

Information Reuse in Dynamic Spectrum Access

Invited Paper

Prashant Krishnamurthy
School of Information Sciences
University of Pittsburgh
Pittsburgh, Pennsylvania 15260
Email: prashk@pitt.edu

Martin Weiss
School of Information Sciences
University of Pittsburgh
Pittsburgh, Pennsylvania 15260
Email: mbw@pitt.edu

David Tipper
School of Information Sciences
University of Pittsburgh
Pittsburgh, Pennsylvania 15260
Email: dtipper@pitt.edu

Abstract—Dynamic spectrum access (DSA), where the permission to use slices of radio spectrum is dynamically shifted (in time and in different geographical areas) across various communications services and applications, has been an area of interest from technical and public policy perspectives over the last decade. The underlying belief is that this will increase spectrum utilization, especially since many spectrum bands are relatively unused, ultimately leading to the creation of new and innovative services that exploit the increase in spectrum availability. Determining whether a slice of spectrum, allocated or licensed to a primary user, is available for use by a secondary user at a certain time and in a certain geographic area is a challenging task. This requires “context information” which is critical to the operation of DSA. Such context information can be obtained in several ways, with different costs, and different quality/usefulness of the information. In this paper, we describe the challenges in obtaining this context information, the potential for the integration of various sources of context information, and the potential for reuse of such information for related and unrelated purposes such as localization and enforcement of spectrum sharing. Since some of the infrastructure for obtaining fine-grained context information is likely to be expensive, the reuse of this infrastructure/information and integration of information from less expensive sources are likely to be essential for the economical and technological viability of DSA.

I. INTRODUCTION

A variety of applications and services make use of radio transmissions, ranging from garage door openers and baby monitors to cellular networks with smart phones and satellite communications. Dynamic Spectrum Access (DSA) refers generally to the process where the permission to transmit in slices of radio spectrum can dynamically change across such services and applications in time and space. This is in contrast to the most common current approaches, which include licensed and unlicensed access to the radio spectrum. In the former, a *license* to a slice of radio spectrum is granted by the Federal Communications Commission (in the US) to a *primary* user of the slice of spectrum for extended periods of time for use for a specific service or application. For example, the band between 869 MHz to 894 MHz is *allocated* for use for downlink (base station towers to mobile stations) transmissions by cellular telephone networks (which also transmit data traffic in addition to cellular voice telephony in this spectrum these days). In this band, certain slices of

radio spectrum are assigned (licensed) in specific geographical areas on a long term basis to a service provider such as Verizon wireless, precluding its use by anyone else for any other purpose in these regions. This current static approach gives exclusive permission to the primary user to transmit in the slice of radio spectrum in a certain geographical area. But it also allows the primary user to leave the radio spectrum idle, thereby reducing the utilization of radio spectrum, whose value and need is only increasing in time. Finally, the time required to obtain a license can be measured in months (if one is taking over an existing license, with its constraints, including geography, permitted service, power levels, etc.) to decades (if spectrum must be reallocated from one purpose to another before a license can be issued). This does not support innovation in wireless particularly well.

The primary existing alternative to licensed use is to use *unlicensed* bands. Spectrum management authorities worldwide (e.g., the FCC in the US) have allocated certain bands (e.g., the 2.4 – 2.483 GHz bands in the US) to “industrial, scientific and medical” (ISM) uses where no license is required to transmit. WiFi, cordless telephones, and Bluetooth devices are all examples of systems that use these bands. Only some restrictions are placed on the transmit power and in some cases the technical details of the transmission schemes in such bands. Unlicensed users must follow such general rules in order to use these bands (devices need certification by the FCC in the US), but they need not obtain a license to use the frequency bands. Unlike licensed bands, these bands have not been allocated to a particular service (e.g., mobile communications, broadcasting). As a result, considerable innovation in devices and services has taken place using these bands. The principal limitations of unlicensed bands are that (1) the power limitations mean that supported services are exclusively short range and (2) the openness to use by anyone means that the performance of systems that use unlicensed bands is highly unpredictable in space and time.

An alternative to both licensed and unlicensed use, DSA allows *secondary* users of spectrum who can use the spectrum allocated to primary users, when such usage occurs when the primary transmissions do not exist in the simplest case or when such usage will not harm the primary communications. DSA

has been an area of interest from technical and public policy perspectives over the last decade¹. The underlying belief is that DSA, if implemented, will increase spectrum utilization, especially since many spectrum bands are relatively unused, ultimately leading to the rapid creation of new, innovative and perhaps inexpensive services that exploit the resulting increase in spectrum availability.

There are several approaches to DSA that have been examined and investigated. Determining whether a slice of spectrum, licensed to a primary user, is available for use by a secondary user at a certain time and in a certain geographic area is a challenging task. This requires “context information” which is critical to the operation of DSA. Such context information can be obtained in several ways, with different costs, and different quality/usefulness of the information. Opportunistic spectrum access with cognitive radios [1] is perhaps the most commonly used example of DSA. In this case, cognitive radios “sense” whether or not a primary user is actually transmitting in a slice of radio spectrum, and if not, the cognitive radios can make use of the radio spectrum becoming secondary users of the spectrum. There are several challenges in correctly sensing the usage of spectrum by cognitive radios as mentioned later. A database of usage of spectrum [2], especially for predictable services, is another approach towards DSA. In this case, a potential secondary user of a slice of radio spectrum consults a database prior to transmitting to determine whether or not the spectrum slice can be used without *harmful interference* to the primary user. A third approach suggests “sensing as a service” [3]. In this case, a network of sensors is deployed to specifically sense the utilization of spectrum and provide this information to secondary users.

Since the infrastructure for obtaining fine-grained context information is likely to be expensive (e.g., sensing as a service), the reuse of this infrastructure/information and integration of information from less expensive sources are likely to be essential for the economical and technological viability of DSA. In this paper, we describe the challenges in obtaining this context information, the potential for the integration of various sources of context information, and the potential for reuse of such information for related and unrelated purposes. We use two particular examples – localization and enforcement of spectrum sharing.

The rest of the paper is organized as follows. In Section II we explain what we mean by context information of spectrum usage, why it is important and how it is not easy to readily obtain this information. Section III classifies various sources of information about spectrum usage and describes their benefits and limitations. We discuss integrating such sources in Section IV and in Section V, we examine two services – enforcement and positioning that can *reuse* the information and infrastructure for context information of spectrum usage

¹Spectrum sharing has been endorsed at the highest levels of government in the US. See, for example this Presidential Memorandum: <http://m.whitehouse.gov/the-press-office/2013/06/14/presidential-memorandum-expanding-americas-leadership-wireless-innovatio>

in DSA. Finally, we conclude the paper.

II. CONTEXT INFORMATION OF SPECTRUM USAGE

Figure 1 shows a simplified concept of DSA for understanding the need for context information related to spectrum usage and the types of information that are needed for successful DSA implementation.




Fig. 1. Scenario for understanding spectrum usage information

In Figure 1, a primary transmitter (base station), using a certain slice of radio spectrum, provides coverage over a geographical area shown by the dashed line. The determination of this geographical area is not trivial. Coverage implies that devices designed to operate with the primary user’s service are able to reliably communicate with the primary base station. This generally implies a minimum level of the received signal strength (RSS) from the base station, which is assumed to be able to reliably detect communications from the primary device in its coverage area. The physical distances at which the RSS is larger than the minimum acceptable RSS P_{min} depend on a variety of factors such as the radio frequency, the operating environment which impacts path-loss (or attenuation in signal strength), the antenna types/directionality/mounting at the transmitter and receiver, the modulation schemes, the data rates, and interference faced by the receiver from other transmissions at the same frequencies, that may be geographically separated. Consequently, the information related to spectrum usage should include these parameters as *metadata* for the benefit of a secondary user (to estimate when and where a secondary transmission may be possible).

Despite the complexities mentioned here, this description is still a simplified view of spectrum usage. The traffic characteristics of the primary user can determine “how much” of the spectrum or radio resources are actually being used (e.g., in CDMA systems, the number of active transmissions determine the amount of capacity left in a certain frequency band). Simple occupancy of the spectrum by the primary user may not necessarily imply that there is “harmful interference” to a primary receiver if a secondary user were to use the

spectrum at the same time. For example, consider a primary receiver at a distance d that is close to the primary base station. The RSS from the base station, in a simple case, may be written as [4]:

$$P_{\text{rx,primary}} = P_{\text{tx,primary}} \times K \times d^{-\alpha}$$

where the constant K depends on the metadata mentioned above and the path-loss exponent α also depends on some of the metadata and the operating environment (antenna heights, urban, rural or indoor areas, etc.). Let us now suppose that a secondary transmitter is located at a distance l from the primary receiver. The interference power from the secondary transmitter at the primary receiver will be:

$$P_{\text{rx,interference}} = P_{\text{tx,secondary}} \times K \times r^{-\alpha}$$

assuming that the characteristics of the secondary-tx to primary-rx link are identical to those of the primary link. In reality, this link is likely to have a different path-loss exponent. The signal to interference ratio for the primary user will be:

$$SIR = \frac{P_{\text{rx,primary}}}{P_{\text{rx,interference}}} = \frac{P_{\text{tx,primary}}}{P_{\text{tx,secondary}}} \times \left(\frac{r}{d}\right)^\alpha$$

We observe here that the SIR is a function of the distances d and r , the path-loss exponent α as well as the transmit powers of the primary and secondary transmitters. In theory, the secondary transmitter could transmit at a very low power and never cause interference to a primary transmitter. Small-scale and large-scale fading effects, the type of modulation schemes employed, and the data rates will determine what is really an acceptable SIR in a given situation. If there are multiple transmitters and receivers all using the same frequency in a specific geographical area, things get still more complicated.

Note that lack of “harmful” interference does not mean that there are not (negative) economic consequences for primary users. As pointed out in [5], the value of spectrum is degraded when (opportunistic) secondary users are present because it decreases the long term potential benefit of the license and demonstrates to a (potential) buyer of the spectrum that it is potentially costly to exclude external users (i.e., to enforce the usage rights granted by the license [6]). Some spectrum management authorities have considered issuing licenses that explicitly enable sharing; in those cases, one would expect the prices at license auctions to be lower because the sharing costs would be priced in.

One approach to simplifying the problem of determining when and where secondary use is possible has been to allow secondary transmissions where primary transmissions do not exist (or cannot be detected by a primary receiver) in the first place. This is sometimes called a “spectrum hole” or a “spectrum white space”. However, as seen in Figure 1, this is not really as simple as it might seem. The secondary transmitter in Figure 1 is outside the coverage area of the primary transmitter, but depending on its transmit power and the directionality of antennas employed, it may still cause interference to a primary user within the dashed coverage area.

Another approach to simplifying the problem has been to use cooperative sharing approaches rather than opportunistic ones. In cooperative sharing, secondary users explicitly coordinate their transmission with the primary user(s) [7]. This “coordination” may take the form of usage agreements for a price. In this situation, the necessary context information would have to be provided by the primary user directly.

Such technical difficulties aside, the context information related to spectrum usage is an essential component for DSA. Such context information should include the usage metadata mentioned above, SIRs, locations of transmitters and receivers, the capabilities of receivers, fading margins, transmission schemes, data rates, and temporal usage data (e.g., traffic characteristics) as well as information related to the regulatory parameters of the spectrum, such as acceptable uses, maximum transmit powers, antenna and tower information, etc. As we will see next, most sources of context information of spectrum usage have far less information. The creation of a *radio environment map* that includes some (but not all) of this metadata has been suggested in [8], [9] to facilitate *cognitive radio* based networks (see below). However, the creation of such radio environment maps is expensive as we discuss in Sections III and IV.

III. SPECTRUM USAGE INFORMATION SOURCES

In this section, we describe the various sources of spectrum usage information. Our objective here is to indicate the differences in granularity in space and time and the information that such sources provide. In general, context information can be acquired through observation (sensing) or by asking (database query, explicit coordination with primary user). Figure 2 summarizes the five types of sources that we consider below.




Fig. 2. Sources of Spectrum Usage Information

A. Spectrum Usage Databases

A straightforward approach to decide whether or not a slice of radio spectrum is available for use by a secondary user is to find such information in a database. The FCC’s Universal Licensing System [10] maintains a list of all the primary spectrum license holders and information related to the restrictions on the spectrum (such as the geographical area where the license applies, the service, and technical aspects related to radio transmission). This database does not have any temporal usage information, nor does it provide any other guidelines for secondary users.

Temporal and geographical usage information can get fairly complex if the usage of spectrum under consideration by a primary user changes randomly in time. Databases of spectrum usage and availability make sense for slices of radio spectrum where the primary user's behavior is mostly predictable. Such spectrum slices have been characterized as *static* or *periodic* in terms of their temporal usage in [2]. TV White spaces, essentially unused channels in the frequency bands allocated for broadcast television is an example of a static radio spectrum slice. The possibility of causing harmful interference to primary users is very small when using a TV white space (where there is really no TV broadcast in that geographical region).

An example of a database that can be polled by secondary devices for information about spatially available TV white spaces is Google's spectrum database [11]. Google's spectrum database has been certified by the FCC to provide information about spectrum availability for approved TV white space devices. Such devices can determine the frequency slices that are available locally (based on their latitude and longitude) for use as secondary users of the white space spectrum. For example, at 135 N. Bellefield Avenue, Pittsburgh, PA (where the University of Pittsburgh's School of Information Sciences is located), this database indicates that there are seven available TV white space channels, each 6 MHz wide, for a total of 42 MHz of spectrum for use by secondary users (based on specific rules set by the FCC). The transmit power allowed on a specific available channel is indicated in this data. Google's database also includes information about primary users in TV channels that are **not** available for secondary use. The location of such primary transmitters (latitude, longitude), their radiated power, the TV channel that is being used, and details of antennas is included in this data enabling prediction of the region covered by such transmissions.

Spectrum measurement databases also exist in universities. The Illinois Institute of Technology's (IIT) Spectrum Observatory in Chicago provides the temporal "occupancy" of various frequency bands [12]. The measurements are from antennas mounted on the top of a 22 story building in Chicago, but they are made by a fairly sophisticated spectrum analyzer in a wide frequency range (30 MHz to 6 GHz). Consequently, the actual radio frequency measurements are quite accurate, but exist only for the particular fixed location of the three antennas. As reported in [12], the measurements in frequency bands larger than 3 GHz are unreliable because of the large attenuations of signals at such higher frequencies. The directionality of antennas imply that occupancy by GPS and satellite transmissions are not very accurate. Further, the time granularity is on the order of tens of seconds. Thus, the database, while indicative of occupancy, is not entirely useful for quantifying the availability of spectrum for sharing or dynamic access. Another example of a study of spectrum occupancy is one conducted in urban Atlanta, suburban Atlanta, and rural North Carolina by researchers at Georgia Tech [13].

B. Cognitive Radios

Cognitive radios [14], [15], first proposed in 1999, have been viewed as the primary enablers of DSA through "opportunistic spectrum access" or OSA. In the case of OSA, the cognitive radios will sense whether or not primary transmissions exist in a given area and decide on using a particular slice of the radio spectrum (with the appropriate transmit power). These *smart, distributed, context-sensitive, and spectrum agile* radios have been envisioned to (a) "learn" the spectrum usage in the environment – hence smart (b) communicate with each other without the need for a centralized architecture – hence distributed (c) develop the spectrum usage context based on sensing and measurements over a wide range of frequency bands, cooperation, and other sources of data – hence context sensitive and (d) be able to communicate over a wide range of frequency slices, switching between them as their availability/quality change – hence spectrum agile.

The progress in hardware development for accomplishing these lofty goals of cognitive radios has remained slow [16]. Cognitive radios have mostly remained in the research/academic domains without any widespread practical or commercial usage [17]. A study of various testbeds using cognitive radios in [16] indicates that most use open source GNU radios. This study also indicates that many testbeds have radios that are not necessarily spectrum agile, use specific transmission schemes (such as orthogonal frequency division multiplexing or OFDM), and use simple energy detection to determine whether a primary transmission exists or not. They are not strictly smart and there is very little "learning" in the devices in such testbeds.

Cognitive radios further suffer from the problems mentioned in Section II and in [17]. Energy detection provides a binary output of whether or not a channel is occupied and does not include the various types of context information that is often needed for DSA. As noted earlier, the absence of a primary transmission need not imply harm-free transmission, note does its existence imply harmful interference if the same spectrum slice is used. Moreover, adding sensing and geo-positioning mechanisms to cognitive radios makes them expensive hindering potential innovations in exploiting DSA. As an alternative, the authors in [3] proposed the unbundling of the sensing functionality into a separate service that we discuss next.

C. Sensing as a Service

Outsourcing the problem of spectrum hole detection and identification has the advantage of reducing the economic and energy costs of cognitive radios. The cost of sensing is spread over *all* DSA systems so they do not have to be replicated. The concept here imagines a large number of networked sensors that are dispersed geographically as shown in Figure 3 (B). These sensors, while stationary, can dynamically capture the characteristics of spectrum usage over time. Since the sensors are geographically stationary, a sensing system can also learn about the long term statistical behavior of *each* band

in the sensing range, which could lead to improved spectrum efficiency and reductions in interference events to primary users. The sensors, in theory, could detect RSS values at various locations allowing secondary users to use the spectrum more efficiently. For instance, in the discussion in Section II, the secondary user typically would be unable to transmit within the coverage area shown in Figure 1. If the sensor network were to determine that the RSS of a primary user is very high, the secondary user may be able to transmit even within the coverage area as long as the application needs of the primary user were met. Similarly, a secondary user may be efficiently precluded from transmitting outside the coverage area if harmful interference is anticipated inside the coverage area. In other words, the sensing as a service could have sufficient granularity in space, time, and information related to spectrum usage to enable DSA in a fine grained manner. Sensing as a service would still not be able to determine the traffic characteristics of the primary user which may leave some spectrum resources unused.

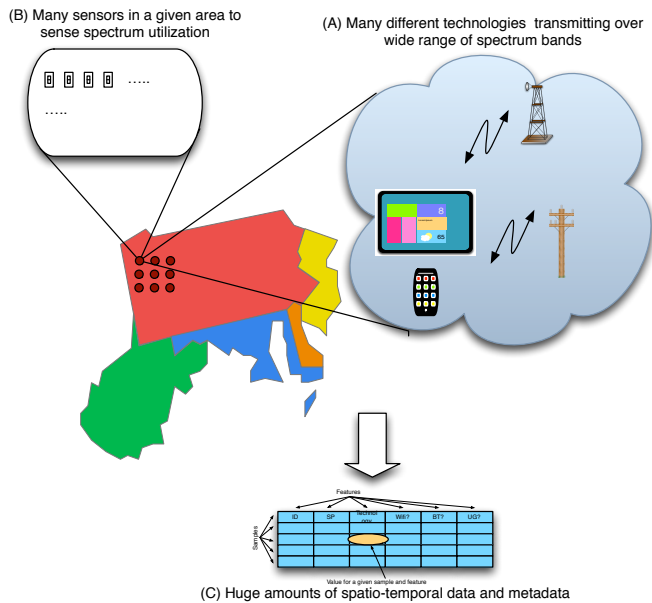


Fig. 3. Challenges in sensing as a service

But such wideband sensors would produce significant amounts of streaming data (one sensor has been estimated to produce around 10 GB of raw data per hour) that must be analyzed and processed into a usable form, perhaps rather rapidly (see Figure reffig:sense). There are tradeoffs in terms of the fidelity of sensed information and the amount of data generated, as well as the duration for which a certain radio spectrum slice is sensed [18]. Further each sensor provides partially overlapping information from nearby sensors, so this data must be integrated/fused/processed to provide a single view of the spectrum context. This will potentially need infrastructure that incorporates machine learning, data mining, and big data processing capabilities.

Yet another challenge related to sensor networks is the

required capital costs. For the network to be effective, a significant geographic region must be outfitted with sensors. The cost of this network depends on the needed density of the sensors, and this, depends on the frequency band being measured (higher frequency bands require higher densities to detect spectrum holes properly due to higher path loss at higher frequencies). These costs have to be spread over a large number of secondary users (and a large number of services) for it to be cost effective.

In summary, the infrastructure and data transmission/processing needs of such a service could be prohibitively expensive unless the infrastructure could be reused for other purposes (which is the focus of Section V).

D. Primary Users

Primary users of spectrum form an extremely important source of context information. The primary users are the most aware of their own transmissions. In the case of unidirectional transmissions such as TV broadcast, the usage or occupancy of the spectrum slice may not entirely help without knowledge of the receivers that are actually active [17]. However, in the case of bidirectional transmissions, metadata such as the traffic load in the primary service will only be available to the primary user and will be useful if shared for DSA purposes.

Different operating contexts exist in which a primary user can share context information. In the case of the Cognitive Pilot Channel (CPC), primary users would transmit occupancy information on a common channel that could be used by all (potential) secondary users [19], [20]. This reduces or eliminates the need (or precision) of other context acquisition techniques. In explicit coordination, primary users authorize secondary users to operate in a location and under certain conditions and provide the necessary operating parameters, eliminating the need for any additional context acquisition. In several types of bidirectional services (such as cellular telephony or wireless local area networks), base stations or access points transmit a variety of information on the air for smooth operation of the network [4]. Some simple examples are the “beacon”-like signals transmitted by specific technologies to enable mobile devices to discover services, to select the best cell site or to perform handoffs. In the case of the 2G cellular service GSM, the broadcast control channel or BCCH carries some useful information. In wireless LANs, a traffic indication map allows sleeping mobiles to wake up and check whether there are packets pending for them at an access point. In general, if such beacon signals are enhanced with additional metadata that a sensor network as a service can exploit, the information carried by the beacon signals would be reused for DSA.

E. Miscellaneous

With the emergence of smartphones (with a variety of integrated sensors) and advanced applications, it is possible to obtain usage information in a different way today. There are several innovative new sources of information that could

add approximate metadata related to the context information of spectrum usage. We consider a couple of examples here.

Crowd sourcing has been applied as a useful technique to complete tasks that are easier for human beings to accomplish, to obtain quick estimates of sentiments in the population (e.g., through Twitter) and so on. Applications in smart phones that can tap into the sensors integrated within can be a valuable source of context information. For example, one could imagine an application running in the background on the smartphones of a large number of users that periodically provides location information and RSS values across a particular range of frequencies to a repository. Although spectrum usage has not been specifically considered, a crowd sourced approach towards performance analysis of cellular networks has been suggested in [21]. Large scale feedback from customers related to performance issues using an application on a smartphone has been reported in [22] in comparison to traditional customer tickets and traditional network troubleshooting approaches. In this case, the application is able to provide the location information, the coverage (essentially RSS related), and connectivity at a given time to the service provider. Similarly, it is conceivable that the application can provide other kinds of spectrum usage information to a repository. In practice, smartphones may be limited as spectrum sensors as their radios are designed to operate in a specific number of bands (those assigned to one or more carriers for mobile communications as well as the unlicensed bands) and are not tunable outside of those bands. But the number and range of such bands has increased substantially, especially with the emergence of 4G services such as LTE.

Location-based social networks [23] provide information about user activity in various locations over time, mostly reported through smartphones through what are known as “Check-ins” in venues (that can include shops, parks, businesses, homes, etc.). The user activity can be correlated with usage of radio spectrum in certain frequency bands. Such check-ins can perhaps be enhanced to include metadata that is relevant to DSA.

IV. INTEGRATING INFORMATION FROM DIVERSE SOURCES

It appears that no single solution currently exists for creating a universal repository of context information related to spectrum usage, at least one that can be easily adopted for DSA. This brings up the intriguing possibility of creating a repository that fuses/integrates data from a variety of sources. We observe that there may be different levels of granularities that may enable a sequential discovery of the spectrum opportunities for DSA in a manner similar to hybrid positioning and localization methods adopted in recent times. The analogy here is as follows: There is a long time to first fix with the global positioning system (GPS) that can be as large as a few minutes. To address this delay, hybrid localization schemes use information about neighboring cell towers or WiFi access points to quickly get a coarse granularity fix on the position of a mobile device while waiting for a more accurate reading with

GPS. In a similar manner, one can conceive of a multi-step repository where a secondary user hones in on a set of potential frequency slices for reuse through a simple database and then polls a different repository or the sensor network service for more granular information, and finally perhaps resort to local sensing (stand-alone or cooperative), towards selecting the best possible bands. However, it is quite clear from the discussion in the previous section that this can be an intricate problem because of the diversity of sources, the types of data they provide, the potential inconsistencies, and missing data.

The number and locations of sensors to obtain the necessary RSS and context information is still an unknown and an open research question. It may be too expensive (and unnecessary) to deploy sensors even every few tens of meters. It may also be impossible to have sensors at all required locations because of networking, power, zoning, and other such issues. The heights of the sensors is also an important factor, especially if DSA deployment will involve wireless systems in buildings and urban areas. In general, we can think of this as a problem of *missing* data, where the required context information is missing at locations in space and over time for a variety of reasons. Handling missing temporal and spatial data is complicated, but research in this area has explored techniques to interpolate RSS values using sampling and kriging². A survey of such methods is available in [24] where the existence of several open problems in such interpolation schemes is stated.

Another challenge in integrating context information into a unified repository is that there are no evaluations of the problems that arise with handling conflicts and for determining application requirements. For the former problem, there are open questions as to how conflicting context information either from sensors or other sources should be handled. Radio propagation is a random phenomenon that depends on the terrain, movement of devices, etc. and the fact that a conflict exists between two sensors located closely in space may not be surprising. Models such as those discussed in [24] may be used here as well. The latter problem needs application requirements to be clearly quantified, and these may change in space and time, with mobile devices employing DSA. For instance, a certain transmit power may be sufficient for a certain application’s throughput needs. This particular transmit power may be within the limits that the context information related to the primary user suggests. However an additional transmitter deployed by the secondary user at the same or different location may cause the interference to exceed that allowed requiring changes to the limits on transmit power for both devices. There are no easy solutions that have been proposed for this problem.

Finally, the “shelf life” of context information (how long the context information should be stored/maintained/consulted) can vary based on the applications. While we do not elaborate on it here, context information, especially if it is to be reused

²Kriging is a method of interpolating values of a random quantity that changes spatially using all the available samples in geographical area

for enforcement (see next section), may need to be maintained for a longer time period.

V. REUSE OF SENSING INFRASTRUCTURE FOR ENFORCEMENT AND POSITIONING

According to [2], the sensor network based approach is the most costly when implemented in a complete matrix. This may be moderated if it can be “amortized” over a large number of uses and users. If the sensor network is used only for DSA, the latter is dependent on the emergence of a robust DSA market. The former relies on information sharing with other applications/services. One such service is enforcement. If the sensor network is useful for other kinds of applications, its cost may be shared between DSA and such applications. Another such application, not entirely dependent on DSA is positioning (localization) of mobile devices. We briefly consider these two services/applications in the following sections.

A. Enforcement

By “enforcement”, we refer to the *technological ability* to ensure compliance of the rules and regulations imposed by a regulatory agency such as the FCC on secondary users in a DSA environment. Enforcement can also concern itself with compliance with explicit agreements between primary and secondary users. For example, secondary users may be required to not transmit at a power that is higher than a certain value in some frequency slice in a certain period of time in a certain geographical region. The question arises as to how one can *detect* or *prevent* violations of such rules.

A brief treatment of the basis of enforcement (especially as applied to DSA) can be found in [6]. In general, for property rights (such as spectrum licenses) to be meaningful, it is necessary to enforce those rights, which generally results in excluding certain users, uses or behaviors. Enforcement may be *ex ante*, that is, aimed at preventing interference, or *ex post*, that is, aimed at remediating and remunerating interference events. *Ex ante* enforcement usually provide less flexibility and less sharing efficiency than *ex post* approaches do (illustrated in [25]). Finally, the cost of enforcement is largely dependent on the precision to which one may want to exclude certain users or behaviors.

Sensor networks can assist *ex post* enforcement by providing records of signal power of signals transmitted by specific users in specific locations at specific times. Whether these records provide sufficient evidence to support adjudication of a dispute or a violation is outside the scope of this paper, but this does highlight how information from a sensor network might be reused for another purpose. But it also requires additional functionality from a sensor network. As originally conceived, a sensor network largely monitored *primary* user activity to identify spectrum holes. In this application, it must also monitor *secondary* user behavior. The technical requirements of this depends very much on the particular sharing arrangement, as was highlighted in [25].

Ex ante enforcement can also be aided by a sensor network. One form of *ex ante* protection in spectrum sharing is the

use of an “exclusion zone”, which is a region in which the secondary user may not transmit. A sensor network would be able to geolocate secondary user activity and signal them when they approach the boundary of the exclusion zone. In some sharing scenarios, it is possible that the primary users could turn off the secondary users transmissions when they cross into the exclusion zone.

This kind of environmental monitoring for communications systems is not without precedent. For example, there has been research work on monitoring of the medium for identifying the causes for packet losses, mostly in ad hoc networks (see for example [26]). Wireless systems will face more challenges than these kinds of approaches because of the heterogeneity of air interfaces and the multipath fading environment.

B. Positioning

While GPS is the most popular positioning infrastructure, it suffers from a variety of problems, such as time to first fix and accuracy in indoor and urban areas because of multipath and low signal strengths from satellites. GPS makes use of time of flight techniques for positioning a mobile device. The idea of *location signatures* or *fingerprints* has been suggested/employed as a reasonable alternative to time and direction of flight schemes for positioning. This has especially been the case with WiFi location fingerprinting [27]. A simple version of this method is as follows. A database of mean RSS values from various WiFi access points that can be heard at various locations is created. A simple database entry would include the (local/global) coordinates of the location and its associated fingerprint which is a vector of access point IDs (MAC addresses) and the mean RSS value measured at that location. The creation of this database is called the *offline* phase. In the *online* phase, a mobile device measures the RSS from various access points that can be heard at its location, and the resulting vector is compared to those in the database. The closest match provides the estimate for the location of the mobile device. There are numerous suggested ways in which this closest match can be obtained. Instead of only the RSS, other metrics such as the multi path profile may also be used for fingerprinting. In fact, only the existence or not of the RSS may also be used for positioning [28]. The characteristics of the RSS in WiFi for positioning has been explored in [29] where the non-stationarity of the RSS, and the variation seen due to the WiFi chipsets has been identified. These variations cause errors in position estimates of mobile devices.

The work in this area is very similar to the needs and context information for DSA in that a repository of radio measurements is used to identify the location. The radio measurement at a location is expected to be *unique*, so that it accurately identifies the location without ambiguity or errors. The measurements/values in the location signature need not only be mean RSS values. They can be “occupancy” as in [28]. They can include vendor specific information to address the differences noted in [29], so that a mobile device, with a knowledge of its chipset may poll a subset of the fingerprint

database. In many research papers, the fingerprint database is referred to as a *radio map* similar to the nomenclature used in [8] for DSA's context information.

Thus, it is very possible for a repository of context information developed for DSA to be exploited for identifying device locations. This may be through a subset of measurements that a mobile device or cognitive radio performs for accomplishing DSA. If there is sufficient granularity in the deployment of the sensor infrastructure, it may be possible for a mobile device to poll these sensors for location information (or for the sensors to sense transmissions from a mobile device and estimate its location as in [28]). In other words, the economics of the "sensing as a service" would change with rapid and accurate positioning services offered in conjunction with DSA related tasks.

VI. CONCLUSIONS

In this paper, we review the importance and intricacy of "context information" for successful emergence of dynamic spectrum access, which can boost radio spectrum availability significantly, potentially resulting in innovations in services and mobile devices. We discuss several open challenges and problems related to the creation of a useful repository of context information for DSA. The intricacy of the context information, which varies in time and space, implies that a costly infrastructure is necessary for compiling and updating this context information, likely through a network of geographically dispersed sensors producing large amounts of data. The reuse of this infrastructure and information is essential for the economic success of DSA. We describe two potential services for reusing the infrastructure and information - enforcement of DSA related rules, policies, or agreements and accurate, yet rapid positioning of a variety of mobile devices.

ACKNOWLEDGMENT

This work was supported in part by the U.S. National Science Foundation under Grant 1149422. We thank our current and former students Aylin Aksu, Mohammed Altamimi, Liu Cui, and Kamol Kaemarungsi who have helped us in our research work in this area and in refining our thoughts related to this topic.

REFERENCES

- [1] Q. Zhao and B. Sadler, "A survey of dynamic spectrum access," *Signal Processing Magazine, IEEE*, vol. 24, no. 3, pp. 79–89, may. 2007.
- [2] M. Weiss, M. Altamimi, and L. Cui, "Spatio-temporal spectrum modeling: Taxonomy and economic evaluation of context acquisition," *Elsevier Telecommunications Policy*, vol. 36, pp. 335–348, 2012.
- [3] M. Weiss, S. Delaere, and W. Lehr, "Sensing as a service: An exploration into practical implementations of dsa," *IEEE Dyspan*, April 6–9 2010.
- [4] K. Pahlavan and P. Krishnamurthy, *Principles of Wireless Networks: A Unified Approach*. Prentice Hall PTR, 2002.
- [5] M. Weiss, "Spatio-temporal spectrum holes and the secondary user," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN 2011)*. IEEE, 2011, pp. 216–222.
- [6] M. Weiss, W. Lehr, L. Cui, and M. Altamimi, "Enforcement in dynamic spectrum access systems," *Telecommunications Policy Research Conference*, September 21–23 2012.
- [7] M. B. Weiss and W. H. Lehr, "Market based approaches for dynamic spectrum assignment," University of Pittsburgh, Working Paper <http://d-scholarship.pitt.edu/2824/>, 2009.
- [8] Y. Zhao, L. Morales, J. Gaeddert, K. Bae, J.-S. Um, and J. Reed, "Applying radio environment maps to cognitive wireless regional area networks," *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySpan), 2007.*, pp. 115–118, apr. 2007.
- [9] Y. Zhao, "Enabling cognitive radios through enabling cognitive radios through radio environment maps," Ph.D. dissertation, Virginia Polytechnic Institute and State University, May 2007.
- [10] J. W. Mayo and S. Wallsten, "Enabling efficient wireless communications: The role of secondary spectrum markets," *Elsevier Information Economics and Policy*, vol. 22, pp. 61–72, 2010.
- [11] Spectrum database – Google. [Online]. Available: <https://www.google.org/spectrum/whitespace/index.html>
- [12] T. Taher, R. Bacchus, K. Zdunek, and D. Roberson, "Long-term spectral occupancy findings in Chicago," *IEEE Dyspan*, 2011.
- [13] A. J. Petrin, "Maximizing the utility of radio spectrum: Broadband spectrum measurements and occupancy model for use by cognitive radio," Ph.D. dissertation, Georgia Institute of Technology, 2005.
- [14] J. Mitola, "Cognitive radio: An integrated agent architecture for software defined radio," PhD Dissertation TRITA-IT AVH 00:01, 2000.
- [15] P. Mahonen and M. Zorzi, "Cognitive wireless networks," *IEEE Wireless Communications*, vol. 14, no. 4, pp. 4–5, 2007.
- [16] P. Pawelczak, K. Nolan, L. Doyle, O. Ser Wah, and D. Cabric, "Cognitive radio: Ten years of experimentation and development," *Communications Magazine, IEEE*, vol. 49, no. 3, pp. 90–100, 2011.
- [17] K. Pelechrinis, P. Krishnamurthy, M. Weiss, and T. Znati, "Cognitive radio networks: realistic or not?" *ACM CCR*, April 2013.
- [18] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Communications Surveys and Tutorials*, vol. 11, no. 1, 2009.
- [19] S. Delaere and P. Ballon, "Multi-level standardization and business models for cognitive radio: The case of the cognitive pilot channel," in *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, 2008, pp. 1–18.
- [20] O. Sallent, J. Perez-Romero, R. Agusti, and P. Cordier, "Cognitive pilot channel enabling spectrum awareness," in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, 2009, pp. 1–6.
- [21] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Cacares, "Obtaining in-context measurements of cellular network performance," *ACM IMC*, November 2012.
- [22] Y. Jin, N. Duffield, A. Gerber, P. Haffner, W.-L. Hsu, G. Jacobson, S. Sen, S. Venkataraman, and Z.-L. Zhang, "Large-scale app-based reporting of customer problems in cellular networks," *ACM W-MUST*, 2011.
- [23] P. Krishnamurthy and K. Pelechrinis, "Location based social networks," in *Advanced Location-Based Technologies and Services*, H. Karimi, Ed. CRC Press, June 2013.
- [24] C. Phillips, D. Sicker, and D. Grunwald, "A survey of wireless path loss prediction and coverage mapping methods," *IEEE Communications Surveys and Tutorials*, vol. 15, no. 1, First Quarter 2013.
- [25] M. Weiss, M. Altamimi, and M. McHenry, "Enforcement and spectrum sharing: A case study of the 1695-1710 mhz band," *8th International Conference on Cognitive Radio Oriented Wireless Networks (Crown-Com)*, July 8–10 2013.
- [26] J. Ning, S. Singh, K. Pelechrinis, B. Liu, S. V. Krishnamurthy, and R. Govindan, "Forensic analysis of packet losses in wireless networks," *IEEE ICNP*, 2012.
- [27] P. Krishnamurthy, "Wifi location fingerprinting," in *Advanced Location-Based Technologies and Services*, H. Karimi, Ed. CRC Press, June 2013.
- [28] A. Aksu and P. Krishnamurthy, "Sub-area localization: A simple calibration free approach," in *13th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2011.
- [29] K. Kaemarungsi and P. Krishnamurthy, "Analysis of WLAN's received signal strength indication for indoor location fingerprinting," *Elsevier Pervasive and Mobile Computing*, vol. 8, no. 2, April 2012.