Achievable Privacy-Performance Tradeoffs for Spectrum Sharing with a Sensing Infrastructure

Matthew Clark*†, Konstantinos Psounis*
*University of Southern California, Los Angeles, CA
{clarkma,kpsounis}@usc.edu
†The Aerospace Corporation, El Segundo, CA

***Abstract*—Motivated by growing demand for radio frequency spectrum, spectrum regulators are laying the groundwork for centralized, dynamic spectrum sharing systems that will enable efficient access to spectrum for new wireless technologies. With a database of spectrum user information, and an infrastructure of spectrum sensors, these sharing systems will leverage cognitive radio concepts to automatically identify suitable spectrum for users. Whether incumbent users should provide information directly to the database or rely on the sensor network for detection remains an open question with implications for the effectiveness of the sharing and the privacy of the users. Both methods present the potential to explore a privacy-performance tradeoff within the sharing system.**

In this work we assess this tradeoff with both sensing and interface obfuscation approaches in a spectrum sharing system. We identify key design parameters in a formal model for the sharing system architecture, and conduct a thorough simulation study of a real-world use case to quantify privacy and performance. While abstract models suggest sensing based solutions should compare favorably with obfuscation heuristics applied to the user interface, the performance of realistic sensor network designs suggests that achieving a favorable privacy-performance tradeoff with sensor networks may be significantly limited by practical considerations.

I. Introduction

Dramatic growth in wireless applications and technologies, such as cellular, Wi-Fi, and the internet of things, demands increased access to radio frequency spectrum. Unfortunately, desirable frequency ranges are limited, and there is no unencumbered spectrum for new services. Further, replacing legacy technologies is time consuming and expensive, meaning that rapid introduction of a new technology requires improvements to how we share spectrum.

How to best employ spectrum sharing technologies remains an open question. Decentralized cognitive radio solutions face challenges such as the ”hidden node problem,” and difficulty with remediation of misbehaving devices [1]. Centralized solutions have been introduced, offering potential efficiency advantages and simplified RF devices [2]. In this setting users interface directly with Spectrum Access Systems (SAS) which maintain databases of spectrum policy and use information.

In the U.S., the Federal Communications Commission has issued rulemakings to create a Citizens Broadband Radio Service (CBRS) managed by dynamic SAS, opening the 3550-3700 MHz band for access to new commercial services [3], [4]. New entrants are expected to share the band with incumbent systems, which will retain priority access. The SAS will interface with spectrum users, and will also be informed by an infrastructure of spectrum sensors, called the Environmental Sensing Capability (ESC). The SAS is expected to identify suitable protections to prevent harmful interference to priority/primary users (PUs) which must be enforced by the SAS when granting spectrum access to secondary users (SUs).

Government entities e.g., military radars, comprise the PUs in CBRS. Incumbent users have raised concerns about maintaining the privacy of their operations, suggesting that some PUs should rely entirely on the ESC instead of sharing information directly with the SAS [5]. The SAS would need information such as locations, frequencies, time of use, and susceptibility to interference, where any of these may be considered very sensitive by the incumbents and should be protected from exposure to a potential adversary. Privacy may also be preserved for PUs communicating directly with the SAS by obfuscating the information they provide. PU privacy depends on the accuracy and precision of the user data provided to the SAS as well as of the ESC estimation capability. Coarse precision will require more conservative access to the spectrum by SUs in order to avoid harmful interference. As a result, there is a potential tradeoff between the privacy of the PUs and the utility of the shared spectrum that can be achieved by the SUs.

In this work we study the privacy-performance tradeoff in terms of design options for the SAS and privacy strategies of the PUs. We focus on inference attacks, i.e., adversaries that attempt to learn information about the users without disrupting the system. We leverage a SAS architecture and analytical model to evaluate the SAS, validating and expanding on our findings with a thorough simulation study of the CBRS use case. As a result, we find that while ESC based solutions can theoretically achieve good performance and privacy relative to interface obfuscation mechanisms, when practical limitations of sensor network implementation are accounted for, even fairly simple interface obfuscation schemes achieve significantly better performance and privacy. Further, the privacy and performance of an ESC based system is determined largely by the density of the sensor network deployment, which should be anticipated to be a major cost factor. With a method to evaluate sensing, interface obfuscation, privacy, and performance in the

This work was supported by the Aerospace Corporation Study Assistance Fellowship program, the NSF under grants CNS-1618450 and ECCS-1444060, and CISCO Systems under a CRC grant.

ISBN 978-3-903176-02-7 © 2018 IFIP

design and operation of a centralized spectrum sharing system, we believe the results of this work will help to enable more effective spectrum sharing, freeing bandwidth for use by new wireless technologies.

The paper is organized as follows. Related work is reviewed in Section II. A system model for the SAS, ESC, and user interface is provided in Section III. In Section IV, we analytically investigate the utility-privacy tradeoff of sensing and database systems. For the specific case of CBRS, we conduct a thorough simulation study in Section V, followed by our conclusions in Section VI.

Throughout this paper, we will use uppercase to denote vectors, arrays, and their elements, and subscripts to index the elements. Lowercase will denote scalar variables, where subscripts are used to distinguish variables that are similar in nature. Similarly, we use superscripts to distinguish related arrays. Calligraphic font will be used to denote sets, and bold face to denote random variables.

II. Related Work

Privacy of spectrum sharing is studied in [6]–[11], but the formulations are limited to SU-centric privacy and cannot be directly extended to issues of PU privacy. PU privacy with a SAS is considered in [12] where the authors assess strategies for the SAS falsely denying SU resource assignments to protect PU privacy. This obfuscation strategy can be considered as a special case of the more general SAS privacy framework we present here. In [13]–[15], PU privacy preserving obfuscation methods are considered against adversary inference attacks based on observation of the SAS assignments, while in [16], a protocol is presented for PUs and SUs to both preserve their location privacy by randomly perturbing the information they report to the SAS. These works do not consider a general SAS that includes a sensing component, nor do the privacy models address adversaries able to hack the SAS directly. This precludes assessment of PU privacy in the CBRS setting and a more general methodology is needed.

Spectrum sensing has received much attention in the literature. Many formulations on the subject of cognitive radio assume sensing is conducted directly by SUs, which make local decisions on how to access the spectrum [17], [18]. The performance of an individual sensor is limited by any fading along the interference path, while uncertainty in the thermal noise floor of the sensors can also degrade performance of an individual sensor [19]. Cooperative sensing techniques have been proposed where multiple networked sensor measurements are used to achieve more accurate detection, and may be well suited for an ESC with an infrastructure of spectrum sensors. Optimal determination from the sensor measurements can provide increased robustness [20], [21], but at the cost of communication overhead and complexity. As a compromise, approximate approaches are employed to fuse measurements, including hard decision voting methods and linear fusion [22], [23]. One limitation of these approaches is that they consider a binary PU state where either the PU is present or it is not. This presents a challenge in extending these approaches to the SAS setting, where multiple PUs may be operating, further motivating the more general system model we use to study the privacy-performance tradeoff.

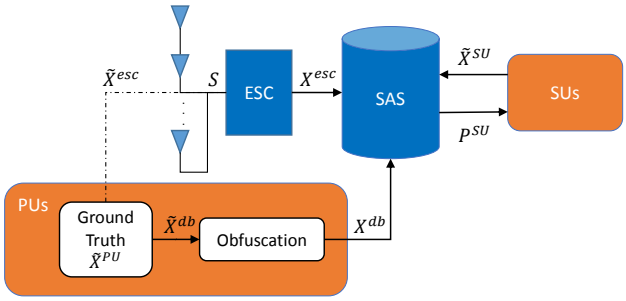

Fig. 1. System Model.

Machine learning techniques, where solutions are learned from a set of training data have been considered in limited cognitive radio and spectrum sharing settings [24], [25]. [26] treated the sensing problem with a reinforcement learning solution, but this requires deployment of sensors at PU boundary locations to provide reliable feedback to support the learning process. Since PUs are mobile and cannot be confined to a fixed area, it is not clear that such a reinforcement learning approach can be applied. Machine learning has also been applied to the sensing problem in [27], [28] where several machine learning algorithms were evaluated for their effectiveness in sensing incumbent users, but the issue of privacy was not specifically included in these models.

In this work, we formulate a general SAS framework that encompasses the ESC, database, direct user interfaces, obfuscation mechanisms, and appropriate metrics for performance and privacy. We identify and study relevant design parameters under an adversary inference attack model. Through analysis and simulation, we show under what conditions different implementations are most effective, and draw insights from the observed characteristics of the performance-privacy tradeoff.

III. System Model

Spectrum sharing with a SAS consists of a direct interface with users, an ESC composed of a network of sensors, and a process that determines spectrum assignments for SUs. Here we describe an extension to the spectrum sharing system model in [14] to include a sensing component. This system is illustrated in Figure 1 and we describe each component in the following sub-sections, deferring discussion of the privacy model to Section IV. While the model is intended to be generally applicable to spectrum sharing scenarios, to ensure concepts are clear, we focus on a specific CBRS example where the PUs operate military radars and the SUs consist of cellular network operators[1]. Suppose the SU transmissions are from the cellular user equipment (UE) to base station (BS) receivers. BS transmissions to UEs are assumed to occur at another frequency or time and can be treated analogously.

[1]CBRS also includes a tier of General Authorized Access users which can be considered as a subset of the SUs in our model.

ISBN 978-3-903176-02-7 © 2018 IFIP

A. User Interfaces

An SU requests an assignment via a connection that does not rely on spectrum from the SAS by sending information on its present location and parameters. We denote this information for n_{SU} SUs with the set $\tilde{\mathcal{X}}^{SU} = \{\mathcal{L}^{SU}, G^{SU}, \mathcal{P}\}$. $\mathcal{P}$ identifies the range of useful transmission powers, frequency tuning ranges, and a range of useful bandwidths determined by hardware limitations and application specific requirements. SU locations are given by the set $\mathcal{L}^{SU} \subseteq \mathcal{L}$, where $\mathcal{L}$ is the set of all discretized locations in the considered region. To allow the access system to take advantage of frequency dependent scheduling, an SU may also send channel state information for the links in the SU network, e.g., estimated channel gains on the UE-to-BS transmission path, denoted by the array $G^{SU} = [G^{SU}_{i\ell}]$, where ℓ is the frequency channel index and i is the index for the ith SU.

We will refer to the set of PU system and operational ground truth parameters as $\tilde{\mathcal{X}}^{PU} = \{\tilde{\mathcal{L}}^{PU}, \tilde{P}^{PU}, \tilde{I}, \tilde{\Lambda}\}$. $\tilde{\mathcal{L}}^{PU} \subseteq \mathcal{L}$ is the true set of PU locations, PUs will operate with transmission powers $\tilde{P}^{PU}$, $\tilde{I}_j$ is a harmful received interference power threshold for the jth PU, and $0 < \tilde{\Lambda}_j < 1$ is a reliability requirement for the PU, i.e., the maximum probability that the threshold given by $\tilde{I}_j$ can be exceeded. This reliability parameter accounts for inherent uncertainty due to the ESC detection process, any obfuscation strategies, and imperfect SAS prediction of aggregate interference from the SUs. Some or all of these parameters may be used as a basis for information communicated directly to the SAS database. We denote $\mathcal{X}^{db} = \{\mathcal{L}^{db}, P^{dB}, I^{db}, \Lambda^{db}\}$ as the set of potentially obfuscated information received by the SAS from the PU.

B. Environmental Sensing Component

We denote the PU information potentially detectable by the ESC with the set $\tilde{\mathcal{X}}^{esc} = \{\tilde{\mathcal{L}}^{esc}, \tilde{P}^{PU}\} \in \mathbb{X}^{esc}$, where $\mathbb{X}^{esc}$ is the set of all possible PU states. Here, $\tilde{\mathcal{L}}^{esc} \subseteq \tilde{\mathcal{L}}^{PU}$ since some PUs may be receive-only and undetectable by the ESC. Other PU transmission characteristics may also be detected by the ESC, e.g., waveform details, but for simplicity and brevity, we will not include these aspects explicitly in our formulation.

Because the operations of the incumbent military radars in CBRS are considered sensitive, SAS stakeholders have proposed that sensing systems should not enable precise geolocation of PUs [29]. Instead, sensors in the system should be limited to detecting received signal strength, and should not include the use of directional antennas to support angle of arrival estimation, nor the use of precise timing information that would support time and frequency difference geolocation techniques. With these restrictions, the ESC in SAS will consist of a network of energy detectors.

A total of n_{esc} energy detectors are deployed in the region to detect PUs. We assume quiet periods of duration ν are scheduled where no SUs are granted access to the spectrum and measurement of the PU transmissions is only affected by the sensor bandwidth b and thermal noise η_{esc}. The energy detected by the sensors, are denoted by the random vector $\mathbf{S} \in \mathbb{R}_+^{n_{esc}}$ where the randomness follows from an additive white Gaussian process for modeling the sensor thermal noise.

The sensor measurements in $\mathbf{S}$ are sent to a centralized node for estimation of the PU state, denoted $\mathcal{X}^{esc} = \{\mathcal{L}^{esc}, P^{esc}\}$. The ESC interpretation problem was studied in [28] and found to be impractical to solve optimally. For our CBRS case study, we will implement and compare machine learning solutions for an ESC, which [27] and [28] found to be effective for spectrum sensing.

We partition the region into discrete sensing cells. For any ESC implementation, we can estimate the probability of missed detection, denoted p_{md}, i.e., the probability that the ESC fails to detect a PU in a particular sensing cell, as well as the probability of false alarm, denoted p_{fa}, i.e., the probability that the ESC detects a PU present in a sensing cell where no PU is actually operating. These error rates will depend on the physical parameters of the sensor network, increasing with thermal noise η_{esc} and decreasing with the number of samples, i.e., the time-bandwidth product $b\nu$, as this allows the effect of the noise to be averaged out. The error rates will also decrease with higher density sensor deployments, as this will tend to result in more sensors nearby the PU, receiving the PU transmission with high signal-to-noise power ratio, γ.

C. SAS Assignments to SUs

The SAS manages n_c frequency channels and assigns discrete power levels over discrete time slots, where the duration of these slots are chosen as a trade between efficiency and complexity of the system. Assignments will protect n_{PU} PU locations identified in $\mathcal{X}^{db}$ and in $\mathcal{X}^{esc}$. The SAS will assume a propagation model with uncertainty when predicting channel gains between SUs and detected PU locations, i.e., $\mathbf{G}^{PU} = [\mathbf{G}^{PU}_{ij\ell}] \in \mathbb{R}^{n_{SU} \times n_{PU} \times n_c}$ is the random array for the channel gains between each PU (e.g., radar) and SU device (e.g., UE) with indices ℓ, i and j corresponding to the frequency channel, the SU and the PU respectively.

For each time slot, the SAS will allocate spectrum to SUs to maximize some utility function subject to robust constraints protecting the PUs from harmful interference. A SAS assignment function $f()$ should return maximum transmit power allocations for each SU-channel pair as an array $P^{SU} = [P^{SU}_{i\ell}] \in \mathbb{R}_+^{n_{SU} \times n_c}$ such that

$$Pr\left(\sum_{\ell=1}^{n_c}\sum_{i=1}^{n_{SU}} P_i^{SU}\mathbf{G}^{\mathbf{PU}}_{\mathbf{ij}\ell} \geq I_j^{db}\right) \leq \Lambda_j^{db} \qquad (1)$$

A power assignment of zero excludes an SU from transmitting in the corresponding frequency channels during this time slot. In this way, $f()$ acts as an admission control, channel assignment, and power assignment function. The selection of $f()$ as well as any obfuscation involved in the reporting of $\mathcal{X}^{db}$ and $\mathcal{X}^{esc}$ will affect the utility of the spectrum for the secondary users and the primary users' privacy.

Identifying solutions to (1) is non-trivial. For the purpose of this paper, we will offer a general methodology, but when a specific form for $f()$ is called for in the following results and ESC evaluations, we leverage the approach in [28], which accounts for uncertainty in both the channel state information and the sensor detection. The utility $U()$ may be left general,

ISBN 978-3-903176-02-7 © 2018 IFIP

addressing considerations including throughput, fairness and multiple access among SUs. In the remainder of this work, we take the sum-rate of the SUs as our metric for SU utility, assuming fairness between SUs is handled external to the SAS, e.g., by the cellular network operator.

IV. Privacy Analysis

Here we define the adversary threat model, identify metrics to quantify privacy, and specify obfuscation mechanisms in the SAS model. This will allow us to formally define the SAS design problem for use both in abstract analysis and for specific application to spectrum sharing in CBRS.

A. Adversary Threat Model

An adversary will observe the spectrum sharing system, with observations denoted by the set $\mathcal{Y}$. Modeling the PU state as random, an adversary makes an inference attack by estimating $p_{\hat{\boldsymbol{\mathcal{X}}}}(\tilde{\mathcal{X}}^{PU} = \mathcal{X}|\mathcal{Y})$, a distribution for the PU state given the observations. Adversaries may have different levels of access to information in the sharing system. In this work, we will specifically consider the case where the adversary has direct observation of the SAS through hacking or other means. Thus the adversary observes $\mathcal{Y} = \{\mathcal{X}^{db}, \tilde{\mathcal{X}}^{SU}, S\}$, i.e., the user communications with the SAS and the ESC measurements.

Other adversary models could be considered, e.g., an adversary might observe the assignments granted by the SAS to SUs. Related works have already shown that SU assignment obfuscation strategies can preserve PU privacy effectively under this threat model [14], and such strategies are equally applicable to sensing and interface obfuscation approaches. Similarly, an adversary that is able to hack the PU systems via the interface with the SAS might also be considered. In this case, privacy will be totally lost for any PU with their true information, $\tilde{\mathcal{X}}^{PU}$, stored on their system, corresponding to a cyber security challenge beyond the scope of the privacy analysis in this paper. For PUs that only store obfuscated parameters on their system or else rely on the ESC instead of reporting their information, the adversary does not gain any additional information by hacking these PU systems, corresponding to a case of the threat model we consider here.

B. Privacy Metrics

We measure PU privacy in terms of the quality of the adversary estimate given its observation of the SAS. We focus on the issue of location privacy for our case study in Section V, though time and frequency use privacy could be treated analogously. Specifically, we apply average distance error and search area metrics from [30].

The average distance error between the true PU locations and the adversary estimate is computed by partitioning the region into $\tilde{n}_{PU}$ sub-regions such that each partition consists of the set of cells from the original region that are closest to one location in $\tilde{\mathcal{X}}^{PU}$. Denote these sub-regions $\mathcal{L}^i$ for $i \in \{1, ..., \tilde{n}_{PU}\}$ corresponding to the indices for the PU entries that generate each sub-region. Let ℓ^i be the true location of the ith PU. Then the average distance error is given by

$$\frac{1}{n_{PU}} \sum_{i=1}^{n_{PU}} \sum_{\ell \in \mathcal{L}^i} \|\ell - \ell^i\| \frac{p_{\boldsymbol{\ell}}(\ell|\mathcal{Y})}{\sum_{\ell' \in \mathcal{L}^i} p_{\boldsymbol{\ell}}(\ell'|\mathcal{Y})}, \tag{2}$$

where $p_{\boldsymbol{\ell}}(\ell|\mathcal{Y})$ is the adversary estimated probability that any arbitrary location $\ell \in \mathcal{L}$ is contained in the true PU set $\tilde{\mathcal{L}}^{PU}$ given the observations $\mathcal{Y}$. The norm is Euclidean distance.

The area the adversary would need to search to reliably find the PUs, given the observations, serves as a second privacy metric. Let each discrete cell in the region have area α. For an estimate $\hat{\mathcal{X}}$, the search area is

$$\sum_{\ell \in \mathcal{L}} \alpha \mathbf{1} \left(\min_{\hat{\ell} \in \hat{\mathcal{X}}} \|\ell - \hat{\ell}\| \leq \max_{\ell^i \in \tilde{\mathcal{L}}^{PU}} \min_{\ell \in \hat{\mathcal{X}}} \|\ell - \ell^i\| \right), \tag{3}$$

where $\mathbf{1}()$ is the indicator function. Both the average distance error in (2), and the search area error in (3) offer an intuitive measure of the PU privacy loss. Either metric can potentially be tied to actual PU operator requirements, depending on the specific threat(s) of concern to that operator.

C. Obfuscation

The effectiveness of adversary inference attacks can be limited by two obfuscation mechanisms in the SAS architecture. First, sensor measurements are noisy. Even an optimal estimator will experience missed detections and false alarms, introducing uncertainty for the adversary. Privacy will depend on the physical parameters of the deployed sensor network. Second, the PU system may apply a random obfuscation function, translating $\tilde{\mathcal{X}}^{PU}$ to $\boldsymbol{\mathcal{X}}^{db}$. This function may add false PU entries, introduce missed detections and false alarms, as well as add random noise to the true PU values.

Let g_{ESC} and g_{PU} be functions modeling these obfuscation mechanisms. Also, let $h()$ be a function to compute a chosen privacy metric, e.g., average distance error. To quantify the effectiveness of the obfuscation methods, we can view the impact on SU utility and PU privacy in the context of a formal optimization over a time horizon of t time slots, i.e.,

$$\max_{g_{ESC}, g_{PU}} \quad \mathbb{E}\{h(\boldsymbol{\mathcal{Y}^1}, ..., \boldsymbol{\mathcal{Y}^t}, (\tilde{\mathcal{X}}^{PU})^1, ..., (\tilde{\mathcal{X}}^{PU})^t)\} \tag{4a}$$

$$\text{subject to} \quad \mathbb{E}\{U((\boldsymbol{P^{SU}})^1, ..., (\boldsymbol{P^{SU}})^t, G^{SU}, \mathcal{P})\} \geq c_t \tag{4b}$$

$$(\boldsymbol{\mathcal{X}^{esc}})^\tau = g_{ESC}((\tilde{\mathcal{X}}^{esc})^1, ..., (\tilde{\mathcal{X}}^{esc})^\tau) \tag{4c}$$

$$(\boldsymbol{\mathcal{X}^{db}})^\tau = g_{PU}((\tilde{\mathcal{X}}^{PU})^1, ..., (\tilde{\mathcal{X}}^{PU})^\tau) \tag{4d}$$

$$(\boldsymbol{P^{SU}})^\tau = f((\boldsymbol{\mathcal{X}^{db}})^1, ..., (\boldsymbol{\mathcal{X}^{db}})^\tau, (\boldsymbol{\mathcal{X}^{esc}})^1, ... \tag{4e}$$

$$..., (\boldsymbol{\mathcal{X}^{esc}})^\tau, \tilde{\mathcal{X}}^{SU}) \tag{4f}$$

$$Pr\left(\sum_{k=1}^{n_c} \sum_{i=1}^{n_{SU}} (\boldsymbol{P_i^{SU}})^\tau \mathbf{G_{kij}^{PU}} \geq \boldsymbol{I_j^{db}} \right) \leq \boldsymbol{\Lambda_j^{db}} \tag{4g}$$

$$1 \leq j \leq n_{PU},\ 1 \leq \tau \leq t. \tag{4h}$$

Problem (4) maximizes the expected PU privacy over the obfuscation functions, subject to (4b), a constraint on the

ISBN 978-3-903176-02-7 © 2018 IFIP

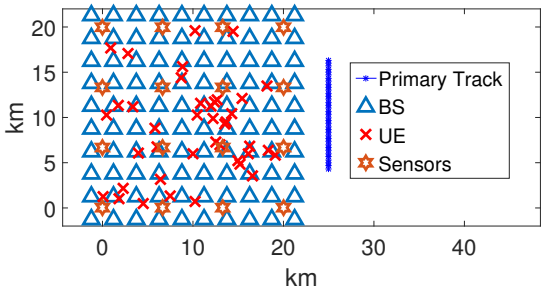

Fig. 2. Sample Topology with Ship-borne PU Radar

expected utility of the SU network exceeding threshold c_t, and to (4g), which ensures the interference to the PUs is held to a sufficiently low level.

Given that the optimization is over two functional spaces, even with limiting assumptions, solving problem (4) analytically is non-trivial and beyond the scope of this paper. However, we can apply known obfuscation heuristics and compare SAS designs rigorously in this framework. Further, the analytical formulation allows us to make a key observation about the role of the ESC. For any solution to (4) that relies solely on the implementation of an ESC to detect location information, i.e., where $\mathcal{L}^{db} = \emptyset$, there exists another solution that achieves equal or better privacy without degrading SU utility which does not employ an ESC at all, i.e., where g_{ESC} always returns an empty set. Thus, choosing to rely on detection by the ESC rather than direct communication with the SAS database cannot offer improved privacy or SU utility.

V. Simulation Results

We now approximate solutions to Problem (4) in a case study of SAS designs for the CBRS setting, quantifying privacy and performance for implementations relying on either an ESC or PU interface obfuscation.

A. System Setup

CBRS regulations specify two kinds of PU radar with parameters in [31]. Ground based PU radars operate within specified protection zones, while ship-borne radars will require interference protection from SUs in coastal areas. We assume a 20 km by 20 km region where a network of cellular SUs operate. SU BSs are deployed on a grid with a 2.5 km inter-site spacing, a 5 MHz receive bandwidth and -101.5 dBm thermal noise. A cellular network of this size could support hundreds of UEs transmitting simultaneously, but for ease of simulation, we deploy 40 UEs randomly in the region, finding this does not impact the relative performance between different SAS designs. ESC sensors are deployed on a grid with variable inter-sensor spacing. The sensors measure the full 5 MHz cellular bandwidth. An SU quiet period of variable duration is scheduled every 30 seconds, where the sensors integrate measured energy and the ESC attempts to estimate the PU state. Figure 2 plots an example topology with a track for a ship-borne PU radar moving due north, 5km off the coast from the protection region.

Ground based PUs are assumed to transmit with +30 dBm, ship-borne PUs transmit with +60 dBm, and SU power assignments from the SAS are in the range -40 dBm to +24 dBm, corresponding to typical UE transmit powers. A breakpoint model [32] is used for mean channel gain with a freespace model out to the breakpoint, and a path loss exponent of 4 applied beyond the breakpoint. The PUs, BSs, UEs and sensors are placed at heights of 15, 15, 2 and 3 meters respectively. We assume log-normal shadowing with a standard deviation of 10 dB. PU receivers are assumed to have a harmful interference power threshold $\tilde{I} = -114$ dBm. For SU utility, we assume sum-rate throughput, computed as the sum Shannon capacity of all SU assignments.

B. Machine Learning for ESC Estimation

For the ESC estimation problem, we leverage the machine learning library in MATLAB to experiment with Support Vector Machines (SVM), ensembles of decision trees, and logistic regression methods as proposed for spectrum sensing in [28]. To identify suitable parameters for our models, train our machine learning classifiers, and compare methods, we generate a set of training data with 12,000 observation-label pairs, $\{S, \tilde{\mathcal{L}}^{esc}\}$, in the training set. The first 10,000 pairs are used to train the models while the remaining 2,000 are used to verify that the models generalize beyond the data they are trained on. Each training sample is generated with an independent, identically distributed topology. Because the number of potential classes described by $\tilde{\mathcal{L}}^{esc}$ is very large, we partition the region into a grid of sensing cells, and separately train a classifier for each sensing cell.

In Figure 3, we plot the performance of ESC estimators based on a sensor grid with 2 km inter-sensor spacing, a 2 km sensing cell resolution, and observations based on the last 8 measurements of the nearest 20 sensors. True detection rates are plotted against false alarm rates as the so-called receiver operating characteristic (ROC), where the achieved rates are adjusted by selection of a decision threshold for each method. The machine learning approaches include logistic regression, two SVMs, one with a linear kernel and one with a Gaussian kernel, and three decision tree ensembles, where ensembles are produced via bagging, adaptive boosting, or random subspace sampling. A hyper-parameter search was conducted for each approach, and the results shown correspond to the hyper-parameters yielding the largest area under the ROC curve.

The linear SVM and logistic regression implementations clearly outperform the other methods. This holds for inter-sensor spacings we examined from 2 km to 10 km. On a PC with a 2.2 GHz processor, the boosted ensemble, random subspace ensemble, and SVM with the Gaussian kernel compute their classifications in hundreds of microseconds, and the remaining methods require only tens of microseconds. With its combination of speed and accuracy, we select the SVM classifier with a linear kernel and implement it as our ESC estimator in the CBRS case study.

ISBN 978-3-903176-02-7 © 2018 IFIP

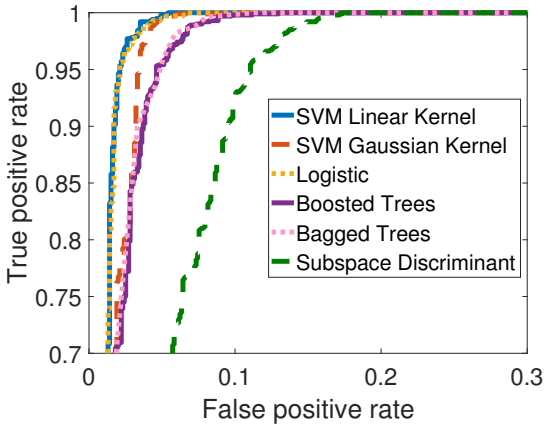

Fig. 3. ESC Receiver Operating Characteristic

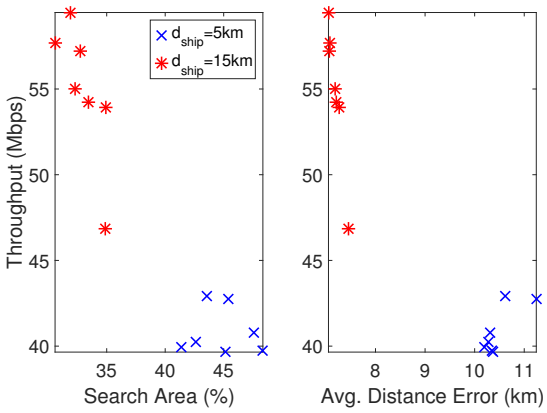

Fig. 5. PU Ship Radar and Sensing Utility vs Privacy

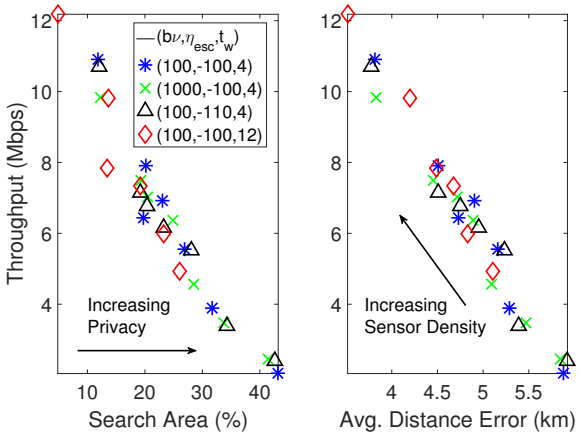

Fig. 4. PU Ground Radar and Sensing Utility vs Privacy

C. Sensing Performance

1) Ground-based PU radar: First we consider a scenario with a ground-based PU radar operating at a random location in the 20 km by 20 km region, which relies on the ESC to detect its location. Figure 4 plots the sum-rate utility of the SU network versus our two privacy metrics. Operating points are plotted for different sensor network designs, where each operating point is averaged over 100 random topologies with 20 minutes of simulated SAS operation. We consider an integration time ν of either 20 or 200 microseconds achieving a time-bandwidth product ($b\nu$) of 100 or 1000 respectively. We select a thermal noise power (η_{esc}) of -100 dBm or -110 dBm, and conduct the ESC estimate based on either the last 4, or last 12 measurements (t_w) of the nearest 20 sensors to each sensing cell. The sensing cells have a 2km resolution, and we consider inter-sensor spacings of 2km, 3.3km, 4km, 5km, 6.6km, 10km, and 20km. Note that the inter-sensor spacings are not explicitly labeled, but increasing sensor density moves along the points in the figure from right to left, increasing the achieved SU utility while sacrificing PU privacy.

Findings: The utility-privacy tradeoff is relatively linear, with achievable utility in the range 2 to 12 Mbps sum-rate. Privacy can be selected in the range of 3 to 6 km average distance error, or, with respect to search area, to between 2 and 45% of the area of the original region. In this scenario, the received radar signal to noise ratio is high at the sensors, such that the impact of thermal noise and time-bandwidth product on the accuracy of the ESC and adversary estimations is small in all cases. In contrast, the effect of the random shadowing on the channel gains is more significant, and is mitigated with more independent observations, either from an increased density of sensors, or more samples in time.

2) Ship-borne PU radar: Considering the case of a ship-borne radar moving north to south along the coast as in Figure 2, we plot the achieved privacy-utility tradeoff in Figure 5, considering sensor designs with -100 dBm thermal noise and a time-bandwidth product of 100. We examine one case where the ship is 5 km off the coast, and another where the ship is 15 km off the coast. The sensors are assumed to be deployed on land in the protection zone where the SU network operates. The ESC attempts to determine the PU state with a sensing cell resolution of 2 km over a 25 km square region of ocean closest to the SU network.

Findings: Linearity of the privacy-utility tradeoff roughly holds here. SU throughput is increased for the scenario with the ship further from the SU network, as expected, but it is somewhat surprising that the privacy is decreased. The large kW transmitter of the closer PU produces high signal-to-noise ratios at all sensors. The adversary can easily estimate the location of the PU, but is highly uncertain about the potential for other PUs that might be operating. Thus there are many potential PU states that may have produced the sequence of observations. PUs could be effectively hidden behind the high signal-to-noise ratios of the close PU. Since our metrics quantify the ability of the adversary to estimate the PU state over the entire region, the closer ship case appears to have relatively good privacy. Applying another metric tailored to the ability of an adversary to estimate any one PU would suggest much poorer privacy. For the ship 15 km off the coast, the system is able to conclude that no PUs can be operating near the coast, and is able to estimate that specifically one PU is operating. Even in this case, the most accurate sensing deployment is only able to estimate the PU location with an

ISBN 978-3-903176-02-7 © 2018 IFIP

average distance error of 7 km, and a search area that is 30% of the size of the original protection region.

D. Interface Obfuscation Performance

Now consider PUs that provide obfuscated location information directly to the SAS. We employ three obfuscation strategies. In the first, we artificially impose missed detections on the input by omitting any PU entries from the information provided to the SAS with a fixed probability. Second, we vary the resolution of the information reported to the SAS, i.e., rather than reporting that a PU is at a specific location, the PU reports that it is contained within a square area. Third, we randomly add a fixed number of false PU entries to the information communicated, and maintain those entries over the course of the simulation.

1) Ground-based PU radar: The privacy-performance tradeoff of the selected obfuscation strategies is shown in Figure 6 for the ground PU radar. Artificial missed detection probabilities and reporting resolution are labeled explicitly, while the number of false entries is not labeled, but correspond to the operating points along each curve, with zero false entries on the far left, and an increasing number of false entries as we observe operating points to the right.

Findings: Notably, there is a knee in the characteristics for search area privacy, where additional false entries have a reduced impact on the throughput of the SUs relative to the first few false entries. This is intuitive since the first few are likely to cause many SUs to switch to an interference constrained mode, while additional entries find SUs already operating at reduced powers due to interference constraints from the earlier entries. Also, artificial missed detections have a negligible effect on privacy, particularly when a larger number of fake entries are included, and the increased risk of harmful interference to the PUs may not be justifiable. In terms of average distance error, the privacy-utility tradeoff is approximately linear out to a 6 km average distance error. At this point, additional false entries have little impact on SU throughput or average distance error. Finally, reporting resolution is found to have a substantial impact on the privacy-performance tradeoff, where a coarse resolution substantially reduces the throughput but also achieves higher privacy, particularly in terms of the search area privacy.

2) Ship-borne PU radar: **Findings:** The characteristics of the privacy-performance tradeoff are shown in Figure 7 where we apply a 2 km reporting resolution. Note here that the ship operating 5 km off the coast does not have better privacy than the ship operating 15 km off the coast as we saw in the sensing case, since the high power transmission does not affect the interface obfuscation strategies. Addition of fake entries does not significantly degrade the throughput in the 5 km case since the random placement of the entries is likely to be further from the shore. The 5 km ship does appear to offer a higher achievable average distance error, but we find this is an artifact of the chosen PU operating region, where the 5 km ship operates near the region edge, and the 15 km ship operates closer to the center.

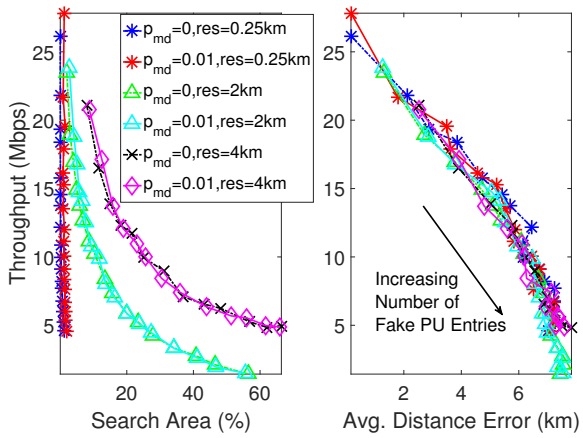

Fig. 6. PU Ground Radar Interface Obfuscation

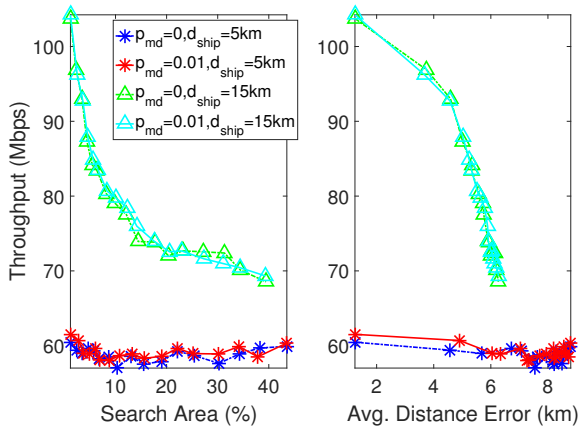

Fig. 7. PU Ship Radar Interface Obfuscation

E. Interface Obfuscation Versus Passive Sensing

In Figure 8, we compare the sensing and interface obfuscation approaches directly, with the ground-based radar scenario on the left, and the ship-borne scenario on the right. Both plots appear consistent with our finding in Section IV, i.e, for any sensing system design, we can find an interface obfuscation strategy that performs at least as well. In fact, in the ground radar scenario, we find that interface obfuscation significantly outperforms sensing, where the same level of privacy can be maintained while offering the SU network nearly double the sum-rate. While not as dramatic, interface obfuscation outperforms sensing in the ship-borne scenario as well.

We find that none of the sensing designs can offer throughput on par with the interface obfuscation strategies. This results from the difficulty in designing and modeling a sensing system that is accurate enough to offer high SU utility while also maintaining interference protection for the PUs. The most effective methods to improve sensing accuracy, i.e., higher densities of sensors and longer histories of observation, will drive the real-world complexity and cost of deploying sensor networks. In particular, deploying very high densities of energy detectors may present a prohibitively high cost since each sensor will require acquisition of hardware, real-estate to host

ISBN 978-3-903176-02-7 © 2018 IFIP

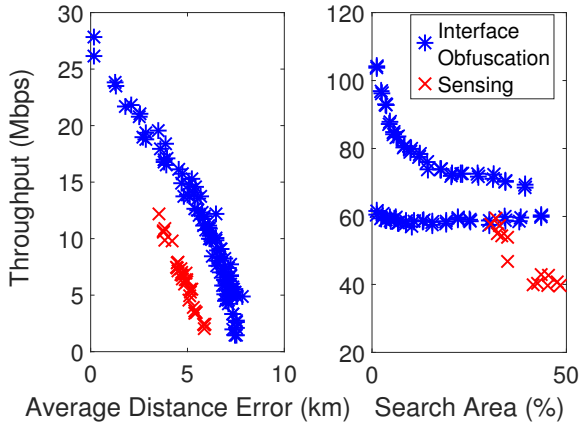

Fig. 8. Interface Obfuscation Versus Sensing

the sensor, electrical power, and a communication link back to the ESC. On the other hand, interface obfuscation strategies are all easily implementable in software, and the primary cost driver is the implementation of the interface itself.

Finally, consider the flexibility of interface obfuscation versus sensing. The performance of a sensing system is largely dependent on hardware implementations, limiting the potential to adapt to evolving technologies and operational requirements employed by PUs and SUs. Machine learning sensing estimation also presents a potential challenge for operational flexibility because such approaches are only effective when the training data accurately reflects the real-world operational environment. Any changes to the PU system hardware or operational behaviors may require retraining of the ESC estimators. Alternatively, interface obfuscation strategies can potentially be adapted on the fly, with independent PUs having the ability to dial in their own level of required privacy.

VI. Conclusions

We have modeled privacy and performance of centralized spectrum sharing systems, encompassing passive sensing and interface obfuscation with spectrum users. With abstract analysis, and through specific assessment of a practical sharing scenario, we found that practical interface obfuscation can perform at least as well as any theoretical passive sensing system based on an adversary that has hacked into the sharing system directly, and may significantly outperform realistic sensing system designs. This work can help to inform and enable the implementation of spectrum sharing systems that jointly satisfy user utility and privacy requirements.

References

[1] J. Carroll et al., "Case study: Investigation of interference into 5 GHz weather radars from unlicensed national information infrastructure devices," *US Dept. of Commerce, NTIA Report TR-12-486*, June 2012.

[2] Y. Han et al., "Spectrum sharing methods for the coexistence of multiple RF systems: A survey," *Ad Hoc Networks*, vol. 53, pp. 53 – 78, 2016.

[3] "Report and order and second further notice of proposed rulemaking," *FCC*, , no. 15-47, GN Docket No. 12-354, Apr. 2015.

[4] "Order on reconsideration and second report and order," *FCC*, , no. 16-55, GN Docket No. 12-354, May 2016.

[5] P. Atkins, "Re: Commercial operations in the 3550-3650 MHz band (GN Docket No. 12-354)," *National Telecommunications and Information Administration Letter to the FCC*, Mar. 2015.

[6] S. Li et al., "Location privacy preservation in collaborative spectrum sensing," in *Proceedings IEEE INFOCOM*, March 2012, pp. 729–737.

[7] W. Wang and Q. Zhang, "Privacy-preserving collaborative spectrum sensing with multiple service providers," *IEEE Trans. Wireless Commun.*, vol. 14, no. 2, pp. 1011–1019, Feb 2015.

[8] Z. Gao et al., "Security and privacy of collaborative spectrum sensing in cognitive radio networks," *IEEE Wireless Commun.*, vol. 19, no. 6, pp. 106–112, 2012.

[9] M. Grissa et al., "Preserving the location privacy of secondary users in cooperative spectrum sensing," *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 2, pp. 418–431, Feb 2017.

[10] Z. Qin et al., "Preserving secondary users' privacy in cognitive radio networks," in *IEEE INFOCOM*, April 2014, pp. 772–780.

[11] Z. Gao et al., "Location privacy in database-driven cognitive radio networks: Attacks and countermeasures," in *INFOCOM*. IEEE, 2013, pp. 2751–2759.

[12] A. Robertson et al., "Spectrum database poisoning for operational security in policy-based spectrum operations," in *MILCOM - IEEE Military Communications Conference*, Nov 2013, pp. 382–387.

[13] B. Bahrak et al., "Protecting the primary users' operational privacy in spectrum sharing," in *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DYSPAN)*, April 2014, pp. 236–247.

[14] M. Clark and K. Psounis, "Can the privacy of primary networks in shared spectrum be protected?," in *IEEE INFOCOM*, Apr. 2016.

[15] N. Rajkarnikar, J. M. Peha, and A. Aguiar, "Location privacy from dummy devices in database-coordinated spectrum sharing," in *IEEE DySPAN*, March 2017, pp. 1–10.

[16] Z. Zhang et al., "Achieving bilateral utility maximization and location privacy preservation in database-driven cognitive radio networks," in *IEEE 12th International Conference on Mobile Ad Hoc and Sensor Systems*, Oct 2015, pp. 181–189.

[17] Y. Gai and B. Krishnamachari, "Decentralized online learning algorithms for opportunistic spectrum access," in *IEEE GLOBECOM*, 2011.

[18] L. Luo and S. Roy, "Efficient spectrum sensing for cognitive radio networks via joint optimization of sensing threshold and duration," *IEEE Trans. Commun.*, vol. 60, no. 10, pp. 2851–2860, Oct. 2012.

[19] R. Tandra and A. Sahai, "SNR walls for signal detection," in *IEEE J. Select. Topics Signal Process.*, Feb. 2008, vol. 2, pp. 4–17.

[20] Z. Quan et al., "Optimal Multiband Joint Detection for Spectrum Sensing in Cognitive Radio Networks," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 1128–1140, Mar. 2009.

[21] G. Xiong et al., "Spectrum sensing in cognitive radio networks: Performance evaluation and optimization," in *Physical Communication*, 2013, vol. 9, pp. 171 – 183.

[22] Z. Quan et al., "Optimal Linear Cooperation for Spectrum Sensing in Cognitive Radio Networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 28–40, Feb. 2008.

[23] M. Zhou et al., "A Reliable Collaborative Spectrum Sensing Scheme Based on the ROCQ Reputation Model for Cognitive Radio Networks," in *IEEE Vehicular Technology Conference (VTC)*, May 2012, pp. 1–5.

[24] C. Clancy, J. Hecker, E. Stuntebeck, and T. O'Shea, "Applications of Machine Learning to Cognitive Radio Networks," *IEEE Wireless Commun.*, vol. 14, no. 4, pp. 47–52, Aug. 2007.

[25] M. Bkassiny et al., "A Survey on Machine-Learning Techniques in Cognitive Radios," *IEEE Communications Surveys Tutorials*, vol. 15, no. 3, pp. 1136–1159, 2013.

[26] A. Galindo-Serrano and L. Giupponi, "Distributed Q-Learning for Aggregated Interference Control in Cognitive Radio Networks," *IEEE Trans. Vehicular Technology*, vol. 59, no. 4, pp. 1823–1834, May 2010.

[27] K. M. Thilina et al., "Machine Learning Techniques for Cooperative Spectrum Sensing in Cognitive Radio Networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 11, pp. 2209–2221, Nov. 2013.

[28] M. Clark and K. Psounis, "Designing sensor networks to protect primary users in spectrum access systems," in *IEEE/IFIP WONS*, 2017.

[29] "CBRS operational security technical specification," *Wireless Innovation Forum*, , no. WINNF-15-S-0071, June 2016.

[30] M. A. Clark and K. Psounis, "Trading utility for privacy in shared spectrum access systems," *IEEE/ACM Trans. Netw. (to appear)*, 2017.

[31] E. Drocella et al., "3.5 GHz exclusion zone analyses and methodology," *US Dept. of Commerce, NTIA Report 15-517*, June 2015.

[32] A. F. Molisch, *Wireless Communications*, John Wiley & Sons Ltd., West Sussex, England, 2009.