

How can polycentric governance of spectrum work?

Martin BH Weiss
(Corresponding Author)
School of Info. Sciences
University of Pittsburgh
Pittsburgh, PA
Email: mbw@pitt.edu

Prashant Krishnamurthy
School of Info. Sciences
University of Pittsburgh
Pittsburgh, PA
Email: prashk@pitt.edu

Marcela M. Gomez
School of Info. Sciences
University of Pittsburgh
Pittsburgh, PA
Email: mmg62@pitt.edu

Abstract—Spectrum policy in the US (and throughout most of the world) consists generally of a set of nationally determined policies that apply uniformly to all localities. However, it is also true that there is considerable variation in the features (e.g., traffic demand or population density), requirements and constraints of spectrum use on a local basis. Global spectrum policies designed to resolve a situation in New York City could well be overly restrictive for communities in rural areas (such as central Wyoming). At the same time, it is necessary to ensure that more permissive policies of central Wyoming would not create problems for NYC (by ensuring, for example, that relocated radios adapt to local policies). Notions of polycentric governance that have been articulated by the late E. Ostrom [16] argue that greater good can be achieved by allowing for local autonomy in resource allocation.

Shared access to spectrum is generally mediated through one of several technologies. As Weiss, Altamimi and Liu [21] show, approaches mediated by geolocation databases are the most cost effective in today’s technology. In the database oriented Spectrum Access System, or SAS, proposed by the FCC, users are granted (renewable) usage rights based on their location for a limited period of time. Because this system grants usage rights on a case-by-case basis, it may also allow for greater local autonomy while still maintaining global coordination. For example, it would be technically feasible for the database to include parameters such as transmit power, protocol, and bandwidth. Thus, they may provide the platform by which polycentric governance might come to spectrum management. In this paper, we explore, through some case examples, what polycentric governance of spectrum might look like and how this could be implemented in a database-driven spectrum management system.

In many ways this paper is a complement to [20]. In that paper, we evaluate emerging SAS architectures using Ostrom’s socioeconomic theory. In this paper, we explore how a SAS-based system could be constructed that is consistent with Ostrom’s polycentric governance ideas.

The approach proposed in this paper aims at addressing spectrum management as an emergent phenomenon rather than a top down system. This paper will describe the key details of this system and present some initial modeling results in comparison with the traditional global model of spectrum regulation. It will also discuss some of the concerns associated with this approach.

I. INTRODUCTION

Spectrum policy in the US (and throughout most of the world) consists generally of a set of nationally determined policies that apply more or less uniformly to all localities. However, the economics, user needs, and technical considerations of spectrum are geospatially very different. The population density of New York City (NYC) is 27016 people

per square mile¹ while the population density of Eugene, Oregon is 3572 people/square mile². The building profile in NYC is very different from that in Eugene, thereby creating very different radio propagation characteristics. These numbers illustrate the stark differences between two cities in the US. When we extend such considerations to rural communities, the differences become more drastic. Inhabited ranches may be separated by miles – applications may require wireless communication technology for herding of cattle across thousands of square miles with very few human beings in contrast to watching HD videos with Netflix in an apartment complex while competing with many similar streams in close proximity. Consequently, it is reasonable to assume that the spectrum policies designed to resolve a situation in NYC could well be overly restrictive for communities in central Wyoming.

Technologies for wireless systems have undergone rapid changes in the past two decades. Mass production of wireless chipsets and the ability to rapidly tweak them are real possibilities. Similarly, spectrum policy has seen rapid changes in the last years. The historical presumption in favor of exclusive license based access has changed. The introduction of unlicensed bands introduced new models for and applications of spectrum use. Beginning in the early 2000s [11], we began to see a serious engineering and policy thrust in support of the development of Dynamic Spectrum Access (DSA) systems that would enable much denser sharing of electrospace³. The emergence of this new policy thrust was enabled by the development of reconfigurable radios, such as software radios and cognitive radios.

The same technologies that have facilitated more intensive sharing of spectrum also allow for a broad new approach to spectrum policy. Instead of a “one size fits all” spectrum policy, it is now possible to have local spectrum policy that is tailored to the needs and spectrum environments of particular communities. The FCC has already taken steps in this direction in the CBRS (3.5GHz) proceeding [9] by allowing higher power levels in rural areas than in urban ones. While this is certainly a positive development, it still relies on a

¹https://en.wikipedia.org/wiki/List_of_United_States_cities_by_population_density

²https://en.wikipedia.org/wiki/Eugene,_Oregon

³This term is used to describe a multidimensional space for spectrum occupancy or use [5], [15]

1	Boundaries	5	Graduated Sanctions
2	Appropriateness	6	Conflict Resolution Mech.
3	Collective Choice	7	Minimal Recog. of Rights
4	Monitoring	8	Nested Enterprises

TABLE I: Critical attributes of long surviving CPR systems

centralized Spectrum Access System (SAS) that is *most likely not governed locally*. Inspired by experiences with polycentric governance of common pool resource systems, we propose that this spectrum policy can be locally determined through automated and semi-automated negotiation between devices. If radio spectrum can be best classified as a common pool resource system, the question that naturally arises is whether (and when) the kind of durable, decentralized governance system characterized by Ostrom and her collaborators can be developed.

Shared access to spectrum is generally mediated through one of several technologies. As Weiss et.al. [21] show, approaches mediated by geolocation databases are the most cost effective in today’s technology. In this paper, we explore how the database technology on which many modern spectrum sharing systems are built (e.g., TV White Spaces, Spectrum Access Systems [8], License Shared Access) can also be used as a foundation for the kind of resource governance system described by Ostrom [16]. Such a system will effectively implement a differential spectrum policy that is designed to optimize locally where possible. In these systems, users are granted (renewable) usage rights based on their location for a limited period of time. Because this system grants usage rights on a case-by-case basis, it may also allow for greater local autonomy while still maintaining global coordination. For example, it would be technically feasible for the database or SAS to include parameters such as transmit power, protocol, and bandwidth. Thus, they may provide the platform by which polycentric governance might come to spectrum management. In the following discussion, we use a very flexible notion of SAS to illustrate the principles behind polycentric governance of spectrum. Later we use the idea of a radio appliance (RA) to consider a more detailed version of polycentric governance.

II. REQUIREMENTS AND SYSTEM OPERATION

In the body of her work, Elinor Ostrom studied frameworks for the governance of common pool resource systems [16]. Several researchers have argued that spectrum meets the definition of common pool resource systems under our current technological endowment [13], [20]. Ostrom identified the list of attributes for the governance of long surviving common pool resource (CPR) systems contained in Table I.

Ostrom and her collaborators studied resource systems such as fisheries, forests and irrigation systems. Unlike these systems, spectrum is a *constructed* resource that does not exist apart from radio technology [19], [20]⁴. Thus, it is necessary

⁴Prior to Marconi, there was no discussion of spectrum rights. The need for spectrum coordination and the emergence of rights associated with spectrum grew with radio technology and its use [12].

	Full owner	Proprietor	Auth. claimant	Auth. xmitter	Auth. Rcvr
Access	X	X	X	X	X
Withdrawal	X	X	X	X	
Management	X	X	X		
Exclusion	X	X			
Alienation	X				

TABLE II: Distribution of rights by user type

to construct a technological foundation for a decentralized governance system for spectrum.

1. **Boundaries:** In Ostrom’s description, two types of boundaries are relevant: those that define the user community and those that define the resource. Resource boundaries are often physical delimitations of what is included in the governance system and what is excluded. For spectrum, the physical delimitation is largely determined by the electrospace of the transmission system. From a system design perspective, this is determined by technical parameters such as transmission power, frequency band, antenna height, antenna type, etc. User boundaries are more challenging. Here, the rights structure and governance systems come into play. To address this, Schlager and Ostrom [18] developed the rights typology defined in Table II (adapted here for application to radio spectrum). “Access” and “Withdrawal” rights are referred to as *usage rights* whereas the remaining are referred to as *collective action rights*. In this context, collective action rights are implemented in the SAS⁵, and these, in turn, govern the usage rights. Locally determined spectrum policy, then, means a (hyper) local SAS. The user community, then, is determined by the the local SAS, which must coordinate with adjacent (and nearby) local SASs or, possibly, a regional SAS.
2. **Appropriateness to local conditions** This item includes (a) “rules restricting time, place, technology, and/or quantity of resource units are related to local conditions” as well as (b) “The benefits obtained by users from a CPR, as determined by appropriation rules, are proportional to the amount of inputs required in the form of labor, material, or money, as determined by provision rules.” (from [3]) This is the attribute that is perhaps of most interest here. Appropriateness to local conditions means that the economic and management relationships are appropriate for the local conditions and the people affected.
3. **Collective Choice Arrangements** This item addresses the extent to which participants in the CPR system have a say in its management; that is, are endowed with and able to execute their collective choice rights. In spectrum sharing, collective choice is embodied in local SASs as well as in the protocol that implements

⁵Since we are envisioning an environment of nested rights, this may refer to the system of SASs in reality.

Type 1	Events due to the routine operation of a sharing ecosystem
Type 2	Events due to “rogue” or malicious users
Type 3	Events due to faulty equipment of authorized spectrum users
Type 4	Events when all users are in compliance with regulations

TABLE III: A typology of interference events

inter-SAS coordination [20].

The next three items address the enforcement of CPR governance relationships. Enforcement has been addressed elsewhere (see, for example, [22]); in brief, it is necessary to enforce both usage rights and collective action rights.

Usage rights in spectrum as a common pool resource are typically expressed in the form of interference protection. Enforceable interference events might be subdivided into four distinct types, as outlined in Table III.

Type 1 interference events might occur due to aggregation of similar devices, propagation anomalies, location errors, etc. In these kinds of events, we expect that the radios are compliant with applicable technical and operational regulations. We cannot say the same about Type 2 interferers, which might be software radios that have temporarily been programmed to operate in a band and may not make an effort at compliance with the appropriate technical and operational requirements. These may or may not have a typical physical characteristic that would allow them to be automatically identified. Type 3 events are due to leaky cables, poor filters, etc. We would expect these to be licensed devices that fit no particular pattern or lack a particular physical characteristic. Type 4 events occur when regulations or licenses are incomplete and/or poorly written or assigned.

For the purpose of this paper, we consider only Type 1 events, which is also generally the focus of the CPR literature. Other types of interference events are also of substantial importance, but that is outside the scope of this paper.

Weiss *et.al.* [20] have made a case for the need to enforce collective action rights as well as interference rights in spectrum sharing systems. To date, collective action rights have been exerted through the NTIA’s CSMAC⁶ process and in the FCC (through the Administrative Procedures Act (APA)). As Cox *et.al.* point out transparency is of importance for sustainable CPR management; in spectrum systems, these collective action is codified in software-based SAS systems which will require transparency as well.

4. **Monitoring** This refers to the ability to observe and audit the use and governance of the CPR system. It also refers to the ability to monitor the overall “health” of the system. The accountability of monitors to the members of the benefactors of the CPR is important as well.

For spectrum, this can imply the need to monitor spectrum use, or have the ability to detect interference that rises above a “harm-claim threshold” [6]. As to the second component, the “health” of the ecosystem can perhaps best be measured by the noise floor [10].

5. **Graduated Sanctions** When violations of the operational rules occur, sustainable CPRs allow for sanctions that are scaled to the seriousness and context of the offense. Sanctions may be assessed by other benefactors and/or officials who are accountable to them.

Some research on appropriate and effective sanctions has been performed by Woyach and Sahai [24], [25]. In their work, they focus on non-monetary “spectrum jails” that penalize misappropriators. Alternately, Malki and Weiss [14] consider penalties in the context of the law and economics literature. In both cases, the penalties and sanctions have been focussed on interference events. Work remains to be performed on how sanctions might be applied when collective action rights are violated.

6. **Conflict Resolution Mechanism** Also of importance is rapid access to low-cost local arenas to resolve conflicts among benefactors or between benefactors and officials. The federal enforcement system currently in place through the FCC’s enforcement bureau does not meet this criterion. It is neither local or low cost. Thus, a local system must include a mechanism for automated or semi-automated negotiations among systems to coordinate their usage of spectrum.

The remaining factors refer to the broad institutional context in which the CPR exists. On the one hand, there is a question of legal context and on the other, one of broader operational context.

7. **Minimal Recognition of Rights** This refers to a formal (or informal) recognition by central authorities of the self-governance of local authorities. The FCC already delegates spectrum management functions to organizations such as the ARRL [4]. As well, the databases for the TVWS and CBRS are effectively managing spectrum on behalf of the FCC, so it seems that this recognition would, in principle, be possible.

8. **Nested Enterprises** In many sustainable CPR systems, the use, provision, monitoring, enforcement, conflict resolution, and governance activities are organized in multiple layers of nested enterprises. This is helpful because (a) CPR systems exist in a broader environment that may require external coordination; (b) not all conflicts and disputes can be resolved locally, so it might be necessary to appeal to an independent authority; (c) it may be more cost effective to use a super-local or regional entity to coordinate usage than to resolve usage or governance conflicts locally.

III. SYSTEM ARCHITECTURE AND CASE STUDIES

The current trend toward database-oriented spectrum sharing systems allows for the implementation of a decentralized spectrum policy. We envision a system based on a geo-located and networked “radio appliance” (RA) that controls the protocols and transmission parameters of the radios associated with it. It is effectively a small scale Spectrum Access System (SAS). This appliance (and the associated radios) might be

⁶<https://www.ntia.doc.gov/category/csmac>

Requirement	System Feature
Boundaries	Operating boundaries determined by transmit power and antenna
Congruence with Local Conditions	Locally determined spectrum assignment and usage
Collective Choice Arrangements	Open source software for radio appliance
Monitoring Users & resource	Identifying users and documenting transmission & spectrum sensing
Graduated Sanctions	Backoff protocol and explicit coordination
Conflict Resolution Mechanisms	Protocols for negotiating interference protection
Minimal Recognition of Rights	Delegation of spectrum control by FCC, NTIA and SAS operators
Nested Enterprises	Ability to self-organize & delegate spectrum control

TABLE IV: Architecture requirements

owned and operated by a landlord, a farmer, an Internet service provider, etc.

When an RA is installed and powered up, it locates itself, registers itself and obtains an operating license and learns about neighboring RAs. If there are no nearby RAs, the operating parameters of the radios, such as power, bandwidth, protocols (e.g. waiting times in WiFi) would gradually become more permissive until (a) it can meet the requirements of the application in question or (b) an enforceable event occurred. If (a) occurs, the system would maintain steady state, making small adjustments to account for environmental variations (e.g., weather). If instead (b) occurs, the RA would automatically cause the parameters to become more restrictive before gradually increasing again to find the maximum possible performance⁷. As additional RAs come on line, the RAs would begin to negotiate with each other, coordinating frequency usage, time of use, transmission protocol, etc. to minimize the occurrence of Type 1 interference events. If the time required to compute an acceptable allocation becomes too large, the affected RAs defer to a super-local RA. This might be a separate, specialized entity or a software process within each RA that could be invoked.

If, upon powerup, the RA learns that there are other RAs nearby (as might be the case in a suburban or urban area), the RA first determines whether coordination will be required (based on the distance to other RAs). If none is required, it proceeds as described above. If it is required, the RA requests an electrospace allocation for its operation either through negotiation with neighboring RAs or from the super-local RA.

In Table IV, we show how this approach could implement a decentralized spectrum policy regime in the context of sustainable CPR systems.

To illustrate operation, we will consider two cases. The first is a macro-level case that uses the 3.5 GHz band Citizens Broadband Radio Service. Here we will not consider protocols, but instead look at the *coverage* and *interference* aspects.

⁷This is generally similar to the “slow start” mechanism of TCP (<https://web.stanford.edu/class/cs244e/papers/rfc2001.pdf>)

The second is a more fine-grained case of 802.11 at the protocol level. In this case, we will examine the impact of locally adapting the protocol parameters. Both cases make assumptions that are not necessarily realistic, but the objective here is to assess the possibilities with polycentric governance, rather than creating realistic simulations.

A. Case I: Citizens Broadband Radio Service

We begin with a very rural situation: a ranch in central Wyoming. This region is outside of the incumbent’s exclusion zone so sensing naval radars is not a concern. Suppose the rancher wishes to monitor the location and health of her cattle and acquires the proposed radio system. The RA is installed centrally with a tower and an antenna. The radio-based cattle monitors are strapped to the animals. The RA begins by registering itself over the Internet and acquiring information about neighboring RAs and the governance structure in place. In our example, we assume no nearby RAs, so the radios and RA begin by operating with “stock” radio parameters as specified by the FCC for the 3.5GHz band. The RA begins permitting the associated radio devices to gradually increase their power until adequate performance is achieved over the entire ranch.

Let us now assume that a nearby ranch also acquires a similar system. When their RA (RA_2) is powered on, it learns of the presence and location of RA_1 . It also learns that a fully decentralized governance system is in place and that it is outside of the incumbent exclusion zone. If these two RAs are sufficiently far apart to not interfere with each other under the “stock” transmission parameters and, like RA_1 , it would choose the best modulation scheme, as much channel bandwidth and channel time as needed for the application and gradually increase transmission power until adequate performance is achieved over their entire ranch.

This might continue unabated if the two ranches are geographically separate. If they are adjacent, it is likely that transmissions associated with RA_1 would eventually interfere with transmissions associated with RA_2 (or *vice versa*). If interference reached a level where the application performance was sufficiently impaired, RA_1 or RA_2 (whichever first detected the problem) would initiate a negotiation with the other. The RAs can coordinate on transmit power, frequency, time or modulation scheme to mitigate the interference; transmit power would affect reach, frequency band would affect bandwidth, time would affect latency (as well as bandwidth), and modulation might strike a different balance between the three. Each RA would prefer to minimize the consequence of the negotiation on the performance of the application. In the case of the cattle monitoring application, the RAs might prefer to maintain higher power levels at the cost of decreased throughput or increased latency so that the entire ranch can be covered.

If these radio systems became popular in the area, then bilateral coordination would no longer function efficiently. In these cases, one of the radios would be elected to be the spectrum coordinator (this could also potentially rotate among the RAs

in the region). The SC's function would be to allocate the regional spectrum resources among the associated RAs so to optimize its use. This SC would need to be aware of other SCs in the region and to coordinate with super-regional SASs (i.e., participate in a nested governance system).

An enforceable event occurs when the transmission of a radio device associated with one appliance diminishes the operation of a radio associated with a different appliance to the point where the function of the radio application is measurably impaired (e.g., latency, throughput). An enforceable event would trigger an automated negotiation protocol whose goal is modifying the radio parameters to minimize the consequence of sharing. The outcome of the negotiation might be decreasing transmission power, coordinating the frequency bands used, time sharing the transmissions, etc. The radio appliances would then gradually increase power until another enforceable event occurred.

It may also involve fusing authenticated data regarding the enforceable event. If the frequency of enforceable events increased above a threshold, a super local (or regional) coordinating device would be invoked to optimize use of the spectrum. In this way, spectrum governance is nested [1], providing resilience against a single point of failure in the governance process.

This represents a means to learn from the local knowledge and policy changes, which results in a system that is more responsive to environment threats at multiple scales. In the end, this system can compensate for the failure of some units with the successful response of others [2]. This can be an advantage with respect to global policy mechanisms. Thus spectrum management would be best modelled as an emergent phenomenon rather than a top down system.

B. Case II: IEEE 802.11 Protocol Tweaks

Urban environments with dense deployment of WiFi access points provide another environment where polycentric governance of radio spectrum regulation may make sense. In apartment complexes and other densely populated urban areas, it is likely that a variety of usage scenarios exist. It is possible that there are low and bursty data rate users that primarily browse e-mails and the web and very high data rate users that stream HD video to their portable devices in close vicinity. In an unplanned WiFi deployment, such users may share the same WiFi channels. The RA in this case may choose to specify different protocol parameters for different users based on their willingness to tolerate latency or other performance parameters.

Specifically, in the case of WiFi, transmissions are based on collision avoidance. The protocol requires devices to sense the radio channel, and after a transmission in a "free" communication medium or after detecting the medium to be busy, devices enter a backoff, where they wait for a random duration within a contention window. This contention window is initially set to a small value for all devices and then doubled if packet losses are observed. One strategy that could be deployed is to specify *different* contention windows for different devices

based on a negotiation with the RA. While this is similar to some of the QoS approaches in WiFi [17], the devices that employ QoS do so on their own, without any coordination with an RA. Thus, the resulting behavior may fail to be as desired. One could imagine an RA being installed in each apartment complex or a specific urban geographic area (e.g., a hotel and its surroundings). With IEEE 802.11ac, it may be possible for an RA to facilitate negotiation of the channel bandwidth as well - where channel bandwidths of 20, 40 or 80 MHz may be negotiated to reduce interference or increase throughput as necessary.

IV. SIMULATION OF CASE I

A. First Simulation Environment

We developed an Agent-based model in NetLogo 5.2.1⁸ to test the first case study. This simulation tool permits to utilize three types of agents: patches, turtles and links⁹. Patches are static agents. In our model, we utilize these agents in order to define ranches, base-stations or radio appliances (RAs), cattle and the communication process between the RA and cattle.

NetLogo patches are utilized for creating ranches. In this way, we have defined each ranch as a square formed by 256 patches. This means that each side of the ranch will have 16 patches. We further assume that each patch represents a small square of 100 meters per side, hence the approximate area of each ranch is 1 square mile. It is important to note that each ranch has specific x-axis and y-axis coordinates, which are useful for defining the ranch boundaries and further determining the required coverage.

We utilized NetLogo turtles for creating base stations or radio appliances (RAs) and cattle. The RA is located at the center of each ranch. In turn, the cows belonging to each ranch are randomly placed within the ranch area. This means that a cow's specific x-axis coordinate is a random number, uniformly distributed between the maximum and minimum x-coordinates of the ranch (e.g., $x = 1$ and $x = 16$). The same applies for its y-axis coordinates.

Links are currently utilized to illustrate the flow of information between the cattle and the corresponding RAs.

Our model supports the creation of up to four ranches, where the communications system of this case study has been deployed. In this way, we envision an environment where an RA periodically monitors the cattle by pinging them in order to determine their location. Additionally, the sensors attached to the cattle can transmit information about the nutrition and hydration levels of each cow. In case these levels fall below a critical threshold, the sensor is expected to send an emergency message that will alert the farmer of the situation.

Given that the cows represent the mobile agents of the system, they have been modeled as moving agents. In an

⁸Information on this programming tool can be found in the NetLogo User Manual [23]

⁹"Turtles are agents that move around in the world. The world is two dimensional and is divided up into a grid of patches. Each patch is a square piece of "ground" over which turtles can move. Links are agents that connect two turtles." [23]

attempt to portray the level of activity of real cows, we have set our cow agents to move only 30% of the time. In this way, at every simulation time unit (i.e., tick) each cow agent generates a random number, uniformly distributed between 0 and 100. When this number is lower than 30 the cow moves. The actual movement of a cow is given by an angle with respect to the initial direction that the cow faces and a number of “steps” in that direction. As such, we have modeled the movement of a cow as a shift to the right with an angle θ_1 and a rotation to the left with an angle θ_2 . The values of θ_1 and θ_2 are defined as uniformly distributed random numbers between 0 and 35. After the angle that the agent faces is determined, the cow moves one step forward in that direction. When a cow reaches the borders of its ranch it turns to face the center of the ranch and continues with the aforementioned movement sequence. We have utilized this movement pattern as a means to avoid the agents moving in one single direction and going beyond the boundaries of their own ranches.

1) *Model Parameters*: For this communication system to operate, we need to define the appropriate transmission, reception, coverage and propagation parameters. Fig. 1 shows an example of the location of one ranch in terms of its x-,y-coordinates and the maximum radius that needs to be covered. Taking into account this maximum coverage distance, we can utilize an applicable propagation model and the link budget formula in order to determine what is the minimum transmission power that an RA requires for this type of coverage.

The propagation model that we have chosen is the Extended Hata model for sub-urban environments, as it has been used by the NTIA for calculating exclusion zones in the 3.5 GHz band [7]. The actual parameters utilized in these calculations are summarized in Table V.

Parameter	Value
Frequency	3500 MHz
Transmitting antenna height [Hb]	10 m - 20 m
Mobile antenna height [Hr]	2 m
Distance	1.2 Km

TABLE V: Parameters utilized for calculation of path loss with Extended Hata Model for suburban environments

After obtaining the applicable path loss values, we consider the link budget formula (1), where P_r is the minimum power required by the receiver, G_t and G_r are the transmitter and receiver antenna gains, L_t and L_r are the transmitter and receiver losses, respectively and L_p is the path loss calculated with the Extended Hata model. The specific values for the aforementioned parameters are presented in Table VI.

$$P_t = P_r - G_t + L_t - G_r + L_r + L_p \quad (1)$$

2) *Finding coordination events*: The coexistence of multiple ranches in close proximity may result in mobile devices (i.e., cow sensors) suffering harmful interference. According to the governance model we study, interference would call for coordination among the interfering entities. Applying this

(1, 16)

(16, 16)




Fig. 1: Graphic representation of one ranch in the simulation environment

Parameter	Value
Minimum Received Power - P_r	-80 dBm
Transmitter antenna Gain - G_t	3.3 dB
Receiver antenna Gain - G_r	3.3 dB
Loss at the transmitter - L_t	2 dB
Loss at the receiver - L_r	2 dB

TABLE VI: Parameters for calculating the minimum required transmission power

to our model, when we consider more than one ranch, the coverage area of each RA may interfere with the coverage area of the remaining RAs. In Fig. 2 we illustrate this interference situation for an area with four adjacent ranches. The circles surrounding each ranch represent the coverage area of the RA at the center of each ranch. As it can be observed, the coverage areas overlap thus creating interference-prone zones, which may call for coordination events among RAs.

For interference assessment, we consider there to be harmful interference when there are cows in the interference area (i.e., mobile devices) and when the power level in the interference zone is above the noise threshold of these devices. Consequently, when these two conditions are met, we consider that a coordination process needs to take place among the interfering parties.

The aforementioned conditions are a consequence of the transmit power chosen by each RA and the noise tolerance of the mobile devices attached to the cattle. To illustrate these factors in our model, each RA responds to a specific profile that determines the transmission power level to utilize. The profiles have been defined as follows:

- Profile 0 : The transmission power of RA_x is the minimum power it requires to cover the entire area of the




Fig. 2: Interference zones

ranch where it operates.

- Profile 1: The transmission power of RA_x is the maximum allowed by the FCC, i.e., 47 dB.
- Profile 2: The transmission power of RA_x is 10 dB above the maximum power allowed by the FCC.

In the case of the mobile devices, three interference tolerance levels are also defined:

- Level 0: -100 dB
- Level 1: -90 dB
- Level 2: -80 dB

Note that all the mobile devices belonging to a particular ranch will be assigned the same level.

3) *Simulation Results*: In this section we present the evaluation of three simulation scenarios, which aim at assessing the interference events that would take place and thus require coordination among RAs. The three simulation scenarios attend to the following characteristics:

- Every RA transmits at the power level allowed by the FCC (i.e., 47dB) and the mobile devices in each ranch have the highest tolerance to noise.
- Every RA transmits at the minimum power it requires to cover the entire ranch where it operates (i.e., 52dB) and the mobile devices in each ranch have the lowest interference tolerance.
- Each RA and group of mobile devices in every ranch choose independently¹⁰ their transmit power profile and interference tolerance level, respectively.

¹⁰To simulate the independence of each RA and group of mobile devices, their profiles and interference tolerance levels have been assigned as a uniformly distributed random number from 0 to 2. Subsequently, the transmit power and tolerance values have been assigned as defined in subsection IV-A2.

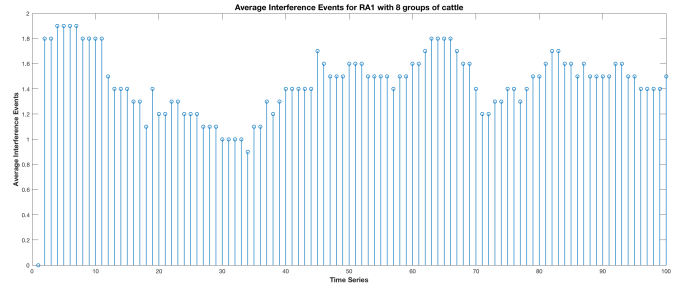


Fig. 3: Aggregate interference incidences for RA_1 with varying number of mobile devices in a worst-case scenario

In each of the aforementioned scenarios we also explore how the number of mobile devices (i.e., groups of cattle) affect the number of interference events. In this manner, we simulate scenarios where the group of cattle varies from three to eight. The results presented in what follows correspond to those obtained while exploring the model behavior in 2000 simulation time units or “ticks” and calculating the average over 10 repetitions.

a) *Scenario 1*: Our simulation results show zero interference events, which means that when the RAs limit their transmission power to the maximum established by the FCC, there is no need for negotiation among them. However, it should be noted that this power level does not allow the RAs to fully cover the area of the ranch. Indeed, for the minimum received power we consider in our model, this power level provides a coverage radius of approximately 900 meters. It follows that, when compared to the required coverage radius of 1130 meters, approximately 20% of the area is left uncovered. In this scenario we are presented with a trade-off between avoiding inter-ranch interference (thus negotiations among RAs) and reaching the required coverage.

b) *Scenario 2*: In the second scenario we find a significant number of interference incidences throughout our simulation. Given that all RAs operate under similar configurations, in Fig. 3 we present the results obtained with RA_1 , which show how the number of interference events when there are 8 groups of cattle in the ranch.

As it can be observed, on average, there are 1 to 2 interference events per time instance, which would require coordination among RAs.

c) *Scenario 3*: This scenario aims at capturing varying behavior of the RAs and sensitivity levels of the mobile devices. This could be a representation of what might happen in a real-world scenario as the owners of the equipment may be able to change the operational parameters of their devices or choose those that best adapt to the environment. Fig. 4 shows the results obtained for RA_1 in this scenario.

B. Second Simulation Environment

We modified the first simulation environment in order to consider a geographical area where RAs are in closer proximity. For this purpose, we now model a 9-cell grid where each RA operates and serves groups of users. The actual

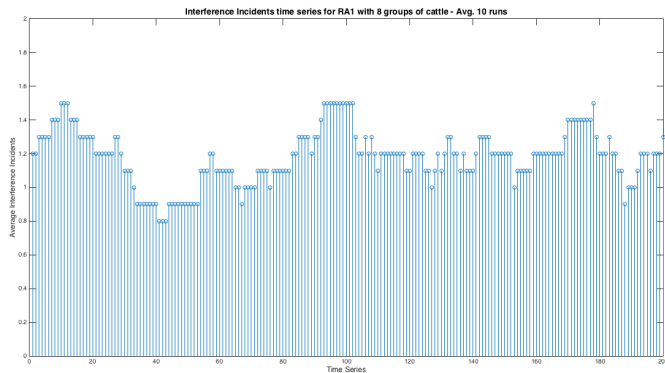


Fig. 4: Interference events when the RAs’ transmit power and mobile device sensitivity are particular to each ranch.




Fig. 5: Configuration of second simulation environment

configuration of this scenario is shown in Fig. 5. Note that the goal of this scenario is to determine how the number of interference events changes as we consider more crowded spaces.

The model parameters and the method for identifying interference events are the same as those utilized for the first simulation environment and presented in subsections IV-A1 and IV-A2, respectively. We ran our simulation considering the situation in which each RA chooses independently its transmission power according to the three profiles explained in subsection IV-A2. Each profile is assigned to every RA as a uniformly distributed random number between 0 and 2.

1) *Simulation Results:* In this section we include the results for the most *critical* case. This corresponds to the interference events generated and perceived by the RA and users belonging to the cell (or ranch) located at the center of the grid. To generate our simulation results, we varied the number of user groups from 3 to 8.

Fig. 6 shows the interference incidences for RA5 during the first 200 simulation time ticks. RA5 is located in the central cell. Note that the aggregate interference events correspond to the average number of events across 10 simulation runs, each lasting 2000 time ticks.

Our simulation results show that as the number of user

groups increases, so do the number of interference events. Note that the total number of interference events is lower than that of the first case. This is due to the lower transmission power required for coverage of the entire cell area as the latter is significantly smaller than the case of four, bigger, ranches. From these results we can infer that a negotiation mechanism that involves reduction of transmission power may effectively reduce the number of interference events in the studied environment¹¹.

C. Implementing Negotiations

After analyzing the interference events and their patterns in the previous subsections, we study how the incorporation of a negotiation mechanisms impacts the number of interference events and thus performance of each RA. The negotiation process operates as follows: Let RA_x be the RA causing interference and RA_y the “interference victim”.

- 1) If the transmit power of RA_x is higher than the minimum required to cover its ranch area, RA_x should reduce its transmit power (preferably to the minimum)
- 2) If the transmit power RA_x is already at the minimum (or below), RA_x switches to a different frequency band (within the 3500 to 3700 MHz range).

We tested this negotiation mechanism in the 4-ranch scenario, considering each RA to choose its power and device interference tolerance independently from one another. The results obtained are presented in Fig. 7

Fig. 8 compares the average interference events that occur in the first 200 time ticks, instead of the aggregate values. We have chosen this period as a sample of what occurs throughout the entire simulation, as there is not a significant change in the results obtained as time progresses. In this figure, the reduction of the interference events is evident when implementing a simple negotiation mechanism. It is also important to note that the power adjustment is done in terms of the power required for ranch coverage instead of what is allowed by regulation. In this manner, some RAs may be transmitting at a power higher than the limited established by the FCC and still is able to avoid causing interference to the neighboring RAs.

V. CONCLUSION

The database driven Spectrum Access System (SAS) was designed to implement ex ante control over a defined electrospacetime to enforce a particular set of usage rights in the form of protection from interference events. As such a SAS is a governance mechanism for spectrum. Though SAS was designed as a centralized control mechanism, in this paper we propose and demonstrate through some toy models the feasibility of using SAS to implement a decentralized, locally driven spectrum policy. This approach to spectrum use could dynamically adapt to the needs of every particular area and thus manage resource access, interference and assignment. Decentralized, polycentric forms of spectrum management

¹¹This situation would hold as long as a lower transmit power still permits to fulfill the coverage requirements.




Fig. 6: Interference Events for the central cell in a 9-cell grid

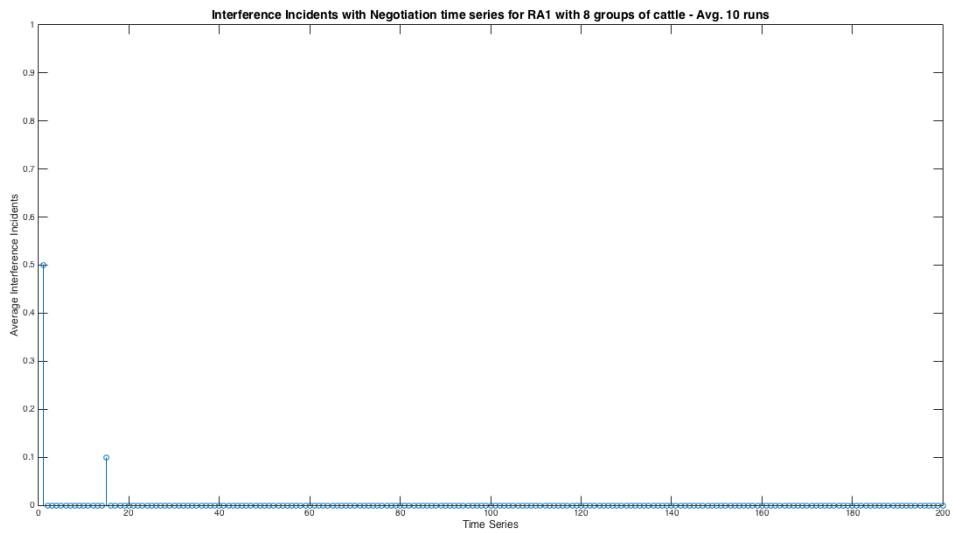


Fig. 7: Aggregate Interference Events for RA1 when negotiation is implemented

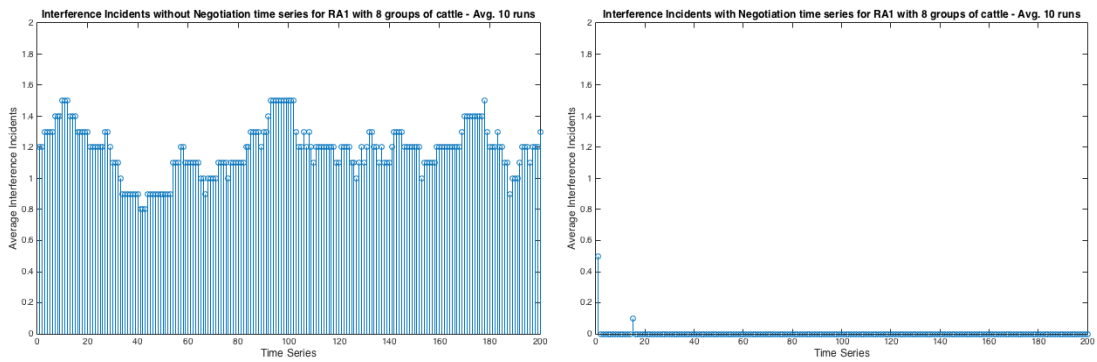


Fig. 8: Comparing average interference events in scenarios with and without negotiation

have properties that can be mapped to the resource systems studied by Elinor Ostrom and her collaborators.

To explore this approach, we examine two case studies: the first explores the need for governance in a communications system developed utilizing the CBRS band, while the second one focuses on alternatives provided by tweaks to the IEEE 802.11 protocol.

For our analysis of the CBRS band, we have developed an Agent-based model that portrays the utilization of a system for monitoring cattle in a rural area, where four ranches are located in relatively close proximity. Through the evaluation of our simulation results, we can first point out that while the power limits established by the FCC avoid the need for negotiation among entities, the associated cost includes a reduction in the coverage area, which is an important factor in the system we explore. Additionally, we find that as radio appliances increase their transmit power, interference events arise; however, the increment of this events is not constant or linear, which permits us to infer that there are occasions in which interference is not present. This would leave us with opportunities to negotiate access in terms of time in addition to frequency. Finally, when we characterize the different agents participating in each ranch independently from one another, we find a lower number of interference events. This suggests that heterogeneity of operation also creates opportunities for negotiation and coordination among devices.

The second case study analyzes the application of polycentric governance to an environment where IEEE 802.11 is the predominant technology. The difference in performance requirements, quality of service and resource usage may make it suitable to deploy a system of local governance rather than a global coordination system.

In both cases, open questions remain regarding various technical and regulatory issues. Nevertheless, through our case studies we point out that incorporating a mechanism that allows entities to learn from local conditions and adjust to particular environments may permit to create enforcement methods that are more suitable than those which have been broadly defined.

REFERENCES

- [1] Arun Agrawal. Sustainable governance of common-pool resources: Context, methods, and politics. *Annual Review of Anthropology*, 32:243–62, 2003.
- [2] Arun Agrawal and Elinor Ostrom. Collective action, property rights, and decentralization in resource use in india and nepal. *Politics & Society*, 29(4):485–514, 2001.
- [3] M. Cox, G. Arnold, and S. Villamayor Tomás. A review of design principles for community-based natural resource management. *Ecology and Society*, 15(4), 2010.
- [4] Liu Cui, Marcela M Gomez, and Martin BH Weiss. Dimensions of cooperative spectrum sharing: Rights and enforcement. In *New Frontiers in Dynamic Spectrum Access Networks*. IEEE, April 2014.
- [5] Arthur S De Vany, Ross D Eckert, Charles J Meyers, Donald J O’Hara, and Richard C Scott. A property system for market allocation of the electromagnetic spectrum: A legal-economic-engineering study. *Stanford Law Review*, pages 1499–1561, 1969.
- [6] J. Pierre de Vries. Optimizing receiver performance using harm claim thresholds. *Telecommunications Policy*, 37(9):757 – 771, 2013.

- [7] E. Drocella, R. Sole Richards, F. Najmy, A. Lundy, and P. McKenna. NTIA Report 15-517. 3.5 GHz Exclusion Zone Analysis and Methodology.
- [8] FCC. Amendment of the Commission’s Rules with Regard to Commercial Operations in the 3550-3650 MHz Band. GN Docket No. 12-354.
- [9] FCC. In the matter of Amendment of the Commission’s Rules with Regard to Commercial Operations in the 3550–3650 MHz Band.
- [10] FCC. Office of Engineering and Technology announces technological advisory council (TAC) noise floor technical inquiry. DA-16-676.
- [11] FCC Spectrum Policy Task Force. Spectrum policy task force report. *ET Docket*, (03-237), 2002.
- [12] Thomas W. Hazlett. The rationality of useregulation of the broadcast spectrum. *The Journal of Law and Economics*, XXXIII(1):133–175, April 1990.
- [13] Christian A. Herter. The electromagnetic spectrum: A critical natural resource. *Natural Resources Journal*, 25:651–663, July 1985.
- [14] Amer Malki and Martin B.H. Weiss. Ex-post enforcement in spectrum sharing. In *Telecommunications Policy Research Conference*. Telecommunications Policy Research Conference, September 2014.
- [15] Robert Matheson and Adele C. Morris. The technical basis for spectrum rights: Policies to enhance market efficiency. *Telecommunications Policy*, 36(9):783 – 792, 2012.
- [16] Elinor Ostrom. Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100:1–33, June 2010.
- [17] W. Pattara-atikom, P. Krishnamurthy, and S. Banerjee. Distributed mechanisms for quality of service in wireless lans. *IEEE Wireless Communications (Special issue on QoS in Next-generation Wireless Multimedia Communications Systems)*, 10(3), 2003.
- [18] Edella Schlager and Elinor Ostrom. Property-rights regimes and natural resources: A conceptual analysis. *Land Economics*, 68(3):249–262, August 1992.
- [19] Howard A. Shelanski and Peter W. Huber. Administrative creation of property rights to radio spectrum. *Journal of Law and Economics*, 41(S2):581–609, October 1998.
- [20] M. B. H. Weiss, W. H. Lehr, A. Acker, and M. M. Gomez. Socio-technical considerations for spectrum access system (sas) design. In *Dynamic Spectrum Access Networks (DySPAN), 2015 IEEE International Symposium on*, pages 35–46, Sept 2015.
- [21] Martin B.H. Weiss, Mohammed Altamimi, and Liu Cui. Spatio-temporal spectrum modeling: Taxonomy and economic evaluation of context acquisition. *Telecommunications Policy*, 36(4):335 – 348, 2012.
- [22] Martin B.H. Weiss, William H. Lehr, Liu Cui, and Mohammed Altamimi. Enforcement in dynamic spectrum access systems. In *Telecommunications Policy Research Conference*. Telecommunications Policy Research Conference, September 2012.
- [23] Uri Wilensky. Netlogo user manual.
- [24] K. Woyach and A. Sahai. Why the caged cognitive radio sings. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 431–442, 2011.
- [25] K.A. Woyach, A. Sahai, G. Atia, and V. Saligrama. Crime and punishment for cognitive radios. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 236–243, Sept 2008.