Efficient QoS Support for Secondary Users in Cognitive Radio Systems

Daniel Willkomm
Telecommunication Networks Group
Technische Universität Berlin, Germany
Email: willkomm@tkn.tu-berlin.de
Adam Wolisz
Technische Universität Berlin, Germany
and University of California, Berkeley, US
Email: awo@ieee.org

***Abstract*—The fundamental condition for the legal approval of Dynamic Spectrum Access (DSA) approaches is the protection of the Primary User (PU). However, for DSA to become an attractive service reality, it is crucial to ensure also some Quality of Service (QoS) support for the Secondary User (SU) communication. In this paper, we discuss sensing based opportunistic spectrum access approaches, in which PU protection is achieved by a properly organized sensing process and SU communication reconfiguration. While the required reliability of the sensing process can be expressed in terms of rarely enough overlooking the PU, we assume that the proper QoS for the SU is given by maintaining — with a given confidence level — a minimum bandwidth availability for the SU in spite of PU dynamics. In this paper we present an overview of approaches which might be used to achieve these objectives. In addition we point out, that both the sensing process and the secondary link maintenance (necessary to keep the required bandwidth in spite of reconfiguration due to detected PUs) require a significant spectrum overhead. We identify and elaborate a fundamental tradeoff in using these overheads in either sensing or link maintenance, and present examples of its optimization.**

I. Introduction

Opportunistic spectrum sharing is a very promising Dynamic Spectrum Access (DSA) approach for more efficient spectrum usage. In opportunistic spectrum sharing so called Secondary Users (SUs) detect through a sensing process parts of the spectrum temporarily not used by the license holders (the Primary Users (PUs)) and utilize them. A set of SUs desiring to communicate with each other and being within communication range is further referred to as Secondary User Group (SUG). Spectral resources to be used by an SUG are to be determined, and have to be periodically re-sensed to assure that they are still not used by the PU.

The fundamental requirement for the sensing based opportunistic spectrum usage is to protect the PU, i.e., ensure non-interference beyond some — very limited — scope. To quantify this scope, each PU has to specify a so called maximum interference time ($t_{\max}$), which specifies the maximum time a re-occurring PU can tolerate interference from an SU before the interference is considered to be harmful. After this period, the PU should be sure that no interference from SUs will take place. Obviously, $t_{\max}$ heavily depends on the service provided by the PU — it is, e.g., set to 2 s for usage of white spaces in the TV bands. Complementary to $t_{\max}$, the probability of false negatives in the sensing process, i.e., the probability of not detecting the PU although it is present, is defined. To ensure the proper protection of the PU a strict — very small — limit on the acceptable probability of these false negatives of the sensing process and $t_{\max}$ have to be specified — it is frequently postulated to fix these parameters by a legal act.

PU protection imposes severe challenges on the communication of an SUG: The detection of PUs possibly re-occurring to the spectrum being in secondary usage requires that no data transmission takes place in the spectrum to be sensed for the duration of the sensing process. Furthermore, if the sensing process reports the appearance of a PU, the affected parts have to be immediately vacated and — possibly — replaced by spectrum not used by PUs (so called link maintenance).

Both these phenomena imply potentially frequent changes of the used spectrum during a secondary communication session. Maintaining the QoS of a secondary communication despite these changes is challenging, and — in general — possible only by using some level of over-provisioning of spectral resources for the secondary communication session. In this paper, we give an overview of how spectrum sensing and secondary link maintenance can be organized to ensure protection of the PU communication **and** maintain a given QoS despite the challenges explained above.

More precisely, we investigate the spectral overhead needed for secondary QoS support under the constraint of a proper protection of the PU communication. The overhead used for sensing (e.g., time of sensing, number of sensors) has an influence on the quality of the sensing results, measured by the probability of false positives, i.e., the probability of reporting the presence of a PU although the PU is not present. The rate of false positives, on the other hand, has an influence on the required link maintenance overhead: the link has to be reconfigured not only if the PU appears, but always when the sensing process *reports* an appearance (which includes false positives). Therefore, there is a possibility to trade the overhead for sensing vs. the overhead for link maintenance while assuring the same level of QoS for the secondary communication.

The goal of this paper is a discussion of the mechanisms for sensing and for link maintenance used in opportunistic spectrum sharing and a consideration of the above mentioned tradeoff.

The reminder of the paper is structured as follows. In Section II we will discuss the organization of the spectrum sensing process, ways to control its accuracy as well as incurred spectral overhead. Section III is devoted to the organization of secondary usage and QoS support in spite of re-occurring PUs. In Section IV we formally introduce metrics for quantification of both the above mentioned spectral overheads and explain in detail the tradeoff between the overhead in sensing and the overhead in link maintenance. In Section V, we present a system design example for which we analyze the influence of selected aspects of the sensing and the link maintenance process, and present a minimization of the required joint spectral overhead. We summerize our findings in Section VI.

II. Spectrum Sensing

Spectrum sensing is the process of deciding — based on measurements — whether or not a PU is present in a given spectrum band, further referred to as PU band. The measurement techniques used range from simple power measurements to matched filter or cyclo-stationary feature detection [1]. Conceptually, there are two different sensing processes to be performed: *initial* sensing and *periodic* sensing.

Initial sensing has to be performed by members of a Secondary User Group (SUG) before any spectrum will be assigned to them for joint secondary usage. The initial sensing process often covers a wide spectrum range trying to determine which PU bands — possibly strongly differing in their features — are available for secondary communication. It is usually not time critical.

Out of the spectrum detected to be available, an SUG forms a so called Secondary User Link (SUL) — the spectrum which is to be jointly used for communication within the group.

After the setup of an SUL, *periodic* sensing has to be done regularly during the whole period of its usage. The SUG has to verify that the used resources are still available, i.e., that no PU has appeared within PU bands used by the SUL. In contrast to the initial sensing process, the periodic sensing process can be constrained to spectrum resources used by the SUL. Periodic sensing is time critical since a PU band reclaimed by a PU has to be vacated within the maximum interference time (t_{max}), which puts a strict upper limit on the sensing time.

A fundamental problem of spectrum sensing can be stated as follows: It is impossible to reliably detect the PU in a certain spectrum range, while at the same time performing data transmission in that range. It is, thus, crucial to ensure that the spectrum range to be sensed is not used for data transmission by the SUG at the same time.

A. Sensing Performance Metrics

The fundamental performance criterion of spectrum sensing is *reliability*. The usually applied measure for reliability is the probability of false negatives (P_{fn}), i.e., the probability of not detecting a PU although it is present. To ensure the protection of the PU, the sensing process has to ensure to keep P_{fn} below a certain threshold (usually defined by regulation). The probability of detecting a PU although it is not present (probability of false positives (P_{fp})) is a good additional performance measure for the sensing process. We refer to this metric as the *quality* of the sensing process.

Meeting the requirement of reliability (false negatives) and improving the quality of the sensing process (false positives) comes at a cost: the *complexity* and the *spectral overhead* of spectrum sensing. We will consider here primarily the spectral overhead. Given the constraint that no data transmission can take place within the spectrum being sensed, the amount of spectrum used for sensing (defined as the product of the spectrum bands used for sensing and the time in which these bands are sensed) requires spectral overhead. In addition, the exchange of sensing information between different sensors also requires spectral resources. Apart from the spectral overhead, spectrum sensing costs are also influenced by the sensing hardware and software complexity, the amount of energy burned for sensing, etc.

B. Diversity Approaches

Due to the nature of wireless communication "one short sensing sample" does not provide reliable information on the presence of a PU. In order to achieve better reliability and quality of sensing the general idea is to reduce the noise of individual measurements by applying diversity. Three possible diversity dimensions are time, frequency, and space.

Diversity in time is achieved by performing sensing over a certain time span. However, as already pointed out, the sensing time is strictly bounded by the maximum interference time (t_{max}). Furthermore, time samples should be spaced by a lag comparable to the coherence time of the wireless channel in question, in order to have independent sensing samples.

Another possibility of improving the quality of sensing is diversity in the frequency domain [1]. Instead of only taking one sensing sample for a whole PU band, the PU band can be split in sub-channels and each sub-channel can be used to get a sensing sample. As for sensing diversity, the achievable diversity in the frequency domain is strictly limited: on one hand by the bandwidth of the PU band and on the other hand by the coherence bandwidth, to get independent samples. Note that diversity in time and in frequency directly increases the spectral overhead of the sensing process.

Due to the above limitations a PU cannot be detected reliably by just one sensing device [2]. In *distributed* sensing, the spatially distributed sensing entities exchange their local sensing results and combine them to a joint decision. Using spatial diversity requires spectral overhead not only for the sensing itself but also for the exchange of the sensing data. In addition, the decision making process takes time and introduces an additional delay. Obviously, the overhead for distributed sensing depends on the number of participating sensors.

A lot of effort has been devoted to the analysis of the sensing process, and it is nowadays pretty well understood how the quality of sensing depends on the *diversity* of the sensing

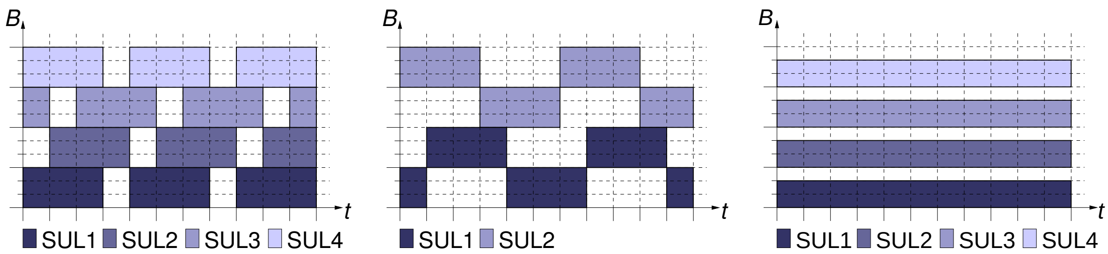

(a) Interrupted sending (b) Dynamic Frequency Hopping (c) Partial sensing

Fig. 1. Approaches for sensing organization: The colored areas represent spectrum used for different communications links (SULs), the white areas spectrum available for sensing.

measurements [3]. The spectral overhead, however, has only recently gained attention in the research community. In [4], the authors show a tradeoff for the secondary throughput between the sensing time and the probability of false positives. In [5] additionally the influence of spatial diversity on the sensing quality is considered. We are not aware of any publications investigating the spectral overhead of distributed sensing approaches.

C. Periodic Sensing Organization

There are basically three different approaches to the organization of the periodic sensing process. The approach most frequently followed in the research community so far is interrupted sending (Figure 1(a)), in which secondary data transmission has to be periodically interrupted to perform sensing. Note that, in general, interruptions of individual SULs do not need to be synchronized. If the sensing process detects a PU in a used PU band, the SUL is shifted to another, unused band.

The second approach, Dynamic Frequency Hopping (DFH) (Figure 1(b)), follows the idea of performing sensing and data transmission in parallel in different PU bands and switching them cyclicly in such a way that — with exception of the very short switching gap — a continuous sending can be assured [6]. One drawback of DFH is that a more complex radio front end — being able to sense and send in parallel — is needed.

The third approach is partial sensing (Figure 1(c)). The idea is that not the whole PU band is used for the SUL, but that some part is always left idle to perform sensing in parallel. If a PU is detected in its band, data transmission in the affected PU band has to be discontinued and the SUL has to be shifted to another, unused band. Similar to the DFH approach, the radio front end has to be capable of performing sending and sensing in parallel.

III. Link and QoS Management

Secondary User Groups (SUGs) create Secondary User Links (SULs) — using temporarily idle spectral resources — to perform data transmission. There are two possibilities for constructing SULs: contiguous SULs and non-contiguous SULs (see Figure 2).

A contiguous SUL consists of a contiguous amount of spectrum. It can span a whole PU band, multiple PU bands, or only parts of a PU band, as shown in Figure 2. The use of contiguous SULs is envisioned, e.g., within the IEEE 802.22 standardization [7].

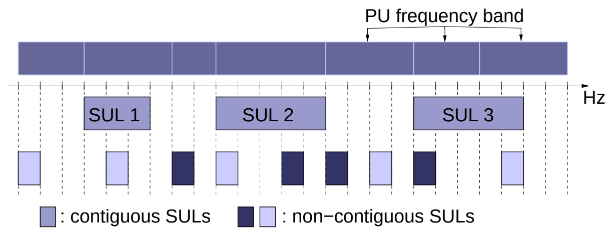

Fig. 2. Examples for contiguous and non-contiguous Secondary User Links (SULs)

Non-contiguous SULs consist of multiple, non-adjacent spectrum bands. One approach, introduced in [8], is to consider the PU bands as a bundle of smaller sub-channels. The idea behind the non-contiguous SULs is to scatter the SUL over multiple PU bands, such that (i) a reappearing PU is less affected (only a few sub-channels are used), and (ii) only a very small number of sub-channels has to be exchanged in the SUL if a PU appears. Note that using this approach partial sensing can be naturally supported. In Figure 2 two non-contiguous SULs are shown.

In this paper we will constrain ourselves to the case of non-overlapping SULs, i.e., the case in which each sub-channel is part of *at most* one SUL. We will outline research challenges for more complex scenarios in Section VI.

A. Link Setup, Maintenance, and Release

Setup of an SUL only happens once at the beginning of the communication. Based on the sensing results the SUL is build using available spectral resources. This can be done in a centralized or distributed manner. In the centralized approach, the central controller of the SUG has to gather the sensing results from the SUG members, decide on which resources to use for the SUL, and distribute the decision back to all SUG members. In the distributed approach the members of the SUG jointly decide on which resources to use for the SUL.

The initial sensing and initial negotiation of parameters requires spectral overhead. The impact on the overall system performance is, however, usually small as compared to the resource usage during the lifetime of the SUL.

The link maintenance process is responsible for the proper operation of an SUL during its whole lifetime. Generally speaking, link maintenance is the process of surveying the availability of the spectral resources used for the SUL and adjusting them in case a PU was detected. Providing a proper QoS of the SUL despite the necessary reconfigurations is an essential part of the link maintenance process. Conceptually, link maintenance can be divided into link reconfiguration, responsible for the release and adding of spectral resources to the SUL, and reconfiguration compensation, responsible to compensate potential temporal performance degradations of the QoS due to reconfigurations.

The spectral overhead required for link maintenance highly depends on the type of QoS to be supported by the SUL. The QoS metric commonly used within the spectrum domain is the minimum bandwidth available for the SUL at *any* time. For an SUL, in order to support this QoS requirement, the link maintenance process has to ensure that even in case of reconfigurations of the SUL the minimum bandwidth is available with a very high probability.

At the end of the communication session the SUG members have to achieve a joint understanding that the periodic sensing process stops and no further usage of the resource of the SUL is authorized. Similar to the setup of an SUL, this process requires spectral resources, which, however, are usually small compared to the resource usage during the lifetime of the SUL.

B. Link Reconfiguration

If the periodic sensing process indicates that a PU appeared within a PU band used for the SUL, the link has to be reconfigured. In order to reconfigure the SUL, the communication peers have to stop data transmission on the sub-channels to be vacated and negotiate / decide which new sub-channels to add to the SUL to compensate for the lost ones and maintain the QoS. The decision which new sub-channels to use is based on sensing results indicating which sub-channels are available. Note that this might require to trigger a new sensing process. In order to add new sub-channels to the SUL (replacing possibly lost ones) it has to be ensured that these sub-channels are not used by a PU. Depending on the capabilities of the sensing process, this can be done proactively during the periodic sensing process, i.e., in parallel to sensing the sub-channels used by the SUL, or reactively upon request of the link reconfiguration process. Whether sensing on potential new sub-channels is done proactively or reactively obviously has a big impact on the time required for link reconfiguration.

Link reconfiguration can be realized centrally controlled or distributed, which has an influence on the required communication overhead. The required communication overhead further depends on the size of the SUG. The bigger the SUG, the more control traffic is required for the negotiation of the new resources and the longer is the duration of the reconfiguration process. Although there are some investigations on protocols for link reconfiguration (e.g. [9, 10]), we are not aware of any investigations of the spectral overhead imposed by such approaches.

C. Reconfiguration Compensation

Link reconfiguration requires time during which the bandwidth of the SUL is reduced. To maintain the QoS of providing a minimum bandwidth availability, these bandwidth reductions have to be compensated for.

Over-provisioning is a common compensation approach, i.e., adding redundancy to the SUL such that even if data transmission on *some* part of the SUL has to be interrupted, there is still enough spectrum available for the SUL to offer the minimum bandwidth.

The basic idea is to add X redundant sub-channels to the N_{sul} sub-channels of the SUL and apply coding in such a way that the receiver can decode the message, if *any* N_{sul} out of the $N_{\text{sul}} + X$ sub-channels are received. For details please refer, e.g., to [11]. Using this approach the SUL can tolerate the concurrent appearance of PUs on up to X sub-channels and still maintain the required link bandwidth (provided that the sub-channels are all located in different PU bands). We refer to this approach as *redundancy* approach. The amount of redundancy added to the SUL has a direct influence on the spectral overhead.

Figure 3(a) shows an example using $X = 1$ redundant sub-channels to support a minimum bandwidth of $N_{\text{sul}} = 2$ sub-channels resulting in a total number of 3 sub-channels used for the SUL (sub-channels 1, 5, and 7). In the second sensing period the PU covering sub-channels 1 and 2 appears, such that sub-channel 1 has to be excluded from the SUL. However, the QoS requirement is nevertheless satisfied since there are still two sub-channels available for the SUL. During the next data transmission period, the link has to be maintained, i.e., the communication peers have to agree on a new (redundant) sub-channel to be added to the SUL (sub-channel 3 in the figure).

An alternative to the redundancy approach is a *resource reservation* approach in which some backup spectrum is reserved for the SUL, which is instantly available if the link needs to be reconfigured. The SUL can then be immediately switched to use the backup spectrum. This obviously implies that, apart from the sub-channels used for the SUL, additional backup sub-channels need to be maintained. Equally to the regular sub-channels of the SUL, the backup sub-channels have to be regularly sensed to ensure that no PU appeared. Reservation of backup spectrum also requires the SUG to ensure that the respective spectrum is not used by other SULs. Obviously, the amount of resources that are reserved for an SUL has an influence on the spectral overhead.

In Figure 3(b) we show an example for the resource reservation approach. Initially, sub-channels 1 and 7 are used for the SUL while sub-channel 3 is maintained as a backup sub-channel. During the second sensing period, the PU covering the first two sub-channels appears and consequently sub-channel 1 has to be vacated. Within the reconfiguration period (t_{reconf}) a replacement for sub-channel 1 has to be selected from the backup sub-channels[1] and in addition a new backup

[1] Although we are only showing one backup sub-channel in our example, there will generally be more than one, so that a selection process has to take place if the SUL needs to be reconfigured.

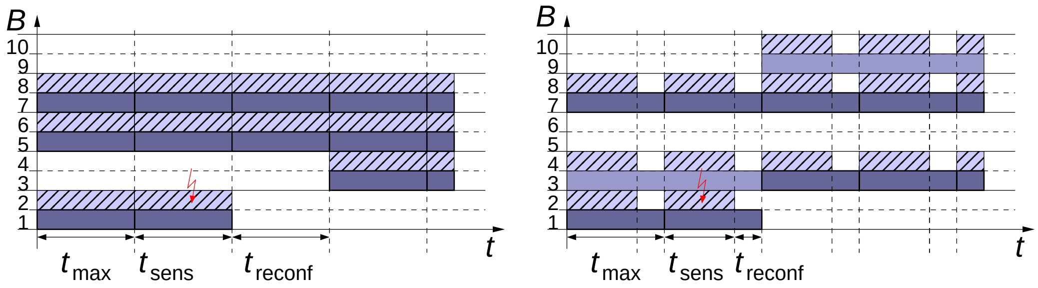

(a) Redundancy approach

(b) Resource reservation approach

Fig. 3. Link maintenance approaches for non-contiguous SULs and partial sensing. The dark areas correspond to data transmission periods, whereas the light, shaded areas indicate spectrum sensing periods. The light unshaded areas mark reserved resources.

sub-channel has to be selected (sub-channel 9 in our example). At the end of $t_{\max}$ the selection process has to be finished such that the SUL can be switched to use sub-channel 3 instead of 1. One assumption in this example is that the time needed to switch the sub-channel usage of the SUL (from sub-channels 1 and 7 to 3 and 7) is negligible and has no influence on the QoS.

While there are investigations on QoS support for secondary communication based on redundancy (e.g. [11]) and resource reservation (e.g. [12]), the spectral overhead required has — to our knowledge — not been investigated yet.

IV. Overall System Design Considerations

In the previous sections we have introduced the spectral overhead involved in both the sensing process and the link maintenance process. Recall, that an SUL has to be maintained even if the sensing process reports a false positive. Thus, we can consider the probability of false positives (P_{fp}) as a parameter linking both the above mentioned overheads. A high P_{fp} allows indeed for a small sensing spectral overhead but imposes potentially a larger overhead for link maintenance. Reducing P_{fp} on the other hand, results in a larger overhead required for sensing but also in a smaller overhead required for link maintenance. There is, thus, a clear tradeoff between the spectral overhead used for sensing and the spectral overhead used for link maintenance.

In order to enable an orderly performance analysis of this tradeoff, we will now introduce some measures for efficiency of spectrum usage considering the sensing and maintenance overhead jointly. For simplicity we consider in the following a spatial area small enough as to assure the *unified* availability or non-availability of spectrum for opportunistic re-usage. For such an "elementary area" we define available spectrum in terms of frequency and time availability. To quantify the available spectrum, we conceptually divide the time axis into slots of length $t_{\max}$ and define available spectrum at a basis of $t_{\max}$. We will differentiate between three notions of available spectrum.

A. Theoretically Available Spectrum

As the theoretically available spectrum ($C_{\max}$) we denote the sum of spectrum which is not used by the PU (within a time slot and the elementary spatial unit under consideration). The theoretically available spectrum is the "ground truth" or the "Gods view" and reflects the real availability of spectrum (based on the spectrum usage of the PU). It can be seen as a benchmark to compare the performance of different DSA approaches. Up to now, there are only very limited publicly available research results analyzing the theoretically available spectrum. Part of the reason is that in-network data is needed, which spectrum owners often are not willing to share. However, in a cooperation with a big U.S. network provider we were able to publish one of the first investigations of the theoretically available spectrum in a cellular voice network [13]. Another example is the assessment of unused TV spectrum based on the evaluation of the TV station database for the U.S. [1].

B. Sensed Available Spectrum

The sensed available spectrum (C_{sens}) is the spectrum *sensed* to be available for secondary usage and is, thus, an estimate of the theoretically available spectrum as discovered using a given sensing approach. How close this estimate comes to the theoretically available spectrum depends on the quality of the sensing process, i.e., the probability of false positives (P_{fp}).

Recall that the reliability of the sensing process is to be guaranteed, i.e. the sensing process has to provide some very limited probability of false negatives. Thus, the sensed available spectrum depends on the probability of false positives. It specifies how much of the idle spectrum can be discovered without declaring spectrum used by the PUs as idle. Note, however, that the sensed available spectrum does not consider the spectral overhead required for sensing. The spectral overhead will be accounted for in the effectively available spectrum (C_{eff}) defined below.

C. Effectively Available Spectrum

In Section II we have explicitly indicated that each sensing approach requires some spectral overhead: keeping some spectrum resources unused in order to make sensing possible and also spectral resources needed for exchanging sensing data in the case of distributed sensing. The definition of the sensed available spectrum ignores this need.

Additionally, as indicated in Section III, also the link maintenance process requires spectral overhead to maintain the secondary QoS: backup spectrum / redundancy to compensate for reconfigurations and additional spectrum for signaling (to negotiate the reconfiguration of the SUL within the SUG).

Therefore, we introduce the additional concept of effectively available spectrum (C_{eff}) being the part of the spectrum which actually can be used for secondary data transmission according to the spectrum sensing results achieved on the basis of the used approach for spectrum sensing and SUL reconfiguration.

In fact, the effectively available spectrum quantifies the efficiency of a given secondary usage approach in terms of how much of the unused spectrum could be effectively recovered for secondary usage. Additionally, the effectively available spectrum enables to quantify the tradeoff between the sensing overhead and the overhead required for secondary link maintenance.

V. System Design Example

In the previous sections we have identified the spectral overhead required for spectrum sensing and link maintenance and introduced performance metrics to quantify this overhead. In this section we give a simple example showing the potential tradeoffs on the effectively available spectrum (C_{eff}) imposed by the different overheads. While the presented analysis only considers some selected aspects, it clearly shows the potential design space and motivates further research.

A. System Model

We assume for our investigation a system based on Orthogonal Frequency Division Multiplexing (OFDM), which means that spectrum sensing can be done *in parallel* for the whole operation range of the secondary network. We consider a system model supporting partial sensing and non-contiguous SULs. The investigated link maintenance approaches are based on the examples of Section III-C.

For the performance analysis we consider a scenario with N_{pu} different PUs each covering a bandwidth of B hertz divided into N_{sub} sub-channels of bandwidth B_{sub}. The maximum interference time of each PU is $t_{\max}$. The probability that a PU is active within $t_{\max}$ is set to P_{pu}. Let us denote by P_{fn} and P_{fp} the probability of false negatives and false positives, respectively, and by $P_{\text{d}} = 1 - P_{\text{fn}}$ the probability of detecting the PU.

The sensed available spectrum (C_{sens}) is determined by the probability of false positives, i.e., the theoretically available spectrum ($C_{\max}$) less the reduction in available spectrum due to false positives. Assuming a slotted system (with a slot length of $t_{\max}$) C_{sens} per slot can be computed as $C_{\text{sens}} = 1 - P_{\text{fp}}$ (defined as a fraction of $C_{\max}$).

As introduced in Section II, the probability of false positives (P_{fp}) depends on the amount of spectrum used for sensing. This can be quantified by the number of sensing samples (N). In the time domain, N depends on the sensing time t_{sens} and the sampling frequency f_{s}. Furthermore, we assume that the SUs use OFDM based sensors, such that N_{sens} power samples are recorded in parallel with N_{sens} being the number of sub-channels used for sensing. We choose both, f_{s} and the bandwidth of a sub-channel such that sensing samples are i.i.d. Given the mentioned constraints $N = N_{\text{sens}} \cdot t_{\text{sens}} \cdot f_{\text{s}}$.

An SUL consists of N_{sul} sub-channels and the time needed for reconfiguration of the link is specified by t_{reconf}. We assume the use of non-contiguous SULs and that the sub-channels of an SUL are distributed over different PUs. We constrain ourselves to non-overlapping SULs for this investigation. We use the following values for our analysis: $N_{\text{pu}} = 100$, $N_{\text{sub}} = 50$, $N_{\text{sul}} = 10$, $P_{\text{pu}} = 0.2$, $P_{\text{d}} = 0.99$, $f_{\text{s}} = 500\,\text{kHz}$, $t_{\max} = 0.5\,\text{s}$, and $t_{\text{reconf}} = 0.1\,\text{s}$.

B. Performance Analysis

For the calculation of the probability of false positives (P_{fp}) we use the formula introduced in [4], which defines P_{fp} as a function of the number of power samples (N) used:

$$P_{\text{fp}}(N) = Q\left(\sqrt{2\gamma+1}\, Q^{-1}(P_{\text{d}}) + \sqrt{N}\,\gamma\right) , \qquad (1)$$

with $Q(x)$ being the Q-function and γ the received Signal-to-Noise Ratio (SNR). For details please refer to [4].

The spectral overhead required for partial sensing depends on the number of sub-channels used for sensing (N_{sens}). Thus, for both, the reservation as well as for the redundancy approach, the effectively available spectrum only considering sensing overhead (C_{so}) can be computed as

$$C_{\text{so}} = C_{\text{sens}}\left(1 - \frac{N_{\text{sens}}}{N_{\text{sub}}}\right) . \qquad (2)$$

In order to calculate the overhead required for the reconfiguration compensation approaches, we need to calculate the number of sub-channels to be reserved for the SUL (for the reservation approach) or the number of redundant sub-channels to be added to the SUL (for the redundancy approach). The number of additional sub-channels needed depends on the number of sub-channels that need to be reconfigured concurrently (within one slot) which in turn depends on the probability that a PU was detected. Recall, that it does not matter whether the PU really appeared or whether the sensing process reported a false positive — the respective sub-channel has to be replaced in either case. The probability that a sub-channel has to be replaced can be computed as

$$P_{\text{rec}} = P_{\text{pu}} \cdot P_{\text{d}} + (1 - P_{\text{pu}}) \cdot P_{\text{fp}} . \qquad (3)$$

Using P_{rec} and the binomial coefficient, and given a number of X redundant / reserved sub-channels, we can calculate the probability P_{int} that there are not enough sub-channels available to support the required QoS [11]. Choosing a target probability of $P_{\text{int}} = 0.01$ we can numerically find the optimal number of reserved / redundant sub-channels X_{opt} satisfying this criteria. Using X_{opt} sub-channels for the two maintenance approaches, C_{eff} can be determined as

$$C_{\text{eff}} = C_{\text{so}}\left(1 - \frac{X_{\text{opt}}}{N_{\text{sul}} + X_{\text{opt}}}\right) . \qquad (4)$$

For details of the calculations above please refer to [14].

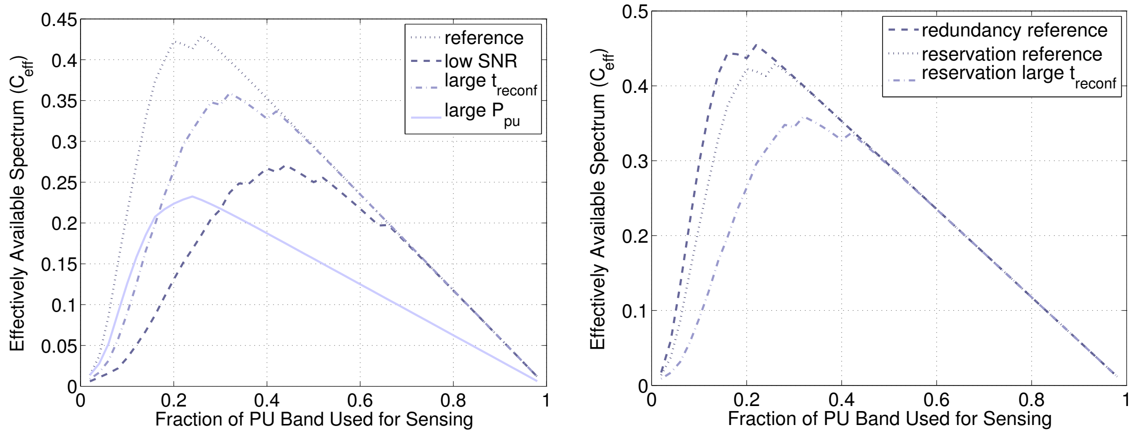

(a) Reservation approach.

(b) Comparison of the redundancy and the reservation approach.

Fig. 4. Spectral efficiency (C_{eff}). The "reference" values used are $t_{\text{reconf}} = 0.1\,\text{s}$, $P_{\text{pu}} = 0.2$, $SNR = -22\,\text{dB}$; for "low SNR" we use $SNR = -24\,\text{dB}$; for "large t_{reconf}" we use $t_{\text{reconf}} = 0.25\,\text{s}$; for "large PU" we use $P_{\text{pu}} = 0.5$.

C. Performance Results

Figure 4(a) shows the performance results for the resource reservation approach and the influence of different system parameters on the results. All graphs clearly show the tradeoff between the spectral overhead used for sensing and for link maintenance resulting in an optimum amount of sensing overhead which maximizes the effectively available spectrum (C_{eff}).

Looking, e.g., at the reference graph (dark, dotted line), the maximum for C_{eff} is reached if 22% of each PU band is used for sensing. Using less spectrum for sensing increases the probability of false positives which results in a higher probability of link reconfiguration and thus requires more sub-channels to be reserved in order to ensure the QoS requirement. Conversely, increasing the amount of spectrum used for sensing reduces the effect of false positives but also results in more spectrum blocked for sensing and, thus, not usable for data transmission. Using less spectrum for sensing than the optimum, false positives are the dominating effect on C_{eff} whereas using more than the optimum the spectrum blocked for sensing is the dominating effect on C_{eff}.

In Figure 4(b) we compare the redundancy with the reservation approach. While the redundancy approach shows the same general behavior as the reservation approach, its performance is slightly better for the used reference values. This results from the different sensing times used in both approaches. In the reservation approach, t_{sens} depends on t_{reconf}: The longer t_{reconf}, the shorter t_{sens}. For the redundancy approach t_{sens} is independent from t_{reconf}. Obviously, the difference between both approaches depends on the length of t_{reconf}.

Note, however, that if the performance of multiple SULs is considered, the reservation approach has the advantage of a potential multiplexing gain. Instead of having each SUL maintaining its own set of backup spectrum, multiple SULs could cooperatively maintain a common backup spectrum resulting in an increase of the overall spectral efficiency.

VI. Conclusions

In this paper, we have investigated mechanisms for spectrum sensing and maintenance of Secondary User Links (SULs). We show by means of a hypothetical example how, given a certain spectral overhead, it is — in principal — possible to assure a certain minimum bandwidth availability for an SUL despite necessary reconfigurations due to the detection of PUs and still assuring a proper protection of the PU communication.

Furthermore, we show that there exists a clear tradeoff between the spectral overhead used for sensing and the spectral overhead used for link maintenance. We have demonstrated that this tradeoff is driven to a big extent by the sensing quality. In practice the cost and complexity of high accuracy sensors will have a big influence on the system design. This paper demonstrates clearly, that the inaccuracy of the sensing process can be to some extent compensated by a very efficient (low overhead) design of the link maintenance process.

Obviously, we have given here only some basics for the efficient design of secondary QoS support in systems with opportunistic spectrum usage. Further research should give more consideration to several important additional aspects, some of which are outlined below.

In order to precisely quantify the overhead tradeoff shown in this paper several additional effects need to be considered and quantified. One aspect having a big influence on the overhead required for both, spectrum sensing and link maintenance is the usage pattern of the PU. A very agile PU frequently using its spectrum but only for short time periods potentially results in more required overhead than sluggish one.

Furthermore, the signaling overhead required by the Secondary User Group (SUG) needs to be quantified. Signaling is required to derive a verdict out of a distributed sensing process. The required overhead depends on many factors such as the size of the SUG or the algorithm used for merging the sensing samples. Signaling overhead is also required by the link maintenance process for link reconfiguration. In this paper we have only considered a generic reconfiguration time to qualitatively access the required overhead. In depth studies are needed for quantifying this overhead.

Policies for the selection of spectral resources to be used for secondary communication are another aspect to consider. Ideally, an SUG will be able to operate in a wide spectrum range covering multiple different types of PU bands. The "spectral footprint" of a PU signal, i.e., the specific electromagnetic characteristic of the signal has an influence on how hard it is for the sensing process to detect such a signal. The same applies to the strength of the signal at the sensing devices. Finally, $t_{\max}$ is also of importance for the overhead. In which PU bands to place an SUL, thus, has an influence on the spectral overhead required by the sensing process.

The presented work was restricted to the case of non-overlapping SULs. With secondary usage becoming more and more popular, there will, however, likely be cases of multiple SUGs operating in the same area resulting in an overlap of their communication ranges. If those SUGs operate in an uncoordinated manner, collisions in spectrum usage are unavoidable. An interesting research topic is, thus, to quantify the spectral overhead required for coexistence mechanisms, supporting a collision-free operation and comparing it to the spectrum lost due collisions in case of uncoordinated operation.

Last but not least, more data about the real performance of sensing systems and reconfiguration mechanisms, as well as usage statistics of PUs, is needed for an exact quantitative performance analysis. Such data will become available as more experimental work on opportunistic spectrum usage will be published.

REFERENCES

[1] R. Tandra, S. M. Mishra, and A. Sahai, "What is a spectrum hole and what does it take to recognize one?" *Proc. IEEE*, vol. 97, no. 5, pp. 824–848, May 2009.

[2] A. Sahai, R. Tandra, S. M. Mishra, and N. Hoven, "Fundamental design tradeoffs in cognitive radio systems," in *TAPAS '06: Proceedings of the first international workshop on Technology and policy for accessing spectrum*. New York, NY, USA: ACM Press, 2006, p. 2.

[3] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Communications Surveys Tutorials*, vol. 11, no. 1, pp. 116–130, 2009.

[4] Y.-C. Liang, Y. Zeng, E. Peh, and A. T. Hoang, "Sensing-throughput tradeoff for cognitive radio networks," *IEEE Trans. Wireless Commun.*, vol. 7, no. 4, pp. 1326–1337, 2008.

[5] Y.-J. Choi, Y. Xin, and S. Rangarajan, "Overhead-throughput tradeoff in cooperative cognitive radio networks," in *Proc. of the IEEE Wireless Communications and Networking Conference WCNC 2009*, Apr. 2009, pp. 1–6.

[6] W. Hu, D. Willkomm, L. Chu, M. Abusubaih, J. Gross, G. Vlantis, M. Gerla, and A. Wolisz, "Dynamic frequency hopping communities for efficient IEEE 802.22 operation," *IEEE Commun. Mag.*, vol. 45, no. 5, pp. 80–87, May 2007.

[7] IEEE, "IEEE 802.22 working group," 2010. [Online]. Available: http://www.ieee802.org/22

[8] R. W. Broderson, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "CORVUS: A cognitive radio approach for usage of virtual unlicensed spectrum," White Paper, 2004, a condensed version of this paper is available in the proceedings of the 14th IST Mobile Wireless Communications Summit 2005 under the title "A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum".

[9] C. Han, J. Wang, and S. Li, "A spectrum exchange mechanism in cognitive radio contexts," in *Proc. of the 17th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC 2006*, Sep. 2006, pp. 1–5.

[10] Q. Shi, D. Taubenheim, S. Kyperountas, P. Gorday, and N. Correal, "Link maintenance protocol for cognitive radio system with ofdm phy," in *Proceedings of the 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, 2007 (DySPAN 2007)*, 17-20 April 2007, pp. 440–443.

[11] D. Willkomm, J. Gross, and A. Wolisz, "Reliable link maintenance in cognitive radio systems," in *Proceedings of the 1st IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN) 2005*, Baltimore (MD), Nov. 2005, pp. 371–378.

[12] S. Subramani, S. Armour, D. Kaleshi, and Z. Fan, "Spectrum scanning and reserve channel methods for link maintenance in cognitive radio systems," in *Proc. of IEEE Vehicular Technology Conference VTC Spring 2008*, May 2008, pp. 1944–1948.

[13] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary user behavior in cellular networks and implications for dynamic spectrum access," *IEEE Commun. Mag.*, vol. 47, no. 3, pp. 88–95, Mar. 2009.

[14] D. Willkomm, "Spectrum efficient QoS support for secondary users in cognitive radio systems," Telecommunication Networks Group, Technische Universität Berlin, Tech. Rep. TKN-09-008, Dec. 2009.