Independent Calculation of Move Lists for Incumbent Protection in a Multi-SAS Shared Spectrum Environment

Michael R. Souryal, *Senior Member, IEEE*, and Thao T. Nguyen

***Abstract*—In the 3.5 GHz Citizens Broadband Radio Service (CBRS), secondary users are managed by spectrum access systems (SASs) to protect incumbents from interference. Current practice requires each SAS to exchange detailed user information with other SASs, and to use a common algorithm to suspend transmissions so that an aggregate interference percentile is below a predefined threshold. We propose a simplified method that utilizes a tight bound on the aggregate interference distribution. Simulation results show that the proposed approach trades off a marginal reduction in spectral efficiency to greatly simplify incumbent protection procedure, allowing each SAS to independently manage its users.**

***Index Terms*—Aggregate interference, CBRS, incumbent protection, spectrum access system, spectrum sharing.**

I. Introduction

In shared spectrum bands with prioritized users, it is common for an incumbent to require protection from harmful interference from lower priority users. Depending on the type of incumbent, the protection criterion can be defined in terms of mean, median, or a certain percentile of the aggregate interference. In the Citizens Broadband Radio Service (CBRS) in the U.S. [1], the federal incumbents are military radar systems. A spectrum access system (SAS)—a centralized spectrum resource allocation service—must ensure that the 95th percentile of the aggregate interference from lower tier users stays below a predefined threshold at the receiver of a federal radar system [2].

Computing the aggregate interference percentile is feasible if the probability distribution has a known closed form. However, in the federal incumbent protection case, the probability distribution has no known closed form, and thus, requiring estimation of the percentile. The challenge is further compounded in scenarios consisting of multiple SASs managing user transmissions in the same geographic area. In such scenarios, the SASs must coordinate to ensure that the percentile protection requirement is met. This coordination involves each SAS exchanging information of its user population (e.g., cell locations, transmission powers, antenna configurations) with every other SAS on a regular basis, so that, together, they meet the aggregate interference limits of incumbents from their collective user population. It also requires that all SASs use a common, standardized algorithm for choosing which of their users to suspend or move to a different channel—known as the "move list" in CBRS standards—when a dynamic incumbent is active on the channel.

To advance development of this band, both SAS administrators and regulators have called for a simplified incumbent protection method to facilitate operation, testing and certification of individual SASs. However, until now, no research has been found in the literature that provides a solution to this problem.

The aim of this study was to develop a simple method that allows each SAS in a multi-SAS environment to *independently* manage its user population without detailed knowledge of other SASs' users while guaranteeing that the aggregate interference percentile from all users does not exceed a predefined threshold at the incumbent receiver. Furthermore, the proposed method does not require all SASs to use a common algorithm for determining which users to suspend. Rather, each SAS is allocated an interference budget and has the flexibility to manage its user population as it sees fit to meet that budget. This letter addresses the fundamental challenge of modeling the superposition of signals having non-identical or even unknown distributions. It leverages a bound on the statistical interference to ensure that the aggregate percentile threshold is met.

The remainder of this letter is as follows. Section II provides a background on the federal incumbent protection in CBRS. Section III describes the bound on interference distribution and derives a simple limit that each SAS must comply with when managing its user population. Numerical examples for CBRS federal incumbent protection areas in the U.S. are presented in Section IV. And finally, a conclusion is drawn in Section V.

II. Federal Incumbent Protection in CBRS

The National Telecommunications and Information Administration (NTIA), which regulates U.S. federal use of spectrum, defined geographic areas where military radars may operate and require protection from harmful interference [3]. These dynamic protection areas (DPAs) are defined offshore for shipborne radars as well as on land for ground-based radars. A DPA may or may not be active (i.e., needing protection) at any given time, depending on whether a radar is active in that area. Designated sensor detects radar's signals and informs SASs of activated DPA on certain channels. SASs are collectively required to manage the usage of spectrum resources such that the 95th percentile of the aggregate interference power of all co-channel users within the "neighborhood" (i.e., a DPA-specific distance) of an active DPA is below a DPA-specific

The authors are with the National Institute of Standards and Technology, Communications Technology Laboratory, Gaithersburg, MD, 20899 USA e-mail: {souryal, thao.t.nguyen}@nist.gov.

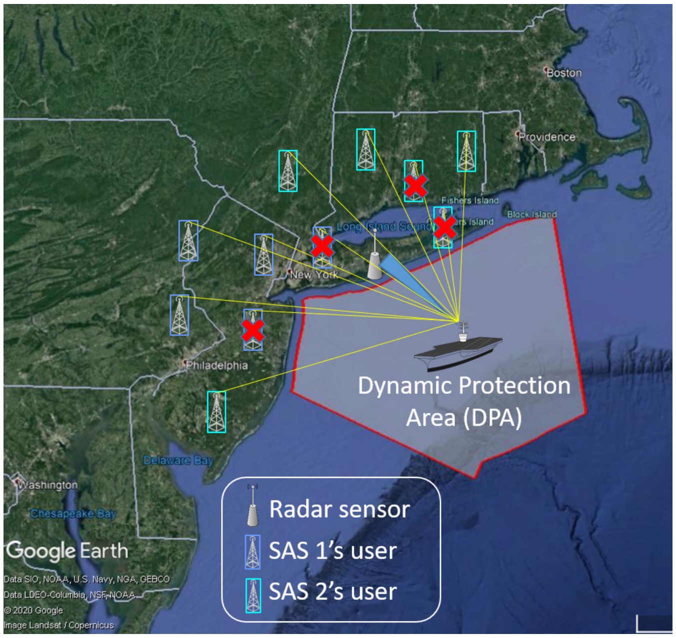

Fig. 1. A graphic illustration of systems deployed within the neighborhood of an offshore dynamic protection area (DPA) near New York.

threshold power level, e.g., −144 decibels (dB) relative to 1 mW (dBm) per 10 MHz, anywhere within the DPA. Fig. 1 illustrates an offshore DPA near New York, a radar sensor, and users managed by 2 SASs deployed near the DPA.

To mitigate interference to an active DPA, each SAS exchanges its user information and uses a standard "move list" algorithm to move users off the protected channel in the vicinity of the DPA. The algorithm considers all protection points in the protected area. For any protection point and channel, it determines all the users within neighborhood and sorts them based on their median interference contribution from smallest to largest. The algorithm finds a subset of the sorted list of users to be removed such that the 95th percentile of the aggregate interference from the remaining users is below the required threshold for all possible receiver azimuths of radar antenna. The overall move list for the DPA is the union of the move lists of all protection points. Due to its complexity, the move list is computed offline on a daily basis and only utilized when the DPA becomes active. Details and reference implementation of the move list algorithm can be found in [2, R2-SGN-24] and [4].

The standard move list algorithm uses the irregular terrain model (ITM) in point-to-point mode with time variability [5], [6] to compute the path loss from each transmitter to the incumbent receiver. The time variability results in a piecewise lognormal probability distribution of the path loss [5]. The effective isotropic radiated power (EIRP) of each transmitter is assumed fixed, therefore the received power at the incumbent from a transmitter is also piecewise lognormal.

The aggregate interference at the incumbent is taken as the power sum of all co-channel transmissions. Given the nature of the propagation model, the probability distribution of the aggregate interference power has no closed form expression and varies in each calculation. Since the standard move list algorithm uses Monte Carlo method to estimate the 95th percentile of the aggregate interference, it causes *uncertainty* in move list results. In [7], we proposed an approach to utilize bounds on the aggregate interference distribution to compute *deterministic* move lists, which can be used in testing and operational of the SASs. However, we did not address the issues that require the SASs to exchange detailed user information and to use a common, standardized algorithm. These requirements were imposed on the SASs to ensure consistency in move list calculation, but they put a strain on the SAS resources and complicate the testing process.

Building upon one of the tightest bounds on the aggregate interference distribution found in [7], we propose a simplification that can remove both these requirements. We will show that if a given pth percentile upper bound of each interference contribution is below its allocated interference budget, then the pth percentile of the aggregate interference is guaranteed to be below the sum of the allocated interference budgets.

III. Proposed Simplification

In this section, we present a tight bound on the distribution of aggregate interference and define interference criterion for each SAS based on that bound.

A. Bound on Interference Distribution

In general, the aggregate interference from multiple transmissions on a wireless channel at a receiver will have a probability distribution which may or may not be known. Even if the distribution of each individual contribution to the aggregate is known, the distribution of the aggregate may not have a closed form (e.g., the sum of lognormally distributed variates). In such cases, bounds on the distribution are useful.

In [7], we analyzed several concentration inequalities, e.g., Markov, Chebychev, Camp-Meidell, and Van Dantzig, for potential bounds on the distribution of aggregate interference. We found the *Van Dantzig inequality* [8] to provide the tightest lower bound on the cumulative distribution of the aggregate as a function of just its first two moments:

$$Pr\left\{I \leq x\right\} \geq \frac{8\left(x-\mu_I\right)^2}{3\sigma_I^2+8\left(x-\mu_I\right)^2}\,, \tag{1}$$

where I is the aggregate interference power, μ_I is its mean, and σ_I^2 is its variance. The bound is valid if the probability distribution of I has a second derivative and its density function is convex at the tail. In fact, the bound can be applied to all unimodal continuous probability density functions in their convex part.

Solving (1) for x, we obtain an upper bound on the pth percentile above the mean as

$$x \leq \sigma_I\sqrt{\frac{3p}{8\left(1-p\right)}}+\mu_I \quad , \quad x>\mu_I\,. \tag{2}$$

For example, the 95th percentile ($p = 0.95$) is no more than $\sigma_I\sqrt{\frac{57}{8}}+\mu_I$.

While the upper bound (2) only depends on the first two moments of a distribution, it is instructive to observe how

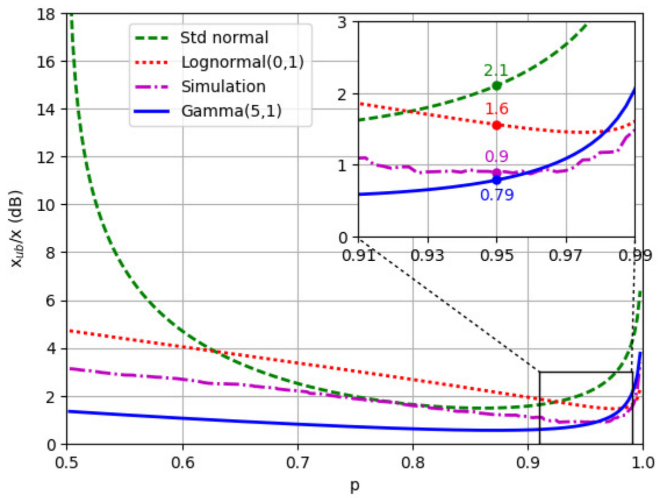

Fig. 2. Ratio of pth percentile upper bound (2), x_{ub}, to the exact pth percentile, x, for some distributions.

loose or tight the bound is at different p values. Fig. 2 plots the ratio in decibels (dB) of the upper bound to the exact pth percentile value as a function of p for four different distributions. They include a standard normal random distribution (in green), motivated by the central limit theorem; a lognormal distribution (in red), which has been used to approximate the sum of lognormal variates; a Gamma distribution (in blue), representing the power sum of signals that are independent and exponentially distributed in power (or Rayleigh distributed in amplitude); and finally, an unknown distribution (in magenta), which was computed as a sum of piecewise lognormal variates in our simulation. We observe that, while the bound can be loose near the median and extreme upper tail, it is quite reasonable around $p = 0.95$, where it ranges from 0.8 dB for the Gamma variate to 2.1 dB for the standard normal variate.

B. Individual SAS Interference Criterion

In CBRS, the interference protection requirements for federal incumbents are stringent [2]. The SASs need to consider the worst case scenario, in which the incumbent is active and must be protected at a predefined interference threshold.

Let the incumbent protection requirement be that the pth percentile of the aggregate interference power of all co-channel users not exceed a threshold, t. Furthermore, assume that the co-channel users are managed by M independent SASs, and that the jth SAS is allocated an interference budget of t_j, such that $\sum_{j=1}^{M} t_j = t$. Now, let N_j be the number of users managed by the jth SAS. Let $\mu_j = \sum_{n=1}^{N_j} \mu_{j,n}$ and $\sigma_j^2 = \sum_{n=1}^{N_j} \sigma_{j,n}^2$ be the total mean power and the variance, respectively, at the incumbent receiver of the N_j users managed by the jth SAS. Note that, $\mu_{j,n}$ and $\sigma_{j,n}^2$ are the mean and variance, respectively, of the received power of the nth user of the jth SAS. Also, we assume the received powers are uncorrelated. Then, in the appendix we prove the following theorem.

Theorem 1: If each SAS manages its users such that

$$\sigma_j \sqrt{\frac{3p}{8\,(1-p)}} + \mu_j \le t_j \quad , \quad 1 \le j \le M, \tag{3}$$

then the cumulative distribution of the overall aggregate interference, I, at the threshold t satisfies

$$Pr\,\{I \le t\} \ge p. \tag{4}$$

In other words, the pth percentile of the aggregate interference from all users managed by all M SASs does not exceed t.

The implication of the theorem is that, provided each SAS knows its interference budget t_j and satisfies the simple rule in (3), the overall aggregate interference protection requirement of the incumbent (4) is guaranteed to be met. The protection criterion is a vast simplification over the current procedure used for federal incumbent protection in CBRS. It removes both the need for global knowledge *and* the need to use a common, standardized algorithm. The tradeoff of using an upper bound for the percentile is that each SAS might be overly conservative in the interference management of its users. The analysis below sheds light on the extent of this spectral reuse tradeoff.

IV. Analysis

We present numerical results of using the proposed interference protection criterion in context of CBRS in the U.S. After describing the modeling assumptions, we present the results in terms of two metrics including the total number of users moved from the channel in order to protect the incumbent (i.e., the size of the move list) and the realized aggregate interference of all co-channel users at the incumbent.

A. Modeling Assumptions

For this analysis, we utilize models for propagation, deployment, and interference management that have either been codified in CBRS standards [2], implemented in test certification software [4], or utilized in the CBRS community for similar analyses (e.g., [9]).

Given M SASs manage a total of N co-channel users in the neighborhood of a DPA, we assume a non-uniform assignment whereby approximately $\frac{N \times j}{\sum_{j=1}^{M} j}$ users are assigned to the jth SAS. We also considered another scenario, by which the users are divided uniformly among the M SASs, approximately $\frac{N}{M}$ users per SAS. However, since the results in both scenarios are similar, we only present detailed results of the non-uniform assignment scenario. The number, locations, antenna heights, and EIRP of transmitters are modeled as in [9].

Finally, we set the interference budget of the jth SAS, t_j, in proportion to the number of co-channel users, N_j, that the jth SAS manages in the neighborhood of the DPA protection point (similar to requirement in [2, R2-SGN-16]) as

$$t_j = \frac{N_j}{N} t \tag{5}$$

where t is the protection level of the DPA in watts. We use the deterministic operational move list method in [7] for calculating the local move list of a SAS, though in practice, a SAS is free to use any method to satisfy criterion (3).

TABLE I
MOVE LIST SIZE AND AGGREGATE INTERFERENCE RESULTS AS A FUNCTION OF THE NUMBER OF SASS FOR DPA EAST7.

Number of SASs	Move List Size	Increase in Move List Size (%)	Max 95th Prctl. of Agg. Interf. (dBm/10 MHz)	Decrease in Agg. Interf. Prctl. (dB)
1	24 527	-	-147.30	-
2	24 867	0.58	-147.45	0.14
3	25 250	1.22	-147.82	0.51
4	25 543	1.72	-147.90	0.60
5	25 891	2.31	-148.04	0.74
10	27 080	4.32	-148.56	1.25

B. Numerical Results

In this section, we first examine the results for a single DPA East7 in detail, and then, we summarize the results for all DPAs along the U.S. coasts. The 95th percentile of aggregate interference threshold for all co-channel offshore DPAs is very stringent, i.e., -144 dBm/10 MHz, at every point in the DPA, as required in [2], [3]. The proposed interference protection criterion is applied to $M \in \{1, 2, 3, 4, 5, 10\}$ SASs. Note that the case of $M = 1$ (a single SAS) is equivalent to the current CBRS protection procedure where each SAS has a global picture of the users and applies a common algorithm for determining the move list. The $M = 1$ case serves as a baseline for comparison with the proposed protection criterion applied to multiple SASs.

For DPA East7, we deployed $N = 59\,120$ users in the neighborhood of the DPA. Table I shows the results for move list size (second column) and maximum 95th percentile aggregate interference (dBm/10 MHz) (fourth column), as a function of the number of SASs. The maximum 95th percentile aggregate interference is computed over all protection points in the DPA and over all radar receiver azimuths. As the number of SASs increases, the number of users on move lists also increases, meaning the proposed protection criterion is increasingly conservative. On the other hand, as the number of SASs increases, the number of remaining users (i.e., users not on move lists) decreases, resulting in a lower realized aggregate interference percentile, and always below the -144 dBm/10 MHz protection threshold.

Table I also shows the increase in move list size (third column) and the decrease in aggregate interference percentile (fifth column), from a single SAS to multiple SASs $M \in \{2, 3, 4, 5, 10\}$. Note that, we computed the increase in move list size as a percentage of the total number of co-channel users in the DPA neighborhood, N. For DPA East7, we observe that the changes for both move list size and aggregate interference percentile are insignificant.

Fig. 3 illustrates a geographic view of the move list in the case of a single SAS for DPA East7. Red and blue markers represent users on and off the move list, respectively. Green markers represent the protection points in the DPA that were used to compute the move list.

Fig. 4 illustrates the results for the 26 DPAs off the east and Gulf coasts (including DPA East7) and the 14 DPAs off the west coast of the continental U.S., for $M \in \{1, 2, 3, 4, 5, 10\}$

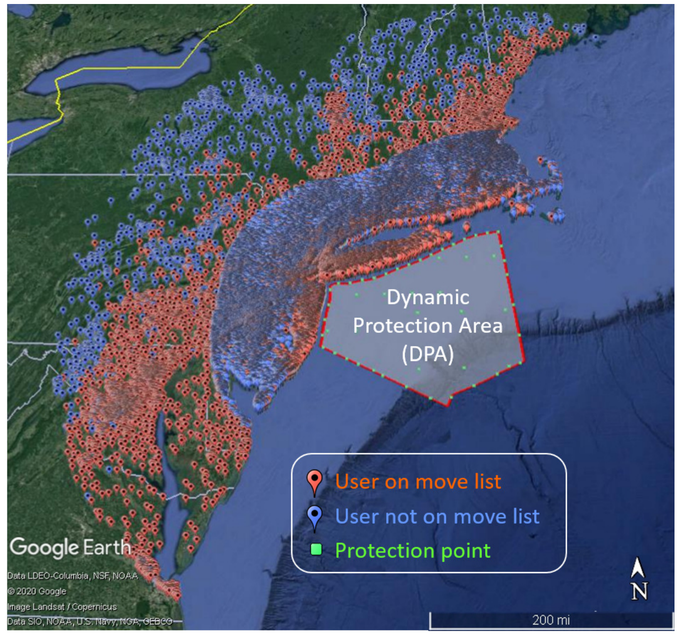

Fig. 3. Simulated move list for DPA East7.

SASs. At the time of writing of this letter, the U.S. Federal Communications Commission has certified five independent SAS providers to provide commercial CBRS service.

The bar charts on the left in Fig. 4 show both the total number of co-channel users in the neighborhood of the DPA and the number who are on move lists. Similar to the DPA East7 results, as the number of SASs increases, the number of users on move lists also increases. The increase in move list size, from one SAS to five SASs, varies in the range of $[0.09, 24.05]$ % across all offshore DPAs. The high values are driven by a few DPAs (e.g., East13 and East25), in which the proposed method is cautiously adding more users to the move lists to protect the incumbent. However, for all 40 DPAs, the median increase in move list size is only 2.69 %, which is relatively small.

The line charts on the right in Fig. 4 show the maximum realized 95th percentile of the aggregate interference in each case. Again, as the number of SASs increases, the realized aggregate interference percentile decreases. We observe that the decrease in the aggregate interference percentile, from one SAS to five SASs for all DPAs, changes within the range of $[0.74, 5.36]$ dB with a median of 2.36 dB. For most DPAs, the maximum realized 95th percentile is just below the -144 dBm/10 MHz threshold for a single SAS and gradually decreases as the number of SASs increases. However, a special case can be observed for the DPA East17, which has the max 95th percentile of aggregate interference much lower than -144 dBm/10 MHz for all values of M. This is because most users in the neighborhood of DPA East17 were put on the move list, regardless of the number of SASs.

V. CONCLUSION

We found that the proposed methodology can significantly simplify the current practice while ensuring interference protection to the federal incumbents in CBRS. It allows for

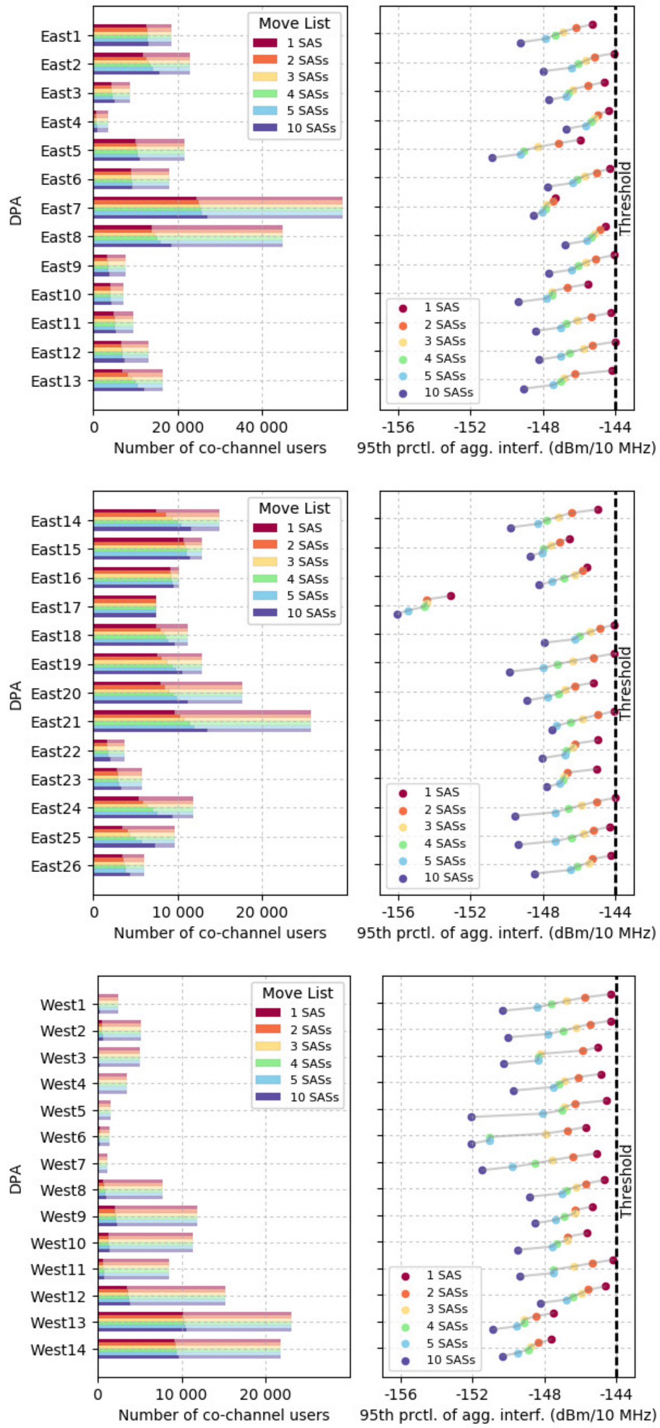

Fig. 4. Move list size and aggregate interference in dB relative to 1 mW (dBm) per 10 MHz, as a function of the number of SASs.

flexibility and innovation by SAS administrators. Using this approach, each SAS can manage its move list independently of others without the need to exchange detailed users information. Furthermore, it can apply its own method for determining its move list without the need to use a common move list algorithm.

However, the proposed method comes with a tradeoff in spectral reuse that tends to grow with the number of SASs managing users in the same area. It may put more users on the move list than a more complex, precise solution. This causes the aggregate interference of all SASs' co-channel users to be lower than necessary to protect the incumbent. For future work, a tighter bound than the Van Dantzig's on the aggregate interference distribution could improve the spectral efficiency. While this work was motivated by federal incumbent protection in the 3.5 GHz band, the proposed method is applicable to similar sharing arrangements in other bands, as well.

Appendix
Proof of Theorem 1

From (3) and $\sum_{j=1}^{M} t_j = t$, we can show that

$$\begin{aligned} \sum_{j=1}^{M}\left(\sigma_j\sqrt{\frac{3p}{8\,(1-p)}}+\mu_j\right) &\leq \sum_{j=1}^{M} t_j \\ \sqrt{\frac{3p}{8\,(1-p)}}\sum_{j=1}^{M}\sigma_j+\sum_{j=1}^{M}\mu_j &\leq t. \end{aligned} \tag{6}$$

By induction that, we get

$$\sqrt{\sum_{j=1}^{M}\sigma_j^2} \leq \sum_{j=1}^{M}\sigma_j. \tag{7}$$

Substituting (7) into (6) gives

$$\sigma_I\sqrt{\frac{3p}{8\,(1-p)}}+\mu_I \leq t, \tag{8}$$

where $\mu_I = \sum_{j=1}^{M}\mu_j$ and $\sigma_I^2 = \sum_{j=1}^{M}\sigma_j^2$ are the total mean power and variance, respectively, of all M SASs' users at the incumbent receiver, assuming they are uncorrelated. Finally, (8) together with (2) guarantees that the exact pth percentile of I (i.e., x in (2)) does not exceed t.

References

[1] "Citizens broadband radio service," 47 C.F.R. § 96, 2016.

[2] "Requirements for commercial operation in the U.S. 3550–3700 MHz citizens broadband radio service band," Wireless Innovation Forum Document WINNF-TS-0112, Version 1.9.1, Mar. 2020. [Online]. Available: https://cbrs.wirelessinnovation.org/release-1-standards-specifications

[3] 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

[4] SAS Testing and Interoperability Repository. [Online]. Available: https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System

[5] 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

[6] Irregular Terrain Model. [Online]. Available: https://www.its.bldrdoc.gov/resources/radio-propagation-software/itm/itm.aspx

[7] T. T. Nguyen and M. R. Souryal, "Deterministic move lists for federal incumbent protection in the CBRS band," *under review for possible publication in IEEE Transactions on Cognitive Communications and Networking*, 2020.

[8] G. Blatman, T. Delage, B. Iooss, and N. Pérot, "Probabilistic risk bounds for the characterization of radiological contamination," May 2017. [Online]. Available: https://arxiv.org/abs/1701.02373

[9] 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.