, Kirk Webb*,
M Basit Iqbal Awan*, Jacob Bills*, Miguel Gomez‡, Aarushi Sarbhai*,
Greg Durgin†, Sneha Kasera*, Neal Patwari*‡, David Schurig‡, Jacobus Van der Merwe*
*Kahlert School of Computing, University of Utah
†School of Electrical and Computer Engineering, Georgia Institute of Technology
‡Electrical and Computer Engineering, University of Utah

***Abstract*—Radio Dynamic Zones (RDZs) are being explored by the research community as an approach to safely test and evaluate spectrum sharing mechanisms and technologies. There is general consensus in the research community regarding the conceptual architecture of an RDZ. In this paper, we present our work on the POWDER-RDZ, a prototype RDZ developed and built on the POWDER platform. We present a practical RDZ architecture and explore a number of end-to-end use-cases. We present the design and implementation of OPENZMS, our prototype RDZ Zone Management System, and evaluate it in the POWDER platform.**

***Index Terms*—Radio Dynamic Zone, Zone Management System, spectrum sharing, spectrum management**

I. INTRODUCTION

INTEREST in wireless applications continues to grow unabated, and with it continued demand for spectrum. Spectrum is, however, a finite resource and it is widely accepted that *spectrum sharing* will be the only viable approach to address this supply-demand mismatch [2]. Spectrum sharing is not a panacea, however, and incumbent spectrum users are rightly deeply concerned about the possible impact of spectrum sharing approaches on their respective wireless applications.

A concept being pursued by the research community to explore, and hopefully allay, these concerns is a *Radio Dynamic Zone (RDZ)*. In essence, an RDZ is a "spatial volume" at a particular geographic location where wireless experimentation, and spectrum sharing approaches in particular, can be performed in a controlled way and specifically in such a manner that the potential impact on incumbent spectrum users can be reasoned about and monitored so as to understand and reduce the associated impact and risk. RDZs will be equipped with the necessary tools, mechanisms, equipment and processes to realize such safe exploration.

Received 30 August 2024; revised 14 January 2025 and 10 April 2025; accepted 11 April 2025. Date of current version 17 April 2025. This work was presented in part at IEEE DySPAN 2024 [1] [DOI: 10.1109/DySPAN60163.2024.10632848]. This work was supported by the NSF through the following grants: #1827940, #2232463, #2232464 and #2232465. *(Corresponding author: Jacobus Van der Merwe)*

David Johnson, Dustin Maas, Alex Orange, Leigh Stoller, Kirk Webb, M Basit Iqbal Awan, Aarushi Sarbhai, Sneha Kasera, Neal Patwari, and Jacobus Van der Merwe are with the Kahlert School of Computing, University of Utah, Salt Lake City, UT 84112 USA (e-mail: johnsond@cs.utah.edu, kobus@cs.utah.edu). (Jacob Bills was with the Kahlert School of Computing at the time of this work, and is now with Georgia Tech Research Institute.)

Serhat Tadik and Greg Durgin are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: serhat.tadik@gatech.edu, durgin@gatech.edu).

Miguel Gomez and David Schurig are with the Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, UT 84112 (e-mail: david.schurig@utah.edu).

By necessity the realization of an RDZ is quite complex, involving regulatory and technical challenges, as well as buy-in from incumbents and other stakeholders. The research community is actively investigating these challenges, examining solutions, and exploring spectrum sharing use-cases that might be enabled by RDZs [3]–[9]. There is general consensus in the research community regarding the conceptual architecture of an RDZ [10]. In this paper, we present our work on the POWDER-RDZ, a prototype RDZ developed and built on the POWDER platform [11]. We make the following contributions:

A practical end-to-end RDZ architecture: Our first contribution is to flesh out the conceptual RDZ architecture into a concrete RDZ design and architecture capable of supporting and experimenting with end-to-end RDZ workflows. A *Zone Management System (ZMS)* (called a Zone Management Engine in [10]) is at the center of an RDZ. We formalize the other key role players in an RDZ, i.e., *Spectrum Providers*, *Spectrum Monitors* and *Spectrum Consumers*, and develop a *Zone Abstraction Layer (ZeAL)* to enable these role players to interact with the ZMS via well-defined interfaces.

End-to-end RDZ use cases: We explore the generality of our RDZ architecture by considering the end-to-end workflows associated with a number of different RDZ use-cases. First, we consider using the ZMS to provide general spectrum awareness and spectrum management in an RDZ: i.e., monitoring incumbent spectrum use and making spectrum allocation/use decisions for experimental transmission systems to ensure safe operation in the RDZ. Second, we examine explicit spectrum sharing in the RDZ with an incumbent mobile wireless provider: e.g., a provider who might temporarily, and for short time-scales, grant use of a part of its spectrum to the RDZ. Finally, we consider spectrum sharing with a sensitive spectrum user: e.g., a weather radar system, or a radio telescope installation used in astronomy exploration.

A federated RDZ architecture: We expand upon the core RDZ architecture and propose a federated zone management approach, and map it to our end-to-end use cases. Our approach is motivated by the need to support both hierarchies of spectrum allocation and peer-to-peer zone coexistence and interference mitigation, and we argue that federated capabilities are critical to the success of the RDZ and ZMS concepts. We develop an extended form of POWDER-RDZ, called *RDZ-in-RDZ*, that serves both as an early example of federated capabilities, as well as a real over-the-air, on-demand environment for prototyping new RDZ capabilities.

An open-source, end-to-end ZMS: The key enabler of our work is OPENZMS, the Open Zone Management System, an open-source prototype implementation with the necessary features to operate a real radio dynamic zone. OPENZMS is built with a containerized, cloud native approach, where ZMS core functions are realized as loosely-coupled services with well-defined APIs. This design enables a broad range of deployment options, and, critically, enables flexibility in core function implementation. For example, spectrum intelligence and decision making in OPENZMS is realized as a digital spectrum twin (DST) [6], [7]. OPENZMS is a standalone software system that can be deployed to manage any RDZ; POWDER integrates with OPENZMS over the ZeAL interface to form POWDER-RDZ.

Prototype RDZ implementation and evaluation: Using OPENZMS, we created POWDER-RDZ, a functional realization of our RDZ architecture, and its evaluation in the POWDER platform. Specifically, we integrate OPENZMS with the POWDER platform and leverage its capabilities to facilitate examples of the RDZ role players (spectrum provider, monitor and consumer). Through OPENZMS, we use two RF propagation analysis tools in POWDER-RDZ: an open source version which builds on the TIREM RF propagation model [8], as well as a commercial RF propagation engine which implements several RF propagation models.

Open source ZMS and end-to-end artifacts: To the best of our knowledge, POWDER-RDZ is the first practical end-to-end realization of an RDZ. Our overall objective is for this work to inform and be used by the community to explore RDZ concepts and experiment with spectrum sharing methodologies. To that end, OPENZMS has been released as open source software [12] to the community. OPENZMS has been deployed and expanded for use at the Hat Creek Radio Observatory [13], [14]. We will release our POWDER-RDZ end-to-end use cases as ready-to-run profiles on the POWDER platform to enable others to replicate and build on our work.

II. ARCHITECTURES FOR RADIO DYNAMIC ZONES

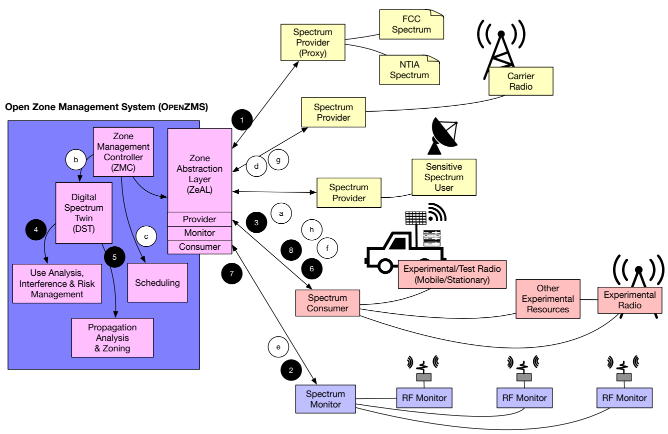

Fig. 1. RDZ architecture.

Based on experience and lessons learned while operating the POWDER platform [11], an outdoor over-the-air wireless testbed and spectrum research facility, and drawing on a set of spectrum-sharing use cases, we propose a new architecture for building and operating radio dynamic zones. This architecture is depicted in Figure 1. The goal of creating a fully automated RDZ is a daunting task: real-time, autonomous services must predict radio operational parameters and enforce spectrum constraints for safe coexistence amongst multiple parties; analyze monitoring reports from a wide variety of RF sensors and mitigate interference; support a wide variety of operational modes informed by spectrum policy; and generalize to manage any band of spectrum delegated by a regulatory body and adapt dynamically to policy changes over time. Based on our use case exploration, we chose to create a new architecture and implementation that is extensible to current and future RDZ uses—but that also can model and integrate with existing spectrum-sharing systems such as CBRS.

At a high level, the architecture consists of a zone management system (ZMS) that ensures safe, dynamic spectrum coexistence for a number of spectrum-consuming organizations (zone participants) whose spectrum use overlaps in time, space, and frequency. In our architecture, the ZMS provides a set of logical interfaces to zone participants. These interfaces allow delegation of spectrum to the ZMS to manage; the dynamic consumption of spectrum inside the zone; and reporting of a variety of RF observations and metrics to assure safe coexistence. We examine these interfaces in more detail below.

The most critical aspect for an RDZ is to have *access to spectrum* that can be used for testing and/or spectrum sharing. The various ways spectrum access might be realized is captured by the *Spectrum Provider* role. As shown at the top of Figure 1, the most basic way an RDZ might obtain access to spectrum is through a licensing process via regulatory agencies, e.g., the FCC and/or NTIA. This might, for example, take the form of special temporary authority (STA), a program experimental license (PEL), or an innovation zone (IZ) designation. Another type of Spectrum Provider involves an incumbent spectrum "owner" (or lessee) who collaborates with the RDZ to enable use/sharing with *their* spectrum. (Note that in these cases there is active spectrum sharing between the Spectrum Provider and the RDZ which is managed/controlled by the RDZ's Zone Management System.) Figure 1 shows two examples of this type of Spectrum Provider. First, a mobile carrier, who, for example during periods of low demand, might temporarily allow use of specific spectrum bands in the RDZ, while retaining the right to dynamically revoke such permission whenever their demand increases. Second, a sensitive spectrum user, such as a radio telescope array, coordinates its spectrum use schedule with the RDZ, thus enabling spectrum sharing in its allocated band; but retaining the right to revoke that permission if an astronomy event changes its planned operation.

Of course, the reason to have an RDZ is to enable testing of radio transmission systems and/or spectrum sharing approaches. As shown in Figure 1, this is captured by the *Spectrum Consumer* role. We anticipate many realizations of this role, spanning active and passive spectrum use cases. For example, a standalone experimental/test radio system might be temporarily brought into an RDZ and the capabilities of its radio transmission system would be (manually) shared with the zone operator to determine a space/volume in which it might safely operate. At the other extreme, the RDZ might have other experimental resources, e.g., other radio systems and antennas, compute and networking capabilities, power and mount points, which could be combined with experimental/test systems.

For general spectrum awareness, and to ensure "safe" operation of the RDZ, a robust *Spectrum Monitoring* capability is needed [4]. We again anticipate a variety of Spectrum Monitor realizations in a practical RDZ. Examples include sensors distributed in and around the RDZ that implement location tracking of test radios [3], inline radio monitoring capabilities that are deployed with radio test equipment [15], aerial spectrum sensors for 3-D spectrum awareness [16], or low-cost, ubiquitous RFID tags for spatially-dense sensing [17].

As shown in Figure 1, the final role is that of the Zone Management System (or, more specifically in our realization, the OPENZMS). The ZMS orchestrates and controls the other role players to realize the RDZ. The role of the ZMS is explored in more detail below by considering a number of example end-to-end use-cases.

A. End-to-end Use-cases

Program experimental license (PEL): Perhaps the most basic RDZ use-case involves the use of spectrum associated with a program experimental license or innovation zone designation. I.e., the RDZ is allowed to use the spectrum range for testing, but should ensure that it does not interfere with incumbent spectrum users and does not cause interference outside the zone. With reference to Figure 1 (steps #1-8): (1) At startup, the ZMS is initialized with information about the RDZ, and specifically with any information about PEL spectrum ranges in which the RDZ might operate. (2) The ZMS requests Spectrum Monitor(s) to perform monitoring of the associated spectrum ranges to provide spectrum intelligence for subsequent RDZ operations. (3) Assume that an experimental 5G system in the RDZ requests spectrum to perform a test. The request might take the form of *XX MHz in the range YY-ZZ MHz*, i.e., to match the capabilities of the 5G system. (4) Using the information provided in the Spectrum Consumer request, the ZMS consults with the DST to determine if the request can be satisfied. In our example the outcome of that query, using data provided by the monitoring system, will be that the request can be satisfied in a specific sub-range of the PEL spectrum. (5) The ZMS will also use the DST to ensure that the 5G system, operating at the power levels indicated in the request, will not cause interference outside of the RDZ. (6) The ZMS then informs the Spectrum Consumer to proceed with testing of the experimental 5G system. (7) Should the Spectrum Monitor report a violation by the 5G system (e.g., transmitting outside the allocated range, or transmitting at a power level that would cause interference outside the RDZ), (8) the ZMS will instruct the Spectrum Consumer to take measures to immediately cease 5G test transmissions.

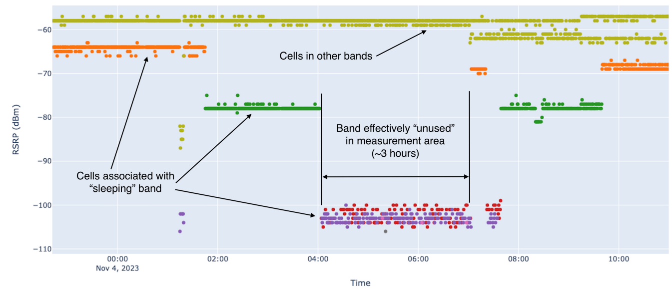

Fig. 2. Example of cells being "idled" resulting in the band in question effectively being unused at the measurement location.

Cooperating mobile carrier: The second use-case depicted in Figure 1 involves spectrum provided by a cooperating mobile carrier. We assume that this spectrum is not a long term sub-lease, but a band that can only be used by the RDZ when not in use by the carrier (i.e., when the carrier explicitly informs the RDZ that the spectrum is not in use). As suggested earlier, this might occur when a provider temporarily idles base stations during periods of low demand to reduce energy usage. Figure 2 shows an example of this behavior in the POWDER area. The figure shows RSRP values reported by a COTS UE in POWDER that are associated with different cells reachable by the UE. The data points at the top are associated with cells operating in bands that are not subject to idling (cells in other bands), while the other data points are all associated with different cells in the band that are subject to being idled (cells associated with "sleeping" band). The effect of this behavior is that there is a period of approximately three hours where the "sleeping band" will *effectively be unused* in the area.[1] For our use-case we assume that the cooperating provider will inform the RDZ when this happens, thus allowing the RDZ to make use of the unused spectrum.[2]

The end-to-end flow associated with this use is depicted in Figure 1 (steps #a-h) and proceeds in a similar manner as the PEL use-case: (a) An experimental 5G system in the RDZ requests spectrum to perform a test, and communicates its requirements to the ZMS. (b) As before, the ZMS consults with the DST to determine if the request can be satisfied. In our example the outcome of that query will be that the request can be satisfied, *but only* when the spectrum is not in use by the carrier. Assume that at the time the ZMS receives this Spectrum Consumer request, the carrier is still using its spectrum. (c) The ZMS notes the outcome of this decision and schedules the Spectrum Consumer request to be satisfied once the spectrum becomes available. (d) At some later time the carrier informs the ZMS that the spectrum range is not in use. The ZMS realizes that the Spectrum Consumer request can be satisfied and, as before, (e) makes a request to perform

[1] I.e., at these low RSRP values (e.g., less than -100 dBm in our example), UEs will switch to using cells with better RSRP values associated with the other bands.

[2] As can be expected and confirmed by our example measurements in Figure 2, these sleeping periods would typically happen during off-peak hours. (In Figure 2 from 4am to 7am local time.) Note that, while likely not generally useful, for an RDZ, access to spectrum for testing will be highly useful even when such access occurs during the night.

monitoring to detect violations, (f) and informs the Spectrum Consumer to proceed with testing of the experimental 5G system. (g) We assume that the carrier anticipates starting to operate in the band again and communicates that to the ZMS. (h) At this point the ZMS system informs the Spectrum Consumer to cease use of the carrier's spectrum band.

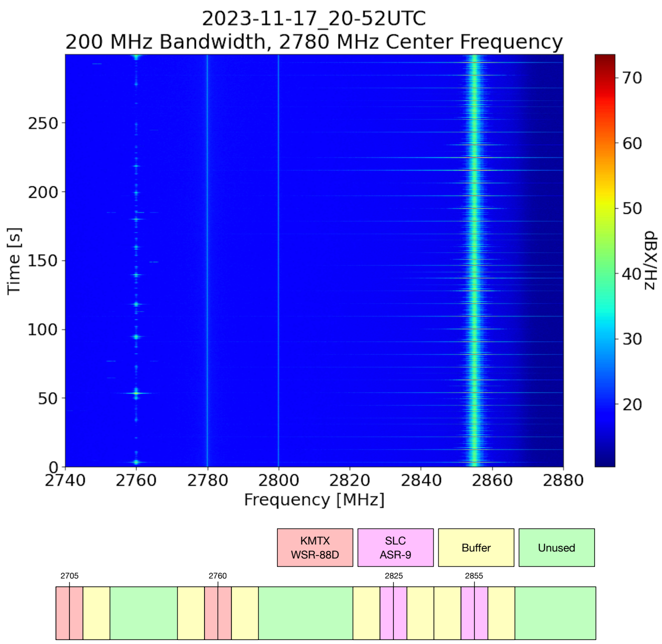

Fig. 3. Radar occupancy in POWDER area (2.74-2.88 GHz).

Cooperating weather radar system: Sharing with weather radar systems presents a use-case that is conceptually similar to the carrier sharing use-case, but from technical, safety and regulatory perspectives provides unique challenges. For example, consider sharing with weather radar systems using the 2.7-2.9 GHz range. This band is used by NOAA's Next Generation Weather Radar (NEXRAD) systems, as well as airport surveillance radar systems (ASR-9), which is used by the FAA to monitor civilian and commercial air traffic [18]. These are safety-of-life systems that are always operational [18], and indeed have been subject to interference from licensed radio communications systems operating *above* the radar band [19]. Nonetheless, assuming safe sharing mechanisms can be realized, this does present a spectrum sharing opportunity. For example, Figure 3 shows a spectrogram of the 2.74-2.88 GHz range in the vicinity of the POWDER platform, as well as the primary/secondary operating frequencies associated with the two radar systems operating in our area. The spectrogram shows the NEXRAD system operating at 2.76GHz and the ASR-9 at 2.855GHz.[3] As suggested by the occupancy representation at the bottom of the figure, and assuming a 10MHz buffer area around the radar operating ranges, a significant portion of the band (*in this case*) is effectively unused. (This may not be true in other geographic areas with additional operational radar systems.) Further, a notification system, interacting with the ZMS, where the secondary operating range of the radar system is shared when the primary is in use (and vice versa), will offer additional sharing opportunities.[4]

Cooperating radio telescope: Our final example use-case involves a cooperating sensitive passive spectrum user such as a radio telescope array. In this case, the radio telescope would provide the ZMS the location of antennas, the spectrum it uses, sensitivity levels, and more. Further, the radio telescope would communicate its anticipated spectrum use (receive) schedule, and the RDZ would be allowed to schedule use of that spectrum in the resulting time-and-frequency gaps. The radio telescope might, however, dynamically preempt use of its spectrum to explore an emerging astronomical event. This use case maps to a similar set of flows between RDZ components as described above. However, because of the sensitivity of the instrument, both the analysis being performed by the ZMS (and its DST), and the sensitivity of monitors deployed to detect violations may differ.

B. Federated Zone Management

We anticipate that federation will be a key enabler of successful radio dynamic zones. The zone management system interfaces we have described extend naturally to supporting hierarchies of flexible spectrum management as well as peer-to-peer inter-zone coexistence. Hierarchies of spectrum allocation, in which trusted root and intermediate zone management systems delegate spectrum authority and usage policy to lower-level zones, enable dynamic spectrum allocation while providing the policy and mechanism to ensure safe coexistence. Root and intermediate zones may also allocate spectrum to consumers within their zones directly via the consumer interface, but should delegate these allocations insofar as possible to leaf zones, to reap the scalability benefits of hierarchy. Moreover, direct allocation of spectrum to consumers is best left to leaf zones, which can leverage detailed, locale-aware propagation simulations to predict use and mitigate conflicts in accordance with delegated spectrum policy. Leaf and intermediate zones can provide higher-level zone management systems with visibility into interference and usage or occupancy via summary reports, as useful or required. Zone peers—whether by time, space, or frequency or interference constraint—may federate via common roots of spectrum authority, and collaborate to manage coexistence at boundaries by analyzing and resolving interference.

Figure 4 shows an example of federated zones in our architecture. Each *zone* is given spectrum authority over sets of $\langle$time, space, frequency$\rangle$ domains. Our concept of a zone is flexible across these domains: a zone may span an entire country in space, but narrowly scoped to a single hypothetical, new dynamic access band; a specific geographic locale with special spectrum authority (e.g., the Green Bank Observatory in the National Radio Quiet Zone [21]); or perhaps the union of a set of radio dynamic zones, each spanning different geographic locales, but with identical spectrum access privileges for select bands. The *root zone* is a top-level spectrum authority granted

[3] The wider signals associated with the ASR-9 system is the result of clipping in our collection system.

[4] More invasive in-band sharing with radar systems has been proposed [20], but it is unclear whether that could be done in practice.

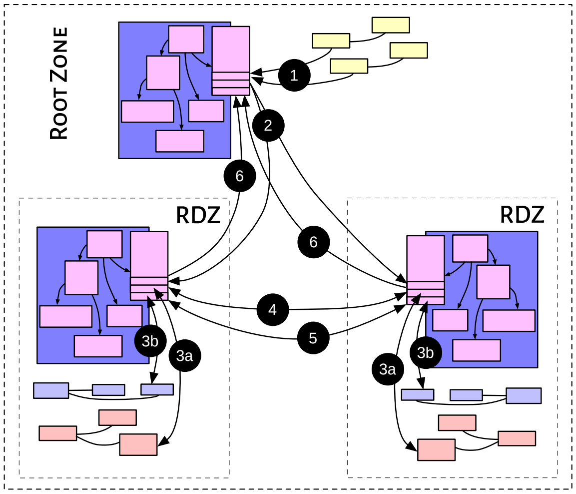

Fig. 4. Federated zone architecture.

by a government spectrum management body (e.g., the US FCC or NTIA). The root zone acts like a certificate authority in a public key infrastructure, and delegates spectrum authority to lower-level zones by signing credentials containing spectrum privileges, policy, interference constraints, reporting requirements, and so on. Intermediate zones may further delegate spectrum to leaf zones as allowed by policy, and leaf zones allocate spectrum to operational radios, allowing each radio's spectrum privileges to be verified.

Specifically, with reference to Figure 4, in Step #1, regulatory agencies would create ⟨time, space, frequency, policy, constraints⟩ spectrum objects via the *Spectrum Provider* interface that a zone management system could further delegate. In Step #2, these allocations are communicated to intermediate or leaf zones, either through requests from lower-level zones via the ZMS's dynamic *grant* protocol, or by manual grant allocation; the recipient zone stores these grants as spectrum objects of its own for further delegation. In Steps #3a and #3b, the leaf zone internally allocates spectrum grants via the *Spectrum Consumer* interface, and monitors usage for compliance via the *Spectrum Monitor* interface, as described earlier. Finally, in Steps #4 and #5, peer zones may exchange monitoring and occupancy reports, and adjust spectrum usage to mitigate interference.

Peer radio dynamic zones are those whose spectrum authority may overlap in ⟨time, space, frequency, constraints⟩: these "adjacent" zones must not interfere with each other according to policy, priority, and interference limits. If interference occurs, the peers may exchange monitoring reports to ensure that each is operating within its spectrum authority. Peer zones may discover each other in two ways. First, zones may be informed of adjacent peers via the delegations of spectrum authority they receive from higher-level zone management systems. Second, zones may observe unknown, adjacent peers via monitoring, and query the higher-level zone (from which their associated spectrum authority originated) to discover matching, authorized peers. We anticipate that peer zones will coordinate overlapping spectrum use by exchanging IEEE Spectrum Consumption Models [22] that summarize zone-wide spectrum use. SCMs model details of per-radio transmission and propagation and sensitivity and reception, may be aggregated to describe an entire zone's use in a specific area of overlap, and therefore offer a mechanism through which peer zones may assess and mitigate inter-zone interference. Each zone may independently predict its usage and refine its own prediction models based on monitoring data, and exchange this information via standardized models that, for instance, do not require each zone to run the same propagation simulations, schedule intra-zone use using the same algorithms, or identically resolve intra-zone interference.

Weather radar. The cooperating weather radar system spectrum sharing use case described earlier may be modeled as two peer radio dynamic zones, as shown in Figure 5a. (Although a non-cooperative version of this use case may also be modeled by describing the weather radar as an incumbent within a single radio dynamic zone, this results in a less-dynamic, coarse-grained version of sharing; and does not properly represent the certified spectrum authority of either participant.) In this example, the radar's zone fully subsumes another zone hosting experimental radio devices, due the potential for even small amounts of radiated power in its coverage area to interfere with the radar. In Step #1, both root zones (in this figure, for instance, the left-hand root zone could be NTIA, whereas the right-hand root zone might be FCC) are configured with spectrum allocation and policy. In Step #2, subsets of this spectrum are delegated to both an intermediate zone (in the POWDER-RDZ PEL use case, this might be NSF) and directly to the weather radar zone. In Step #3, the intermediate zone further delegates a subset of its authority to the radio dynamic zone (in this use case, POWDER-RDZ). In Step #4, the zone peers exchange information to ensure that the radar experiences no interference from the RDZ. In Steps #5a and #5b, each zone operates normally, consuming spectrum, and possibly monitoring for interference. In Step #6, the peer zones might exchange monitoring information to resolve any unexpected interference.

Commercial carrier. The commercial carrier spectrum sharing use case described earlier may be modeled as two peer radio dynamic zones, as shown in Figure 5b. This use case strongly resembles the federated weather radar use case, and most steps are identical in flow and function. In Step #1, the carrier obtains spectrum authority from the FCC, not NTIA. However, the key difference is that the two leaf zones do not overlap. This reflects the case where the zones may potentially interfere with one another in the absence of coordination: their allocations are intended to coexist, enabling same-time, same-frequency coexistence; whereas in the weather radar use case, the RDZ is effectively operating fully within the radar's area of operation.

III. OPENZMS

To manage POWDER-RDZ, we are building OPENZMS, an open-source, end-to-end zone management system. In this section, we describe the design of the services, interfaces, and data models that comprise OPENZMS, and provide a

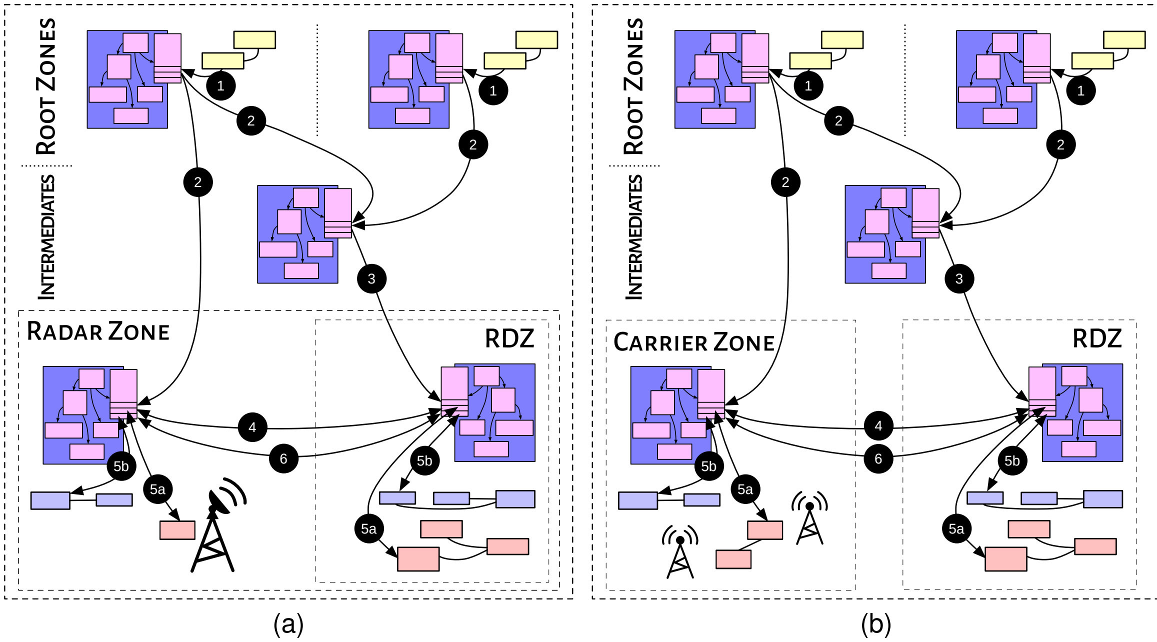

Fig. 5. Federated zone-based spectrum sharing. (a) Weather radar and RDZ. (b) Commercial carrier and RDZ.

brief description of our implementation. OPENZMS employs a centralized spectrum management approach, where spectrum consumers operate under explicit, dynamic delegations of authority from OPENZMS services. This design ensures that OPENZMS can remain in control of an RDZ: it can implement flexible policy-based sharing while still providing protections from harmful interference.

A. Zone Abstraction Layer (ZeAL)

The Zone Abstraction Layer (ZeAL) realizes the APIs that implement the "logical" *Provider*, *Consumer*, and *Monitor* interfaces presented in Section II. At a more granular level, the ZeAL APIs provide the external, publicly-available, "northbound" interface through which organizations and their members participate in spectrum sharing activities within an RDZ. We refer to a participating organization as an *Element*; Elements are the unit of teaming and collaboration within an organization. Elements provide and update OPENZMS with the configuration of resources that they use within the RDZ (e.g., radio transmitters and receivers), and resources that they offer to the RDZ (e.g., spectrum, radio monitors, infrastructure)—either through an automated Element-side service, or via manual ZeAL API invocations or the OPENZMS user interface, when necessary. *Users* are members of one or more Elements, and may be granted one or more *Roles* in a number of Elements to use spectrum or observe its use; or to perform administrative and operational activities within an Element, such as changing an antenna associated with a radio. We expect that many users of RDZ spectrum will not themselves be OPENZMS Users: instead, these Element members may access RDZ resources through user accounts within their organization (Element), and the Element will offer RDZ resources to its own users, via its own abstractions, by consuming the ZeAL APIs.

OPENZMS is designed to automate all aspects of spectrum management in a radio dynamic zone. For Elements to most effectively participate in dynamic spectrum sharing and access, they will interact with OPENZMS through automated Element-side services. To aid these services in responding immediately to changing conditions within the radio dynamic zone, the OPENZMS ZeAL APIs provide a filterable publish/subscribe subsystem, which sends notifications of key events impacting an Element and its radio and spectrum resources (e.g. interference, grant approval or revocation) and control requests (e.g. requests to change a monitor's configuration or task list, or to shut down a radio transmitter).

B. Radio Data Models

To effectively manage an RDZ, OPENZMS defines a data model intended to capture operating characteristics of RDZ radios that consume and monitor spectrum. OPENZMS does not *operate* these radios nor consume spectrum itself; it captures this information to enable effective prediction (e.g., propagation simulation) and analysis of measurements (e.g., interference, model accuracy, and more).

OPENZMS models *Radios* as devices with one or more *RadioPorts*, each of which is connected to an *Antenna*. Radios model the higher-level, generic properties of a device, such as location, description, identifiers, model, and more. RadioPorts define additional details of operation, such as operating bands, tx/rx modes, maximum power level, and attached Antenna azimuth, elevation angle, and location. Antennas define the properties of a specific model of antenna, and can be described via simple properties like gain and beam width; or by detailed radiation patterns (e.g., the widely-used MSI format). Antennas are associated with RadioPorts, and this allows OPENZMS to model the entire radio tx/rx chain during predictive simulations and observation analysis. OPENZMS generically models location and area information to support many different coordinate reference systems, although we expect primary use via WGS84 or UTM due to wide availability of associated data sources.

C. Spectrum

To delegate spectrum to OpenZMS to manage, Elements create *Spectrum* objects. Each Spectrum object consists of start and end constraints, radio constraints (frequency start and end, power, area of operation, maximum leakage allowed outside areas of operation, etc), and per-Element usage policies. These policies allow or restrict Elements from using delegated spectrum; assign maximum priority levels that grants can request; require per-grant approval by the spectrum-delegating Element; and declare usage policy, such as *use-when-unoccupied*. Elements may revoke Spectrum delegations at any time, which will result in revocation of current and future grants allocated within the revoked Spectrum object.

D. Grants

To obtain spectrum from OpenZMS, Elements create *Grant* objects, which consist of the same kinds of start, duration, and radio constraints, and may be associated with one or more RadioPorts. Grant requests may be under-specified across these constraints, allowing OpenZMS to determine a "best fit" allocation of spectrum to the Grant.

OpenZMS grants are stateful and their lifecycle is managed by a grant controller. A newly-created grant that cannot be synchronously *approved* or *denied* will be placed into a *scheduling* state. Once scheduled, if policy dictates that automatic approval is not possible, the Element that delegated the spectrum must approve the grant as well as an OpenZMS administrator, and the grant is placed into an *approving* state until approvals are provided. Once an *approved* grant is accepted by the grantee, it will be placed into the *active* state when its start time arrives. The controller will monitor its state until its expiration time arrives, and then pause the grant and notify the grantee to cease use and wait for acknowledgement. The grant controller implements the server-side of the heartbeat protocol state machine that ensures grant liveness, which we discuss in detail below.

Grants may be associated with two kinds of interference-oriented constraints: *static* constraints that are intended to be evaluated by the scheduler to ensure safe sharing and interference protection; and *runtime* constraints that specify levels and rates of unexpected interference that may occur while the grant is active. A static constraint associated with a grant defines a maximum allowable power level within an area (either a custom area, or the area defined by the operating parameters of the grant) caused by another grant or a third-party. A runtime constraint declares a *metric* whose value may be provided by Observations marked with Annotation values, and tracked while the grant is active, and compared against a threshold to see if the constraint is in violation. Runtime constraints include typical alarm policy knobs, such as declaration of *comparator* function, statistical aggregation function, and period, allowing runtime constraint violation (alarm) state to be tracked over time. Finally, runtime constraints include an optional *criticality* field, which allows the grant requester to communicate to OpenZMS an urgency of resolution. Violations of runtime constraints due to OpenZMS grants are likely to result in grant revocation or replacement (e.g. a new grant with different operational characteristics that may ameliorate the unexpected interference).

Grants may be assigned arbitrary priorities in conformance with delegated spectrum policy limits, allowing OpenZMS to support traditional primary-secondary use cases, where the secondary user may use spectrum in absence of use by the primary. However, this model allows implementation of significantly more complicated scenarios with hierarchies of priorities. If a higher-priority grant is paused, revoked, or deleted, non-conflicting lower-priority grants that were pending may be immediately activated, until the higher-priority grant is activated again.

E. Scheduling

Grant scheduling occurs synchronously when possible, but may require asynchronous responses due to the potentially large range of complex computations required to produce a solution. The OpenZMS scheduler performs simple calendaring and conflict-resolution search to find available spectrum possibilities, but then begins to query its digital spectrum twin, e.g. to look for the best-available unoccupied spectrum, and to ensure that the operating parameters of the grant will not result in transmissions outside the coverage area of the grant's transmitting radio ports. The OpenZMS scheduler does not yet account for all constraints that may be expressed in the model (e.g. geospatial and static interference constraints). An output of successful scheduling includes, for each grant G, a "conflict set" of other grants that cannot be active simultaneously with grant G. If a higher-priority grant G is paused during its active period, the OpenZMS grant controller will activate lower-priority grants in grant G's conflict set.

F. Heartbeats

To ensure grant usage compliance, OpenZMS implements a grant heartbeat protocol through which grantees periodically update the operational status of their grants. These operational status updates allow OpenZMS to quickly confirm that the grantee is properly using the grant or has ceased use when violations occur. Heartbeats may be optionally enforced globally or per-Grant, allowing flexibility for a mix of use cases. Periodic heartbeats are required when a grant is *active* or *paused*. Singleton heartbeats are required when the grant has been *approved* but not yet started; or when the grant is being revoked or replaced. Heartbeats communicate operational status to OpenZMS and ensure both OpenZMS and the grantee know the status of the radios operating via the grant. If the grantee's heartbeat messages are not successfully received and acknowledged without error by OpenZMS, the grantee must cease use of the grant until it can reestablish connection to OpenZMS, and OpenZMS will have either paused or revoked the grant in the absence of successful heartbeat.

G. Claims

Because radio dynamic zones may overlap with incumbent operators whose spectrum allocations are managed by an external spectrum management system, OPENZMS provides a *Claim* abstraction to model incumbent operators and radios. After delegating spectrum to OPENZMS to manage, spectrum providers can create claims that take priority over conflicting grants that the OPENZMS scheduler determines cannot co-exist. Internally, claim operating and interference constraints are modeled similarly to grants, which allows the OPENZMS scheduling algorithms to analyze coexistence properties between grants and claims—so that higher-priority claims block (pend) lower-priority grants. We expect that claims will typically be assigned higher priority than OPENZMS-scheduled grants, but spectrum providers are free to set arbitrary priority values for their claims and grants to implement specific use cases, even allowing claims to take priority over each another.

H. Monitors

OPENZMS is designed to support a wide set of monitoring systems, whether they produce IQ samples or processed and reduced Observations (e.g. power-spectral density, radar operation indications, interference). Monitoring systems may operate in a fixed-deployment, continuously-reporting mode and produce one kind of Observation the entire time they are in operation; or they might be deployed on flexible infrastructure, such software-defined radios that run different host compute software monitoring systems at different times—and are therefore customizable to the needs of a ZMS, which vary over time with spectrum and grant constraints. Monitors may be single radios, or they may aggregate data from multiple radios into individual processed Observations (e.g., to perform TDOA-based localization, monitoring-as-a-service, and more). OPENZMS's monitoring abstractions seek to accommodate any of these monitoring systems.

Just as OPENZMS does not operate the Radios provided by Elements, it does not operate Monitors. For simple OPENZMS deployments, fixed, non-configurable Monitors may provide enough data to manage spectrum—for instance, by simply monitoring occupancy via power-spectral density sweep scans. However, if the Monitors are part of flexible infrastructure or a service that can run different types of monitoring tasks, or can be reconfigured according to different common parameters (e.g., can sweep a different frequency band, change gain), OPENZMS may be able to more efficiently use the monitoring system.

OPENZMS's monitoring abstractions build upon the base radio model abstractions. Elements define *Monitors*, associated with particular RadioPorts, which provides detailed reception characteristics to OPENZMS analysis services. Monitors report *Observations*, which contain radio measurement data and metadata. The OPENZMS design does not dictate specific measurement and data formats, and accepts processed, indexed, and analyzed information (e.g., occupancy data, power-spectral density), or raw sample data (e.g., SigMF). Our goal is to enable many kinds of analysis services within OPENZMS via horizontally-scalable Observation processing pipelines. Analysis services may attach additional processed or learned information to Observations as *Annotations* (e.g., an interference and risk analysis). Finally, Observations may be grouped as *Collections*, which may be useful in cases where a monitoring device cannot capture in one tuning a desired wideband frequency range, and must combine multiple scans at several center frequencies to implement a "sweep" of the wideband range.

I. Deployment Scenarios

OPENZMS is designed and anticipated to manage campus to muncipal- or regional-scale deployments. While a full evaluation of the scaling properties of OPENZMS is beyond the scope of this work, certain high-level deployment properties imply scaling challenges. A highly-dynamic environment with many short-lifetime grants will add to scheduling load and conflict analysis computations, as could a deployment with many densely-packed operational radios. High-fidelity interference analysis and protection will require increased monitor vantage points, which will increase data processing requirements. Total geographic area covered by a deployment may significantly increase the cost of high-fidelity propagation simulations, or require subdividing the covered area into separate yet overlapping propagation simulation domains. Analyzing these properties in depth is the subject of future work.

J. OPENZMS *components/services*

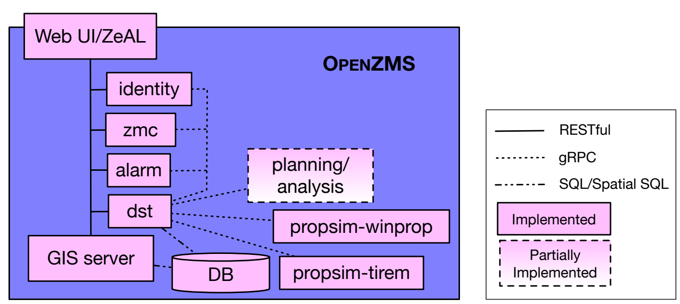

Fig. 6. OPENZMS services.

The OPENZMS ZeAL API is implemented by several services, shown in Figure 6, that provide secure, RESTful endpoints. All services communicate internally over a trusted, RPC-based messaging layer, designed to facilitate high-throughput, low-latency messaging and reactive, event-based analysis pipelines.

1) Core services: The *identity* service provides the Element, User, Role, and Token abstractions and operations. Users can scope Tokens with subsets of their Roles to restrict the set of authorized operations. All OPENZMS services register with the identity service, which maintains a directory to support service-to-service communication.

The *zmc* (zone management controller) service provides the Radio, Monitor, Spectrum, Grant, and related abstractions—it is the primary API endpoint for Elements to populate

OPENZMS with radio device information; to provide spectrum for OPENZMS to manage; and for Elements to reserve spectrum via the Grant abstraction. The *zmc* service currently hosts the spectrum scheduler, which allocates Spectrum to Grants, querying other services as required by policy (e.g. the *dst* service's occupancy data and propagation simulation maps). The *zmc* service pauses, revokes, or replaces Grants when notified by the *alarm* service that they are in violation of operating constraints, or likely to be causing reported interference.

The *dst* (digital spectrum twin) service provides the Observation, Collection, Annotation, and Propsim (propagation simulation) abstractions. Conceptually, it operates as a dual-purpose system, functioning both as a predictive engine and a data analytics hub. This service indexes records of RF measurements, spectrum usage, and radio device data, and provides a predictive query interface. These queries serve two purposes: 1) determining occupancy, such as verifying the availability of X MHz of spectrum for radio Y at power level Z, at a specific time T, and for a duration D; and 2) estimating a transmission power range for radio Y at time T, ensuring emissions remain below power P outside the RDZ. The *dst*'s analysis capabilities empower it to extrapolate intelligence from short-term observations to identify and understand long-term, spatial spectrum usage trends as exemplified by Figure 7. The *dst* service stores Observation data in persistent file storage and indexes within appropriate databases (e.g., geospatial, relational, time-series) to support fast queries. It can invoke propagation simulation services on-demand, but can also simulate in expectation of future usage as new Radio and Spectrum objects are created and updated by Elements.

The *dst* service acts as a client to *propsim* (propagation simulation) services, which generate expected received signal strength maps and other geospatial features for one or more transmitters operating in a given band and power level. These predictions are the basis to facilitate simultaneous, deconflicted shared spectrum use. OPENZMS defines an extensible RPC-based *propsim* job service, and through this interface, the *dst* service can run parameterized propagation simulations to obtain maps with received signal strength data. The *dst* service caches and indexes these maps in its PostGIS database to facilitate geographic sharing under a variety of constraints.

2) Analysis services: The *alarm* service analyzes and responds to unexpected interference reports that can occur due to spatial and temporal changes not captured even with sophisticated planning and a multi-dimensional DST. OpenZMS uses monitor observations, along with historical data from the DST, to manage spectrum access for consumers, revoking access from probable interferers. Credible reports from interference events and monitoring data will be used to update the DST for future risk assessments. The primary concern for incumbents is minimizing interference events. This is especially critical for sensitive applications like radio astronomy that have a very small window of acceptable interference [23]. The *alarm* service provides a real-time response to unexpected interference reports caused by other RDZ consumers. It attempts to identify the source of unexpected interference, and if the source may be another grantee, the *alarm* service will recommend corrective actions to the *zmc* service, such as pausing or revoking the suspected offending grant, or replacing it with a grant with different operating characteristics (e.g. reduced maximum allowed power level). The *zmc* service will consider these recommendations in coordination with spectrum and grant policies and choose a final action.

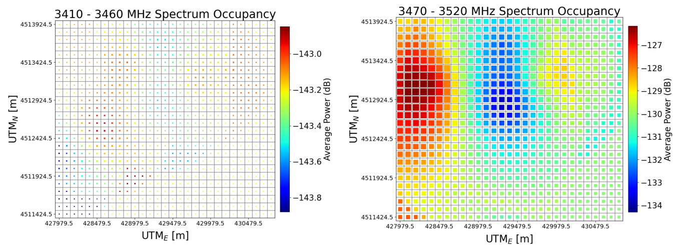

Fig. 7. Spatial spectrum occupancy: 3.41-3.46, 3.47-3.52 GHz.

The DST also provides *spectrum planning* services by analyzing long term spectrum monitoring data to provide spatial spectrum occupancy maps. These maps reveal historical usage patterns and availability of specific frequency bands at targeted locations. To illustrate this functionality, we have analyzed spectrum monitoring data from seven different monitors in the RDZ, spanning from June 2022 to November 2023. Our analysis involved identifying a power threshold that distinguishes noise from spectrum usage. We then calculated two key metrics for spectrum occupancy at all monitored sites: the duty cycle and the average occupancy power. After calculating these metrics, we employed spatial interpolation techniques to estimate values for each *proxel*—or propagation picture element—within the RDZ. Specifically, we used a radial basis function interpolation for duty cycle estimations and a Kriging interpolation for average occupancy power. We then synthesized these spatial estimates, correlating the ratio of the area of a square within each proxel to the proxel's total area with its duty cycle. Additionally, the color of these squares visually represents the average occupancy power, with a corresponding colorbar indicating power levels. The result is depicted in Figure 7. Notably, our observations reveal a significantly lower duty cycle and average occupancy power in the 3410-3460 MHz band compared to the 3470-3520 MHz band in the POWDER deployment area, aligning with insights from instantaneous monitoring data and reinforcing the data's utility in strategic decision-making processes.

K. OPENZMS *Implementation*

OPENZMS is built as a set of containerized, cloud-native services to support horizontal scalability for large analysis and prediction workloads, and to facilitate extensibility via third-party services. Each service that provides parts of the ZeAL API does so over a RESTful JSON "northbound" interface, on which OPENZMS users invoke typical CRUD operations (Create, Read, Update, Delete) across the data model. Users authenticate via API tokens and are authorized via role-based access control (RBAC) (e.g., admin, operator, user, viewer). Internally, services communicate over trusted gRPC service APIs and event streams.

The core *identity*, *zmc*, *dst*, and *alarm* services are implemented in Golang; the *propsim-tirem* and *propsim-winprop* propagation simulation services are implemented in Python. The core services each store data in a relational Postgres database, and the *dst* service further uses the PostGIS extensions to support raster storage of propagation simulation maps and geospatial indexes and queries. We built a web UI that exposes OpenZMS's core abstractions, using the `Vue.js` and `Nuxt.js` frameworks. The UI displays zone status, live measurement graphs, and propagation simulation maps as web map tiles generated and cached by a Geoserver instance, attached to the *dst* services's PostGIS database.

OpenZMS's *propsim-tirem* service wraps the Terrain Integrated Rough Earth Model (TIREM) [24]. TIREM is a tried-and-true theoretical propagation model that considers physical phenomena such as free space path loss, ground reflection, and diffraction. It has also been shown that TIREM's predictions can be enhanced with measurements collected in the RDZ [8]. OpenZMS's propsim-tirem service exposes a variety of TIREM parameters to control the simulation; parallelizes its execution; and outputs maps as GeoTIFF images that are added to the dst service's database for use in geospatial queries. By default TIREM operates under the assumption of isotropic antenna radiation, an approximation that often diverges from real-world conditions. In contrast, the dst service takes into account the specific radiation pattern of the antenna at the terminal. This approach more accurately represents the antenna's gain distribution across different regions, thereby minimizing the discrepancies caused by relying on the generalized assumption of isotropic radiation.

To illustrate OpenZMS's modularity, and to benefit from other RF analysis tools, we have also implemented the *propsim-winprop* service, based on Feko Winprop, a commercial product from Altair. Winprop provides a range of wireless planning and analysis tools, including RF signal propagation analysis in varied environments. Figure 8 shows a sample RF propagation map produced by this service.

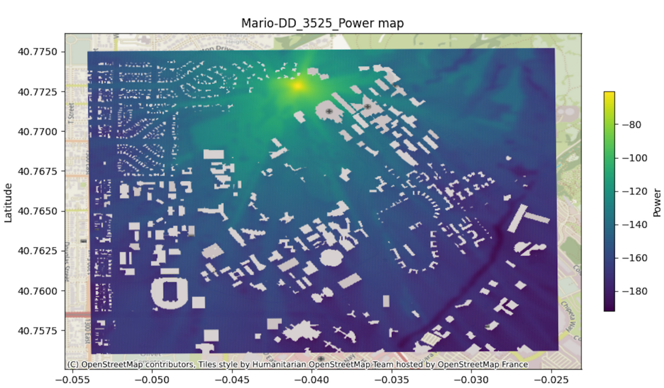

Fig. 8. Example propsim-winprop RF propagation map.

IV. Powder-RDZ Implementation

While we are realizing Powder-RDZ on the Powder platform, our goal with its design, and specifically the design of OpenZMS, is to produce a generic RDZ architecture and a reference ZMS implementation that will be applicable to a broad range of use cases and RDZ deployments. We used and enhanced the Powder platform to achieve this goal. Specifically, we implemented the OpenZMS *ZeAL interfaces* as a thin abstraction layer that maps to existing Powder functions, and used and enhanced other platform features to provide or emulate Spectrum Provider, Spectrum Consumer and Spectrum Monitor capabilities and to execute variations of the use-cases described in Section II. Figure 9 depicts this implementation and the setup we used for evaluation.

A. Spectrum Consumer

Spectrum use in the Powder platform is based on requests associated with our FCC PEL and our FCC IZ designation. The platform also treats spectrum as any other platform resource that can be reserved, allocated, and used by experimenters. For our Powder-RDZ exploration we modified Powder to *cede spectrum control* for a configured spectrum range to the OpenZMS. Specifically, when instantiating an RDZ-related test/experiment in the platform, the normal Powder workflow pauses to request a spectrum grant from the OpenZMS via the ZeAL Spectrum Consumer interface, and resumes once it receives the approved grant from the OpenZMS. Requests from the OpenZMS to stop using spectrum previously granted maps to existing Powder mechanisms to terminate a test/experiment.

The Powder platform is a highly flexible mobile and wireless testbed, and is therefore a good realization of other Spectrum Consumer functions. The platform has the necessary mechanisms to reserve and allocate a variety of wireless equipment (e.g., software-defined-radios (SDRs), commercial-off-the-shelf (COTS) user equipment (UEs) and radio units (RUs)), and to combine that with other platform components (compute nodes, network switches, software stacks) to realize a broad range of functions useful for RDZ testing. For instance, we use the platform's 5G capabilities as our canonical RDZ testing system as described in Sections II and V.

B. Spectrum Monitor

As shown in Figure 9 (light purple elements), we have realized three example implementations of the ZeAL Spectrum Monitor. First, the Powder platform uses *inline spectrum monitoring*, using an RF coupler and specialized software on a dedicated monitoring SDR, to detect any RF transmission violations by experimenters [15]. We implemented the ZeAL Spectrum Monitor interface on this existing monitoring system to report spectrum violations to the OpenZMS. Second, Powder performs regular (a number of times per day) over-the-air monitoring using all available SDRs in the platform. We added the ZeAL Spectrum Monitor interface to this spectrum monitoring system to report spectrum observations to the ZMS. Third, we used the NTIA/ITS spectrum characterization and occupancy sensing (SCOS) [25] open source software, combined with Powder SDRs, to realize an IEEE 802.15.22.3-compliant OpenZMS monitor.

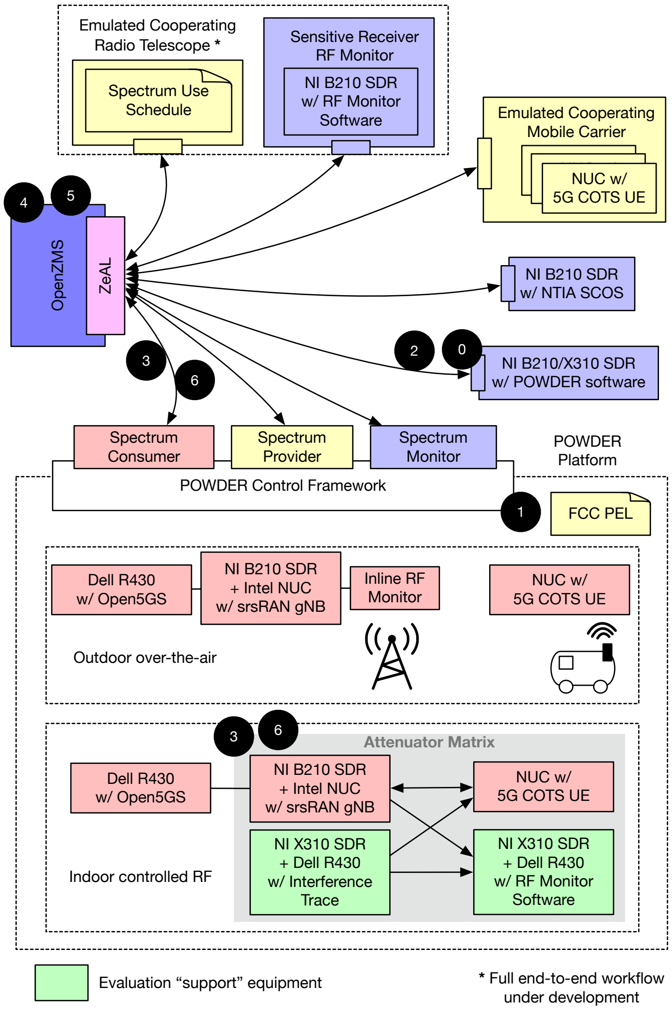

Fig. 9. Implementation and evaluation setup.

C. Spectrum Provider

We have also implemented three Spectrum Provider examples (yellow elements in Figure 9). First, as noted above, spectrum use in POWDER is coupled with our FCC PEL/IZ requests. The platform reports these, or at least the subset associated with RDZ testing, to the OPENZMS via the ZeAL Spectrum Provider interface.

Second, we have developed a *carrier sleep/idle cell detector* to *emulate* a mobile carrier Spectrum Provider. We built an application which uses the standard “cell search procedure” available on COTS UEs to record the availability of cell towers operating on specific cellular bands. This application executes on COTS UEs distributed across the POWDER platform, and reports back to a centralized detector, which determines when cells are idled and uses that information to emulate a cooperating mobile carrier via the ZeAL Spectrum Provider API. The data collected includes the cell identifier, the frequency the cell is transmitting, and the RSRP. The detector determines the maximum RSRP value observed for all cells in the band of interest as observed by associated observation points. If this system-wide maximum RSRP value is less than a threshold, the detector decides that the band is effectively unused in the area covered by the detector, and signals the availability of spectrum to the ZMS. Whenever the system-wide maximum RSRP value exceed this threshold, the spectrum use is revoked.

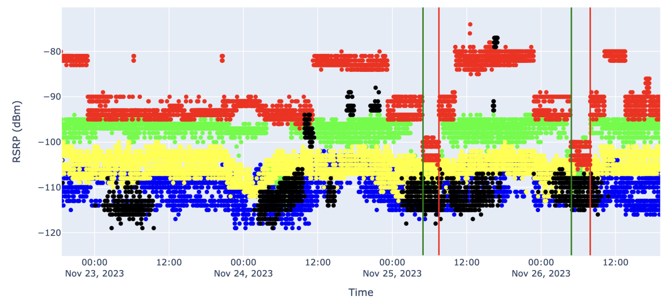

Fig. 10. Idle cell detector emulating cooperating mobile carrier.

Figure 10 shows example output of the detector superimposed on the underlying data from the COTS UEs used by the detection system. Specifically, the figure shows RSRP values, associated with cells operating in the “sleeping band”, as observed by the UEs used in the detector as a function of time. Our detector collected data at five different COTS UEs deployed at fixed endpoints. Cells observed from the same detector location are displayed with the same color. The superimposed vertical green and red lines respectively represent when the monitor decides to provide and revoke spectrum, thus emulating the mobile carrier notifications. (The maximum RSRP threshold used in Figure 10 was -99dBm.)

Finally, as shown at the top of Figure 9, we implemented an *emulated* radio telescope Spectrum Provider. There are two modes of operation for sensitive spectrum providers: exclusive access or shared access. The ZMS allows spectrum sharing on a temporal basis by allowing other consumers to operate when the sensitive spectrum provider is not operating. In the shared access case, the ZMS decides which consumers can operate in the same frequency range as the sensitive spectrum provider without causing harmful interference. The sensitive spectrum monitors are used by the ZMS to enforce the interference constraints and respond to any interference reports by revoking access from potential interferers. We implemented a variant of this spectrum provider where the radio telescope uses its own RF monitors to detect interference and report that to the OPENZMS to take corrective action.[5]

D. Hierarchical Federation: Nested RDZes

Having built POWDER-RDZ and enabled POWDER experiments to access spectrum through a real zone management system with RDZ-based safety and coexistence properties, we created an initial example of the federated RDZ capabilities presented earlier in Section II-B. We refer to this operational mode as *RDZ-in-RDZ*, and its extended architecture is shown in Figure 11.

In this mode, POWDER-RDZ users may instantiate their own private RDZ—the “inner” RDZ—as a POWDER experiment that runs its own OPENZMS, whose spectrum is

[5]We first implemented this as an end-to-end standalone RDZ use case and, at the time of writing, are porting that functionality onto OPENZMS.

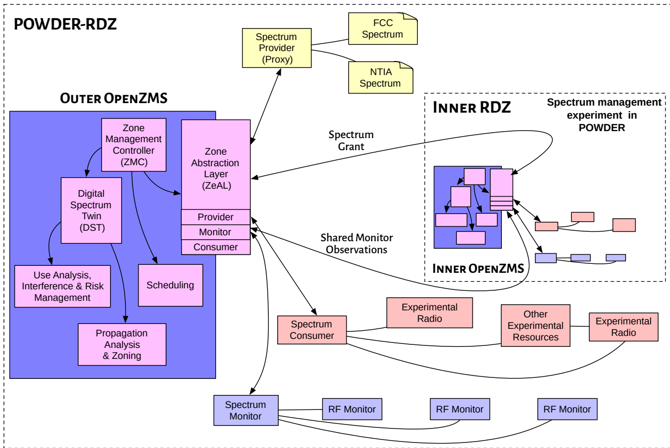

Fig. 11. RDZ-in-RDZ architecture.

allocated by POWDER-RDZ—the "outer" RDZ. Users have full control of their "inner" RDZ, allowing them to replace components and test new workflows and algorithms—but they must provide the same safe coexistence guarantees as the outer RDZ. The user may then allocate spectrum for regular POWDER experiments *via their inner RDZ* instead of the outer RDZ. This functionality is specific to POWDER-RDZ, since it relies on the same experiment automation and radio infrastructure building blocks provided by the POWDER platform that were used to build POWDER-RDZ. However, it builds on the generic federated RDZ concept and uses the existing OPENZMS interfaces (spectrum provider, consumer, monitor) developed for standalone RDZ deployments—and therefore represent a prototype realization of hierarchical federation of two radio dynamic zones.

When an inner RDZ is created, it requests an allocation of spectrum just like another POWDER-RDZ experiment would, but with the intent to further delegate this spectrum as grants to other POWDER experiments, using the Spectrum Provider interface (labeled "Spectrum Grant" in Figure 11). Thus, the inner RDZ receives a Grant of spectrum from the outer RDZ, and represents this as a Spectrum object (a delegation of spectrum with constraints) within its inner RDZ. The inner RDZ must provide safe coexistence for the experiments operating under its spectrum grants, so it requires similar monitoring data as the outer RDZ. To share resources efficiently, monitor data streaming into the outer RDZ is forwarded to the inner RDZ (again, using the existing Spectrum Monitor interfaces; labeled "Shared Monitor Observations" in Figure 11). The inner RDZ operator can also allocate their own inner SDR-based monitors, possibly running different monitoring software and data analysis pipelines.

To ensure that safe coexistence is guaranteed inside POWDER-RDZ, whether an experiment's spectrum is being managed by the inner or outer RDZ, the outer RDZ will terminate both grants and experiments allocated via the inner RDZ, if a transmission violation is detected via the monitors. This is a necessary feature for the virtualized aspect of POWDER-RDZ, although it is not part of the federated, hierarchical RDZ pattern. In a decentralized hierarchy of (sub)delegation of spectrum and policy, each level's zone management system would be responsible to operate safely within its delegated spectrum.

V. EVALUATION

POWDER-RDZ combines numerous (complex) components into a coherent prototype RDZ. A thorough evaluation of each of these components is beyond the scope of this paper. Instead we focus on demonstrating the generality of the POWDER-RDZ architecture and specifically its utility in realizing end-to-end RDZ workflows. We therefore focus our evaluation on an end-to-end functional illustration/evaluation of one of the RDZ workflows realized in the POWDER-RDZ. We specifically use the PEL use-case described in Section II-A.

A. Evaluation setup

Figure 9 depicts the two experimental setups used for our evaluation, as well as steps used in the end-to-end workflow. (These steps are the same as those for the corresponding use-case described in Section II-A and shown in Figure 1.)

These two setups are functionally equivalent, involving an experimental end-to-end 5G system made up of a 5G COTS UE (a Quectel RM520N) connecting to an SDR-based gNodeB (an NI B210 and Intel NUC compute node executing srsRAN software), which connects to a 5G Core instance (Open5GS software executing on a Dell R430 compute node). One version of this setup is deployed in the POWDER outdoor environment and enables over-the-air functionality, while the other is deployed in our controlled-RF environment consisting of equipment in RF-shielded enclosures interconnected through a programmatically controllable RF attenuator matrix. The outdoor over-the-air variant allows us to test POWDER-RDZ under real-world RF channel conditions but doesn't allow us to show the effects of interference from incumbents if the center frequency of operation is poorly chosen, as doing so would generate interference toward licensed transmitters. The controlled-RF variant allows us to show the effect of operating without using the spectrum intelligence provided by OPENZMS, and attempt to transmit in the presence of real-world interference signals without impacting incumbents. As shown in Figure 9, the attenuator matrix setup included two additional NI X310 SDRs with Dell R430 compute nodes (in green). We used one of these SDRs to "play back" an RF interference trace collected in the POWDER-RDZ outdoor setting, and the other to monitor the RF environment within the setup to visualize its functionality. Other than this "visualization" monitoring, all other monitoring and RDZ decision-making in our evaluation involves data from the POWDER outdoor environment. The outdoor evaluation took place weeks after the controlled-RF evaluation, leading to differences in monitor data and existing spectrum grants, and resulting in different spectrum allocations for the 5G test system.

B. End-to-end functional evaluation

In this section, we show representative "snapshots" of results obtained in the end-to-end PEL use-case depicted in

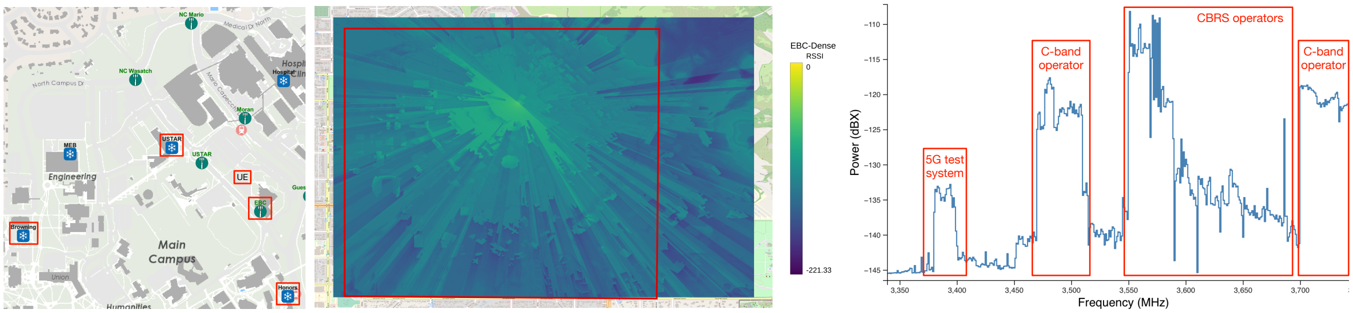

Fig. 12. Results from end-to-end functional evaluation. (a) Partial POWDER map. (b) Check TX power outside RDZ (step #5). (c) RF monitor data with 5G test system (step #6).

Figure 9. Figure 12a shows a part of the POWDER deployment with relevant nodes marked. For the outdoor evaluation, the EBC dense site was used as the gNodeB, and a COTS UE was deployed approximately 125 m from the site, where it attached to the cell and collected link metrics.

To facilitate RF interference playback in the controlled-RF environment, we collected IQ samples using an X310 SDR at the POWDER USTAR rooftop location, with a 200 MHz sample rate and centered at 3.5 GHz. (Step #0 in Figure 9.) This trace showed similar incumbent activity to that seen in Figure 12c in that range, i.e., a C-band operator (the 40 MHz signal from 3.47-3.51 GHz) and a few 20 MHz CBRS operators (from 3.55-3.6 GHz). During the indoor controlled-RF evaluation, this trace was played back to emulate the expected real-world interference. In the outdoor evaluation, similar activity was still present.

In the end-to-end PEL workflow OPENZMS is informed about the PEL spectrum range (3.35-3.6 GHz) (step #1) and it requests monitoring data from an RF monitor (step #2). In our example flow this monitor executed on the POWDER Browning rooftop node for the controlled-RF evaluation, and the Honors rooftop node for the outdoor evaluation. At this point the 5G test setup requests spectrum from OPENZMS to operate at the EBC dense-deployment location in the POWDER outdoor environment with 20 MHz bandwidth in the PEL range (step #3). Figure 12c (excluding the 5G test system activity, which was unoccupied and therefore eventually assigned for the outdoor evaluation) shows a snapshot of the data reported by the monitor and the resulting frequency ranges as determined by the OPENZMS for the outdoor evaluation (step #4).

Before informing the 5G test system about this operating range, OPENZMS performs a geospatial query of propagation simulations via the *dst* service, to ensure that the test transmitter will not interfere outside the RDZ (step #5). Figure 12b shows the result of this query, confirming that operating at the EBC dense-deployment location will not be a problem.

OPENZMS informs the 5G test system to operate with a center frequency at 3.39 GHz for the outdoor evaluation, and 3.45 GHz for the controlled-RF evaluation (step #6). Figure 12c shows the outdoor 5G test system operating at the designated center frequency (and well separated from the incumbents). In addition, we used the controlled-RF environment to *manually* execute the 5G test system at two other center frequencies, i.e., at 3.47GHz and 3.49GHz, representing the scenario where, without the spectral intelligence provided by the OPENZMS, a user might simply pick a center frequency within the PEL-allowed range, thereby generating and receiving varying degrees of interference. Table I shows performance metrics for the controlled-RF evaluation under these different scenarios and the outdoor evaluation with no spectral overlap with incumbent transmitters. Specifically, the table shows the average channel quality indicator (CQI) reported by the srsRAN gNodeB as an indicator of the quality of the wireless channel, as well as the average downlink throughput (measured with iPerf3) as an indicator of application level performance. As expected, the scenarios with overlapping spectrum perform significantly worse than the OPENZMS selected case. The outdoor results are commensurate with the expected path loss and other channel impairments. While we measured the impact of poor frequency selection on the *5G test system* in a controlled-RF environment, in an RDZ such selections will of course similarly impact the incumbent operators, i.e., the exact thing the OPENZMS is preventing.

TABLE I
5G TEST SYSTEM EVALUATION DATA.

Scenario	Center freq. (MHz)	Avg. CQI	Throughput (Mbps)
Outdoor over-the-air:			
OPENZMS selected - no overlap	3390	11.2	20.8
Indoor controlled-RF:			
OPENZMS selected - no overlap	3450	14.1	39.6
Manual - partial overlap	3470	10	10.8
Manual - complete overlap	3490	6.7	4.88

VI. RELATED WORK

To the best of our knowledge, the idea of a national radio dynamic zone (NRDZ) was first conceived by Thomas Kidd [26], as somewhat of an "opposite" to a national radio quiet zone (NRQZ). I.e., while a NRQZ is an area with special rules to protect sensitive receivers inside the zone from "normal" transmitters on the "outside", a NRDZ as conceived by Kidd would protect normal receivers *outside* the zone from special transmitters *inside* the zone. In the US, funding from the NSF's *Spectrum Innovation Initiative* has established a community of researchers that are investigating various aspects of RDZs. A recent paper by Mariya Zheleva et. al. [10] serves as a current community consensus "snapshot" of what an RDZ is, the need for RDZ(s), the features, capabilities and challenges associated with realizing an RDZ, as well as the key required functional components. In this more recent work, the RDZ concept has been generalized to "regional-scale experimental testbeds that can enable spectrum research into – and provide real-world validation of – the coexistence of disparate active and passive [spectrum using] technologies". While clearly not regional-scale, POWDER-RDZ aligns with the vision described in Zheleva's paper, and we believe is the first practical (prototype) realization of an RDZ.

The (N)RDZ concept has also been explored in the context of autonomous aerial and ground spectrum sensors in the AERPAW testbed [3]. They present early results related to spectrum compliance monitoring in the AERPAW platform. The AERPAW team has also published more recent work on formulating an approach for out-of-zone signal detection [4]. These efforts are complementary to our POWDER-RDZ efforts, and we envision detection systems like these to map to the OPENZMS Spectrum Monitor interface to provide spectral intelligence to the ZMS.

Another related effort is the NRDZ project being conducted by the National Radio Astronomy Observatory [5]. The focus of their project is on developing a high-fidelity advanced spectrum monitoring (ASM) device and exploring RDZ concepts in the context of radio astronomy use cases. Our POWDER-RDZ/OPENZMS efforts are focused on exploring a broad range of use-cases and specifically on creating building blocks for the realization of an eventual NRDZ.

Our work builds on various related efforts within our own group. Notably, the POWDER-RDZ DST concept builds upon our earlier digital spectrum twin efforts [6], [7]. Our TIREM-based RF propagation service benefits from ML-based propagation modeling enhancements [8]. We have also performed an initial exploration of automating mobility management of test transmitters in an RDZ [9].

Our work is related to other spectrum sharing approaches, notably the CBRS ecosystem [27], [28]. In that context the CBRS spectrum access system (SAS) bears resemblance to an RDZ ZMS. However, the SAS ecosystem is a single-purpose system built for a well-defined use case: to share the CBRS band in the United States. As such, the CBRS SAS ecosystem lacks the interfaces and mechanisms necessary to ensure deconflicted, parallel spectrum use; offers limited visibility into competing use within the region; lacks a model for spectrum agility and policy, e.g., spectrum providers cannot delegate a new band for the SAS to manage and use; and aims to support an emerging business model, as opposed to enabling broader spectrum sharing testing and exploration in an RDZ. For these reasons, in the OPENZMS architecture, we deliberately chose to develop a novel set of interfaces (i.e., ZeAL) and services that will support a wide variety of spectrum-sharing use cases in varied RDZ realizations.

Another related spectrum management effort is a framework [29] built on the COSMOS testbed [30] in which collaborative wireless networks exchange IEEE Spectrum Consumption Models [22] to coordinate shared spectrum use, assisted by service and monitoring planes. Our work takes a regionally-centralized approach so that RDZ participants may only consume and observe spectrum via authority granted through central mechanism, policy, and optimization decisions; and so that the ZMS can act authoritatively to remove harmful interference when violations are detected. We are investigating the application of SCMs to OPENZMS scheduling algorithms and inter-RDZ cooperation at zone boundaries.

VII. DISCUSSION AND CONCLUSION

In this paper we described our work on POWDER-RDZ in which we presented the first end-to-end realization of a radio dynamic zone (RDZ) architecture and implementation. We illustrated the generality of our design by describing how POWDER-RDZ enables a variety of spectrum sharing use cases. We presented the design and implementation of OPENZMS, an open source zone management system (ZMS), and showed how its cloud-native modular realization eases the integration of different ZMS services and components, including commercial tools. We presented an illustrative end-to-end functional evaluation of our work on the POWDER platform, by showing the importance of ZMS-based spectrum intelligence in managing spectrum in a mobile and wireless testbed.

While "running code" in an operational testbed is a significant step forward, we realize that there is a long way to go towards the ultimate goal of a national radio dynamic zone (NRDZ). We also realize that reaching that goal will require collaboration and cooperation of the broader RDZ community. Thus, it was important for us to design the POWDER-RDZ with the objective of supporting RDZ facilities other than *our* testbed (POWDER), and specifically to make OPENZMS open source. We anticipate that the OPENZMS framework will enable others to explore RDZ concepts without the need to develop everything from the ground-up. Further, as we have shown in this paper, the flexibility of the POWDER platform serves as an ideal sandbox for RDZ related exploration.

We will continue to build on the work described here. Specifically, we will continue to develop, test and validate end-to-end RDZ use-cases and the OPENZMS components that enable them. This includes using OPENZMS as the authoritative spectrum manager in POWDER, evaluating the robustness and scalability of OPENZMS, evaluating the fidelity and accuracy of POWDER-RDZ monitoring systems, and more.

REFERENCES

[1] D. Johnson, D. Maas, S. Tadik, A. Orange, L. Stoller, K. Webb, M. B. I. Awan, J. Bills, M. Gomez, A. Sarbhai, G. Durgin, S. Kasera, N. Patwari, D. Schurig, and J. Van der Merwe, "Powder-rdz: Prototyping a radio dynamic zone using the powder platform," in *2024 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2024, pp. 69–78.

[2] M. Parvini, A. H. Zarif, A. Nouruzi, N. Mokari, M. R. Javan, B. Abbasi, A. Ghasemi, and H. Yanikomeroglu, "A comprehensive survey of spectrum sharing schemes from a standardization and implementation perspective," 2022.

[3] S. J. Maeng, I. Güvenç, M. L. Sichitiu, B. Floyd, R. Dutta, T. Zajkowski, O. Ozdemir, and M. Mushi, "National radio dynamic zone concept with autonomous aerial and ground spectrum sensors," in *2022 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2022, pp. 687–692.

[4] S. Joon Maeng, I. Güvenç, M. L. Sichitiu, and O. Ozdemir, "Out-of-zone signal leakage sensing in radio dynamic zones," in *ICC 2022 - IEEE International Conference on Communications*, 2022, pp. 5579–5584.

[5] C. De Pree, V. Bogan, D. Bordenave, K. Shoemaker, and A. Beasley, "The NRAO national radio dynamic zone (NRDZ) project: Current status," in *American Astronomical Society Meeting Abstracts*, vol. 55, no. 2, 2023, pp. 365–05.

[6] 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," in *2022 IEEE 2nd International Conference on Digital Twins and Parallel Intelligence (DTPI)*, 2022, pp. 1–6.

[7] S. Tadik, K. M. Graves, M. A. Varner, C. R. Anderson, D. Johnson, S. K. Kasera, N. Patwari, J. Van der Merwe, and G. D. Durgin, "Digital spectrum twins for enhanced spectrum sharing and other radio applications," *IEEE Journal of Radio Frequency Identification*, pp. 1–1, 2023.

[8] S. Tadik, M. A. Varner, F. Mitchell, and G. D. Durgin, "Augmented rf propagation modeling," *IEEE Journal of Radio Frequency Identification*, vol. 7, pp. 211–221, 2023.

[9] A. Gottipati and J. Van der Merwe, "FlexRDZ: Autonomous mobility management for radio dynamic zones," in *IEEE Future Networks World Forum*, 2023.

[10] 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*, vol. 61, no. 6, pp. 156–162, 2023.

[11] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, A. Ghosh, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas, C. Martin, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L. B. Stoller, A. Todd, J. Van der Merwe, N. Viswanathan, K. Webb, and G. Wong, "POWDER: Platform for open wireless data-driven experimental research," *Computer Networks*, vol. 197, October 2021.

[12] Flux Research Group, "OpenZMS," https://www.openzms.net/, 2025.

[13] A. Aradhya, D. DeBoer, O. Collaco, W. Farah, C. Forrester, K. Gifford, D. Johnson, S. Llosa, A. Pollak, B. Pearce, M. Ruzindana, A. Sarbhai, B. Stover, S. Tschimben, J. Van der Merwe, K. Webb, and G. Weihe, "An experimental test-bed for investigating spectrum sharing strategies between passive and active users at a prototype National Radio Dynamic Zone (NRDZ)," in *2025 National Radio Science Meeting*, 2025.

[14] A. Sarbhai, D. Johnson, K. Webb, L. Stoller, O. Collaco, A. Orange, S. Zachary, B. Pearce, S. Tadik, S. Llosa, A. Aradhya, A. Pollak, W. Farah, J. Van der Merwe, K. Gifford, N. Patwari, G. Durgin, D. DeBoer, and S. Kasera, "Reactive interference management for radio astronomy in Radio Dynamic Zones using ASTRA," in *2025 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, May 2025.

[15] B. C. Terry, A. Orange, N. Patwari, S. K. Kasera, and J. Van der Merwe, "Spectrum monitoring and source separation in POWDER," in *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, September 2020, pp. 25–32.

[16] S. J. Maeng and O. Ozdemir, "Kriging-based 3-d spectrum awareness for radio dynamic zones using aerial spectrum sensors," 2023, https://doi.org/10.48550/arXiv.2307.06310.

[17] G. D. Durgin, M. A. Varner, M. A. Weitnauer, J. Cressler, M. M. Tentzeris, A. Zajic, S. Zeinolabedinzadeh, R. Zekavat, K. Pahlavan, U. Guler, and K. Van der Merwe, "Digital spectrum twinning and the role of rfid and backscatter communications in spectral sensing," in *2021 IEEE International Conference on RFID Technology and Applications (RFID-TA)*, 2021, pp. 89–92.

[18] U.S. Department of Commerce, "Quantitative Assessments of Spectrum Usage," https://www.ntia.gov/report/2016/quantitative-assessments-spectrum-usage, November 2016.

[19] F. H. Sanders, R. L. Sole, J. E. Carroll, G. S. Secrest, and T. L. Allmon, "Analysis and Resolution of RF Interference to Radars Operating in the Band 2700–2900 MHz from Broadband Communication Transmitters," https://www.ntia.gov/report/2012/analysis-and-resolution-rf-interference-radars-operating-band-2700-2900-mhz-broadband, October 2012.

[20] Z. Khan, J. J. Lehtomaki, R. Vuohtoniemi, E. Hossain, and L. A. Dasilva, "On opportunistic spectrum access in radar bands: Lessons learned from measurement of weather radar signals," *IEEE Wireless Communications*, vol. 23, no. 3, pp. 40–48, 2016.

[21] National Radio Astronomy Observatory, "National Radio Quiet Zone (NRQZ)," https://info.nrao.edu/do/spectrum-management/national-radio-quiet-zone-nrqz-1, 2024.

[22] C. E. C. Bastidas, J. A. Stine, A. Rennier, M. Sherman, A. Lackpour, M. M. Kokar, and R. Schrage, "IEEE 1900.5.2: Standard method for modeling spectrum consumption: Introduction and use cases," *IEEE Communications Standards Magazine*, vol. 2, no. 4, pp. 49–55, 2018.

[23] D. R. DeBoer, S. L. Cruz-Pol, M. M. Davis, T. Gaier, P. Feldman, J. Judge, K. I. Kellermann, D. G. Long, L. Magnani, D. S. McKague, T. J. Pearson, A. E. E. Rogers, S. C. Reising, G. Taylor, A. R. Thompson, and L. van Zee, "Radio frequencies: Policy and management," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 51, no. 10, pp. 4918–4927, 2013.

[24] D. Eppink and W. Kuebler, *TIREM/SEM Handbook*. IIT Research Institute, Tech. Rep., Sept 1986.

[25] "NTIA/ITS SCOS Sensor," https://github.com/NTIA/scos-sensor.

[26] T. Kidd, "National radio quiet and dynamic zones," *CHIPS, the Department of the Navy's Information Technology Magazone*, April–June 2018.

[27] M. M. Sohul, M. Yao, T. Yang, and J. H. Reed, "Spectrum access system for the citizen broadband radio service," *IEEE Communications Magazine*, vol. 53, no. 7, pp. 18–25, July 2015.

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

[29] D. Stojadinovic, P. Netalkar, C. E. C. Bastidas, I. Kadota, G. Zussman, I. Seskar, and D. Raychaudhuri, "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*, 2021, p. 77–84.

[30] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, H. Krishnaswamy, S. Maheshwari, P. Skrimponis, and C. Gutterman, "Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking (MobiCom)*, September 2020.