

Deterministic Move Lists for Federal Incumbent Protection in the CBRS Band

Thao T. Nguyen and Michael R. Souryal
Communications Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, Maryland, U.S.
{thao.t.nguyen, souryal}@nist.gov

Abstract—The 3.5 GHz citizens broadband radio service (CBRS) band in the U.S. is a key portion of mid-band spectrum shared between commercial operators and existing federal and non-federal incumbents. To protect the federal incumbents from harmful interference, a spectrum access system (SAS) is required to use a common, standardized algorithm, called the move list algorithm, to suspend transmissions of some CBRS devices (CBSDs) on channels in which the incumbent becomes active. However, the current reference move list implementation used for SAS testing is non-deterministic in that it uses a Monte Carlo estimate of the 95th percentile of the aggregate interference from CBSDs to the incumbent. This leads to uncertainty in move list results and in the aggregate interference check of the test. This paper uses upper and lower bounds on the aggregate interference distribution to compute deterministic move lists. These include the reference move list used by the testing system and an operational move list used by the SAS itself. We evaluate the performance of the proposed deterministic move lists using reference implementations of the standards and simulated CBSD deployments in the vicinity of federal incumbent dynamic protection areas.

Index Terms—3.5 GHz, aggregate interference, citizens broadband radio service, incumbent protection, radar, spectrum sharing, uncertainty, upper and lower bounds.

I. INTRODUCTION

The 3.5 GHz citizens broadband radio service (CBRS) is known as an innovation band, since it was the first mid-band spectrum opened up by the Federal Communications Commission (FCC) to commercial operators on a shared basis with existing federal and non-federal incumbents [1]. This band not only provides 150 MHz of spectrum (from 3550 MHz to 3700 MHz) to new users but also proves that spectrum sharing in the mid-band is feasible, and its model can be applied to other, similar bands.

In the CBRS architecture, the band is governed by a three-tiered spectrum authorization framework. The first tier includes the federal incumbents (e.g., Navy shipborne radars), existing fixed satellite service (FSS) earth stations, and grandfathered wireless broadband licensees. These users will be protected from harmful interference from lower tiers. The second tier consists of priority access licensees (PAL) within the first 100 MHz portion of the band. The third tier is general authorized access (GAA), which allows open and flexible access to all 150 MHz of the band. The GAA users are permitted to operate on unused channels by higher tiers without causing interference to those tiers.

In the ecosystem of the CBRS band, the CBRS devices (CBSDs) are defined as fixed base stations/access points operating as PALs and GAA users to provide mobile broadband services to end user devices. The spectrum access system (SAS) is responsible for allocating spectrum resources (i.e., frequencies, power, etc.) to CBSDs as well as managing interference of lower tiers to higher tiers. The environmental sensing capability (ESC) is a sensor network that monitors the CBRS band for federal incumbent signals and notifies the SAS upon detecting an incumbent signal. The SAS then re-configures CBSDs to mitigate potential interference to the incumbent within 300 s.

The National Telecommunications and Information Administration (NTIA) has defined ESC-monitored dynamic protection areas (DPAs) that must be protected from harmful interference when an incumbent signal is detected within their limits. Most of the DPAs are coastal and typically begin 10 km offshore. However, there are some smaller, inland or port DPAs and even single-point DPAs at some sites. The DPA database for both the contiguous U.S. (CONUS) and outside of CONUS can be found in [2].

The Wireless Innovation Forum (WInnForum), an industry driven forum, has been developing CBRS standards [3] and test software [4] to foster successful deployment in the CBRS band. The standards specify a move list algorithm [3, R2-SGN-24] to pre-compute a list of CBSD transmissions that must be moved off a protected channel in the vicinity of a given DPA when an incumbent signal is detected in the DPA on the channel. The move list is designed to meet the required interference protection level, such that the 95th percentile of the aggregate interference into the incumbent radar receiver antenna does not exceed a predefined protection threshold at every point in the DPA.

The current WInnForum test harness reference implementation of the move list algorithm calculates a Monte Carlo estimate of the 95th percentile of the aggregate interference. The problem with this approach is that the result is non-deterministic, leading to uncertainties in move list size and in the aggregate interference check of the keep list. To mitigate this problem, uncertainty margins have been used in the WInnForum SAS test code for DPA protection pass/fail criteria. However, these margins were pre-computed based on a predetermined deployment and, hence, do not apply, in general, to other deployment scenarios. Another issue with

this approach is the high computational complexity of Monte Carlo estimation, especially for large deployments.

In this paper, we propose alternatives to the current move list reference implementation. Specifically, instead of relying on a Monte Carlo estimate of the aggregate interference, we use upper and lower bounds on the aggregate interference distribution to obtain deterministic move lists. One of them is a reference move list, which can be used for testing the SAS. The other is an operational move list that the SAS can use. We compare the performance of the proposed approach, in terms of move list size and aggregate interference to the federal incumbent, against the current, non-deterministic implementation. The results are discussed in detail for a single point DPA near Pensacola first, and in compact form for all coastal DPAs along the CONUS.

The remainder of this paper is organized as follows. In Section II, we discuss related work found in the literature. Section III describes the deployment model, propagation model, and current move list reference implementation and its shortcomings. In Section IV, we present applicable bounds on the aggregate interference distribution and propose alternatives to the current implementation. We analyze the performance of the proposed approach in Section V. And finally, conclusions are drawn in Section VI.

II. RELATED WORK

Calculation of aggregate interference, or co-channel interference, in wireless communications is a very important subject that has been studied extensively in the literature. The aggregate interference is simply the accumulation of interference power from several sources. Since the attenuation due to shadowing in wireless channels is often modeled by the lognormal distribution, it is common to assume that the interference contribution from a single source is lognormally distributed. Therefore, most prior work focuses on computing the aggregate interference as a sum of lognormal random variables. Since there is no closed-form expression for the lognormal sum probability density function (PDF), several analytical approximations have been derived. Most of these approaches approximate the sum of lognormal random variables by another lognormal random variable. Some of the popular methods include Fenton-Wilkinson [5] and Schwartz-Yeh [6], Beaulieu-Xie [7], and Mehta et al. [8]. Other authors propose the log-shifted Gamma approximation [9] or the log-skew normal distribution [10] to approximate the lognormal sum distribution. However, as we will show in the next section, the distribution of the interference from an individual CBSD is not simply a lognormal distribution, and hence, approximating the aggregate interference with another lognormal distribution may not be the best option for our application.

Another technique is presented in [11] to compute the PDF of a sum of two random variables on a logarithmic scale. The method can be applied recursively for more than two random variables, and it will give exact results for arbitrary distributions. A practical drawback of this method is that the results are not in closed form and extensive numerical integration is required. Given that thousands or millions of

CBSDs may be deployed in the CBRS band in the near future, using this technique for aggregate interference calculation can be a burden for the SAS operators.

Recognizing the limitations of the above methods, we turn our focus to other alternatives that estimate strict upper and lower bounds on the cumulative distribution function (CDF) of the sums. Farley's method [12] provides an upper bound on the CDF (or a lower bound on the complementary CDF) of a sum of independent, identically distributed lognormal random variables. Slimane extended Farley's upper bound method to include non-identical independent random variables as well as proposed a lower bound on the distribution function [13]. These upper and lower bounds were later generalized for correlated lognormal random variables by Tellambura [14].

In our application, interference contributions are independent but non-identically distributed, so Slimane's upper and lower bounds are suitable. And because these random variables are not lognormally distributed, generalizations of Slimane's equations, applicable to non-lognormal distributions, are used. As mentioned in [13], the upper bound is expected to be tight especially for random variables with large standard deviations. On the other hand, the lower bound is expected to be loose. Although the author proposed tighter upper and lower bounds, these tighter bounds require complicated numerical integration, and we do not see significant improvement in performance from their use.

To search for tighter lower bounds on the CDF, we consider classic concentration inequalities [15]–[17]. The concentration inequalities relate the tail area probabilities of a random variable to its moments. Among these inequalities, we found that the inequalities, which were derived by Markov, Chebyshev, Camp-Meidel, and Van Dantzig, to be the most appropriate for our application.

III. INCUMBENT INTERFERENCE PROTECTION FRAMEWORK

In this section, we describe the commercial deployment model, propagation model, and standard move list algorithm currently being used for interference protection of the incumbent.

A. Commercial Deployment Model

The simulated deployments used in this study were derived from a model used by the NTIA. The NTIA model, described in detail in [18, Section III-A], generates the locations, antenna heights, and transmission powers of a simulated deployment of CBSDs around a given DPA. The numbers of CBSDs and their locations are a function of population, land classification, and many other factors including daytime traveling factor, market penetration factor, and channel scaling factor.

There are two categories of CBSDs in the CBRS band. Category A CBSDs are lower power devices with a maximum effective isotropic radiated power (EIRP) of 30 dBm/10 MHz (i.e., 30 dB relative to 1 mW (dBm) in a 10 MHz channel) and are typically installed indoors. Whereas, Category B CBSDs are higher power devices (47 dBm/10 MHz maximum EIRP) and are professionally installed outdoors [1]. For this study,

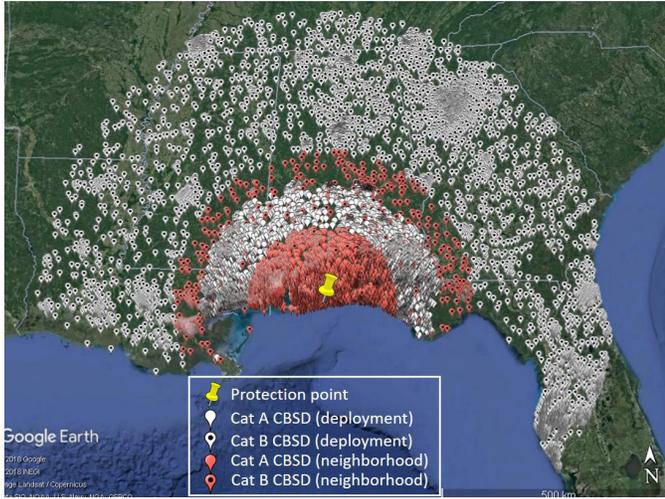


Fig. 1. Simulated deployment of CBSDs near Pensacola.

Category A CBSDs and Category B CBSDs were placed as far as 250 km and 600 km, respectively, from the DPA boundary. Furthermore, all CBSD antennas were configured to be omnidirectional. Sample deployments generated with the NTIA model can be found at [4].

Fig. 1 shows an example of a simulated deployment of CBSDs near Pensacola DPA, which is a single point protection area. The yellow pin indicates the protection point, markers without dots represent Category A CBSDs, and markers with dots represent Category B CBSDs. Even though there are 14 409 CBSDs (in white markers) deployed around the protection point, only a subset of these CBSDs within the “neighborhood” of the protection entity will be included in the aggregate interference calculation, per CBRS SAS standards [3, R2-SGN-24]. In this example, given the neighborhood distances of 150 km and 304 km for Category A and Category B CBSDs, respectively, only 5161 CBSDs (in red markers) are counted to be within the neighborhood. DPA-specific neighborhood distances and protection criteria can be found in DPA keyhole markup language (KML) files provided by NTIA [2].

B. Propagation Model

CBRS standards [3, R2-SGN-03] require a SAS to use the irregular terrain model (ITM) (also known as the Longley-Rice model) [19] in point-to-point mode for the calculation of aggregate interference to federal incumbents. The ITM model does not include clutter loss, hence, it is regarded as a conservative model for interference protection. However, to account for building attenuation, which is also absent in the ITM model, 15 dB is added to the loss if the CBSD is located indoors. Other parameters used in the ITM model are provided in [3, R2-SGN-17]. An open-source reference implementation of the ITM model is available at [4], as are the terrain and other data used by the model at [20].

According to [21], the output of the ITM model is a quantile

$$A = \begin{cases} A', & \text{if } A' \geq 0 \\ A' \frac{29-A'}{29-10A'}, & \text{otherwise} \end{cases} \quad (1)$$

and

$$A' = A_{ref} - V_{med} - Y_T - Y_L - Y_S \quad (2)$$

where A_{ref} is a reference attenuation, V_{med} is an adjustment from the reference attenuation to the all-year median, and Y_T , Y_L , Y_S are deviations due to time, location, and situation variables, respectively.

While A_{ref} and V_{med} are deterministic for a specific path, the values of Y_T , Y_L , Y_S vary and depend on three standard normal deviates, z_T , z_L , z_S , as shown in [21, (5.6), (5.9), (5.11)]. These deviates are defined as

$$z_T = z(q_T), \quad z_L = z(q_L), \quad z_S = z(q_S) \quad (3)$$

where q_T , q_L , and q_S are the desired fractions of time, locations, and situations, respectively, and $z(q) = Q^{-1}(q)$ is the inverse function of the complementary normal distribution.

In the ITM point-to-point mode, since there is no location variability, q_L is set equal to 0.5, and thus, $z_L = 0$ and $Y_L = 0$. Because broadcast mode is used in the model (mode of variability (MDVAR) = 13 in [3, R2-SGN-17(a)(v)]), time variability and situation variability are measured by reliability and confidence, respectively. Fixing the confidence parameter to 0.5 as required in [3, R2-SGN-17(a)(iv)] results in $q_S = 0.5$, and thus, $z_S = 0$ and $Y_S = 0$ [21, (5.11)]. Therefore, the time deviation Y_T is the only remaining non-zero variate in (2).

As shown in [21, (5.6)], the time deviation Y_T is piece-wise linear in z_T as follows

$$Y_T = \begin{cases} \sigma_{T_{minus}} z_T, & z_T \leq 0 \\ \sigma_{T_{plus}} z_T, & 0 \leq z_T \leq z_D \\ \sigma_{T_{plus}} z_D + \sigma_{TD}(z_T - z_D), & z_D \leq z_T \end{cases} \quad (4)$$

where $\sigma_{T_{minus}}$ and $\sigma_{T_{plus}}$ are the slopes and can be computed following the instructions in [21]. The constants z_D and σ_{TD} are related to ducting effects, and their values depend on the climate as shown in [21, Table 5.1]. Given Y_T is piece-wise linear in z_T , which is the inverse function of the complementary normal distribution of q_T , it is clear that Y_T is not simply a normal distribution of q_T . Hence, the path loss quantile output from the ITM model as a function of the reliability parameter, q_T , does not follow a normal distribution.

To illustrate the quantiles of the attenuation distribution in the ITM model, we compute the path loss from a Category A CBSD to the protection point in Pensacola. The CBSD is located indoors and is about 99 km away from the protection point. The climate value, computed using the reference implementation in [4], is maritime temperate overland. We use 10 000 reliability values, evenly spaced in the interval [0.001, 0.999].

Fig. 2 shows the path loss variation of the selected path. The first subplot shows the path loss vs. reliability and the second subplot shows the path loss histogram. Although it is not apparent in the first subplot, the second subplot clearly shows that the path loss distribution has three distinct regions with transitions at reliability values of 0.1 (green dotted line) and 0.5 (red dashed line) directly associated with the three quantile regions in (4). This can be explained by the fact that reliability values of 0.1 and 0.5 are associated with the deviate values z_T of 1.282 and 0, respectively, and because z_D for

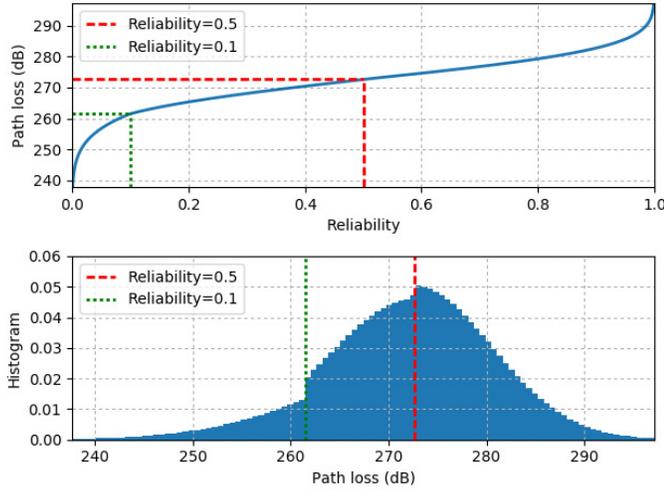


Fig. 2. Path loss example from a Category A CBSD to a single protection point in Pensacola DPA.

maritime temperate overland climate is equal to 1.282 [21, Table 5.1].

C. Standard Move List Algorithm

Regulatory rules require that the SAS operators manage their CBSD transmissions (tier 2 and tier 3) to protect the operations of existing incumbents (tier 1) in the band. To fulfill the requirement, CBRS standards specify the move list algorithm to be executed by all SASs. Details of the algorithm and its reference implementation can be found in [3, R2-SGN-24] and [4].

Given a set of CBSD transmissions that overlap in frequency with a protected frequency range, the move list algorithm identifies which transmissions must be suspended (and possibly relocated to a different channel) to avoid excessive interference in a protected federal incumbent area. In the CBRS specifications, an authorization to transmit is called a “grant.” Hence, a move list is a list of grants that must be suspended when a federal incumbent protection area becomes active. Reasons for activation of a protection area on a given channel include detection of a federal incumbent signal within the protection area on that channel.

To obtain the move list, the algorithm computes the path loss from each CBSD to a point in the protected area and, using a stochastic model for the loss on each link, computes the 95th percentile of the aggregate interference at that point. The algorithm then chooses a subset of the grants that must be suspended (relocated) such that the 95th percentile of the aggregate interference is below a threshold at any point in the protected area.

Pseudocode for the standard move list algorithm is given in Algorithm 1. For any protection point p within a given protection area DPA and for any protected frequency range ch , the standard move list algorithm first determines a set of N_c grants that are within a neighborhood of the protection point p and having or requesting a grant that includes any portion of the protected frequency range ch (Line 2). It then sorts the grants by their median interference contribution to

Algorithm 1: Standard move list algorithm

Input: Set of protection points \mathcal{P} , protection frequency range ch , set of grants \mathcal{G} , protection threshold t , receiver azimuth range \mathcal{A}

Output: DPA move list on ch , $\mathcal{M}_{DPA,ch} \subseteq \mathcal{G}$

```

1 for  $p \in \mathcal{P}$  do
2    $\mathcal{G}_{p,ch} \leftarrow \text{Neighborhood}(\mathcal{G}, p, ch)$ ; //  $\mathcal{G}_{p,ch} \subseteq \mathcal{G}$ 
3    $S_{N_c} \leftarrow \text{Sort}(\mathcal{G}_{p,ch})$ ;
   //  $S_{N_c} = [Grant_1, Grant_2, \dots, Grant_{N_c}]$ 
4   for  $a \in \mathcal{A}$  do
5      $n_c \leftarrow \text{largest } n \text{ s.t.}$ 
     Calc95thPrcntl( $S_n, a$ )  $\leq t$ 
6   end
7    $\mathcal{M}_{p,ch} = \{Grant_{n_c+1}, Grant_{n_c+2}, \dots, Grant_{N_c}\}$ 
8 end
9  $\mathcal{M}_{DPA,ch} = \bigcup_p \mathcal{M}_{p,ch}$ 

```

the protection point from smallest to largest (Line 3). The median interference contribution, $I_{i,median}(\text{dBm})(p, ch)$, of the i^{th} grant to the protection point p on frequency range ch (dBm) can be computed as follows:

$$I_{i,median}(\text{dBm})(p, ch) = P_i(ch) + G_{tx,i}(p) - L_{i,median}(p) \quad (5)$$

where $P_i(ch)$ is the conducted power of the i^{th} grant on frequency range ch in dB relative to 1 mW (dBm), $G_{tx,i}(p)$ is the transmit antenna gain in the direction of point p in dB relative to isotropic (dBi), and $L_{i,median}(p)$ is the median path loss from the transmitter to point p (dB). It is important to note that $I_{i,median}(\text{dBm})(p, ch)$ does not include the receive antenna gain, which might lead to a sub-optimality of the standard algorithm; but the advantage is that the sort need only be done once per protection point and all the subsequent calculations can be parallelized [18].

However, when computing the statistical interference contribution of an individual grant, and then, the aggregate interference, the algorithm takes into account all possible azimuth directions of the incumbent receiver antenna. The azimuth angles are computed by using increments of half beamwidth over the azimuth range of the given DPA, where the beamwidth and azimuth range are defined in [2] for each DPA. For each possible receive antenna azimuth, it must apply the gains of the transmit and receive antennas accordingly, depending on the bearing of each transmitter relative to the protection point to compute the interference contribution $I_i(\text{dBm})(p, ch)$ (dBm) as:

$$I_i(\text{dBm})(p, ch) = P_i(ch) + G_{tx,i}(p) - L_i(p) + G_{rx,i}(p, a) \quad (6)$$

where $L_i(p)$ is a sample of the random path loss from the transmitter to point p (dB) computed using Monte Carlo simulation with a minimum of 2000 trials as required in [3, R2-IPM-03], and $G_{rx,i}(p, a)$ is the receive antenna gain given the azimuth direction, a .

For the sake of notational simplicity, let $I_i(\text{dBm})$ represent $I_i(\text{dBm})(p, ch)$. Then, we define $\{I_1(\text{dBm}), \dots, I_i(\text{dBm}), \dots, I_{N_c}(\text{dBm})\}$ as a set of N_c independent, but not necessarily identical, random variables,

each representing the interference contribution from a CBSD to protection point p on frequency range ch . The associated interference contribution in linear scale (mW) can be computed as $I_i = 10^{I_i \text{ (dBm)}/10}$, for $1 \leq i \leq N_c$. Consequently, the aggregate interference I (mW) of a subset of n sorted grants, $n \leq N_c$, is the sum of the interference contribution of grants $I_1, \dots, I_i, \dots, I_n$ as follows

$$I = \sum_{i=1}^n I_i = \sum_{i=1}^n 10^{I_i \text{ (dBm)}/10} \quad (7)$$

and I (mW) can be converted to log scale by $I \text{ (dBm)} = 10 \log_{10} I$.

For all potential receiver azimuths, the algorithm finds the largest keep list (portion of the sorted list that can be kept), i.e., largest n , so that the 95th percentile of the aggregate interference, I , of these grants does not exceed the protection threshold (Line 5). This step is analogous to that of finding the largest keep list for each receiver azimuth, and then, taking the intersection of these keep lists. Note that the reference implementation of the algorithm uses Monte Carlo techniques to compute the 95th percentile of the aggregate interference. The grants that must be removed to meet the protection threshold are placed on the move list (Line 7). This process is repeated for every protection point in the protection area, and the move list for the protection area is the union of the move lists of the points (Line 9).

To better understand the algorithm, let us reconsider the example of the Pensacola DPA. Recall that this is a single point DPA and there are 5161 CBSDs within the neighborhood of the protection point.

Fig. 3(a) shows histograms of 2000 Monte Carlo samples of the individual interference contribution, $I_i \text{ (dBm)}$, computed using (6), at the receiver azimuth of 261° . To avoid overcrowding the figure, we only show representative histograms at grant indices $i = [1, 200, 400, \dots, 4800, 5000, 5161]$. Because grants are sorted by their median interference contributions (without considering the receive antenna gain), most of the histograms are shown to be gradually shifted to the right as the grant index i increases. Some of the histograms are out of order indicating that these grants are inside the main beam of the receive antenna. This is because the grants are ordered based on the median interference without the receive gain, but the plotted histograms do incorporate the receive gain. Furthermore, it can be seen from the figure that not only the median but also the variance can vary considerably from one grant to another.

Fig. 3(b) depicts the histogram of the aggregate interference, $I \text{ (dBm)}$, of the first n sorted grants at the protection point. To be consistent with Fig. 3(a), n is selected to be in $[1, 200, 400, \dots, 4800, 5000, 5161]$. As the number of grants n increases, the median increases. Although the variance in linear scale (mW) also increases with the increase of n , the variance in log scale (dB) (as shown in the figure) does not necessarily increase. However, it is clear that the aggregate interference distribution depends heavily on the distribution of the largest interference contribution among these grants.

Next, we executed the standard move list algorithm to obtain the move list for the Pensacola DPA. Note that the

Pensacola DPA has only one protection point with a protection threshold of $-139 \text{ dBm}/10 \text{ MHz}$ [2]. Out of 5161 grants in the neighborhood of the protection point, the algorithm generates a move list of size 2320 grants, leaving 2841 grants on the keep list. To check the performance of the algorithm, we calculated the aggregate interference at the protection point at the worst receiver antenna azimuth of 261° at which the incumbent receives the maximum aggregate interference. Fig. 4 shows the normalized histogram and CDF of the aggregate interference of the keep list. The 95th percentile of the aggregate interference was computed to be $-138.79 \text{ dBm}/10 \text{ MHz}$, which is slightly higher than the required protection threshold but within the uncertainty margin of the Monte Carlo process.

The uncertainty of the aggregate interference percentile, inherited from the Monte Carlo process, is the key issue addressed by this paper. The aggregate interference sample distribution changes, even with the same keep list size, every time we repeat the calculation. Fig. 5 shows the aggregate interference CDFs of the same keep list of size 2841 computed by 100 different Monte Carlo processes. The uncertainty is especially large at the head and tail portions of the distributions. We observe that the 95th percentile of aggregate interference varies in the range of $[-139.43, -138.64] \text{ dBm}/10 \text{ MHz}$ around the protection threshold of $-139 \text{ dBm}/10 \text{ MHz}$. To deal with this issue, currently, uncertainty margins are added to the aggregate interference check in the WinnForum SAS test code for DPA protection pass/fail criteria.

The uncertainty in the aggregate interference calculation also affects the move list size. Fig. 6 shows the variation in move list size for 100 trials, ranging from 2283 to 2344 grants. The inconsistency in move list calculation is undesirable in an operational setting, especially in a multi-SAS environment.

IV. DETERMINISTIC MOVE LIST ALGORITHMS

Because of the uncertainty in the current test harness reference aggregate interference calculation, it is desirable to develop an alternative that can give deterministic results. Approximations of the aggregate interference distribution with a lognormal distribution have been studied extensively [5]–[8]. However, as shown in Fig. 4, the aggregate interference distribution on a log scale is not a standard normal distribution, and hence, it is not lognormally distributed on a linear scale. Therefore, approximating the aggregate interference with a lognormal random variable does not seem to be a promising solution. Another alternative is logarithmic convolution to compute the PDF of a power sum of two random variables [11]. This method gives exact results for arbitrary distributions, but numerical integration is needed and it is computationally expensive. For these reasons, we propose using strict upper and lower bounds on the distribution of the aggregate interference. This method gives deterministic results and is computationally inexpensive. The drawback is that it cannot give exact results, and the bounds are only applicable to certain types of distributions.

A. Bounds on the Distribution of Aggregate Interference

In this section, we introduce upper and lower bounds on the distribution function of the aggregate interference. In other

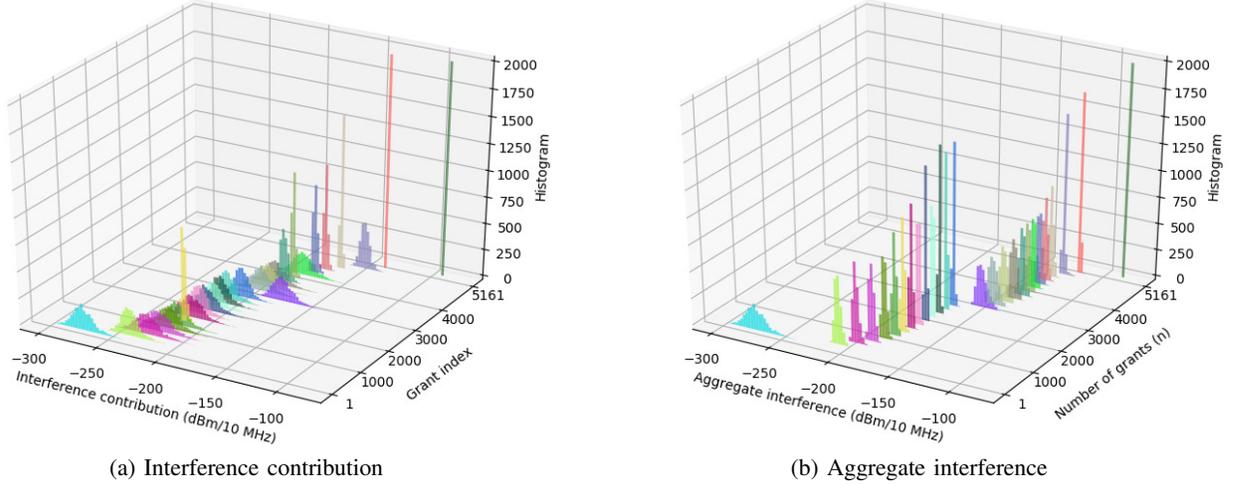


Fig. 3. Example of interference histograms from grants to a single protection point in Pensacola DPA.

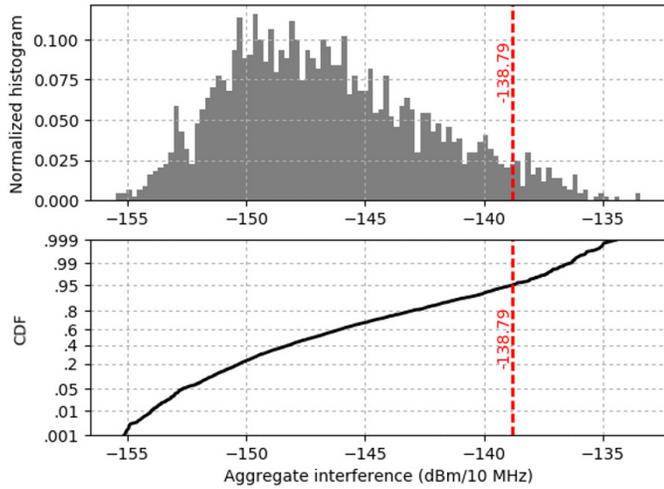


Fig. 4. Aggregate interference of keep list, Pensacola DPA.

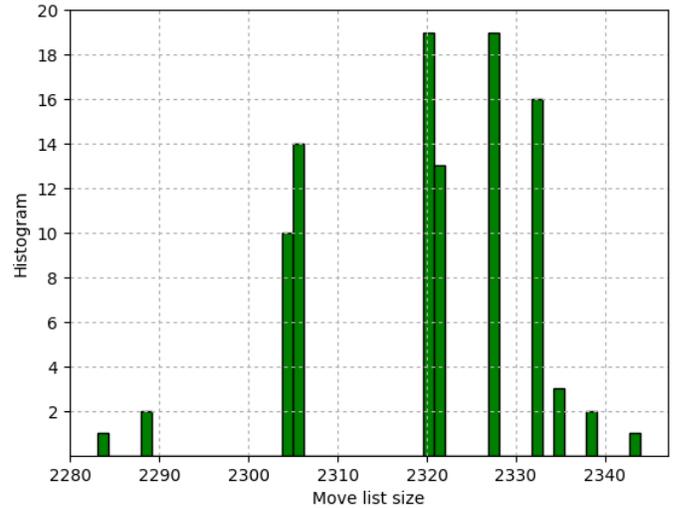


Fig. 6. Uncertainty in move list size, Pensacola DPA.

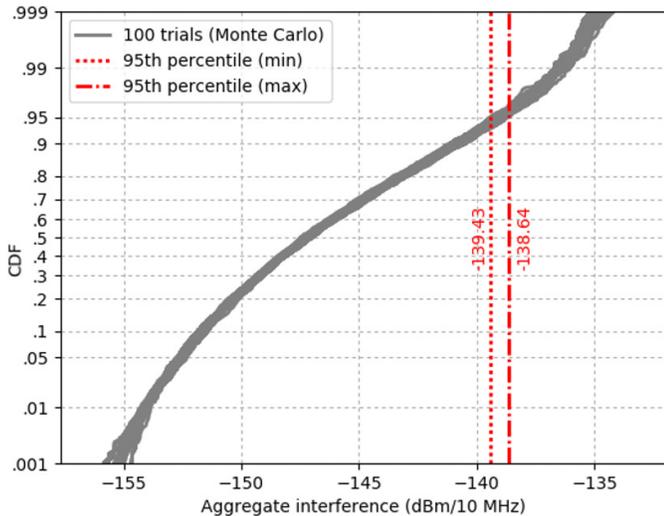


Fig. 5. Uncertainty in aggregate interference of keep list, Pensacola DPA.

words, we estimate upper and lower bounds of the CDF of I , which is a sum of N independent, but not necessarily identical, random variables I_i , where $1 \leq i \leq N$, as shown in (7). These bounds are given in closed form and can be computed easily without the need for Monte Carlo simulation.

1) *Order Statistics Based Bounds*: The upper and lower bounds of the aggregate interference distribution, which rely on the order statistics, were derived by Slimane in [13]. Let $I_{max} = \max(I_i)$ denote the maximum of the N random variables for every outcome. Then, the upper and lower bounds of the random variable I can be obtained as:

$$I_{max} \leq I \leq NI_{max} \quad (8)$$

a) *Upper Bound*: From the first inequality in (8), i.e., $I_{max} \leq I$, and $\forall x > 0$, it can be shown that

$$P(I \leq x) \leq P(I_{max} \leq x) \quad (9)$$

where,

$$\begin{aligned} P(I_{max} \leq x) &= P(I_1 \leq x, \dots, I_i \leq x, \dots, I_N \leq x) \\ &= P(I_1 \leq x) \cdots P(I_i \leq x) \cdots P(I_N \leq x) \\ &= \prod_{i=1}^N P(I_i \leq x) \end{aligned} \quad (10)$$

Let $F_X(x) = P(X \leq x)$ be the CDF of any random variable X . Then, the upper bound of the CDF of I can be obtained as follows:

$$F_I(x) \leq \prod_{i=1}^N F_{I_i}(x) \quad (11)$$

b) *Lower Bound*: Similar to the upper bound, from the second inequality in (8), $I \leq NI_{max}$, we can obtain the lower bound for the aggregate interference distribution I as follows:

$$F_I(x) \geq \prod_{i=1}^N F_{I_i}(x/N) \quad (12)$$

As mentioned in [13], the upper bound indicates that the sum is dominated by the maximum of the N random variables whereas the lower bound is obtained when the N random variables have the same outcome. Therefore, we would expect a tight upper bound especially for random variables with large standard deviations. The lower bound, on the other hand, is expected to be loose since the occurrence of its event is quite low.

To improve these bounds, the author introduced tighter bounds on the distribution by adding the contribution of the minimum and second maximum values of the N random variables to the upper bound and lower bound, respectively [13]. However, these tighter bounds require the computation of the joint distribution functions, which can be inconveniently obtained through numerical integration. Nevertheless, we found that, in our case, the tighter upper bound is similar to the previous upper bound, while the tighter lower bound is tighter than the previous lower bound but not significantly. Hence, we only make use of the upper bound in (11) in our analysis.

2) *Moments Based Lower Bounds*: We use concentration inequalities to compute lower bounds of the CDF. In probability theory, the concentration inequalities relate the tail probabilities of a random variable to its statistical moments. Therefore, they provide bounds of the deviation of a random variable away from a given value (e.g., mean value). This is a classic research topic in the field of statistics and probability [15]–[17]. Here, we focus our work only on a few inequalities that seem to be mostly applicable to our problems. Specifically, these inequalities were derived by Markov, Chebychev, Camp-Meidell, and Van Dantzig.

Let $\mu_I = E[I]$ and $\sigma_I^2 = E[(I - \mu_I)^2]$ be the mean and variance of I , respectively. Since I_i are independent random variables, the mean and variance of the sum can be computed as the sum of means and variances, respectively. In other words, $\mu_I = \sum_{i=1}^N \mu_{I_i}$ and $\sigma_I^2 = \sum_{i=1}^N \sigma_{I_i}^2$, where μ_{I_i} and $\sigma_{I_i}^2$ are the mean and variance of each random variable I_i .

a) *Markov's Inequality*: The Markov inequality [16], [17] is a fundamental inequality from which other inequalities, e.g., Chebychev's inequality, can be derived. Markov's inequality depends only on the mean of the variable. If I is a random variable taking only non-negative values, then $\forall x > 0$,

$$P(I \geq x) \leq \frac{\mu_I}{x} \quad (13)$$

b) *Chebychev's Inequality*: Since we only need to compute the lower bound of the CDF, we will only focus on the one-sided Chebyshev inequality, which is also called Cantelli's inequality [16]. Let I be a random variable with finite expected mean and variance, then $\forall x > 0$,

$$P(I \geq \mu_I + x) \leq \frac{\sigma_I^2}{\sigma_I^2 + x^2} \quad (14)$$

c) *Camp-Meidell's Inequality*: The Camp-Meidell inequality [15] shows that, $\forall x \geq 0$:

$$P(I \geq \mu_I + x) \leq \frac{4\sigma_I^2}{4\sigma_I^2 + 9x^2} \quad (15)$$

The Camp-Meidell inequality requires unimodality of the PDF of I , which is possessed by many continuous distributions such as uniform, Gaussian, lognormal, Weibull, etc. In its two-sided version, the Camp-Meidell's inequality justifies the so-called "three-sigma rule," which states that 95 % of the values are in the interval $[\mu_I - 3\sigma_I, \mu_I + 3\sigma_I]$.

d) *Van Dantzig's Inequality*: The Van Dantzig inequality [15] shows that, $\forall x \geq 0$:

$$P(I \geq \mu_I + x) \leq \frac{3\sigma_I^2}{3\sigma_I^2 + 8x^2} \quad (16)$$

This inequality requires existence of the second derivative of the probability distribution of I and convexity on the density of I . It can be applied to all the unimodal continuous probability laws in their convex part. The tail of most of the classical PDFs is convex as in our case. Although we have only shown the equations for the complementary CDFs of I , i.e., $P(I \geq x)$, the corresponding CDFs can be easily computed as $F_I(x) = P(I \leq x) = 1 - P(I \geq x)$.

To evaluate these bounds, we compute their CDFs and plot them against the aggregate interference distributions of the keep list in Pensacola DPA. Fig. 7 shows Slimane's upper bound (in red), Slimane's lower bound (in cyan), Markov's lower bound (in magenta), Chebyshev's lower bound (in green), Camp-Meidell's lower bound (in yellow), and Van Dantzig's lower bound (in blue). Since the protection requirement is based on the 95th percentile of the aggregate interference, we focus on this portion of the tail of the CDFs. Slimane's upper bound CDF is very tight to the Monte Carlo distributions and it gives a deterministic 95th percentile aggregate interference of -140.06 dBm/10 MHz. On the other hand, among all the lower bound CDFs, Van Dantzig's CDF has the tightest bound and gives a deterministic 95th percentile aggregate interference of -138.07 dBm/10 MHz.

Besides the Pensacola DPA, we also evaluate these bounds at some coastal DPAs (e.g. East5 and East7). We find that the Van Dantzig lower bound CDF might not be very tight to the Monte Carlo distribution as shown in Fig. 7, but it is the tightest bound among the lower bounds.

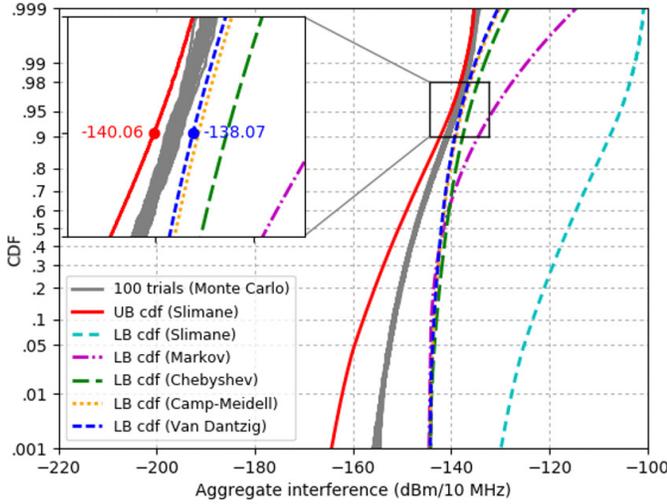


Fig. 7. Bounds of aggregate interference distribution, Pensacola DPA.

B. Deterministic Move Lists

Since Slimane’s upper bound and Van Dantzig’s lower bound provide the tightest bounds among the others on the aggregate interference distribution, we use these bounds to compute two deterministic move lists. The first is a reference move list for use by the test harness that is based on the tightest upper bound of the aggregate CDF. The second is an operational move list for use by the SAS that is based on the tightest lower bound of the aggregate CDF.

The former move list would be slightly smaller than the exact solution (to minimize false positives in testing), and the latter move list would be slightly larger than the exact solution (to minimize true positives, i.e., where the SAS computes a move list that is too small.)

1) *Deterministic Reference Move List*: The deterministic reference move list is based on the upper bound of the aggregate interference distribution. The algorithm is identical to the standard move list algorithm (Algorithm 1). The difference is in the method to compute the largest value, n_c , (to obtain the largest keep list possible) such that the 95th percentile of the aggregate interference does not exceed the protection level (Line 5). Instead of using Monte Carlo techniques to estimate the 95th percentile of the aggregate interference, we employ the upper bound of the aggregate interference distribution as follows.

To meet the protection criteria, $F_I(t) \geq 0.95$ is required. Since $F_I(x) \leq \prod_{i=1}^N F_{I_i}(x)$ in (11), we know that the 95th percentile of the aggregate interference, $F_I^{-1}(0.95)$ is lower bounded by $F_{I,UB}^{-1}(0.95)$, where $F_{I,UB}(x) = \prod_{i=1}^N F_{I_i}(x)$ and $F_{I_i}(x)$ is the CDF of the interference contribution of the i^{th} grant. As a result, the deterministic reference move list is obtained by calculating the 95th percentile in Line 5 of Algorithm 1 with $F_{I,UB}^{-1}(0.95, S_n, a)$:

$$n_c \leftarrow \text{largest } n \text{ s.t. } F_{I,UB}^{-1}(0.95, S_n, a) \leq t \quad (17)$$

Recall that S_n is a set of the first n grants sorted by their median interference contribution to the protection point from

smallest to largest, a is the azimuth direction of the receive antenna, and t is the protection threshold.

2) *Deterministic Operational Move List*: The deterministic operational move list is based on the lower bound of the aggregate distribution. After some manipulations with the Van Dantzig’s inequality (16), the corresponding CDF of the lower bound distribution can be obtained as:

$$F_I(x) \geq \frac{8(x - \mu_I)^2}{3\sigma_I^2 + 8(x - \mu_I)^2} \quad (18)$$

By simply setting

$$\frac{8(x - \mu_I)^2}{3\sigma_I^2 + 8(x - \mu_I)^2} = 0.95 \quad (19)$$

we can ensure that $F_I(x) \geq 0.95$, and hence, the protection criteria is met. From (19), and assuming $x \geq \mu_I$, we can show that the 95th percentile aggregate interference associated with the lower bound CDF can be computed as

$$x = \sqrt{\frac{57}{8}\sigma_I^2} + \mu_I \quad (20)$$

where $\frac{57}{8}$ is a simplified ratio of $\frac{3p}{8(1-p)}$ with $p = 0.95$.

Finally, the deterministic operational move list can be developed by replacing Line 5 in Algorithm 1 with the following:

$$\begin{aligned} \mu_I, \sigma_I^2 &\leftarrow \text{AggMeanVar}\{S_n, a\} \\ n_c &\leftarrow \text{largest } n \text{ s.t. } \left(\sqrt{\frac{57}{8}\sigma_I^2} + \mu_I\right) \leq t \end{aligned} \quad (21)$$

Note that in (21), the aggregate mean μ_I and variance σ_I^2 , are computed as $\mu_I = \sum_{i=1}^n \mu_{I_i}$ and $\sigma_I^2 = \sum_{i=1}^n \sigma_{I_i}^2$, where μ_{I_i} and $\sigma_{I_i}^2$ are the mean and variance of the interference contribution of the i^{th} grant in the sorted list, S_n , at the receive antenna azimuth, a .

V. RESULTS

In this section, we show examples of the proposed deterministic move lists. First, we examine results for the Pensacola DPA in detail, then, we summarize the results for the forty offshore protection areas along the U.S. coasts.

A. Single Protection Point DPA

We computed the proposed move lists in Section IV on simulated deployments of CBSDs around a protection point near Pensacola, Florida. As mentioned earlier, the Pensacola DPA is a single point, inland DPA, and it has a protection threshold of -139 dBm/10 MHz. Out of 14 409 CBSDs deployed in the vicinity of the DPA, there are only 5161 CBSDs in the neighborhood of the protection point.

1) *Move List and Keep List*: The primary output is a deterministic DPA move list, which is a list of grants that must be suspended and moved to another channel to protect the incumbent within the DPA from potential interference. The keep list, on the other hand, is just the complement of the move list, or the grants that may remain active on the protected channel.

Fig. 8 shows the sizes of the deterministic move lists and keep lists computed using the proposed formulae as well as

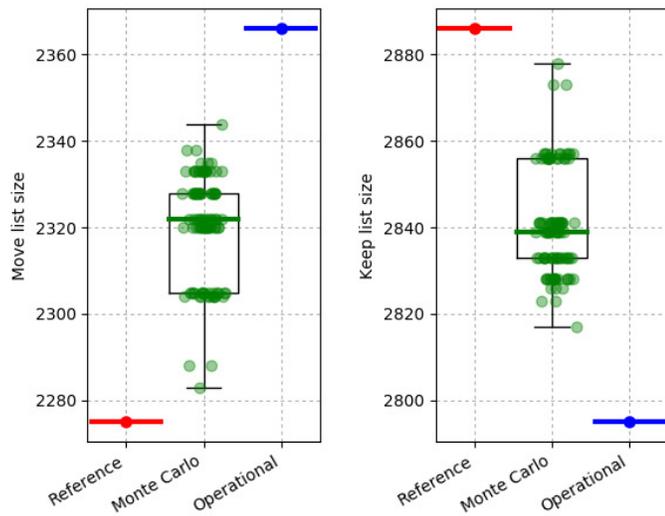


Fig. 8. Move list and keep list sizes, Pensacola DPA.

those of the Monte Carlo approach. The left subplot shows the move list sizes, whereas the right subplot shows the keep list sizes. The deterministic reference move list, computed utilizing the upper bound of the aggregate interference distribution, contains 2275 grants, and, thus, the keep list has the remaining 2886 grants (indicated in red). On the other hand, the deterministic operational move list, employing the lower bound CDF, has a slightly larger size of 2366 and, therefore, a smaller keep list size of 2795 (in blue). The box plots with jitters (in green) show the results of 100 trials of the Monte Carlo approach. As shown in the figure, the proposed deterministic move lists bound the uncertainty of the Monte Carlo outcomes. For the Pensacola DPA, since we could obtain tight upper and lower bounds of the aggregate interference distribution (as shown in Fig. 7), the bounds on the move list and keep list sizes are tight, as well.

The deterministic reference move list can be used by the test harness for testing the SAS. Since the reference move list is smaller than the exact solution, false positives in testing can be avoided. In contrast, the deterministic operational move list can be used by the SAS under test and in commercial operations. Because the operational move list is larger than the exact solution, the 95th percentile of the aggregate interference of the keep list is smaller than the protection threshold. Therefore, its use would satisfy the test requirement and provide an additional margin of protection of federal incumbents.

A geographic view of the difference between the deterministic operational and reference move lists is shown in Fig. 9. Markers without dots are Category A CBSDs, and markers with dots are Category B CBSDs. There are 91 grants in total, accounting for only 1.76 % of the total grants in the neighborhood. These are the grants that have indices between 2796 and 2886 in the sorted list. All Category A CBSDs are within 70 km of the protection point, whereas the Category B CBSDs can be anywhere extending from 115 km to the neighborhood distance of 304 km from the protection point.

2) *Aggregate Interference Check*: To examine the performance of the proposed bounds in terms of interference

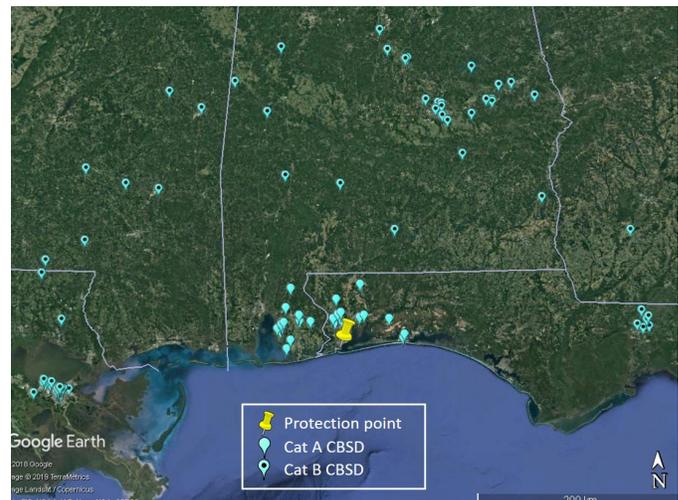


Fig. 9. Difference between operational and reference move lists, Pensacola DPA.

protection, we calculated the 95th percentile of the aggregate interference of each keep list at the protection point for every possible receiver antenna azimuth. Fig. 10 depicts the aggregate interference results. For each keep list, we applied the inequalities used for computing the upper bound CDF (i.e., Slimane’s inequality) and lower bound CDF (i.e., Van Dantzig’s inequality) to obtain the minimum and maximum of the aggregate interference, respectively.

Fig. 10(a) shows the aggregate interference results of the keep list using the upper bound CDF. All the values are below the protection threshold of -139 dBm/10 MHz for all azimuths. The strongest interference level received near the azimuth of 261° is -139.10 dBm/10 MHz for the reference keep list and -141.30 dBm/10 MHz for the operational keep list. The reference keep list has the largest size, thus creating the greatest interference at all azimuths (red line). On the other hand, the operational keep list has the smallest size, thus generating the smallest aggregate interference (blue dashed line). As expected, the aggregate interference of the 100 Monte Carlo keep lists (in green) are bounded by the aggregate interference of the deterministic keep lists.

Fig. 10(b) shows the aggregate results using the lower bound CDF. Because the lower bound CDF upper-bounds the 95th percentile, the aggregate interference values in this plot are higher than those in Fig. 10(a) at every receiver azimuth. Some of the values for the reference and Monte Carlo keep lists exceed the threshold near the azimuth of 261° . This is expected because the interference percentile is overestimated by the lower bound of the CDF.

The results in Fig. 10 are helpful to evaluate the tightness between the upper and lower bound CDFs. However, for the interference protection check, we only need to apply the upper bound CDF to the reference keep list and the lower bound CDF to the operational keep list, separately. The protection requirement is met if the 95th percentile of the aggregate interference in each case does not exceed the threshold of -139 dBm/10 MHz for all azimuths.

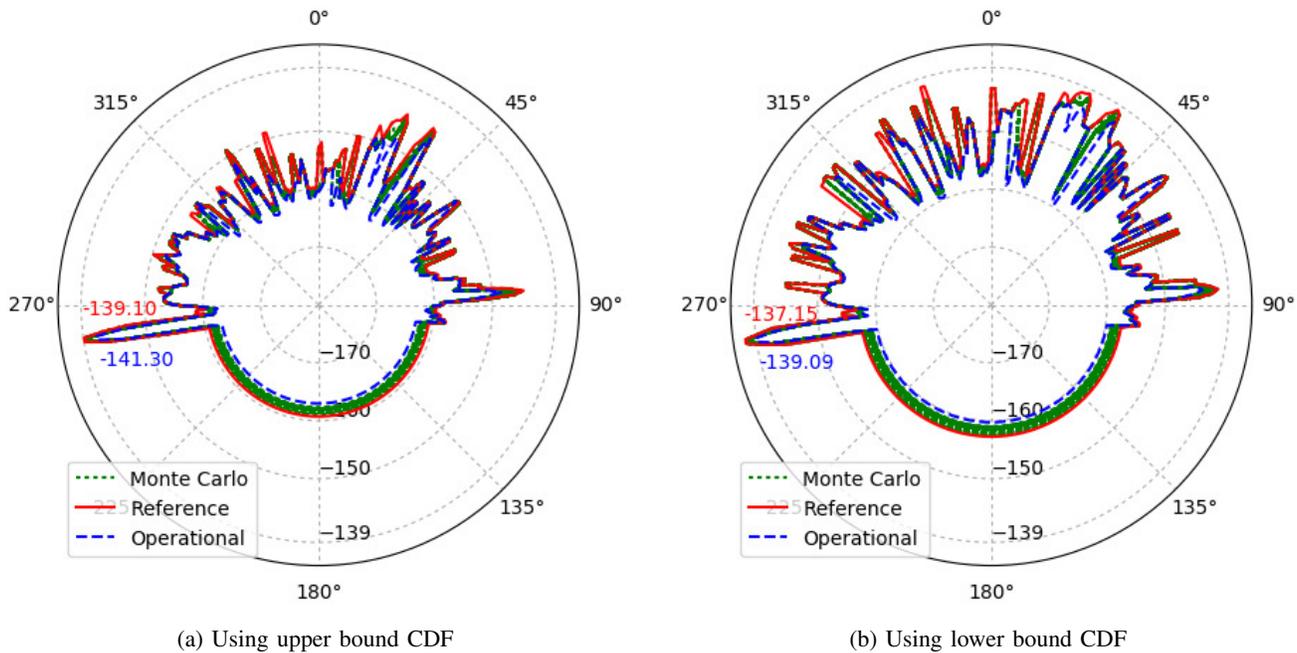


Fig. 10. 95th percentile of the aggregate interference (dBm/10 MHz) of the keep list by receiver azimuth, Pensacola DPA.

B. Results for Offshore DPAs

In this section, we present results for the 40 coastal DPAs surrounding the CONUS. There are 26 DPAs stretching along the East and Gulf coasts (East1 to East26) and 14 DPAs along the West coast (West1 to West14). Unlike the Pensacola DPA having a protection threshold of -139 dBm/10 MHz, these coastal DPAs must be protected at a lower threshold of -144 dBm/10 MHz. Also, unlike the Pensacola DPA which consists of only a single protection point, these DPAs are all protection areas sampled with multiple protection points. In this analysis, we used the “default(25,10,10,5)” protection points builder developed in [4] to generate a total of 50 protection points for each protection area.

Fig. 11(a) shows the reference move list and keep list sizes (in red and light red) and the operational move list and keep list sizes (in blue and light blue). The number of grants within the vicinity of a given DPA varies from thousands to tens of thousands, depending on the population data and geographical area surrounding the DPA. And because of the differences in terrain, some DPAs have most of their neighbor grants put on the move lists (e.g., DPAs East16 and East17 near southern Florida), whereas other DPAs have only a few grants on their move lists (e.g., DPAs West1 to West8 off the coasts of Washington, Oregon, and northern California). Regardless, as expected, the reference move list is always smaller than the operational move list for all DPAs. The ratio between the difference in the two move lists and the neighbor list size can be as small as 0.24 % for DPA West3 and as large as 14.18 % for DPA East5.

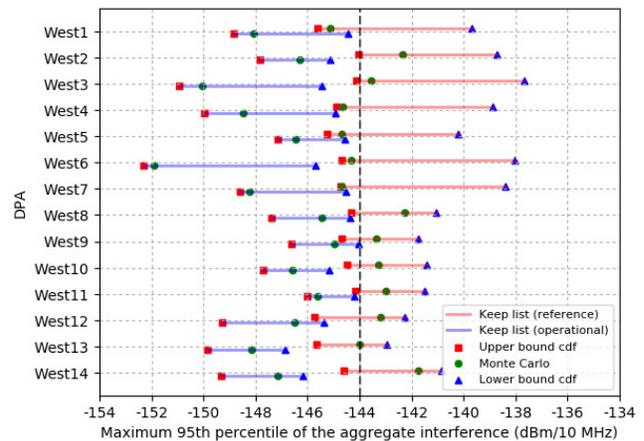
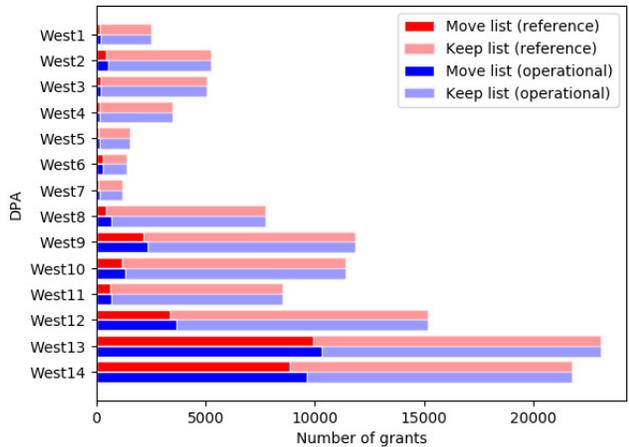
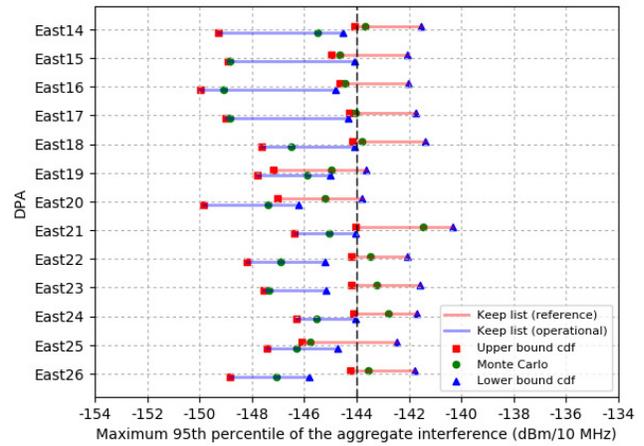
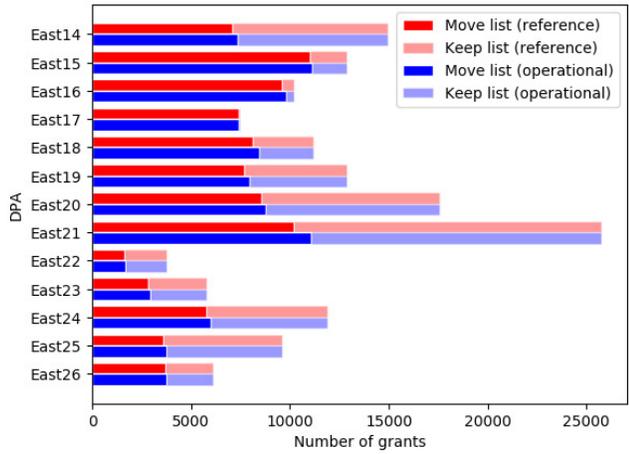
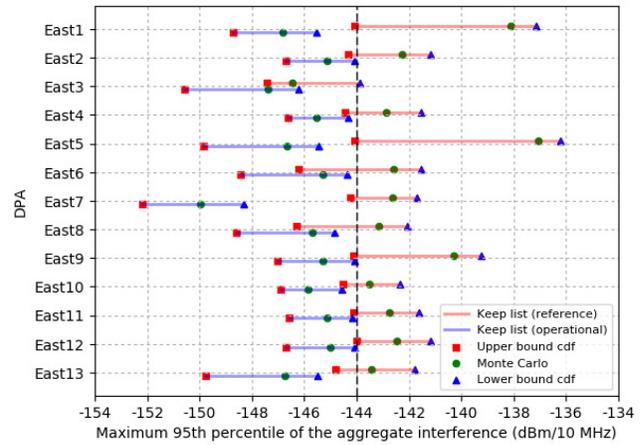
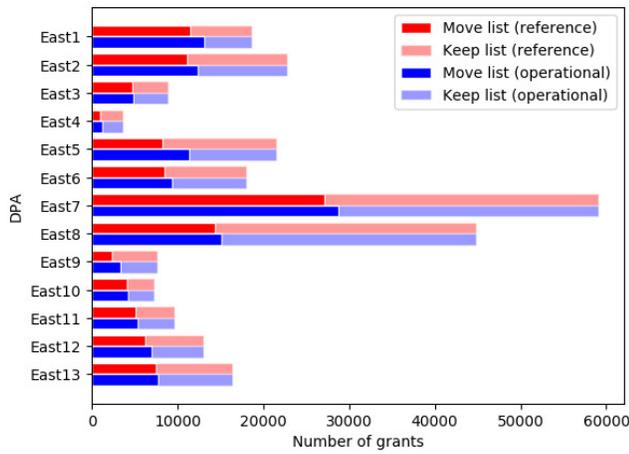
Fig. 11(b) shows the maximum aggregate interference 95th percentile of the keep list over all protection points and all receiver azimuths. Light red lines and light blue lines show the aggregate ranges of the reference keep lists and operational

keep lists, respectively. Red squares represent the deterministic aggregate interference 95th percentiles computed using the upper bound CDF, and blue triangles represent the deterministic aggregate interference 95th percentiles computed using the lower bound CDF. We observe that the difference between the two limits varies from 2 dB (e.g., for the operational keep list in DPA West11) to 8 dB (e.g., for the reference keep list in DPA East5). Large differences reveal that the bounds are not tight at some protection points within the DPA or the azimuth. Green dots show the random aggregate interference 95th percentiles computed using the standard Monte Carlo method. We notice that the green dots stay between the two limits. But in some cases (e.g., DPAs East17 and West7), the green dots are very close to the red squares, indicating the upper bound CDF is too tight to the aggregate interference distribution and might be sensitive to the accuracy of the individual interference contribution CDFs.

An important observation is that the maximum aggregate interference 95th percentile (blue triangle) of the operational keep list (blue line) is always below the -144 dBm/10 MHz protection threshold for all coastal DPAs. Therefore, if all SASs apply the same operational move list, and if the lower bound CDF is used to conservatively check the 95th percentile of the aggregate interference of the remaining keep list, we can ensure that the incumbent protection criteria are met without the need for heuristically obtained margins or the use of Monte Carlo reference move lists. In effect, the margin is built into the statistical bound used to calculate the operational move list.

VI. CONCLUSION

Current federal incumbent protection requirements in the 3.5 GHz CBRS band require the calculation of a percentile of



(a) Move list and keep list sizes

(b) Maximum aggregate interference 95th percentile

Fig. 11. Coastal DPA results: (a) Reference and operational move list and keep list sizes, (b) Maximum 95th percentile of the aggregate interference (dBm/10 MHz), over all protection points and all receiver azimuths, of the reference and operational keep lists.

the aggregate interference, which has a probability distribution with no known closed form. As a result, the current reference implementation of the federal incumbent protection (move list) algorithm resorts to the Monte Carlo method for calculating the 95th percentile of the aggregate interference. The Monte Carlo method computes the path loss from each CBSD to a protection point in the protected area using a stochastic model. This causes uncertainty in the 95th percentile of the aggregate interference, and thus, in the calculation of move list (the list of CBRS transmissions that must be moved off a channel requiring interference protection).

The inherent uncertainty in the standard move list algorithm (which uses the Monte Carlo method) requires that uncertainty margins be used when testing spectrum access systems for compliance with federal incumbent protection requirements. In practice, these margins are calculated based on assumed deployments of CBRS devices. An inadequate choice of margin can either lead to underreported test failures (missed detections) or overreported failures (false alarms).

Using bounds on the distribution function of the aggregate interference, we have presented deterministic, computational alternatives to Monte Carlo estimates of the 95th percentile of the aggregate interference. We proposed two deterministic move lists, a reference move list appropriate for testing SAS compliance and an operational move list for real-time incumbent protection. These move lists inherently build in the margins needed to avoid false alarms in testing and to conservatively protect incumbents in operation. We presented results of these move lists for all 40 DPAs along the coasts of the continental U.S. and for one inland point DPA.

One issue that we did not address in this paper is the computational load of the proposed methods. For future work, we should compare the computation performance of the proposed methods against the current Monte Carlo method. However, the proposed methods have immediate practical value for CBRS testing and implementation. They remove the uncertainty in both the 95th percentile of the aggregate interference and the move list calculation.

DISCLAIMER

The identification of any commercial product or trade name does not imply endorsement or recommendation by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

REFERENCES

- [1] "Title 47 - Part 96 - Citizens broadband radio service," Jun. 2015. [Online]. Available: <https://ecfr.io/Title-47/Part-96>
- [2] "NTIA letter to FCC on commercial operations in the 3550-3650 MHz band," Apr. 2015. [Online]. Available: <https://www.ntia.doc.gov/fcc-filing/2015/ntia-letter-fcc-commercial-operations-3550-3650-mhz-band>
- [3] "Requirements for commercial operation in the U.S. 3550-3700 MHz citizens broadband radio service band," Wireless Innovation Forum Document WINNF-TS-0112, Version V1.9.1, Reston, Virginia, U.S., Mar. 2020. [Online]. Available: <https://cbrs.wirelessinnovation.org/release-1-standards-specifications>
- [4] SAS Testing and Interoperability Repository. [Online]. Available: <https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System>

- [5] L. F. Fenton, "The sum of log-normal probability distributions in scatter transmission systems," *IRE Transactions on Communications Systems*, 1960.
- [6] S. C. Schwartz and Y. S. Yeh, "On the distribution function and moments of power sums with log-normal components," *The Bell System Technical Journal*, 1982.
- [7] N. C. Beaulieu and Q. Xie, "An optimal lognormal approximation to lognormal sum distributions," *IEEE Transactions on Vehicular Technology*, 2004.
- [8] N. B. Mehta, J. Wu, A. F. Molisch, and J. Zhang, "Approximating a sum of random variables with a lognormal," *IEEE Transactions on Wireless Communications*, 2007.
- [9] C. L. J. Lam and T. Le-Ngoc, "Log-shifted gamma approximation to lognormal sum distributions," *IEEE Transactions on Vehicular Technology*, 2007.
- [10] X. Li, Z. Wu, V. D. Chakravarthy, and Z. Wu, "A low-complexity approximation to lognormal sum distributions via transformed log skew normal distribution," *IEEE Transactions on Vehicular Technology*, 2011.
- [11] J. B. Punt and D. Sparreboom, "Summing received signal powers with arbitrary probability density functions on a logarithmic scale," *Wireless Personal Communications*, 1996.
- [12] G. L. Stuber, *Principles of Mobile Communication (2nd edition)*. Kluwer Academic Publishers, 2001.
- [13] S. B. Slimane, "Bounds on the distribution of a sum of independent lognormal random variables," *IEEE Transactions on Communications*, 2001.
- [14] C. Tellambura, "Bounds on the distribution of a sum of correlated lognormal random variables and their application," *IEEE Transactions on Communications*, 2008.
- [15] G. Blatman, B. Iooss, and N. Perot, "Probabilistic risk bounds for the characterization of radiological contamination," in *HAL multi-disciplinary open access archive*, May 2017.
- [16] R. Savage, "Probability inequalities of the Techebycheff type," 1961. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/jres/65b/jresv65bn3p211_a1b.pdf
- [17] S. Bouchron, G. Lugosi, and P. Massart, *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford, 2013.
- [18] M. R. Souryal, T. T. Nguyen, and N. J. LaSorte, "3.5 GHz federal incumbent protection algorithms," in *Proc. IEEE Dynamic Spectrum Access Networks (DySPAN)*, Oct. 2018.
- [19] Irregular Terrain Model. [Online]. Available: <https://www.its.bldrdoc.gov/resources/radio-propagation-software/itm/itm.aspx>
- [20] SAS Data Repository. [Online]. Available: <https://github.com/Wireless-Innovation-Forum/SAS-Data>
- [21] G. Hufford, "The ITS Irregular Terrain Model, version 1.2.2 - The algorithm." [Online]. Available: https://www.its.bldrdoc.gov/media/50676/itm_alg.pdf



Thao T. Nguyen received her B.S., M.S., and Ph.D. degrees, all in Electrical Engineering, from the George Mason University, Fairfax, VA, in 2004, 2006, and 2013, respectively.

She is currently with the Communications Technology Laboratory (CTL) at the National Institute of Standards and Technology (NIST). Prior to joining NIST, she was with Shared Spectrum Company, Vienna, VA, as a Senior Communications Systems Engineer for 8 years. Her main topics of interest are wireless communications and signal processing. Her focus is on research and development of spectrum sharing technologies.



Michael R. Souryal (M'89) received the B.S. degree in electrical engineering from Cornell University, Ithaca, NY, in 1990, the M.S. degree in information networking from Carnegie Mellon University, Pittsburgh, PA, in 1991, and the D.Sc. degree in electrical engineering from The George Washington University, Washington, DC, in 2003.

He is with the Communications Technology Laboratory at the National Institute of Standards and Technology (NIST) conducting research in wireless communication systems. Dr. Souryal came to NIST by way of an NRC Postdoctoral Fellowship in 2004. He was formerly with Telcordia Technologies where he was engaged in new service development for public network providers. Dr. Souryal has also held an adjunct appointment as Professorial Lecturer at The George Washington University. His research interests are in communication theory, spectrum sharing, and public safety communications.