

FlexRDZ: Autonomous Mobility Management for Radio Dynamic Zones

Aashish Gottipati
University of Texas Austin
agottipati@utexas.edu

Jacobus Van der Merwe
University of Utah
kobus@cs.utah.edu

Abstract—FlexRDZ is an online, autonomous manager for radio dynamic zones (RDZ) that seeks to enable the safe operation of RDZs through real-time control of deployed test transmitters. FlexRDZ leverages Hierarchical Task Networks and digital twin modeling to plan and resolve RDZ violations in near real-time. We prototype FlexRDZ with GTPyhop and the Terrain Integrated Rough Earth Model (TIREM). We deploy and evaluate FlexRDZ within a simulated version of the Salt Lake City POWDER testbed, a potential urban RDZ environment. Our simulations show that FlexRDZ enables up to a 20 dBm reduction in mobile interference and a significant reduction in the total power of leaked transmissions while preserving the overall communication capabilities and uptime of test transmitters. To our knowledge, FlexRDZ is the first autonomous system for RDZ management.

Index Terms—Radio Dynamic Zone, Mobility Management, Network Control, AI Planning

I. INTRODUCTION

Radio technology has progressed tremendously in recent years, as seen by the proliferation of software-defined radios, 5G networks, and the advent of terra hertz technology. To continue to drive innovation, researchers require access to adequate radio test facilities, enabling them to develop, benchmark, and validate their test transmitters without worrying about potential impacts on nearby radio infrastructure. In accordance, the notion of developing a universal radio test facility led to the proposed concept of a National Radio Dynamic Zone (NRDZ) [1]. Radio Dynamic Zones (RDZ) are isolated from the outside world spectrally, enabling the deployment of new test transmitters. These zones prevent internal transmissions from escaping, freeing test operators from the responsibility of worrying about interfering with nearby infrastructure.

While an RDZ seeks to coexist with existing infrastructure, an RDZ also seeks to enable users to deploy and test new transmitter technology. In many cases, these test transmitters may be inimical towards nearby infrastructure, motivating the RDZ operator to employ oversight over test transmitters. While operators may be able to exercise oversight over stationary test transmitters due to their static position, mobile transmitters further tax an RDZ operator’s ability to supervise the RDZ properly. For example, many spectrum allocation techniques assume that test transmitters are stationary, failing to model

This material is based upon work supported by the National Science Foundation under Grant Numbers 1827940 and 2232463.

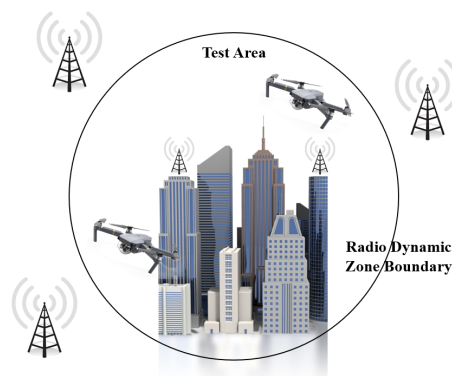


Fig. 1: Mobile RDZ Deployment.

mobile entities accurately [2], [3]; however, in an RDZ, test transmitters may or may not be stationary (e.g., naval and air radar systems). In these cases, utilizing traditional techniques may not be sufficient to prevent test transmitters from impacting nearby infrastructure. We, therefore, choose to explore RDZ mobility based on the increase in management complexity and the assumption that many future test transmitters will be mobile [4].

To simplify our model of transmitter mobility, like others [5], we assume a complete control framework, meaning that all transmitters cede control to a management entity, enabling a higher level of flexibility within the RDZ (see Figure 1). For example, the area in which a mobile test transmitter can operate becomes more flexible through techniques such as spectrum sharing and real-time transmit power adjustments, since leakage and interference can be handled dynamically.

To realize a more flexible RDZ, we propose FlexRDZ. FlexRDZ is a closed-loop, autonomous RDZ manager that seeks to enable the safe operation of an RDZ through real-time control of deployed test transmitters. FlexRDZ is primarily defined by its Hierarchical Task Network (HTN), a graph-based planning approach, and its RDZ digital twin. These two components enable FlexRDZ to dynamically model the RDZ environment and generate a plan to maintain the “health” of the RDZ in the event a compromise occurs, e.g., internal transmissions are detected beyond the RDZ boundary. FlexRDZ utilizes its programmable control framework to execute its generated plan in order to preserve the health of the RDZ. To validate our implementation of FlexRDZ, we deploy and evaluate FlexRDZ within a simulated version of the

Salt Lake City POWDER testbed [6], a potential urban RDZ environment. Our simulations show that FlexRDZ enables up to a 20 dBm reduction in mobile interference and a significant reduction in the total power of leaked transmissions while preserving the overall communication capabilities and uptime of test transmitters.

To realize an autonomous RDZ management system, we make the following contributions:

- The design of FlexRDZ, a near real-time, autonomous RDZ manager that aims to maintain an RDZ in an online fashion.
- A prototype implementation of FlexRDZ that reduces RDZ leakage and preserves transmitter communication.
- Quantitative results that demonstrate the benefits of FlexRDZ by highlighting its performance and use cases.

The remainder of this paper is laid out as follows. Section II discusses key background information related to planning. Section III surveys related work in the space of dynamic spectrum environments. Section IV discusses the design and implementation decisions pertaining to FlexRDZ. Section V highlights FlexRDZ’s capability within the context of autonomous RDZ management. Lastly, section VI comments on the limitations of our approach, proposes future areas of research, and provides concluding remarks.

II. BACKGROUND

AI planning refers to the problem of finding a set of actions that, if executed, will transition the environment from an initial state to a goal state. More formally, given a set of states \mathcal{S} and actions \mathcal{A} , we require a mapping between states and actions $\pi : \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ where the actions push the agent in the direction of the goal state. Furthermore, the transition model defines the probability of a taking an action given the current state: $\pi(a, s) = Pr(a_t = a | s_t = s)$. We call the combination of the mapping and transition model the policy, i.e., a distribution of actions conditioned on the current environment state. This problem is canonically known as a Markov decision process (MDP).

Modern approaches to AI planning utilize a mix of deep reinforcement learning (RL), statistical techniques, and integer programming [7] with the former becoming the dominant planning approach given its application in robot planning and manipulation [8]. At a high level, RL seeks to learn a policy for acting in an environment where the environment is represented by an MDP. In contrast to other deep learning disciplines, RL provides an agent with a reward function, which aims to guide the agent toward favorable states in the state space. Upon taking an action, the agent receives a reward according to a predefined reward function, which provides a metric for gauging the performance of an action. Long-term rewards are penalized via a discount rate, which prioritizes events in the immediate future over the distant future. In essence, the agent seeks to learn a policy by maximizing the returned reward, which, in turn, guides the agents toward favorable states.

However, there are two issues with this formulation. One, while deep learning-based techniques perform well in practice, many of these techniques suffer from the lack of ex-

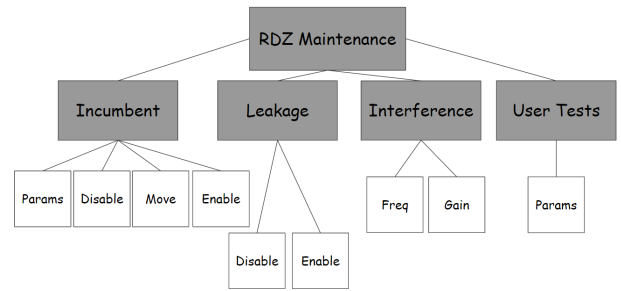


Fig. 2: FlexRDZ’s AI Planner.

RDZ maintenance is divided into four subtasks. For example, to satisfy an incumbent’s request to adjust the area of the RDZ, the HTN returns the following atomic tasks: update the RDZ parameters (e.g., boundary coordinates), disable non-compliant transmitters, move mobile transmitters to the new zone, and enable all compliant transmitters.

plainability [9]; hence, robust safeguards are required when deploying in environments with risky tail-end events. Two, in the real world, state observations are noisy, i.e., states are not fully observable. In this case, the problem becomes a partially-observable MDPs (POMDP), a well-studied problem in robotics [10], but requires more assumptions for safe deployments.

In the context of RDZs, transparency and safety are crucial since an incorrect plan could broadly impact surrounding telecommunications services [5]. Given this, we focus on a classical planning approach known as HTNs [11]. HTNs are dependency graphs that are constructed with domain knowledge, meaning that the policy π is hand-crafted rather than learned. In contrast to other modern methods, HTNs are human interpretable, offering more transparency into understanding why certain decisions or plans were generated [12]. HTNs follow a hierarchical structure where the root node corresponds to the initial state or the start task. The starting task is then recursively broken down into subtasks until an atomic action can be executed (see Figure 2). The root node connects to its descendants via edges, which represent dependencies. To transition from the initial state (root node) to the goal state, the root’s dependencies must be satisfied by carrying out the corresponding atomic actions. When called upon, an HTN generates a sequence of atomic actions to satisfy an overarching goal. We seek to leverage HTNs to dynamically generate plans to mitigate RDZ violations and autonomously maintain the RDZ environment.

III. RELATED WORK

A. Citizens Broadband Radio Service

Interference Management. To mitigate interference within CBRS, the Wireless Innovation Forum (WINNFoRM) suggest three different graph-based approaches to spectrum allocation [13]–[15]. Approaches one and two utilize a graph coloring algorithm to allocate non-overlapping spectrum to nearby transmitters, while approach three employs a recursive clustering algorithm. Gao et al. evaluate all three approaches using a suite of propagation models and GIS map data of Virginia Beach and San Diego [2], [3], showing that all

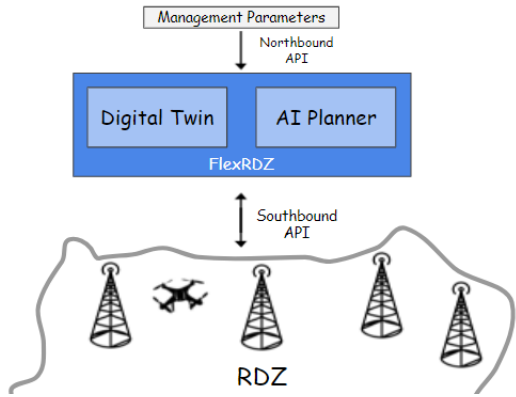


Fig. 3: FlexRDZ Architecture.

three approaches are indeed effective at reducing potential interference. In contrast to WINNFORM, Abbass et al. explore the application of Q-learning for spectrum allocation in CBRS [16]. Specifically, the authors investigate opening up idle access priority access license (PAL) channels to general authorized access (GAA) users. Numerical results demonstrate improvements in spectrum utilization and data rate per unit cost; however, real-world evaluations are necessary.

Although, while the WINNFORM approaches excel at mitigating interference, there are many drawbacks in practice. For example, since these approaches utilize the overlap in deployment area as an estimate for interference between transmitters, the estimates of interference may underestimate or overestimate the actual interference based on the deployed environment [17]. In addition, graph coloring is an NP-complete problem, and, as such, most algorithms are subject to non-polynomial growth, meaning that, with dense transmitter deployments, the computation layer becomes the bottleneck [17]. Lastly, these interference management techniques yield their highest performance when transmit nodes are stationary, which may not hold in an RDZ environment where test transmitters may be mobile, e.g., drones. Thus, to better model dynamic RDZ environments, we supplement our approximations with simulations, leading to the addition of FlexRDZ’s digital twin. We note that we are not solving the interference management problem within CBRS; instead, we build upon techniques used in CBRS and generalize them to RDZs.

Mobility Management. The most relevant work to handling mobility within an RDZ-like environment is related to detecting naval incumbents in CBRS. To detect naval incumbents, the National Telecommunications and Information Administration (NTIA) proposed the use of an environmental sensing capability (ESC) network [18]. An ESC network comprises multiple sensors employed to detect an incumbent’s presence and trigger protective measures upon detection. Nguyen et al. formulated the ESC deployment problem as a set cover problem to compute the minimum number of sensors to cover an area of interest while minimizing the overlapping area between sensors, since overlapping areas may lead to false positives [19]. As opposed to preemptively turning off nearby equipment upon detecting an incumbent, Kang et al. offer a different approach: using a management entity to oversee

interactions between transmitters [20]. Instead of proactively pausing communications in inflection areas (e.g., where naval incumbents are detected), Kang suggests employing more dynamic techniques such as spectrum sharing and virtualization. In our formulation of an RDZ, we implicitly cover both of these cases; however, we emphasize that we are not bound by the same assumptions as CBRS.

B. Radio Dynamic Zones

We now cover related work on RDZs. Maeng et al. propose a spectrum monitoring approach for out-of-zone signal leakage detection and explore spatial correlation-based estimation techniques for signal prediction [21]. The authors constrain their RDZ formulation geographically and assume sensor nodes that cover and monitor the boundary of the RDZ. Through simulation, the authors demonstrate that their spatial correlation-based algorithm enables a larger RDZ radius with sparsely deployed sensor nodes in comparison to propagation-loss techniques. In addition, Maeng et al. present an RDZ concept that relies on both autonomous aerial and ground sensor nodes for radio environment monitoring, enabling real-time radio environment maps of all relevant frequencies and locations [22]. Lastly, [5] discusses key challenges for real RDZ deployments and details a Zone Management Engine for RDZ supervision. Specifically, the Zone Management Engine is composed of a decision engine and three subsystems: spectrum, experiment, and policy management systems. In concert, these three subsystems provide the decision engine with experiment information, spectrum policy rules, user information, and resource allocations to enable dynamic real-time coordination of systems within the zone coupled with user interference protection. In other words, FlexRDZ is a realized prototype of the previously proposed Zone Management Engine [5].

We diverge from previous work in fundamental ways. First, we primarily focus on the problem of streamlining RDZ maintenance through dynamic, flexible planning practices rather than spectrum monitoring. Second, we assume infrastructure for spectrum monitoring and leverage real-time environment maps to estimate and plan for future transmitter behavior. Lastly, we approach RDZ maintenance through a systems perspective, prototyping and evaluating a system for real-time RDZ maintenance and control.

IV. DESIGN AND PROTOTYPE

A. Overview

FlexRDZ is an operational tool that seeks to mitigate RDZ violations through swift resolution and maintain the environment in an autonomous and online fashion. RDZ violations encompass situations that compromise the state of the RDZ (e.g., transmissions are detected beyond the environment boundary or the RDZ fails to accommodate local infrastructure) and cases where user test objectives are conceded.

As seen in Figure 3, FlexRDZ receives management parameters via its northbound interface. These management parameters encompass RDZ maintenance variables (e.g., bound-

ary parameters, leakage thresholds, and interference thresholds) and user test objectives (e.g., reserved areas and frequencies). Upon failing to maintain these parameters, `FlexRDZ` utilizes its planner to generate a plan and resolve the situation dynamically.

Lastly, `FlexRDZ` functions in an online, autonomous fashion. These characteristics are a byproduct of `FlexRDZ`'s southbound interface, which routinely pulls the state of the RDZ from the environment and enables direct control over test transmitters. Illustrated in [22], real-time RDZ state can be monitored via ground-based and aerial sensing infrastructure. By retrieving real-time updates, `FlexRDZ` can, in near real-time, generate solutions and resolve environmental issues with its planner. The combination of `FlexRDZ`'s southbound and northbound interface enables it to function within a closed-loop, online manner. The architecture of `FlexRDZ` is in-line with previously envisioned RDZ supervisors, e.g., a Zone Management Engine [5].

B. Intelligent Control

While planning may deceptively appear straightforward, it is important to note that RDZ environments quickly become too convoluted to manage due to the sheer number of observable environment states over time, making it infeasible to iterate over all potential control solutions. Although this problem is challenging, a great deal of related work has sought to overcome similar complexities in adjacent problems by efficiently searching over the planning space through the use of symbolic methods, artificial intelligence (AI), and various optimization techniques [7]. Like other planning techniques, symbolic methods such as HTNs are utilized to search over the planning space efficiently [11]. We leverage HTNs for their increased decision transparency and reliability—two characteristics essential for maintaining an urban RDZ. We note that the design of `FlexRDZ` does not preclude the use of other planning techniques.

C. Digital Twin Modeling

In combination with GIS map data and transmitter parameters, `FlexRDZ` leverages an internal radio-frequency (RF) model to “digitize” the RDZ environment and generate a radio-environment map in real-time. `FlexRDZ`'s RF model serves to model the RF interactions among transmitters, estimate the coverage for a transmitter, and derive RDZ-specific, key performance indicators. Note that the design of `FlexRDZ` does not rely on any one modeling technique and can generalize to various methods such as path-loss or AI-based. By encoding the environment, `FlexRDZ` can estimate potential states of the RDZ by simulating future actions through its RF model, enabling more robust control updates.

D. Implementation

To prototype `FlexRDZ`, we utilized the Terrain Integrated Rough Earth Model (TIREM) propagation model [23]. TIREM is a set of physics-based algorithms used to estimate the coverage for mobile land radios and point-to-point distances. We




Fig. 4: Simulated RF Map vs. Generated RF Map.

The path-loss contour of our simulation tool appears on the left, while our model’s estimate appears on the right. Note that the estimator closely matches the simulation; however, residuals are still present, e.g., the region near (300, 250).

opted to utilize TIREM as it is the standard propagation model utilized by the United States government. However, through preliminary evaluations, we observed that modeling the entire RDZ environment via TIREM can be extremely costly. In our case, we observed estimation times of approximately 30 seconds for large areas, e.g., the downtown Salt Lake City area, preventing real-time control. Thus, similar to the idea proposed in [24], to alleviate this bottleneck, we trained a neural network to approximate the RF maps produced by TIREM, cutting inference time to approximately 70 ms, and enabling real-time control. We opted to utilize a fusion-based network for their ability to learn more intricate feature representations [25].

As input, the RF model accepts GIS map data that corresponds to the RDZ deployment area and the transmit parameters of the mobile transmitter. The GIS map data is processed through a series of ResNet blocks [26], while transmit parameters are encoded via fully-connected layers. The two embeddings are concatenated and used as input to another fully-connected layer to generate a fusion embedding. From this embedding, the model learns a decoding function to output the predicted RF map.

To train our model, we first generated 10000 training instances. A training instance consisted of the GIS map data, the mobile transmit parameters, and the ground truth TIREM RF map of the downtown Salt Lake City area. After generating the dataset, we randomly split the dataset into train, validation, and test set splits following an 80/10/10 ratio. We empirically chose the fusion embedding size of 2048 as it resulted in the fast training times and an overall storage footprint of approximately 0.8 Mb. Furthermore, we trained our model for 150 epochs via Adam with an initial learning rate of 0.001, using mean-squared error (MSE) as our loss function, until we observed model convergence on our validation set. We verified that our model generates reasonable RF maps, producing an average empirical error of -0.045 dBm per cell (see Figure 4). We emphasize that our learned model is unlikely to generalize to new settings as our training data is drawn from the Salt Lake City area. We leave learning a general model as an area for future work.

Furthermore, to realize `FlexRDZ`'s HTN planning component, we leveraged the open-source HTN planner known as `GTPyhop` [27]. `GTPyhop` provides a basic framework for constructing HTNs, enabling users to specify objectives, sub-tasks, and action primitives. `GTPyhop` utilizes a modified depth-first search to search over the planning space efficiently. Nau et al. provide an in-depth analysis of the `GTPyhop` planning algorithm [27]. Moreover, we constructed `FlexRDZ`'s HTN planner (see Figure 2) based on a small set of atomic primitives: idle, disable transmitter, enable transmitter, and round-robin frequency reassignment, which are routinely utilized to manage dynamic spectrum environments to foster generalizability. Like other works [28], we intertwined the digital twin with the HTN to produce a more robust management policy [28]. For example, suppose that a transmitter has been disabled due to leaked signals. To re-enable the transmitter, `FlexRDZ` must ensure that the newly enabled transmitter does not violate the terms of the RDZ. By leveraging its digital twin, `FlexRDZ` can simulate the impact of its future decision; hence, the transmitter is only re-enabled when its digital counterpart does not leak, which leads to a more compliance-oriented policy.

V. EVALUATION

A. Setup

We utilized our internal RDZ simulation tool to simulate an RDZ deployed within the POWDER testbed. We opted to simulate the POWDER testbed environment as POWDER may provide RDZ functionality to researchers in the future. In addition, the testbed is situated near the downtown Salt Lake City urban area, which directly reflects the urban RDZ deployment scenario laid out in Section I.

The simulated environment consisted of the Salt Lake City area projected into a 400×400 matrix. Each entry in the projected matrix corresponded to the summation of predicted TIREM signal strengths of the deployed transmitters at that given point. Additionally, we leveraged GIS map data of the Salt Lake City area and approximations of the POWDER testbed's campus buildings to model the terrain of the POWDER testbed. The parameters of the digital twin are fixed; hence, we leave dynamic digital twin modeling—utilizing real-time updates to adjust the parameters of the digital twin—for future work. The evaluation seeks to evaluate the performance of `FlexRDZ`'s AI planner (see Figure 2) rather than the performance of `FlexRDZ`'s digital twin. Note that we present a suite of planning-based techniques and seek to demonstrate the efficacy of `FlexRDZ` to bolster RDZ integrity against non-management based environments, e.g., an environment with no planning agent.

To evaluate our HTN planner, we compare with the following planning approaches.

- 1) *Stochastic HTN* The stochastic HTN generates an identical plan to our HTN implementation; however, the agent executes an action with probability $1 - \epsilon$. We set ϵ to 0.1, 0.2, 0.3 (1, 2, and 3 respectively).

- 2) *Model-free RL*. An agent trained via Proximal Policy Optimization (PPO), a state-of-the-art RL algorithm.
- 3) *Random*. The agent adopts a random policy and executes an action at random.
- 4) *Naive*. No planning agent. In this scenario, the RDZ operator adopts a trust-based policy (users will seek to not violate the terms of the RDZ); hence, no control framework is utilized. This is the primary baseline for comparison.

Furthermore, the simulated RDZ was defined by a rigid geographical boundary, which specified where signals above a given power must not be detected, mimicking our desired urban RDZ deployment. The deployment consisted of 10 fixed endpoints deployed at a height of 1.8m with an antenna gain of -2 dBi, 9 rooftop stationary test transmitters deployed at a height between 20m and 40m with an antenna gain of 4.9 dBi, and 8 densely deployed, stationary test transmitters placed at a height of 8m with an antenna gain of 4.9 dBi. The deployed transmitters were evenly split among the frequencies of 3600, 3610, and 3620 MHz. The placement and parameters of transmitters were selected to mimic the current POWDER testbed deployment. A single mobile transmitter was deployed at a height of 30m with an antenna gain of 4.9 dBi, on one of the previously mentioned frequencies, depending on the simulated evaluation. We justify our limited evaluation with only one mobile node since in an early stage POWDER RDZ deployment, we expect the number of mobile transmitters to be constrained by the POWDER testbed deployment area, the density of existing stationary transmitters, and the restricted set of operating parameters.

To measure the performance of our system, we track the following metrics across trials.

- 1) *Leakage Points*. The number of locations outside the RDZ boundary with a value above the designated RDZ power threshold.
- 2) *Total Strength of Leaked Signals*. The total power of induced signals outside the RDZ boundary with a value above the designated RDZ leakage threshold.
- 3) *Induced Interference*. The total mobile interference observed within the RDZ boundary for a given trial.
- 4) *Mobile Signal-to-Interference-plus-Noise Ratio*. The observed SINR during each evaluation step.
- 5) *Mobile Transmitter Uptime*. The percentage of valid time steps in which the transmitter is active. A time step is considered valid if the mobile transmitter is not violating the policies defined by the RDZ operator (e.g., mobile transmissions are limited to the area of the RDZ).

B. RL Training Procedure

To benchmark our planner, we leverage Open AI Stable Baselines and Open AI Gym to train an agent for RDZ maintenance. We adopt the following RL formulation.

State. A state observation corresponds to a 2048 dimensional fusion vector that encodes information about the current state of the simulated RDZ environment as well as the mobile

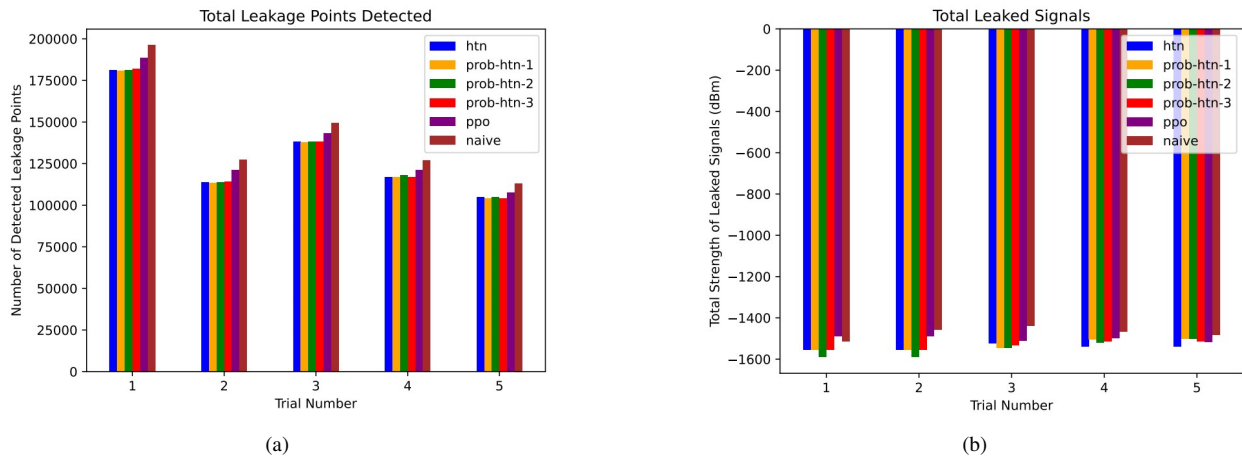


Fig. 5: Observed Mobile Leakage. A smaller number indicates less RDZ leakage, while a larger number indicates more leakage.

transmit parameters. We leverage the fusion backbone of our RF model to encode the state of the simulation.

Action. A discrete value between 0 and 3. We map each action of our HTN to a number between 0 and 3; specifically, 0 corresponds to idle, 1 disable transmitter, 2 enable transmitter, and 3 round robin frequency assignment.

Reward. Our normalized reward is defined per time step as

$$10 \times U + \frac{S}{30} - \frac{I}{I_T} - \frac{P}{A} - \frac{L}{L_T} \quad (1)$$

where U corresponds to the number of time steps where the mobile transmitter is enabled, S is the SINR of the mobile transmitter in dB, I is the induced mobile interference in dBm, I_T is the interference threshold in dBm, P is the number of leakage points detected, A is the total area of the RDZ in units, L is the total power of leaked signals in dBm, and L_T is RDZ power threshold for leaked signals in dBm. While simple, our reward function encourages uptime and open communication channels, while minimizing mobile interference and leakage, all of which are critical to maintaining an RDZ. Note that we clip each reward component by an empirically chosen constant to better shape the reward function and reduce asymmetries between goals.

We represent the learned policy with a multi-layer perceptron (MLP). We trained our agent in simulation for 8000 episodes via PPO [29], a state-of-the-art on-policy training approach, with each episode consisting of 10 time steps. After each time step, we apply the agent’s action where the resulting reward is calculated according to equation 1 until policy convergence is observed.

C. Mobile Leakage

For this evaluation, we set the RDZ leakage threshold to -95 dBm, indicating that if any signals were detected beyond the RDZ boundary above -95 dBm, then this constituted an RDZ violation, i.e., leakage. Since our problem formulation considered an urban RDZ with nearby existing infrastructure, and signals with a strength of -90 dBm are likely to drown

within the noise floor, signals at -95 dBm are likely to be too weak to impact nearby systems.

The evaluation sought to evaluate `FlexRDZ`’s HTN against other planning approaches in terms of leakage mitigation. The naive approach relies on “good faith” that the mobile transmitter will not violate the trust agreement of the RDZ. The evaluation procedure was as follows. A partition of POWDER nodes (9 nodes) was deployed on 3600 MHz. A single mobile transmitter was then randomly deployed on 3600 MHz. The mobile transmitter was then simulated for 50 discrete time steps, with the transmitter moving 5 units in a random direction at each time step. The number of leakage points detected and the total strength of leaked signals were recorded and summed at each step. We repeated the experiment for 5 trials, reusing the same transmit and movement parameters for all planning methods.

We summarize the results of the evaluation in Figure 5. The results show that the HTN-based planners consistently led to less leakage in comparison to the RL-based and naive planning approaches. We omit the results from the random policy as the transmitter uptime was significantly lower than the other planning approaches (see Table I). The decrease in leakage can be observed by the fact the first four bars in Figure 5 are consistently lower than the last two, which correspond to the PPO-based agent and the naive approach. It is important to note that the stochastic planners outperform `FlexRDZ`’s deterministic planner in some cases, e.g., trials 4 and 5 in terms of leakage points and trials 1-3 in terms of leaked signals. This is likely due to the exploration vs. exploitation trade-off commonly discussed in RL literature [29]. In this case, due to the ϵ hyperparameter, the stochastic planners are able to travel to new states that are not reachable by the deterministic HTN. For example, since the deterministic HTN is reactive and not preemptive, it is likely that these stochastic planners are preemptively disabling a transmitter that is about to leak, leading to enhanced leakage mitigation.

On the other hand, while the planners reduce the total number of observed leakage points in comparison to the




Fig. 6: Observed Mobile Interference. A lower value indicates reduced interference, while a higher value indicates increased interference.

naive approach, the amount of leakage points detected is still relatively high. One explanation is that these trials encompass some adversarial cases in which a transmitter attempts to straddle the boundary of the RDZ, i.e., the transmitter continues to oscillate between leaking and compliant behavior. Since `FlexRDZ` does not record the history of a transmitter, the periodic behavior circumvents `FlexRDZ`'s basic leakage mitigation policy. To alleviate these adversarial cases, one could incorporate a strike policy into their RDZ, i.e., after x amount of strikes, the offending transmitter may be penalized (e.g., disabled for a period of time).

Nonetheless, based on the results, the HTN-based planners outperform both the PPO-based policy and the naive approach, while offering greater transparency than their RL counterparts as we can easily trace the planning path through the HTN. We observe improved performance among the stochastic HTN methods for certain trials, indicating that stochastic behavior can potentially enhance our planning policies. Furthermore, we observe a significant reduction in leakage points (nearly 10% in Trial 1) and leaked signal strength (approximately 100 dBm in Trial 1) across trials. Note that the signal results in Figure 5 are presented in log scale (base 10). Most importantly, we emphasize that `FlexRDZ`'s oversight greatly reduces amount of leakage in comparison to non-agent based environments.

D. Mobile Interference

We set the mobile interference threshold to -70 dBm for this evaluation. Therefore, if the mobile interference exceeded -70 dBm, this constituted an RDZ violation. While weak, signals at -70 dBm are likely to impact nearby transmitters, users, and test experiments. We collected and recorded the aggregate mobile interference observed and the SINR for the mobile transmitter. In this evaluation, we aimed to measure `FlexRDZ`'s ability to preserve communication (i.e., maintaining adequate SINR and reducing mobile interference). The procedure and setup were nearly identical to the previous evaluation; however, two POWDER partitions (18 nodes total) were deployed and split among 3600 and 3610 MHz respectively.

The results of the evaluation are summarized in Figure 6 and Figure 7. We omit the results from the random policy as the transmitter uptime was significantly lower than the other planning approaches (see Table I). Based on the results in Figure 6, the stochastic HTN planning methods significantly outperformed all other planning methods; additionally, all planning methods consistently outperformed the naive approach. The consistent decrease in the interference is apparent by the smaller (more negative) interference values obtained by the various planning methods in comparison to the naive approach. Like in the previous evaluation, the stochastic HTN methods' improved interference policy appears to stem from greater state space exploration, with the agent landing in states that have preemptive properties; however, the benefits of the stochastic approaches over their deterministic counterpart are more blatant in this case. Given this, adopting a small perturbation (e.g., $\epsilon = 0.1$) can lead to a noticeable boost in interference reduction and communication preservation. In essence, based on the results, in most trials, `FlexRDZ` outperforms the naive approach with a decrease in total interference of nearly 20 dBm in some cases.

As for the observed mobile SINR, we observe a small difference between the `FlexRDZ` and the naive case. To highlight the performance differences, we plot the SINR results for the mobile transmitter as a CDF (see Figure 7). We assume that the SINR values are normally distributed according to the sample mean and the sample standard deviation of the collected results. For all approaches, most of the collected SINR measurements of the mobile transmitter are above 25 dB, which indicates that the transmitter's transmissions are strong and communication is preserved. Note that the PPO-based agent outperforms all other planning methods significantly. One reason for this may be due to skewed credit assignment. In this case, the provided reward function may be asymmetrically promoting the mobile SINR above all other metrics; hence, the agent learns a policy that prioritizes the mobile SINR, e.g., always adjusting the mobile transmit frequency to promote the mobile communication channel. Although more evaluations may shed light on this discrepancy, the motivations behind the planning trajectory are opaque, unlike the HTN which is human interpretable. Like the previous evaluations, these results indicate that `FlexRDZ`'s planning-based approaches outperform the naive approach by preserving a higher level of communication for mobile transmitters. We emphasize that `FlexRDZ`'s oversight greatly reduces the magnitude of interference in comparison against non-agent based environments.

	HTN	HTN 1	HTN 2	HTN 3	PPO	Random	Naive
Uptime	1.0	1.0	1.0	1.0	1.0	≈ 0.51	1.0

TABLE I: Average Mobile Uptime across Trials.

VI. CONCLUSION

We present `FlexRDZ`, an autonomous RDZ manager, and validated its design through a proof-of-concept prototype. While promising, further research is necessary. For example, leveraging online samples to update the internal propagation




Fig. 7: Observed Mobile SINR. Larger SINR values correspond to improved communication performance, while smaller values indicate degraded communication capabilities.

model’s parameters would improve FlexRDZ’s generalizability. Additionally, real-world benchmarks are crucial, motivating further research into enhancing testbed infrastructure. Despite its limitations, FlexRDZ represents an intersection between a growing interest in autonomous control, digital twin modeling, and dynamic spectrum environments. We argue that systems such as FlexRDZ will be critical towards realizing an autonomous RDZ amidst the growing sophistication and complexity of wireless environments.

REFERENCES

- [1] T. Kidd, “National radio quiet and dynamic zones,” 2018.
- [2] W. Gao and A. Sahoo, “Performance study of a gaa-gaa coexistence scheme in the cbrs band,” in *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2019, pp. 1–10.
- [3] W. Gao, A. Sahoo, and E. Bradford, “Gaa-gaa coexistence in the cbrs band: Performance evaluation of approach 3,” in *2020 IEEE 45th LCN Symposium on Emerging Topics in Networking (LCN Symposium)*, 2020, pp. 39–47.
- [4] W. Chen, J. Liu, H. Guo, and N. Kato, “Toward robust and intelligent drone swarm: Challenges and future directions,” *IEEE Network*, vol. 34, no. 4, pp. 278–283, 2020.
- [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*, 2023.
- [6] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L. B. Stoller, J. Van der Merwe, K. Webb, and G. Wong, “POWDER: Platform for open wireless data-driven experimental research,” in *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, Sep. 2020.
- [7] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, “Dream to control: Learning behaviors by latent imagination,” *arXiv preprint arXiv:1912.01603*, 2019.
- [8] A. Pagoli, F. Chappelle, J. A. Corrales, Y. Mezouar, and Y. Lapusta, “A soft robotic gripper with an active palm and reconfigurable fingers for fully dexterous in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7706–7713, 2021.
- [9] A. Heuillet, F. Couthouis, and N. Díaz-Rodríguez, “Explainability in deep reinforcement learning,” 2020.
- [10] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [11] Y. Wang, H. Sun, G. Chen, Q. Jia, and B. Yu, “Hierarchical task planning for multiarm robot with multiconstraint,” *Mathematical Problems in Engineering*, vol. 2016, 2016.

- [12] K. Chen, N. S. Srikanth, D. Kent, H. Ravichandar, and S. Chernova, “Learning hierarchical task networks with preferences from unannotated demonstrations,” in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1572–1581. [Online]. Available: <https://proceedings.mlr.press/v155/chen21d.html>
- [13] Wireless Innovation Forum, “Operations for Citizens Broadband Radio Service (CBRS); GAA Spectrum Coordination - Approach 1,” <https://winnf.memberclicks.net/assets/workproducts/Recommendations/WINNF-TR-2004-V1.0.%20GAA%20Spectrum%20Coordination-Approach%201.pdf>, May 2019, online.
- [14] —, “Operations for Citizens Broadband Radio Service (CBRS); GAA Spectrum Coordination - Approach 2,” <https://winnf.memberclicks.net/assets/workproducts/Recommendations/WINNF-TR-2004-V1.0.%20GAA%20Spectrum%20Coordination-Approach%202.pdf>, May 2019, online.
- [15] —, “Operations for Citizens Broadband Radio Service (CBRS); GAA Spectrum Coordination - Approach 3,” <https://winnf.memberclicks.net/assets/workproducts/Recommendations/WINNF-TR-2004-V1.0.%20GAA%20Spectrum%20Coordination-Approach%203.pdf>, May 2019, online.
- [16] W. Abbass, R. Hussain, J. Frnda, I. L. Khan, M. A. Javed, and S. A. Malik, “Optimal resource allocation for gaa users in spectrum access system using q-learning algorithm,” *IEEE Access*, vol. 10, pp. 60790–60804, 2022.
- [17] Y. Hsuan, “Impacts of propagation models on cbrs gaa coexistence and deployment density,” 2018.
- [18] Wireless Innovation Forum, “Requirements for Commercial Operation in the U.S. 3550-3700 MHz Citizens Broadband Radio Service Band,” <https://winnf.memberclicks.net/assets/CBRS/WINNF-TS-0112.pdf>, 2020, online.
- [19] S. Joshi, S. Kapuruhamy Badalge, M. Jokinen, T. Hänninen, P. Pirinen, H. Posti, and M. Latva-aho, “Esc sensor nodes placement and location for moving incumbent protection in cbrs,” in *Proceedings of WInnComm 2016*, 03 2016.
- [20] D. H. Kang, K. Balachandran, and M. Buchmayer, “Coexistence performance of gaa use cases using lte-tdt technologies in 3.5ghz cbrs spectrum,” in *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2018, pp. 1–7.
- [21] S. J. Maeng, I. Guvenç, M. L. Sichertiu, and O. Ozdemir, “Out-of-zone signal leakage sensing in radio dynamic zones,” 2021. [Online]. Available: <https://arxiv.org/abs/2111.09490>
- [22] S. J. Maeng, I. Guvenç, M. Sichertiu, B. A. Floyd, R. Dutta, T. Zajkowski, O. Ozdemir, and M. J. Mushi, “National radio dynamic zone concept with autonomous aerial and ground spectrum sensors,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.09014>
- [23] D. Eppink and W. Kuebler, “Tirem/sem handbook.” ELECTROMAGNETIC COMPATIBILITY ANALYSIS CENTER ANNAPOLIS MD, Tech. Rep., 1994.
- [24] T. T. Nguyen, R. Caromi, K. Kallas, and M. R. Souryal, “Deep learning for path loss prediction in the 3.5 ghz cbrs spectrum band,” in *2022 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE Press, 2022, p. 1665–1670. [Online]. Available: <https://doi.org/10.1109/WCNC51071.2022.9771737>
- [25] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, “Multimodal deep learning,” in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [27] D. Nau, Y. Bansod, S. Patra, M. Roberts, and R. Li, “Gtphop: A hierarchical goal+ task planner implemented in python,” in *ICAPS Workshop on Hierarchical Planning (HPlan)*, 2021.
- [28] A. Rasheed, O. San, and T. Kvamsdal, “Digital twin: Values, challenges and enablers from a modeling perspective,” *IEEE Access*, vol. 8, pp. 21980–22012, 2020.
- [29] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” 2017.