

A Study on the Design Space for Harm Claim Thresholds

Janne Riihijärvi, Andreas Achtzehn, Petri Mähönen
Institute for Networked Systems, RWTH Aachen University
Kackertrasse 9, D-52072 Aachen, Germany
email: {jar, aac, pma}@inets.rwth-aachen.de

Pierre de Vries
Silicon Flatirons Center
University of Colorado, Boulder
email: pierredv@hotmail.com

Abstract—Harm claim threshold policies will offer more flexibility in radio regulations. By defining bounds on the acceptable field strength over time and space, these policies will introduce quantitative measures to inter-system coexistence questions. However, this requires tangible and sufficient means of proving claims of interference through spectrum measurements. In this paper we study drive test requirements in terms of necessary sample set sizes and decorrelated sampling. We present order statistics and extreme value theory as powerful mathematical tools for describing necessary confidence intervals, and test their practical viability with data from an extensive measurement campaign. Furthermore, we discuss by means of example necessary extensions to planning tools for wireless networks operating under an interference threshold policy. Our results emphasize that future regulations will need to be accompanied by a rigorous specification of evidence collection requirements in order to compensate for bias and correlation structure in the spectrum data.

I. INTRODUCTION

Existing radio regulation often relies on qualitative and subjective statements when it comes to specifying coexistence rules such as interference tolerance levels between wireless systems. For example, the ITU defines harmful interference as “interference which endangers the functioning of a radionavigation service or of other safety services or seriously degrades, obstructs, or repeatedly interrupts a radiocommunication service” [1]. Such statements are problematic since they are difficult to reason about objectively, and lack clarity when it comes to the rights and responsibilities of wireless operators regarding harmful interference. Because of this there has recently been significant interest in reforming inter-system interference policies in terms of *interference limit policies*. Such policies aim to describe the radio environment a receiver must be able to cope with and the limits to emissions of a transmitter using *objective* and *measurable* criteria instead.

A key example of an interference limit policy is the *harm claim threshold* approach [2] described in the 2012 PCAST report on spectrum, and developed by an FCC TAC Working Group in 2012. A harm claim threshold essentially consists of a field strength profile defined over a frequency range, associated with a percentage of locations and times where the field strength must be exceeded at some confidence level to qualify as harmful interference. The threshold, which defines the interference a wireless operator needs to accept from others, would be low within his own band, higher for

adjacent bands, and highest for bands far away in frequency. Under a harm claim threshold policy, disputes on coexistence issues would become a matter of conducting spectrum use measurements to determine whether the parties involved cause interference exceeding the given threshold.

While interference limit policies in general and the harm claim threshold approach specifically are conceptually attractive and would seem to offer numerous benefits when compared to current regulatory approaches, several open questions remain on how they should be applied in practice. The key design space for harm claim threshold consists of (1) the field strength (or received interference power) threshold at a given frequency range, (2) the percentage of locations and times this threshold must be exceeded, and (3) the confidence level at which the exceedance must be shown. Whereas those parameters are ultimately a matter of policy, the design of measurement campaigns to reliably and sufficiently substantiate claims under harm claim threshold policies will necessarily need to build on knowledge of the statistical principles and practical constraints of extensive spectrum data collection. In this paper we first discuss by means of data analysis and simulations minimum requirements for the design of such campaigns. After introducing mathematical tools from *order statistics* and *extreme value theory*, we use data from a large measurement campaign in the UK to assess their feasibility and prediction capabilities in the process of deriving the actual interference statistics. We found that high thresholds require a substantial sampling density, and that bias from low randomization of measurement locations significantly affects the measurement results.

In the second part of the paper, we take the role of an operator deploying a network subject to a harm claim threshold policy. We discuss simulation requirements of relevant real-world propagation effects (in particular correlated shadowing) and quantify their effects on the convergence rate of the presented estimators. We show that, due to sampling correlations, significantly lower confidence is reached for the same sampling set sizes, i.e., deployment studies will become more complex. These results are highly relevant for the policy community, because they highlight the necessity for a discussion on establishing best practice rules for measurement campaigns for interference limit claims, and for operators wishing to incorporate these policies into their network planning tools.

II. A PRIMER ON ORDER STATISTICS

In order to reason statistically about the exceedance of a signal level threshold, we adopt the following simple probabilistic model. Let X be a random variable, yielding the received power or field strength at a randomly picked location, measured over a given range of frequencies. Denote then by $F_X(x)$ the cumulative distribution function (CDF) of X . If F_X were somehow known, computing different statistics of interference power for harm claim thresholds becomes trivial. For example, $1 - F_X(-40\text{dBm})$ would yield the fraction of locations at which the interference exceeds -40dBm , whereas $F_X^{-1}(0.95)$ would give the interference threshold that is exceeded only at 5% of locations. Unfortunately in any realistic scenario the “true” CDF $F_X(x)$ is unknown, and different properties related to it must be estimated through measurements (in case of, say, policy enforcement) or simulations (in case of planning a network roll-out under given harm claim threshold constraints). We shall focus in this and the following section on the former problem, and discuss simulation-based inference further below in Section IV.

We assume that we have conducted n measurements of the interference power X , denoted by X_1, \dots, X_n with the purpose of establishing a claim on exceedance of a given claim threshold. The *order statistics* of these measurements would simply be the indexed measurement values sorted in ascending order, and denoted by $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(n)}$. Thus $X_{(1)} = \min\{X_1, \dots, X_n\}$, $X_{(n)} = \max\{X_1, \dots, X_n\}$, with the rest of the order statistics carrying information about the *percentiles* $F_X^{-1}(p)$ of X . Note that due to the limited sample size, $X_{(n)}$ is *not* in general the true maximum of X , but some necessarily smaller value we may use to estimate the true maximum. Similarly, for large enough sample we use $X_{(p \times (n+1))}$ as the estimate for the p th percentile of interference power. The probability distribution associated to the inherent error of this estimate is called the *sampling distribution* of the corresponding percentile.

Given the limited sampling size, we need to state the uncertainty of $X_{(p \times (n+1))}$ for any claim. The measure of the estimation error $X_{(p \times (n+1))}$ is, when used as a proxy for $F_X^{-1}(p)$, given by the associated *confidence interval*. Following the result of Conover [3], [4], for unknown F_X the confidence interval for the p th percentile can be constructed from n measurements as follows. First, let Z_p denote the p th percentile of the standard normal distribution $N(0, 1)$. Then the upper limit of the $100(1 - \alpha)\%$ confidence interval is given by $X_{(u)}$, where

$$u = p(n + 1) + Z_{1-\alpha/2} \sqrt{np(1 - p)}, \quad (1)$$

and the lower limit by $X_{(l)}$, where

$$l = p(n + 1) - Z_{1-\alpha/2} \sqrt{np(1 - p)}. \quad (2)$$

If u or l turns out to be non-integer, the corresponding endpoint of the confidence interval can be determined by interpolating between the nearest adjacent order statistics. In addition to

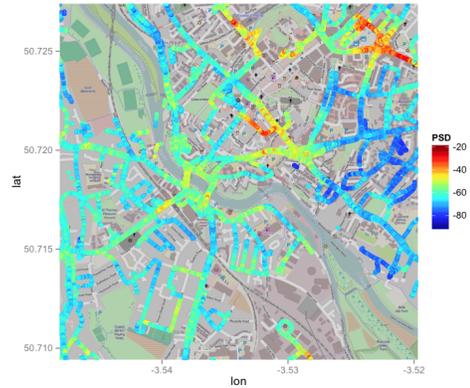


Fig. 1. Example drive test data set used for studying the order statistics and extreme values of received power profiles.

these *two-sided* confidence intervals we will need the *one-sided* confidence intervals, giving only the upper or lower bound for the error. These are obtained simply by replacing $Z_{1-\alpha/2}$ by $Z_{1-\alpha}$ in the corresponding formula above. These upper and lower bounds yield a necessary extension to the regulation policy, as they define the certainty level of the measurement results, and thus the likelihood that the harm claim threshold indeed has or has not been exceeded for measurement results near the threshold. Most importantly, for these bounds to hold, measurements must be gathered in a manner that is *unbiased* in time and space, which must be taken into account when specifying the allowable *measurement plans* in the regulation.

III. ESTIMATING ORDER STATISTICS FROM DRIVE TESTS

We shall now apply the methods outlined in the previous section to extensive drive test data set, and study in detail the properties of both true extreme values of power levels in the data set, as well as those of the different estimators for them based on smaller subsets of the entire measurement data set. These results can be directly used to shed light on adjacent channel interference statistics for real-world wireless systems.

A. Measurement Data Set

As a basis for our study we use drive test data gathered by CRFS on behalf of the UK regulator Ofcom as a part of the “not-spot” study for coverage of mobile broadband. The measurements were conducted using CRFS RFeye nodes, which were set up on rooftop boxes of measurement vehicles together with omnidirectional antennas. We refer the reader to [5] for more details on the measurement campaign.

Since the majority of the geographic region covered by the measurements is sparsely populated and thus with low deployment density of wireless services, we limit the study in this paper on a small region corresponding to the urban downtown area of the city of Exeter. Further, we focus on the UMTS downlink band, as this removes the impact of time domain dynamics which are present in TDMA, OFDMA and random access based systems. The resulting data set shown in Figure 1 consists of over 35 000 samples, covering a square

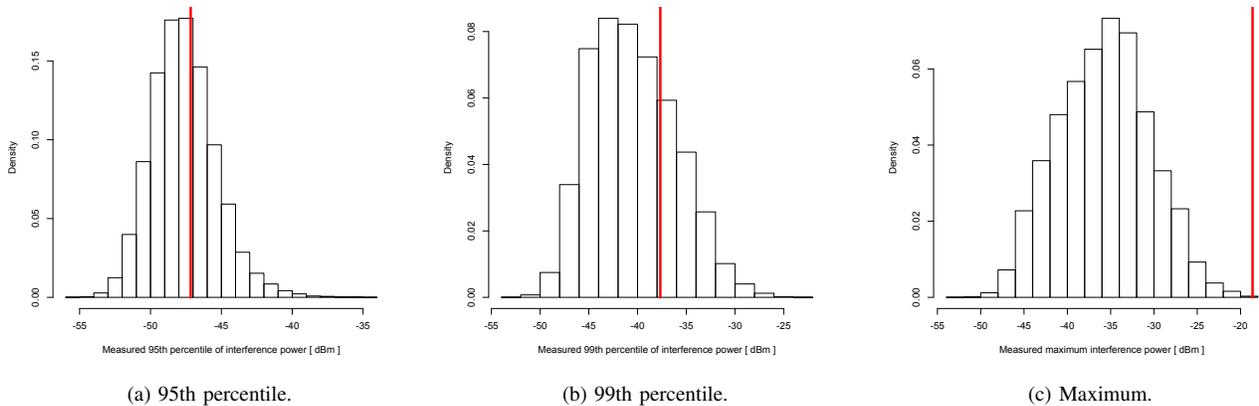


Fig. 2. The sampling distributions (histograms) for 95th and 99th percentiles and the maximum with 100 randomly selected measurement locations as well as the underlying true values (red line).

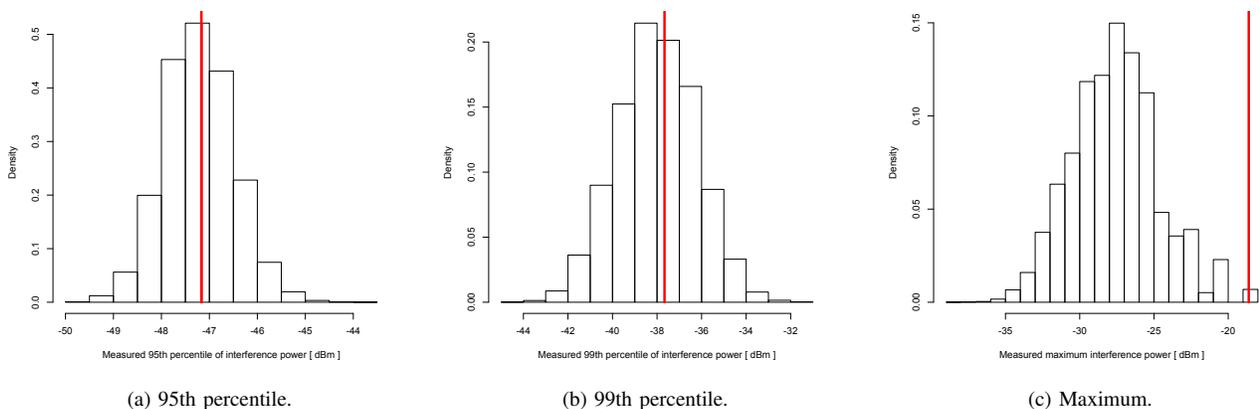


Fig. 3. The sampling distributions and true values for 95th and 99th percentiles and the maximum with 1000 randomly selected measurement locations. Notice that the variability of the estimates has reduced significantly compared to the 100 measurement case shown in Figure 2.

of $2\text{ km} \times 2\text{ km}$. Reported powers are the sum over a selected carrier of a local 3G operator. Such a data set would e.g. be used for an interference claim for a service operating in an adjacent band relative to the UMTS downlink frequencies.

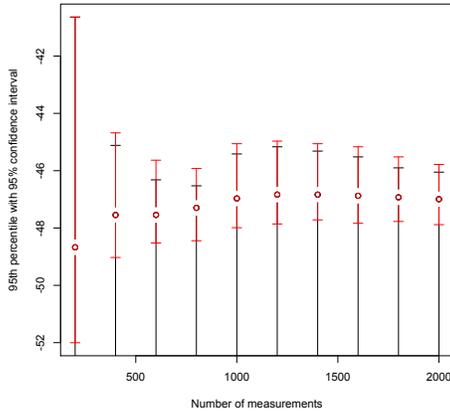
B. Sampling Distributions and Confidence Intervals

We begin by studying the sampling distributions for the different percentiles of the received power levels for different number of measurement points. Figure 2 shows these distributions for $n = 100$ measurements together with the “true” values of these percentiles computed over the entire data set. These plots were obtained by repeatedly selecting random measurements from the entire data set, computing the given percentage, and finally plotting the histogram of all the percentage estimates obtained. We see that in all the cases the sampling errors are substantial as can be expected for such a small measurement count. However, we note that the sampling errors increase rapidly as the percentile is increased, and in the case of maximum ($p = 1$) the errors are significantly biased, resulting in major underestimation of the maximum observable interference level. From Figure 3 we see that increasing the measurement count to $n = 1000$ significantly reduces the errors for the 95th percentile and to some extent also for the

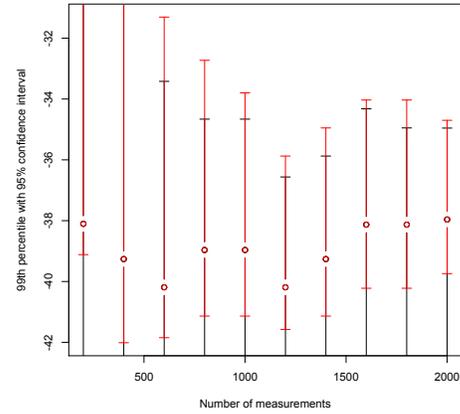
99th percentile. However, the estimates for the maximum are still highly biased and variable.

Several important conclusions can be drawn from these results. First is that regulating for the *worst case interference* is untenable. The high estimation errors for the maximum interference power would force carefully acting operators to vastly restrict transmit powers, and make definite conclusions in harm claim threshold disputes difficult to achieve due to the high level of uncertainty. Second is that also for the lower percentages, the number of measurements conducted has major impact on the estimation accuracy. Further, the choice of percentile in a harm claim threshold entitlement has a significant impact on the number of measurements that the holder has to make to demonstrate a harm claim at a given confidence level. While such statements are qualitatively trivial, we will now see that they can be quantified successfully by using the appropriate confidence intervals for the estimates as outlined in the previous section.

Figure 4 shows the convergence of the estimated received power percentiles together with the one-sided and two-sided confidence intervals. We see that while a relatively small number of measurements suffices to obtain stable estimates

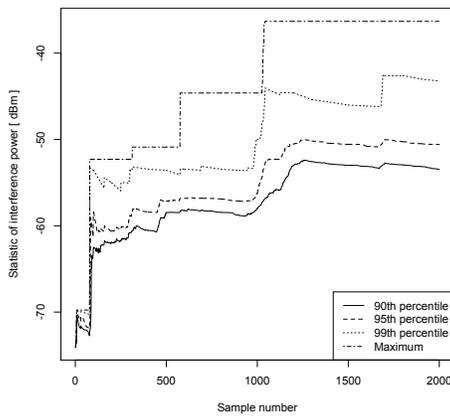


(a) 95th percentile with 95% CI.

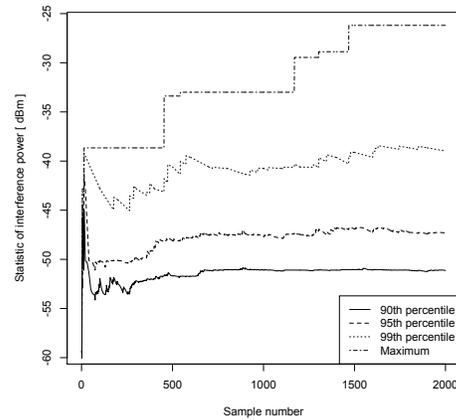


(b) 99th percentile with 95% CI.

Fig. 4. The convergence of estimated order statistics and their confidence intervals for the UMTS data set.



(a) Sampling along a drive path.



(b) Randomized sampling.

Fig. 5. The convergence of order statistics with and without randomisation. In the latter case the convergence to true values is substantially slowed down by spatial correlations in the data.

with bounded confidence intervals for the 95th percentile, for the 99th percentile a much larger number of measurements is needed to obtain finite confidence intervals to begin with, and in any case convergence is very slow. Note that these results are not contradicting those shown in Figures 2 and 3, as the true underlying sampling distributions require the entire data set for their estimation. We conclude that, for a given number of measurements, regulators will need to be very conservative and cautious on their expressiveness in terms of high percentiles of interference levels.

C. Impact of Spatial Correlations

We conclude this section by highlighting the impact of the used *measurement plan* on the results. Measured field strength or received interference power values are highly correlated spatially, resulting in measurements at nearby locations mainly yielding redundant information on the region-wide statistics. This makes it important to “spread out” the measurements used to draw conclusions on harm claim threshold violations.

The confidence intervals used above were also derived assuming uncorrelated data, which can only be achieved through keeping the measurement points sufficiently far apart from each other.

The impact of spatial correlations on the convergence of the different statistics considered until now is illustrated in Figure 5 for two extreme cases of measurement plans. First we look at a non-randomized plan where 2000 measurements are taken consecutively along a path covered by the drive tests. Second is a completely randomized plan, in which the same number of measurements are selected completely at random among all the locations available in the data set. We see that in the first case the convergence of all the estimates is very slow, whereas in the second case all the percentiles except the maximum converge quickly. This highlights the necessity of randomization of the measurement locations: A mischievous operator may select drive paths as to cover particularly high-level emission regions (e.g. close to known base stations), thereby ignoring the smoothing effects of lower power levels

at the coverage edges of a region. Similarly, measurement campaigns only in low coverage areas will yield fundamentally lower interference levels. Policy makers will thus need to decide how to balance between results from heterogeneous coverage regions.

IV. IMPACT OF MODEL ERRORS ON NETWORK DESIGN WITH INTERFERENCE LIMITS

The discussion in the previous section focused on the measurement-driven determination of interference limits for already deployed systems. However, such approaches are not feasible for the network deployment stage, where an operator must decide on the feasibility of deploying a system under an agreed upon harm claim threshold regime. In this section we show through a simulation example that many of the features seen in measurements, such as the high variability of maximum interference power and highest percentiles, are present in the network planning problem as well. Lower percentiles on the other hand result in more predictable behavior under uncertainties in the propagation environment.

Our simulated network is composed of 13 UMTS base stations in the previously discussed $2 \text{ km} \times 2 \text{ km}$ square region in downtown Exeter. We have acquired transmitter locations $\nu_{\text{TX},i}$ and transmit powers $P_{\text{TX},i}$ from a public database of UK regulator Ofcom [6]. Base stations in the scenario transmit at power levels between 40.6–60.3 dBm in the 2.1 GHz downlink UMTS band. Whereas in reality the base stations are owned by different operators and consequentially operate on different carrier frequencies, noting that the transmitter information reflects the initial rollout situation of UMTS in the UK, we have decided to combine all operator base stations to approximate a presumed intermediate network densification.

Received signal strength values $P_{\text{RX},j}$ for the different measurement locations $\gamma_{\text{RX},j}$ in the region are calculated for a fixed receiver antenna height of $h_m = 1.5 \text{ m}$ as

$$P_{\text{RX},j} = 10 \log_{10} \left(\sum_i 10^{(P_{\text{TX},i} - L_{i,j} + \chi_{i,j})/10} \right), \quad (3)$$

where $L_{i,j}$ is the mean pathloss at distance $d_{i,j} = |\nu_{\text{TX},i} - \gamma_{\text{RX},j}|$ from the transmitter. For calculating $L_{i,j}$ we employ the Okumara-Hata model [7] for urban areas.

For a single transmitter setup, the order statistics may be easily derived analytically, assuming that the overall selection of measurement locations is uniform. For multiple transmitters, this approach is infeasible, because the *combined* geometry of link distances will need to be taken into account. Furthermore, local structures such as a buildings, trees, etc. obstruct the individual propagation path and result in deviations from the mean pathloss, modelled through a correction term $\chi_{i,j}$ in (3). It is precisely this term through which uncertainties in network planning under an interference limit policy are captured. Large-scale measurement campaigns, e.g. [8], have shown that this *shadowing term* is locally correlated, i.e. nearby measurement locations are likely to be similarly affected by obstacles. The authors in [9] provide an extensive survey of models of these shadowing correlations, and we recommend their paper for further discussion on the

topic. For space reasons we limit ourselves in this paper to only the most popular model for shadowing, namely Gudmundson’s exponential distance model [10]¹, which defines the correlation coefficient between any two locations $\gamma_{\text{RX},j}$ and $\gamma_{\text{RX},k}$ as $\rho_{j,k} = \exp(-d_{j,k}/(\ln(20) \times d_c))$. The parameter d_c is called the *decorrelation distance* and describes the distance at which the correlation approaches 0.5.

A. Simulation Results

We have developed a custom MATLAB toolchain to create realizations from the correlated shadowing model and to compute the aggregate power levels over the Exeter area. Each simulation was repeated 5 000 times. Our results for different percentile values are shown in Figure 6, where we compare the uncorrelated shadowing and the correlated shadowing case with a decorrelation distance of 50 m.

Figure 6a presents the results over the iterations of the simulation for the uncorrelated case. We see that the percentiles exhibit low variations, e.g. the minimum and maximum deviate by only 0.5 dB, which means that already very few iterations suffice in order to find tangible signal strength thresholds. Only between 39–44 iterations (less than 1 percent) produced statistical outliers, i.e. values that were either larger than $\zeta_u = Q_3 + 1.5 \times (Q_3 - Q_1)$ or smaller than $\zeta_l = Q_1 - 1.5 \times (Q_3 - Q_1)$, where Q_1 and Q_3 are the first and third quartile of the samples, respectively. These results come as little surprise, because for multiple transmitters the variability over repetitions of the simulations in each measurement location solely originate from a sum of i.i.d. lognormal variables. While there is no closed-form analytical solution to describe this distribution, several approximations thereof exist [11]. The maximum encountered value shows a larger spread over the different runs, which complies with our earlier discussion on its statistics. Furthermore, its long right tail becomes apparent from the larger number of outliers (approximately 2.5 percent).

Our simulations with correlated shadowing produce significantly different results, see Figure 6b. Whereas the median values for the various percentiles approximate the uncorrelated case, the spread of the outcomes is significantly larger. These results originate from the decreased “randomness” of the shadowing process, where, due to the interdependencies of the local shadowing terms, deviations from mean signal strength values do not equally even out. We note that for high percentiles, a significant bias of positive outliers becomes apparent, and approximately 2% of all runs produced such irregular results.

Finally, we show the probability density functions (PDFs) of realizations of the 90th percentile and the maximum value for the uncorrelated and correlated case, respectively, in Figure 6c. As discussed earlier, the uncorrelated case results in a strong mode at the mean estimated percentile value, whereas the pdf of the correlated case is significantly flatter and exhibits a heavy right tail. We emphasize again that the uncertainties

¹The authors in [9] suggest an earlier origin of this model, see their discussion.

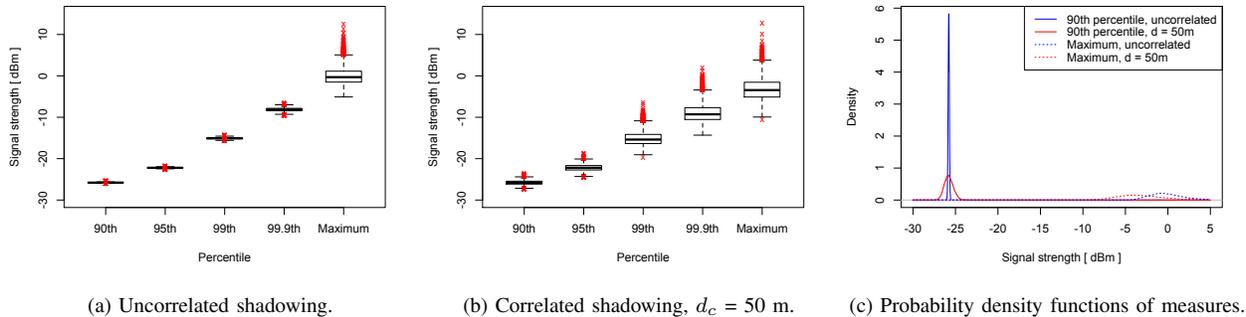


Fig. 6. Simulation results for the correlated shadowing environment.

involved in these estimates need to be absorbed by the network operator in their power budget. Thus, high degrees of estimation error, will require high safety margins.

V. DISCUSSION AND CONCLUSIONS

In this paper we have studied the use of rigorous statistical methods for measurement and estimation problems related to operating wireless systems under interference limit policies, and in particular harm claim threshold policies. We have seen that the percentiles involved in the definition of the harm claim threshold play a crucial role in both reliability with which harm claims can be verified or refuted through measurements, as well as in the uncertainties in network planning an operator must take into account when operating or deploying a wireless system. For enforcement, a policy *must* specify not only the harm claim threshold itself, but the confidence at which exceedances must be measured to occur. Further, since the confidence intervals are always derived under simplifying assumptions, additional safety margins are needed to account for deviations from these assumptions.

Random sampling would in theory be a powerful tool for conducting measurements on harm claim threshold violations. This is needed to overcome the impact of spatial correlations in interference as seen in Section 3. Since conducting measurements at randomly chosen locations is logistically difficult, alternative requirements need to be implemented. One option would be to use stratified sampling to obtain a simple and practical measurement plan design that has most of the benefits of the random design. In this approach the drive test is overprovisioned to obtain larger number of measurements than would be needed in the completely random case, after which the measurement data is culled to retain only measurements that are separated by a certain minimum distance. We plan to advance our research to study the feasibility of this approach.

Finally, we note that the shadowing effects in wireless propagation are still ill-understood. Designing a wireless network under the harm claim threshold approach requires extensive simulations because of the uncertainties in radio propagation and in particular the correlation statistics of shadow fading. Again, if the percentile in the policy is chosen to be too high, the uncertainty associated with the simulation results

will force the network operator to vastly reduce the transmit powers of their network, potentially severely hampering the deployment of new services. However, this paper is not to be misunderstood as arguing against the use of interference limit policies. On the contrary, the authors believe that with careful design, a more efficient spectrum allocation and a more flexible policy environment will be created. A further study into the required policy tools under real-world constraints, e.g. the measurement campaigns setups, will be highly beneficial for the community.

ACKNOWLEDGMENT

We thank RWTH Aachen University and DFG for providing financial support through the UMIC research centre. We acknowledge CRFS Ltd. for technical collection of data, sharing data, and support on the analysis. Ofcom is acknowledged for sharing data and for original funding of data collection. The analysis and conclusions expressed in this paper solely express the authors opinion, and no endorsement or opinions by Ofcom or CRFS should be assumed.

REFERENCES

- [1] M. J. Marcus, "Harmful interference: the definitional challenge," *Wireless Communications, IEEE*, vol. 15, no. 3, pp. 4–7, 2008.
- [2] J. P. de Vries, "Optimizing receiver performance using harm claim thresholds," *Telecomm. Policy*, vol. 37, no. 9, pp. 757 – 771, 2013.
- [3] W. J. Conover, "Practical nonparametric statistics," 1998.
- [4] R. O. Gilbert, *Statistical methods for environmental pollution monitoring*. Wiley, 1987.
- [5] CRFS, "Mobile Not-Spot Measurement Campaign Final Report," Tech. Rep., 2011. [Online]. Available: www.ofcom.org.uk/static/research/CRFS_report.pdf
- [6] Ofcom, "Sitefinder mobile phone base station database." [Online]. Available: <http://sitefinder.ofcom.org.uk>
- [7] M. Hata, "Empirical formula for propagation loss in land mobile radio services," *IEEE Trans. Veh. Technol.*, vol. 29, no. 3, pp. 317–325, 1980.
- [8] N. Perpinias, A. Palaios, J. Riihijärvi, and P. Mähönen, "Measurements of Shadow Correlations in a Suburban Environment on the 485 MHz Band," in *Proc. of VTC Fall 2013*, 2013, pp. 1–6.
- [9] S. Szyszkwicz, H. Yanikomeroğlu, and J. Thompson, "On the feasibility of wireless shadowing correlation models," *IEEE Trans. Veh. Technol.*, vol. 59, no. 9, pp. 4222–4236, 2010.
- [10] M. Gudmundson, "Correlation model for shadow fading in mobile radio systems," *Electronics Letters*, vol. 27, no. 23, pp. 2145–2146, 1991.
- [11] N. Beaulieu, A. Abu-Dayya, and P. McLane, "Estimating the distribution of a sum of independent lognormal random variables," *IEEE Trans. Commun.*, vol. 43, no. 12, pp. 2869–2873, 1995.