Impact of Primary User Activity Patterns on Spatial Spectrum Reuse Opportunities

Janne Riihijärvi, Jad Nasreddine and Petri Mähönen
Institute for Networked Systems, RWTH Aachen University
Kackertstrasse 9, D-52072 Aachen, Germany
email: {jar, jad, pma}@inets.rwth-aachen.de

***Abstract*—Activity patterns of primary users have significant influence on the opportunities for secondary use of spectrum. In this paper we explore this impact in terms of spatial spectrum reuse opportunities. More specifically, we model the primary user activity as a Semi-Markov ON/OFF process with varying distributions of holding times of ON and OFF states. We have used realistic holding time distributions obtained from extensive measurement campaigns. These activity patterns are combined with a collection of primary user transmitter locations, based on a real deployed cellular network. Based on these, we study the impact of activity model parameters on opportunities of spatial spectrum reuse. The results indicate that for the scenarios studied here the overall duty cycle defined by the activity pattern plays the key role also in determining spatial spectrum opportunities. We also discuss applications of the obtained results, including development of lightweight models for occurrences of spectrum opportunities over both time and space.**

I. Introduction

Dynamic Spectrum Access (DSA) by cognitive radios has become an intensively studied approach for alleviating the spectrum scarcity problem caused by current rigid spectrum regulation regime [1]–[3]. Much of the work carried out in cognitive radio research has been targeted towards solving the key technical obstacles to reliable DSA, offering the primary users some degree of guarantee against harmful interference. Most of these efforts have been targeted towards studying and solving the spectrum sensing problem [4], [5], that is, trying to infer the state of the primary user based either on local measurements or data fusion amongst a number of collaborative cognitive radios. Somewhat less emphasis has been placed on studying the actual spectrum reuse opportunities. While several spectrum use measurements have gathered data on prevalence of spectrum opportunities over time, at a fixed location, study of spectrum opportunities over time *and* space has received very little attention.

In this paper we take first steps towards understanding how the activity patterns of primary users impact also the spatial availability of reused spectrum. We adopt a simulation-based approach, using realistic spectrum use models for primary users [6] developed based on our extensive spectrum occupancy measurements [7], [8]. We then combine these activity models with a set of primary user transmitter locations and a simple energy detector model designed to ensure that certain interference probability is not exceeded. This model results in a *protection zone* around each *active* primary transmitter in which a cognitive radio with a given transmit power is not allowed to operate. We then study the structure of the areas outside the protection zones, especially focusing on their dynamics over time and dependency on the precise activity model adopted. This paper extends our earlier work [9] in which spectrum availability over time was considered in scenario with multiple primary transmitters active in the vicinity of the cognitive radio.

The rest of the paper is structured as follows. In Section II we discuss our system model in detail, especially focussing on how the protection zones around the primary transmitters are determined. We also briefly discuss the adopted Semi-Markov ON/OFF model for the activity patterns of individual primary transmitters, and typical parameter values found in different wireless systems. In Section III we outline the simulation environment used, illustrate the chosen scenario in terms of primary user transmitter locations, and discuss the parameter values for the chosen activity models. Results obtained from the simulation campaign are then given and discussed in Section IV, before drawing conclusions and outlining future work in Section V.

II. System model

A. Model for Protection Zones

We shall focus here on a scenario in which a primary network that is willing to share its spectrum with a secondary network tolerates some small loss in service for some of its users or an increase in the transmitted power [10]. The loss in service can be either reflected by a reduction of the coverage area [11] or by an acceptable interference probability [12], [13]. Given a sufficiently advanced cognitive radio technology these losses could be made extremely small, or regulatory and contractual arrangements could be made for situations in which more spectrum is desired for secondary use. We shall also focus on interference on the downlink of the primary network only. We define the interference probability with respect to primary user, p, as the probability that the experienced interference I by p due to secondary activity exceeds a predefined threshold $\iota_{\max}$:

$$\Pr_{\text{int}} \equiv \Pr\left\{I \geq \iota_{\max}\right\}. \tag{1}$$

In this paper, we assume that the primary will accept that any of its users can experience unacceptable interference with a probability ε in the worst case. The experienced interference

I due to the activity of a secondary user s at distance d and transmitting with power P is given by

$$I = P - L(d) + \xi, \tag{2}$$

where $L(d)$ is the path loss between the secondary transmitter and the primary terminal that are separated by a distance d and $\xi \sim \mathcal{N}(0, \sigma_s)$ is the normal shadowing factor with zero mean and standard deviation σ_s. All the variables in (2) are in dBm. Therefore, the constraint on interference probability can be reflected by a constraint on the distance separating the cognitive transmitter from the primary receiver depending on the power of the former. Since the shadowing factor follows a normal distribution, we can rewrite condition (1) using (2) as

$$\frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{\iota_{\max} + L(d) - P}{\sigma_s\sqrt{2}}\right)\right] = \varepsilon. \tag{3}$$

In this paper we consider we use Xia-Bertoni propagation model [14]. Given a frequency f in GHz and distance d between the transmitter and the receiver, path loss $L(d)$ is given by

$$L(d) = K + \beta \log_{10}(f) + \alpha \log_{10}(d), \tag{4}$$

where K, β and α are constants computed using the Xia-Bertoni model and are set respectively to 122.1 dB, 21 dB and 37.6 dB. We consider shadow fading with zero mean and standard deviation of 7 dB. By combining (3) and (4), we obtain

$$d = 10^{\left[\operatorname{erf}^{-1}(1-2\varepsilon)\sigma_s\sqrt{(2)} - \iota_{\max} + P - K - \beta\log(f)\right]/\alpha}, \tag{5}$$

which will be used in the following to define the protection zone around the primary transmitter.

In this paper, we assume that any required approximation should be biased to a more conservative approach with respect to the protection of the primary. Since the positions of the primary terminals are not known to the secondary nodes, the worst case situation is considered. This means that the secondary node should guarantee that an active primary terminal that might be at the position where the interference generated by secondary activity is at maximum will be satisfied (i.e. $\Pr_{\text{int}} \leq \varepsilon$). Therefore, the secondary node will consider the so-called Worst Case primary terminal Position (WCP) for its transmission decision. The WCP is the closest possible primary terminal position in the primary coverage areas that have the lowest path loss toward the secondary node. However, the coverage areas of wireless network cells do not, in general, have known and regular shapes. Therefore the secondary will consider the circle that contains the coverage area as shown in Figure 1, which is again a conservative assumption to protect the primary. This circle will be considered in the following as the estimated coverage area.

The WCP is the intersection of the estimated primary coverage area and the line connecting the primary and the secondary transmitters. This is true when we consider a monotonic distance-dependent path loss where the shadow fading is considered separately in the computation. This conservative approach is chosen in order to protect primary users especially in the scenarios where no information about the shape of the coverage area is available. However, if such information is available, as in the case of the implementation of a Radio Environment Map (REM) [15], more precise estimation of the WCP can be implemented.

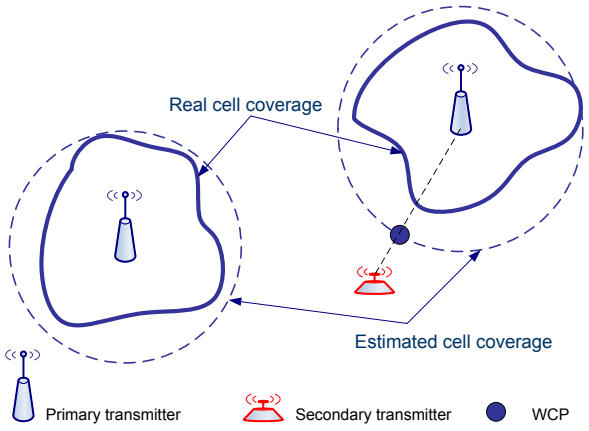

Fig. 1. Estimated Coverage Areas and the Worst Case primary terminal Position (WCP).

The distance given in (5) should be respected for all the primary terminals including the WCP. Therefore we consider that a secondary node with power P will be forbidden to transmit in all circles with a center at an active primary base station and radius Δ given by

$$\Delta = d + R, \tag{6}$$

where d is the distance computed in (5) and R is the radius of the estimated coverage area of the primary base station.

B. Primary User Activity Model

The model described in the previous subsection yields the exclusion zones around the *active* primary transmitters. In order to obtain the area in which spectrum reuse opportunities arise we need additionally to fix the locations of the primary transmitters and their activity patterns over time. We shall discuss the chosen location in the next section as in the context of this paper no model selection is involved there (actual locations from a cellular network are used instead), and focus in the following on the chosen activity model.

We assume that the primary user activity can be described in terms of two states, ON and OFF. We further assume that the times the system is in either of the states follows some fixed but arbitrary distribution, and that durations of successive ON and OFF periods are independent of each other. Thus we arrive at the general two-state Semi-Markov model, also known as an *alternating renewal process*. The case in which both ON and OFF periods are assumed to be exponentially distributed is commonly used in the literature. As shown in [6], this assumption is valid on a number of frequency bands, but not for all of them. Instead distributions with heavier

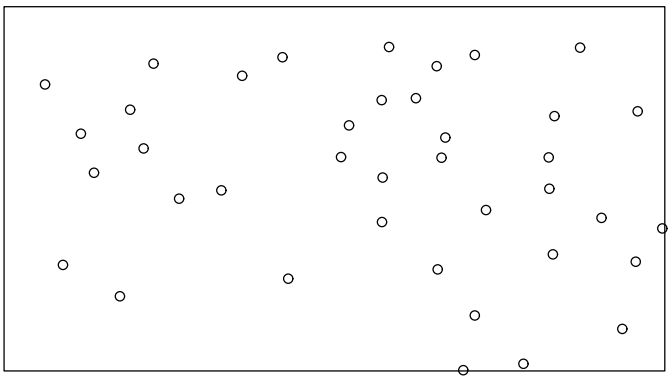

Fig. 2. The primary user transmitter locations used in simulations.

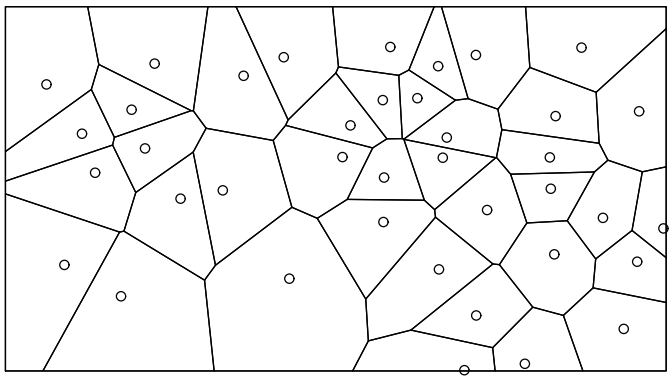

Fig. 3. The approximate service areas for each transmitter obtained from the Voronoi tessellation.

tails, such as the log-normal distribution, were observed for some frequency bands. The measurements also in general validate the assumption of successive ON and OFF periods being independent, although in some cases persistent long-term correlations were observed. In [9] it was observed that such correlations can be qualitatively reproduced by assuming that the observed process is not the activity of a single primary transmitter, but a superposition of several ones, each following a Semi-Markov model of its own. Thus we believe the Semi-Markov assumption is quite appropriate when considering an individual primary transmitter only.

III. Simulation Scenario and Settings

In this section we shall outline the simulation model and settings used in our study. Simulations themselves were implemented in the R environment [16], in part using the spatstat package [17] which contains useful numerical routines for computing areas of arbitrary polygons and for other required tasks. A discrete-event approach was used, with the state transitions of the individual Semi-Markov ON/OFF models supplying the events driving the simulation. The key inputs to the simulator are the models described above, and the set of transmitter locations annotated with the service areas, described below. The output of the simulation is a pair of time series. First of these consists of the areas of the regions in which secondary use of spectrum is possible. The second one contains the lengths of the time intervals these areas are usable. We shall refer to these as the *area sequence* $A(k)$ and *availability sequence* $T(k)$ in the following. For convenience, we normalize the area sequence so that if all the simulation area is usable at time t, we have $A(t) = 1$.

Figure 2 shows the collection of transmitter locations used in the simulations. Originally these correspond to T-Mobile GSM tower locations in a 185 km^2 region in downtown Los Angeles [18], and thus should be representative of planned network deployments in terms of their spatial structure. In order to obtain the protection zones according to the worst case primary receiver position we additionally need an estimate of the service areas of the individual primary transmitter. We use the Voronoi tessellation [19] for this purpose, resulting in the approximate service areas shown in Figure 3. This is naturally just a rough approximation, but we do not expect the adoption of a more realistic coverage model to significantly alter the results. Finally, for each setting of the secondary user transmit power the protection zones around the primary transmitter are precomputed and stored.

We shall apply two distributions as the durations of the ON and OFF periods for the primary transmitters, namely the exponential distribution with probability density function

$$f_{\exp}(t) = \lambda e^{-\lambda t}, \tag{7}$$

where $\lambda > 0$, and the log-normal distribution defined by

$$f_{\ln}(t) = \frac{1}{\sigma t \sqrt{2\pi}} e^{-(\ln t - \mu)^2/(2\sigma^2)}, \tag{8}$$

where $\mu > 0$ and $\sigma > 0$. The mean of the exponential distribution is simply $1/\lambda$, whereas for the log-normal random variable the mean is given by $\exp(\mu + \frac{1}{2}\sigma^2)$. In our measurements we observed for cases where the log-normal distribution resulted in a good fit that typical values of μ being around one in terms of the inter-sampling time used in the detection process, and σ ranging from 0.7 to 2.0. Also values of λ close to unity appeared typical for both ON and OFF period lengths for cases in which exponential distribution produced a better fit than the log-normal case. Thus in the simulations we chose to use the parameter settings $\lambda \in \{1, 2, 3, 5\}$ and $\mu \in \{1, 1/2, 1/3, 1/5\}$, with $\sigma = 1$ in the latter case. Additionally, simulations in which the mean values of the two types of distributions were fixed to be the same were carried out in order to study the extent the results are dominated by first-order statistics of the ON and OFF periods.

Finally, we set the transmit power of the secondary to 16.4 dBm for the bulk of the simulations. Such a value, slightly below usual Wi-Fi maximum transmit power, can be expected to be of correct order of magnitude for several secondary spectrum use scenarios. In the end of the next section we shall also study the impact increasing the transmit power has on the results.

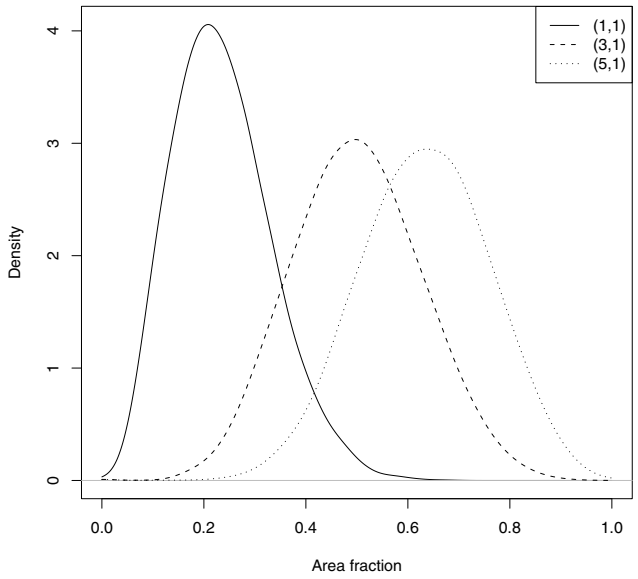

Fig. 4. The probability density functions of the marginal area distributions for different combinations of $(\lambda_{\text{on}}, \lambda_{\text{off}})$ parameters, with both ON and OFF durations being exponentially distributed.

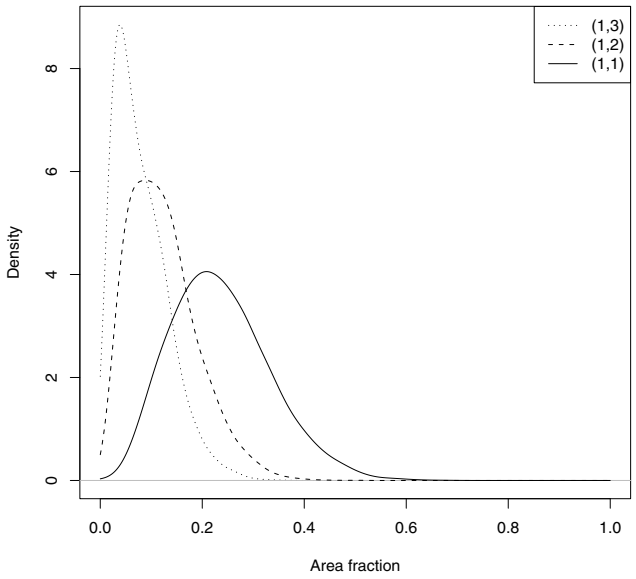

Fig. 5. Further results on density of the marginal area distribution in high duty cycle scenarios. Here $\lambda_{\text{on}} = 1$ throughout, and different values of λ_{off} used are shown in the legend.

IV. Results

We shall now discuss the results obtained from the simulations, beginning from marginal distributions of the area sequences and their dependency on the chosen activity model, and the proceeding with cross-correlation analysis of the time series, and finally exploring the impact of the secondary transmit power.

A. Marginal Distributions of Areas

We shall first study the marginal distribution of the areas in which spectrum can be reused by a cognitive radio with the given transmit power. This corresponds to the study of the distribution of the values $\{A(1), A(2), \ldots\}$ of the area sequence (ignoring the order of the values). As the tool to study these distributions we shall study the kernel density estimates [20] of the area sequence. First of these is shown in Figure 4, illustrating the results for the case of exponentially distributed ON and OFF durations. For the cases shown in the figure the *duty cycle* $\text{DC} = \lambda_{\text{on}}^{-1}/(\lambda_{\text{on}}^{-1} + \lambda_{\text{off}}^{-1})$ giving the proportion of ON-time of the individual transmitters ranges from 50% to 16.67%, corresponding to rather low spectrum utilization by the primary system. We see that even in the 50% duty cycle case (with $\lambda_{\text{on}} = \lambda_{\text{off}} = 1$) spectrum reuse is still possible over substantial proportions of the surface area. In the low duty cycle cases secondary use is possible on over half of the region, and occasionally over 90% of the region is simultaneously suitable for reuse.

Naturally, when the duty cycle is increased, the spectrum opportunities are reduced as well. Figure 5 shows the distributions of the area sequence for high duty cycle scenarios, with $\text{DC} = 66.67\%$ and $\text{DC} = 75\%$, with the 50% case also shown for reference. The density functions clearly become concentrated near the origin, but even in these cases reuse

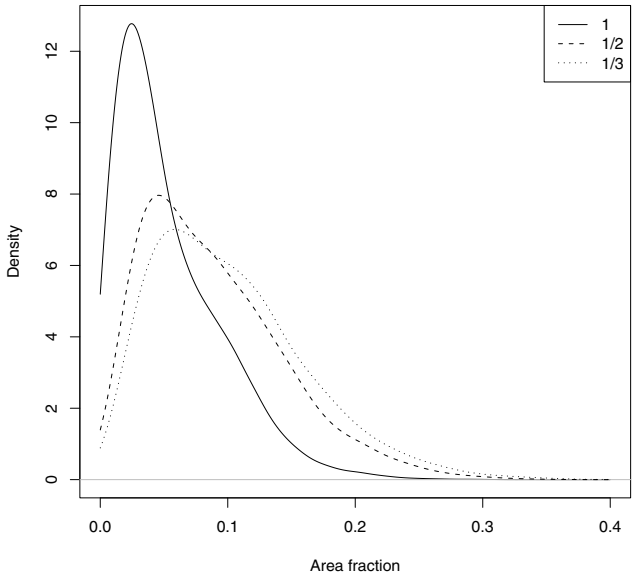

Fig. 6. Marginal area distribution when ON periods have log-normal distribution with μ as shown in the legend, and OFF periods are exponentially distributed with $\lambda_{\text{off}} = 1$.

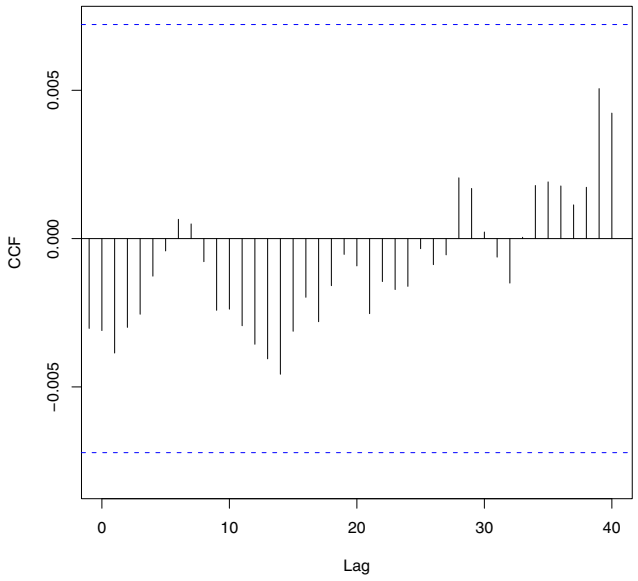

Fig. 7. Cross-correlation function between reusable area and duration of availability for the exponential case with $(\lambda_{\text{on}}, \lambda_{\text{off}}) = (1, 1)$. The dashed lines correspond to the confidence interval for the two series *not* to be correlated.

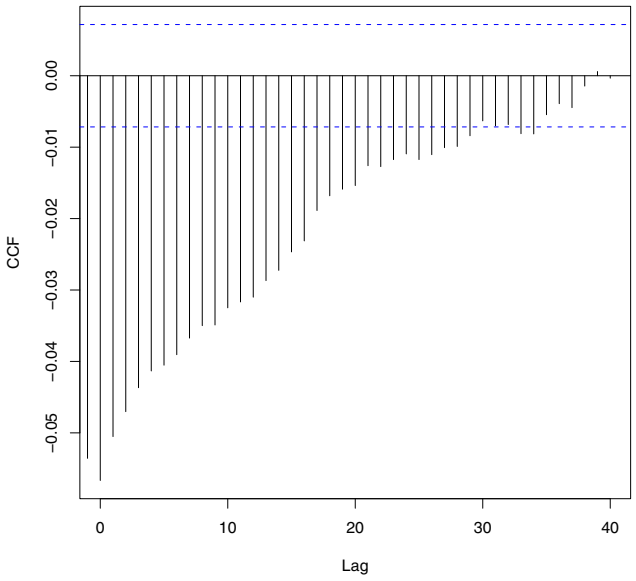

Fig. 8. Cross-correlation function between reusable area and duration of availability for the exponential case with $(\lambda_{\text{on}}, \lambda_{\text{off}}) = (1, 3)$.

opportunities remain.

From Figure 6 we see that the behavior of the area sequence remains qualitatively the same when the ON-periods are log-normally distributed, even though the tail of the log-normal distribution is heavier than of the exponential one. We shall see later on that the shape of the distribution does have a subtle effect on the behavior of the area and the availability sequences, but this effect is not very apparent on the marginal distributions shown here.

B. Dependency between Area and Availability in Time

The next groups of statistics we shall consider characterize the behavior of the area and availability sequences taking into account their temporal structure. We shall first study the dependencies between these two sequences. More concretely, we study if the knowledge of the surface area in which secondary use is possible also yields information on the length of time this area stays available, and vice versa. The linear component of the dependency between the area and time it is available is measured by the *cross-correlation function* of the two sequences, defined by

$$\rho_{A,T}(k) \equiv \frac{1}{\sigma_A \sigma_T} \mathbb{E}\big\{ \big(A(t) - \mu_A\big)\big(T(t+k) - \mu_T\big)\big\}, \quad (9)$$

where μ_A and σ_A are the mean and standard deviation for the sequence $A(t)$, and $\mathbb{E}$ denotes expectation over time. The argument k defines the *lag*, that is, the difference counted in events over which the correlation is measured. Figure 7 shows the cross-correlation for the case of exponentially distributed

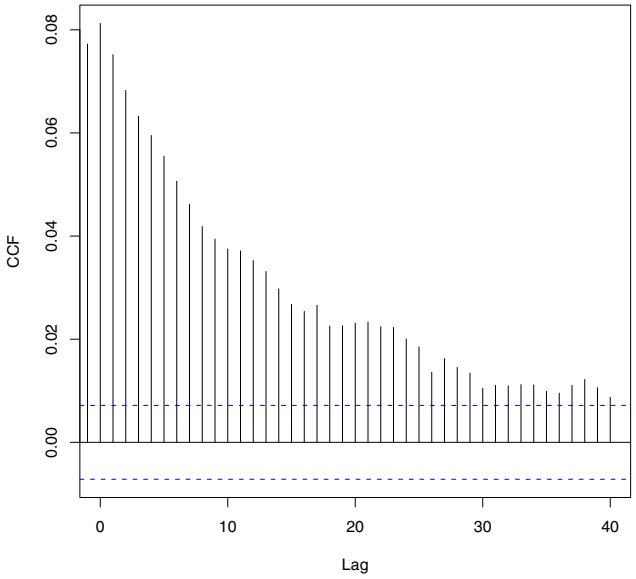

Fig. 9. Cross-correlation function between reusable area and duration of availability for the exponential case with $(\lambda_{\text{on}}, \lambda_{\text{off}}) = (3, 1)$.

ON and OFF period durations with duty cycle of 50%. As can be expected the estimated values of the cross-correlation function are significantly below the confidence intervals for the uncorrelated case, indicating that elements of the area and availability sequences can be considered as independent. This

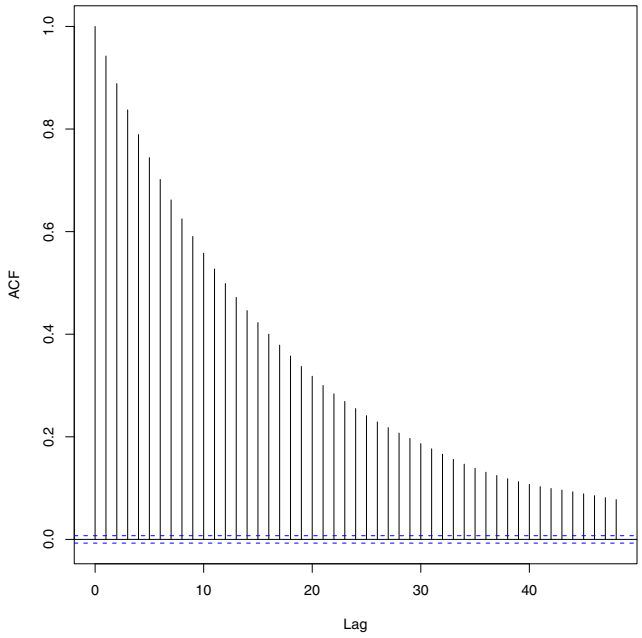

Fig. 10. The autocorrelation function of the area sequence for the exponential case with $(\lambda_{\text{on}}, \lambda_{\text{off}}) = (1, 1)$.

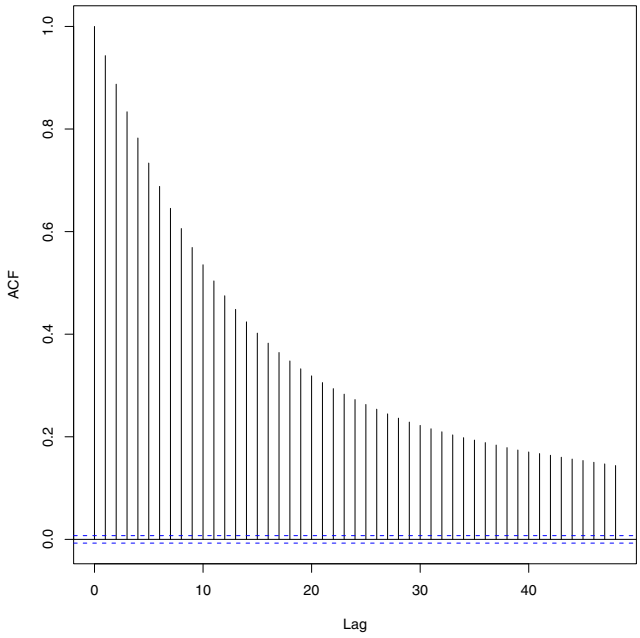

Fig. 11. The autocorrelation function for the case of ON and OFF periods being log-normally distributed, both distributions having mean of unity.

observation has potential applications when, for example, need arises to generate such sequences numerically without full system-level simulation as used here.

Situation changes when the duty cycle is either increased or decreased even for the exponential case. Figures 8 and 9 show the estimates of the cross-correlation functions for the high and low duty cycle cases, respectively. The results show that even though the observed correlations are small, they do appear to be statistically significant. A model for such correlation structure could be again used for simulating area and availability sequences with "correct" statistical characteristics without performing complex simulations, provided that the results of interest are not sensitive to possible higher order correlations in these sequences.

C. Autocorrelations in Area Sequences

We shall continue with the approach used in the previous subsection, but this time study the *autocorrelation function* $\rho_{A,A}(k)$ of the area sequence. The autocorrelation function essentially measures similarity of "nearby" values in the area sequence, and also can be used to predict unknown values of A by means of linear interpolation or extrapolation and simultaneously obtaining confidence intervals for the results. Intuitively one would expect rather heavy autocorrelations in the sequence, since it is defined by state transitions of the individual transmitters. This is indeed confirmed by Figures 10 and 11 showing the autocorrelation functions for the exponential and log-normal cases, with duty cycle of 50% in each case. The autocorrelation function in the log-normal case decays slightly more slowly than in the exponential case, indicating that the distributions of ON and OFF periods have an impact on the structure of area and availability sequences beyond first order characteristics, but at least for the parameter values studied here these impacts are rather small. This indicates that approximating the primary user activity with exponential distributions should not cause a large approximation error, as long as the parameters of the distributions are chosen to yield the correct duty cycle.

D. Influence of Secondary Transmit Power

The last results covered in this section illustrate the impact secondary user transmit power has on the results. Recall that we derived the protection zones around each active primary transmitter by considering the interference probability with a client or a receiver of the primary at the worst case position. This interference probability is influenced heavily by the transmit power of the secondary, thus increasing the required protection zone for higher transmit powers. The impact of this on the area in which spectrum can be reused by the secondary is illustrated in Figure 12. Notice that even though the duty cycle in the simulated scenario is rather low (16.67%), increasing the transmit power of the secondary drastically reduces the spatial extent of the arising white spaces.

V. Conclusions

In this paper we studied the influence the activity patterns of the primary user transmitters have on the area in which cognitive radios have opportunities for spectrum reuse. We

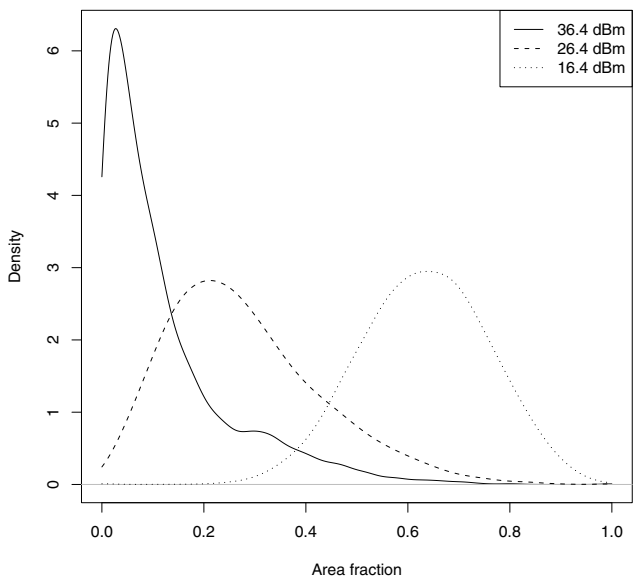

Fig. 12. The influence of transmit power of the secondary on the distribution of the area in which secondary use is possible; exponentially distributed ON and OFF period lengths with $(\lambda_{\text{on}}, \lambda_{\text{off}}) = (5, 1)$.

adopted a system model in which certain interference level is tolerated by the primary user, and service areas of the primary transmitters are used to derive conservative protection zones around them. The sizes of these protection zones are dependent on the local density of the primary transmitters, as well as on the transmit power of the cognitive radios. The results indicate that the duty cycle of the primary plays the key role in determining the distribution of the area in which spectrum can be reused. Higher order statistics reveal that distributions of ON and OFF periods do have impact beyond the duty cycle, but at least for the scenarios considered here these influences are quite small. Example applications of the results include development of lightweight models for reuse opportunities, facilitating the evaluation of the capacity of the secondary system without complete system-level simulations.

ACKNOWLEDGMENT

The authors would like to thank RWTH Aachen University and the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) for providing financial support through the UMIC research centre. We would also like to thank the European Union for providing partial funding of this work through the ARAGORN and NEWCOM++ projects. We would also like to thank Huawei Technologies Co., Ltd. for providing partial funding of this work through the CIRION project.

REFERENCES

[1] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "Next generation/dynamic spectrum access/cognitive radio wireless networks: a survey," *Elsevier Computer Networks*, vol. 50, no. 13, pp. 2127–2159, September 2006.

[2] Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 79–89, May 2007.

[3] D. Čabrić, S. M. Mishra, D. Willkomm, R. W. Brodersen, and A. Wolisz, "A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum," in *Proc. of IST Mobile Wireless Communications Summit*, Dresden, Germany, June 2005.

[4] T. Yücek and H. Arslan, "A Survey of Spectrum Sensing Algorithms for Cognitive Radio Applications," *IEEE Communications Surveys & Tutorials*, vol. 11, no. 1, pp. 116–130, first quarter 2009.

[5] J. Ma, G. Y. Li, and B. H. Juang, "Signal Processing in Cognitive Radio," *Proceedings of the IEEE*, vol. 97, no. 5, pp. 805–823, May 2009.

[6] M. Wellens, J. Riihijärvi, and P. Mähönen, "Empirical Time and Frequency Domain Models of Spectrum Use," *Elsevier Physical Communication, Special issue on Cognitive Radio: Algorithms & System Design*, vol. 2, no. 1–2, pp. 10–32, March–June 2009.

[7] M. Wellens and P. Mähönen, "Lessons Learned from an Extensive Spectrum Occupancy Measurement Campaign and a Stochastic Duty Cycle Model," *Springer Mobile Networks and Applications*, published online: http://dx.doi.org/10.1007/s11036-009-0199-9, August 2009.

[8] M. Wellens, J. Wu, and P. Mähönen, "Evaluation of spectrum occupancy in indoor and outdoor scenario in the context of cognitive radio," in *Proc. of International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom)*, Orlando, FL, USA, August 2007, pp. 420–427.

[9] M. Wellens, J. Riihijärvi, and P. Mähönen, "Modelling Primary System Activity in Dynamic Spectrum Access Networks by Aggregated ON/OFF-Processes," in *Proc. of IEEE Workshop on Networking Technologies for Software Defined Radio Networks (SDR), held in conjunction with IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, Rome, Italy, June 2009.

[10] R. Tandra, M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one?" *Proceedings of the IEEE*, pp. 824–848, 2009.

[11] E. G. Larsson and M. Skoglund, "Cognitive radio in a frequency-planned environment: some basic limits," *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 4800–4806, December 2008.

[12] A. Nasif and B. Mark, "Collaborative opportunistic spectrum access in the presence of multiple transmitters," in *IEEE Global Telecommunications Conference (GLOBECOM 2008)*, 30 2008-Dec. 4 2008, pp. 1–5.

[13] J. Nasreddine, O. Sallent, J. Pérez-Romero, and R. Agustí, "Positioning-based framework for secondary spectrum usage," *Physical Communication*, vol. 1, no. 2, pp. 121–133, 2008.

[14] L. Maciel, H. Bertoni, and H. Xia, "Unified approach to prediction of propagation over buildings forall ranges of base station antenna height," *IEEE transactions on vehicular technology*, vol. 42, no. 1, pp. 41–45, 1993.

[15] Y. Zhao, B. Le, and J. H. Reed, *Cognitive Radio Technology*. Elsevier, 2006, ch. Network Support: The Radio Environment Map, pp. 337–363.

[16] R Development Core Team, *R: A Language and Environment for Statistical Computing*, http://www.R-project.org [Last visited: 28th of February, 2010], R Foundation for Statistical Computing, Vienna, Austria, 2009, ISBN 3-900051-07-0.

[17] A. Baddeley and R. Turner, "Spatstat: an R package for analyzing spatial point patterns," *Journal of Statistical Software*, vol. 12, no. 6, pp. 1–42, 2005, ISSN 1548-7660. [Online]. Available: www.jstatsoft.org

[18] http://www.t-mobiletowers.com/TowerSearch.aspx?mode=new [Last visited: 28th of February, 2010].

[19] F. Aurenhammer, "Voronoi diagramsa survey of a fundamental geometric data structure," *ACM Computing Surveys (CSUR)*, vol. 23, no. 3, p. 405, 1991.

[20] B. Silverman, *Density estimation*. London: Chapman and Hall, 1986.