This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

The Effectiveness of Opportunistic Spectrum Access: A Measurement Study

Vinod Kone, Lei Yang, Xue Yang, Ben Y. Zhao, and Haitao Zheng

***Abstract*—Dynamic spectrum access networks are designed to allow today's bandwidth-hungry "secondary devices" to share spectrum allocated to legacy devices, or "primary users." The success of this wireless communication model relies on the availability of unused spectrum and the ability of secondary devices to utilize spectrum without disrupting transmissions of primary users. While recent measurement studies have shown that there is sufficient underutilized spectrum available, little is known about whether secondary devices can efficiently make use of available spectrum while minimizing disruptions to primary users. In this paper, we present the first comprehensive study on the presence of "usable" spectrum in opportunistic spectrum access systems, and whether sufficient spectrum can be extracted by secondary devices to support traditional networking applications. We use for our study fine-grain usage traces of a wide spectrum range (20 MHz–6 GHz) taken at four locations in Germany, the Netherlands, and Santa Barbara, CA. Our study shows that on average, 54% of spectrum is never used and 26% is only partially used. Surprisingly, in this 26% of partially used spectrum, secondary devices can utilize very little spectrum using conservative access policies to minimize interference with primary users. Even assuming an optimal access scheme and extensive statistical knowledge of primary-user access patterns, a user can only extract between 20%–30% of the total available spectrum. To provide better spectrum availability, we propose *frequency bundling*, where secondary devices build reliable channels by combining multiple unreliable frequencies into virtual frequency bundles. Analyzing our traces, we find that there is little correlation of spectrum availability across channels, and that bundling random channels together can provide sustained periods of reliable transmission with only short interruptions.**

***Index Terms*—Cognitive radio, measurement, software radio, wireless communication.**

I. Introduction

RADIO spectrum is perhaps the wireless industry's most valuable asset. The deployment and growth of any wireless network depend on the amount of spectrum it can access. Despite its recognized value, current policies on spectrum distribution are highly inefficient. Spectrum frequency ranges are assigned statically to wireless carriers in long-term leases, generally ignoring market demands that vary significantly over time. Over the years, the large majority of frequency ranges have been assigned, leaving little room for new technologies or growth. Meanwhile, demands for previously assigned frequencies have dropped significantly, leaving most ranges woefully underutilized at an average of 5% of capacity [1].

Opportunistic and dynamic spectrum access is a new access model designed to "extract" unused spectrum from allocated but underutilized frequencies, supporting newcomer traffic without affecting existing owners. In this model, wireless devices that need spectrum locate and "opportunistically (re)use" unused frequencies ranges. These "secondary" devices take great precaution to avoid disrupting original or "primary" users and immediately exit the frequency whenever they detect traffic from primary users. Through this carefully planned access model, secondary devices can increase spectrum utilization with zero or bounded disruptions to existing owners. Note that compared to more liberal spectrum access rules [2], this "conservative" access model is easier to implement and much more likely to gain acceptance with regulators and primary users.

The success of the dynamic spectrum access model depends heavily on both the availability of unused spectrum and whether secondary devices can efficiently extract and utilize them. While a number of measurement studies have measured and modeled the availability of unused spectrum [3]–[8], the community has generally overlooked the second factor and optimistically assumed that secondary devices can always efficiently utilize available spectrum. Despite its importance, little is known about whether secondary devices can efficiently make use of available spectrum, given the hard constraints of avoiding disruptions to primary users. This is understandable since such a study requires access to a fine-grained measurement trace of spectrum usage, which has not been available until recently.

In this paper, we present the first comprehensive study of performance in opportunistic spectrum access systems that limit disruptions to unpredictable primary users. Our goal is to understand whether dynamic spectrum access can provide reliable spectrum to secondary users while respecting hard disruption limits that protect primary-user transmissions. Our study can address key concerns about the feasibility of supporting traditional network applications in this new model by performing a deep analysis of a large collection of spectrum usage measurements. These measurements are taken from four locations across the globe: two in Germany, one in the Netherlands, and one in Santa Barbara, CA. Each measurement uses a spectrum

Manuscript received March 23, 2011; revised December 01, 2011; accepted February 14, 2012; approved by IEEE/ACM Transactions on Networking Editor M. Liu.

V. Kone was with the Department of Computer Science, University of California, Santa Barbara, CA 93117 USA. He is now with the Runtime Systems Group, Twitter, San Francisco, CA 94107 USA (e-mail: vinod@cs.ucsb.edu).

L. Yang was with the Department of Computer Science, University of California, Santa Barbara, CA 93117 USA. He is now with the Wireless Communications Lab, Intel Labs, Hillsboro, OR 97124 USA (e-mail: leiyang@cs.ucsb.edu).

X. Yang is with Intel Labs, Hillsboro, OR 97124 USA (e-mail: xue.yang@intel.com).

B. Y. Zhao and H. Zheng are with the Department of Computer Science, University of California, Santa Barbara, CA 93117 USA (e-mail: ravenben@cs.ucsb.edu; htzheng@cs.ucsb.edu).

Digital Object Identifier 10.1109/TNET.2012.2191571

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

analyzer to sweep a range of radio frequencies between 20 MHz and 6 GHz for a period of 2–7 days, capturing the raw energy level observed on each of the 200-kHz frequency channels at a periodic interval of 0.65 or 1.8 s. These results capture, at a very fine granularity, when specific radio frequencies are occupied by primary users in the measurement area. This dataset is unique in its combination of wide frequency coverage (20 MHz–6 GHz), measurement length (one week for three of the locations), and measurement frequency (one sweep per 1.8 or 0.65 s compared to 75 s of prior studies [5]). We extract from them spectrum occupancy traces (occupied or free) across a large set of frequencies, covering 5922 wireless channels and a total of more than 5 billion data points for analysis. While four locations are in no way representative of spectrum usage in general, these measurements do provide initial insights into whether opportunistic spectrum access has the potential to support traditional networking applications.

Our analysis of spectrum availability (Section III) confirms that most assigned frequencies are heavily underutilized. Out of 5922 channels analyzed, an average of 26% (or 1267 channels) were partially occupied (5%–95% occupancy). We are primarily interested in evaluating dynamic spectrum access on these channels since other channels are either fully occupied (20% of our dataset, or 1317 channels) or can be statically allocated as free channels (54%) of our dataset, or 3338 channels. We also observe that spectrum availability varies significantly based on the frequency range and measurement location. More importantly, short-term availability varies significantly across time, and both idle duration and busy periods show high variance. This highly variable spectrum availability poses significant challenges to secondary devices, making it harder to access and utilize a channel while respecting a fixed limit of disruptions to primary users.

In Section IV, we use these spectrum traces to compare the performance of two “optimal” opportunistic access mechanisms: one scheme where secondary devices have zero knowledge of primary-user patterns, and one where secondary devices have accurate statistical knowledge of the primary-user accesses [9]. We are shocked to find that, even with accurate statistical knowledge of primary-user accesses, secondary devices can only extract 20%–30% of the available spectrum under a reasonable disruption limit of 10%, and less than 10% of spectrum if the disruption limit drops to 1%. In addition, spectrum extracted from each channel is heavily fragmented and scattered across time. As a result, the equivalent channels available to secondary devices are highly unreliable—spectrum access on each channel is frequently interrupted and often takes 10–100 s before being restored.

However, there is hope. We propose and evaluate *frequency bundling*, where secondary devices build reliable transmission channels by combining together multiple unreliable frequencies, essentially utilizing frequency diversity to compensate for the lack of reliability on individual channels. To evaluate different bundling strategies, we analyze correlation between availability patterns of different 200-kHz channels and find little or no correlation (Section V). This availability independence means that we can significantly improve overall reliability by simply bundling random channel pairs together. Experimental results from our datasets are promising. Using a random bundling strategy, the improvement in channel reliability scales exponentially with the size of the bundle. For example, bundling 5–10 randomly selected channels together will reduce the secondary device’s blocking time by two orders of magnitude. The resulting new channel enjoys average transmission periods of 120–1300 s while being occasionally interrupted by 2–4 s.

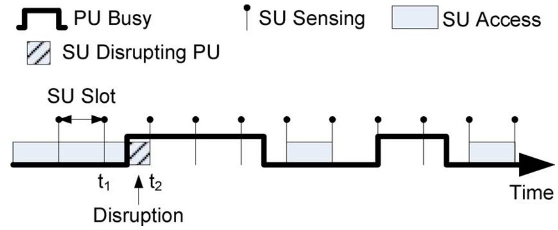

Fig. 1. Illustrative example of opportunistic spectrum access. The bold line shows the primary user (PU)’s channel occupancy. A secondary user (SU) periodically senses channel to detect primary user and determine whether to access the channel. A disruption occurs if the primary user returns in the middle of secondary user transmissions.

In summary, our study provides a first look into the feasibility of accessing spectrum opportunistically while respecting hard limits to disruptions to primary users. We show that statistical knowledge of spectrum occupancy can improve the performance of opportunistic access by a factor of 2–3. Nonetheless, we also show that given the unpredictable nature of primary-user access, current spectrum access methods cannot provide usable channels to secondary devices. Only by bundling multiple unreliable channels together can we provide reasonable levels of reliability to network applications on these devices.

II. Overview

In this section, we first provide background information on opportunistic spectrum access. We then describe the objectives of our investigation and the datasets we use.

A. Opportunistic Spectrum Access

Opportunistic spectrum access involves two entities: primary users or original owners of allocated but underutilized frequencies, and secondary users who seek to make use of unused spectrum, under the hard constraints of avoiding disruptions to primary users at all costs [10], [11], [9].

Fig. 1 shows a representative example of opportunistic spectrum access on a partially used primary-user channel. A secondary user x accesses the channel using a slotted sensing-then-access mechanism. At the start of each slot, x senses the channel to detect whether any primary user is present, often using a RF energy detection [12]. If the channel is occupied, x does nothing and waits until the next slot. If the channel appears to be unused, x will decide whether to access the channel in the current slot. In order to satisfy hard primary-user disruption limits, x must carefully access the risk of using the channel because the primary user can potentially return in the middle of its transmission slot. When necessary, x will give up using an idle channel to avoid disrupting the original owner.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

TABLE I
15 ORIGINAL SPECTRUM OWNERS AND THEIR FREQUENCY RANGES (MHz) MEASURED BY THE RWTH DATASET

Original owner	TV1	Aviation	Marine	TV2	TV3	GSM900 UL	GSM900 DL	DAB
Freq. Range (MHz)	41-67	109-136	157-173	175-229	471-861	890-915	935-960	1453-1491
Original owner	Meteo	GSM1800 UL	GSM1800 DL	DECT	UMTS UL	UMTS DL	ISM	
Freq. Range (MHz)	1675-1710	1710-1785	1805-1880	1882-1897	1920-1980	2110-2170	2400-2500	

B. Goals

By analyzing real-world measurements on primary-user spectrum usage patterns, we have three key goals. First, we wish to understand the feasibility and effectiveness of opportunistic spectrum access. More specifically, we seek to examine the availability of both completely unused and intermittently used spectrum. For intermittently used channels, we also seek to examine the amount of spectrum actually accessible to secondary devices, given the hard constraints of avoiding disruptions to primary users.

Second, we seek to examine the role of various design decisions and network factors in opportunistic spectrum access, including the disruption limit set by the original owners, the time granularity of spectrum access, and the type of information available to secondary devices about the original owners.

Finally, we are interested to examine practical issues in utilizing extracted spectrum to support today's wireless services. Because the extracted spectrum is fragmented across time and frequency, we seek to identify ways to build reliable wireless transmission from scattered spectrum pieces.

C. Datasets

We use two datasets in our analysis. They are unique in their combination of wide frequency coverage, extensive measurement length, and fine-grained measurement frequency.

The first dataset, used for most of the analysis, records the received signal strength across 20 MHz–6 GHz at three locations over a period of one week. Table I lists some of the original owners and their frequency ranges. The measurement was performed by the Mobnets group of RWTH Aachen University, Aachen, Germany [13]. An Agilent E4440A high-performance spectrum analyzer with a resolution bandwidth of 200 KHz and frequency span of 1500 MHz was used [14]. The three measurement sites were: 1) on a balcony of a residential building in Germany (GER1); 2) inside an office building in Germany (GER2); and 3) on a rooftop in the Netherlands (NED). At each location, a spectrum analyzer repeatedly swept the 20-MHz–6-GHz frequency range, measuring signal energy on each of the 200-kHz frequency channels. The measurement uses a 1.8-s sweep time. That is, any two subsequent measurements on a single channel were 1.8 s apart. Using this dataset, we analyzed 5622 channels corresponding to the service bands listed in Table I.

The second dataset came from our own measurements at the University of California, Santa Barbara (UCSB), over a period of two days at the end of April 2010 when school was in session. The goal of these measurements is to sample primary-user access patterns at a finer granularity than the first dataset. It contains the received energy strength in the 1925–1995-MHz GSM frequency band, observed in an office trailer. We configured a GSM1900 digital receiver (Agilent E6454C) as a spectrum analyzer that swept the GSM frequency with a resolution of 200 kHz. Unlike a wideband spectrum analyzer, our digital analyzer only tunes to GSM frequencies. However, since it covers a much smaller frequency range (300 channels), we can increase the sweep frequency to once every 0.65 s.

Preprocessing: We preprocess our datasets to convert the received signal strength traces to spectrum occupancy patterns (busy or idle) on each measurement channel. To do so, we use the energy-detection method [5], [7] and select (for each 200-kHz measurement channel) an energy threshold of -107 dBm that is specified by the IEEE 802.22 standard for TV bands [15]. We declare a frequency channel as occupied (or busy) at a given time if its measured signal strength is above the threshold. While service bands could use different thresholds to protect their transmissions, there are no reasonable guides on what those individual thresholds should be. Thus, we apply this known threshold uniformly across different service bands. For the NED location in the RWTH measurement as well as our own UCSB measurement, we use a slightly higher threshold of -100 dBm. This is to compensate for the presence of stronger noise floor due to the proximity to a railway station in the case of NED (also recommended by [7]) and the presence of metal walls and obstacles in the case of UCSB measurements.

D. Assumptions

We make a few assumptions in order to perform analysis on the measurement datasets.

First, because both measurements sweep the frequency band sequentially to measure a wide frequency range, they do not capture usage activities at time granularity smaller than the sweeping time. Thus, we set the secondary user's access slot size to be the same as the sweeping time. Having said that, our results from the finer-grained UCSB data set [16] show that the granularity of sensing interval does not impact our conclusion on the low extraction and blocking time. In addition, we note that the sweeping times of our datasets (1.8 s for the RWTH dataset and 0.65 s for the UCSB dataset) are two orders of magnitude smaller than previous measurements of 75 s [5].

Second, we capture the effect where a primary user returns to the channel in the middle of a slot in our calculations of the primary-user disruption rate. Specifically, if an idle slot is followed by an occupied slot, then the primary user is likely to arrive in the middle of the first slot. If the secondary user decides to transmit in the first slot, we flag this slot as creating a disruption to both the primary user and the secondary user. We compute the primary-user disruption rate as the ratio of primary-user busy blocks that suffer any disruption [9].

Finally, we assume that secondary users' sensing is accurate, and that multiple secondary users coordinate their access to avoid transmission collision. Since our focus is on studying the impact of spectrum usage patterns of original owners, we abstract multiple coexisting secondary users into a single

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

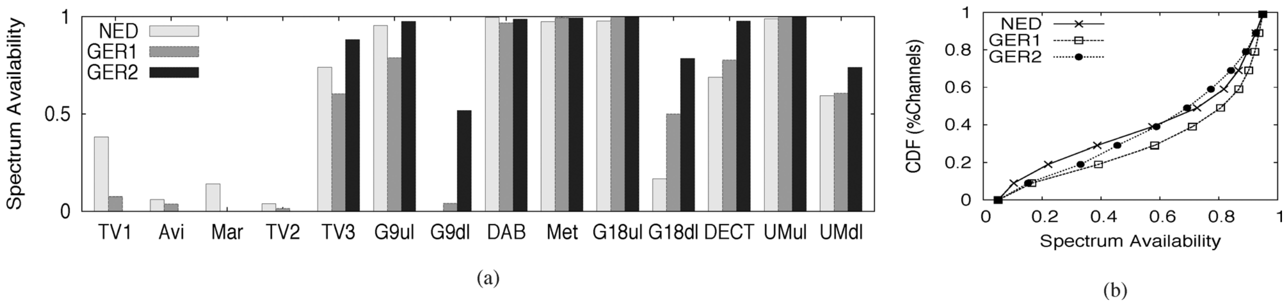

Fig. 2. (a) Average spectrum availability of various service bands over the entire measurement period. The services are ordered in the ascending order of their operating frequencies. Ample unused spectrum exists at all three locations, but the availability varies across locations and frequencies. (b) Cumulative distribution of spectrum availability of all partially used spectrum channels, which is evenly distributed between [0.05, 0.95].

secondary link. The design and overhead of optimal spectrum sensing and coordination protocols, although important, are out of the scope of this paper. We refer the reader to [10], [11], and [17]–[19] for more details on cooperative spectrum sensing and sharing.

III. Spectrum Availability Analysis

The performance of opportunistic spectrum access depends heavily on the sustained availability of unused spectrum. In this section, using the RWTH data set, we examine in detail the availability of spectrum, its dependency on frequencies and locations, as well as its temporal dynamics. In total, we analyzed a one-week spectrum usage patterns (busy or idle) on each of the 5622 frequency channels. In the following, we first describe our findings on overall spectrum availability across frequencies and locations, and then present observed temporal dynamics on instantaneous spectrum availability.

A. Overall Spectrum Availability

We define *spectrum availability* (SA) as the percentage of measured intervals where a channel is not occupied by existing owners in a given time frame. While each service has its own operating channel width, in this study we treat each 200-kHz measurement band as a single spectrum channel.

Our study of 5622 spectrum channels corresponding to 15 selected services (Table I) shows that many spectrum channels are either completely free or partially used. Interestingly, for some of the services (e.g., TV3, GSM1800DL, and UMTSDL), the spectrum availability varies significantly across channels within the same service. To examine the impact of measurement location, Fig. 2(a) shows the spectrum availability measured at the three locations (NED, GER1, and GER2), averaged over a period of one week and across channels within each service band.

We make two key observations from these results. First, for all three locations, a significant portion of allocated spectrum is available for secondary devices. Second, the availability varies significantly across frequencies. Very low frequencies (TV1, Aviation, Marine, TV2) are heavily occupied, while others experience only light and moderate usage. The cellular uplink bands (GSM900UL, GSM1800UL, UMTSUL) are mostly idle because their signals are significantly weaker than those of downlink transmissions, and are thus harder to detect even using high-end spectrum analyzers. Nevertheless, we use these uplink measurements to examine opportunistic access, assuming that secondary users take extra precautions on these bands to avoid disrupting primary users, e.g., by lowering their transmit power.

After examining each channel in detail, we found that out of 5622 channels analyzed, 1176 channels are *partially occupied*, i.e., whose average spectrum availability is within [0.05, 0.95], and 3181 channels are idle, i.e., whose availability is greater than 0.95. In Fig. 2(b), we plot the cumulative distribution of the spectrum availability across these partially occupied channels and see that their availability is evenly distributed between 0.05 and 0.95. In the rest of the paper, we will focus on these partially occupied channels for which we must rely on opportunistic spectrum access to extract unused spectrum.

B. Dynamics of Available Spectrum

To understand both long- and short-term spectrum availability trends, we analyze the dynamics at two different granularity levels. To understand day-to-day trends, we start from dividing traces into half-hour segments and compute for each segment the average spectrum availability. Fig. 3(a) plots the resulting spectrum availability observed over 6 days on three selected GSM1800DL channels with intermediate spectrum availability, one for each location. In this case, spectrum availability varies significantly over time and displays a weak 1-day periodicity.

A more precise view of the channel idle/busy durations is shown in Fig. 3(b), for NED. It represents a randomly selected GSM1800DL channel for a period of 1 h between 11 AM and noon. In this example, the channel busy duration varies between 1.8 and 20 s, while the idle duration varies significantly between 1.8–100 s. The large variance in idle durations, however, poses significant challenges to secondary devices, making it harder to access and utilize a channel while respecting a fixed limit of disruption to original owners. We examine this challenge and its impact in greater detail next in Section IV.

IV. Performance of Opportunistic Spectrum Access

Our analysis of real-world measurements has demonstrated the ample scope for opportunistic spectrum access. In this section, we investigate its performance in terms of "extracting" the unused spectrum without disrupting original owners. As illustrated in Fig. 1, secondary devices sense and access spectrum in a slotted manner. Without knowing exactly when the primary user will return, secondary devices must take great precaution

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

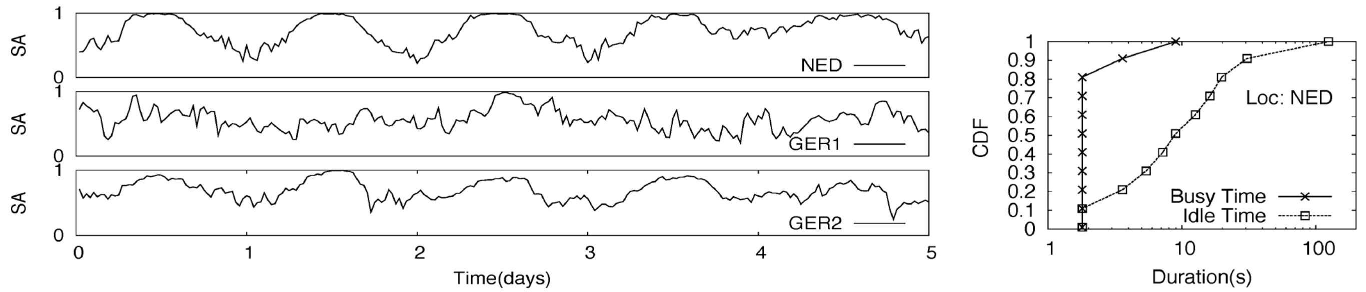

(a) (b)

Fig. 3. (a) Availability, averaged every 30 min, varies significantly over 5 weekdays on randomly selected GSM1800DL channels (one per location). (b) Cumulative distribution of busy and idle periods over 1 h. Large variation in idle duration poses significant challenges for opportunistic access. The GER1 and GER2 results are similar to the NED result and thus omitted.

and occasionally give up using an idle channel. As a result, they cannot extract all the available spectrum. Using the RWTH dataset, we seek to understand how much spectrum a secondary device can actually obtain.

Specifically, our analysis answers three key questions.

- What is the rate of spectrum extraction? Can statistical knowledge on primary-user spectrum usage patterns improve the performance and, if so, by how much?
- Is the average spectrum availability a reliable predictor of the amount of spectrum extracted?
- What is the usability of the extracted spectrum? How long must a secondary user wait to access a channel and how long does the access last?

In the following, we first describe the access strategies used in our analysis, and then address these questions.

A. Access Strategies

Given the primary-user disruption limit η and the probability density function of primary-user idle duration, prior work has developed optimal access strategies for opportunistic spectrum access [9]. At a high level, a secondary user x senses the channel at the start of an access slot t. If the channel is busy, x does nothing and waits until the next slot. If the channel is idle, x estimates the risk of accessing the current slot, using its past channel observations, the primary-user idle duration statistics $f(\cdot)$ and the primary-user disruption limit η. It has been proved that, using small access slots, the above strategy is optimal and satisfies the primary-user disruption limit. Further details can be found from [9].

We apply this optimal strategy to create two practical opportunistic access schemes.

- *No knowledge-based access (NKA):* This scheme requires no knowledge about primary-user usage patterns. Secondary devices will access a channel with a probability η (the primary-user disruption limit) when sensing it idle, leading to an extraction rate around η. This is the optimal result if the primary-user idle time follows the exponential distribution [9].
- *Statistical knowledge-based access (SKA)*. It assumes that secondary devices have the exact statistical distribution of primary-user idle time, $f(\cdot)$. Such knowledge is either provided by original owners or a third party or built by secondary devices via online/offline learning.

We note that secondary users can schedule channel access to utilize all available spectrum if and only if they can completely predict each primary user's spectrum usage events. This ideal scheme, however, is only feasible when the primary user displays a deterministic access pattern, which we did not find in our measurement datasets. Thus, we did not consider it in our analysis.

The SKA scheme requires an accurate statistical distribution of primary-user idle time. Results in Section III show that the distribution varies significantly over time, especially within the same day. To make a fair evaluation, we apply time-series analysis to segment traces of each frequency channel into multiple time segments, each displaying stable availability [20]. Our segmentation results show that the spectrum availability of segments are stable only for a short time (80% of segments are <3 h in length).

B. Spectrum Extraction Rate

For each partially used channel, we measure the spectrum extraction rate as the ratio between the amount of spectrum actually obtained by secondary devices and the amount of available spectrum. By default, the primary-user disruption limit $\eta = 0.1$.

SKA Versus NKA: Fig. 4 plots, for each of the 15 services, the one-week average of the spectrum extraction rate. Without any knowledge on primary-user idle time, NKA's extraction rate is roughly 10% (due to $\eta = 0.1$). SKA, on the other hand, improves the extraction rate by 2–3 times. This demonstrates the benefits of having statistical knowledge of the primary-user access patterns.

A disappointing observation is that even with accurate statistical knowledge on primary-user access patterns, the average extraction rate is only 15%–35%. To further explore this problem, we also plot in Fig. 5 the cumulative distribution of SKA's extraction rate among all the segments of partially occupied channels. Across all locations, the median extraction rate is 19%, and 80% of the segments can produce no more than 37% extraction rate.

The low effectiveness can be attributed to two factors: 1) the spectrum usage patterns are highly random and hard to predict, so without a reliable estimation on channel idle duration, secondary devices are forced to be overly conservative; or 2) the access slot used by secondary devices is too large, forcing them

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

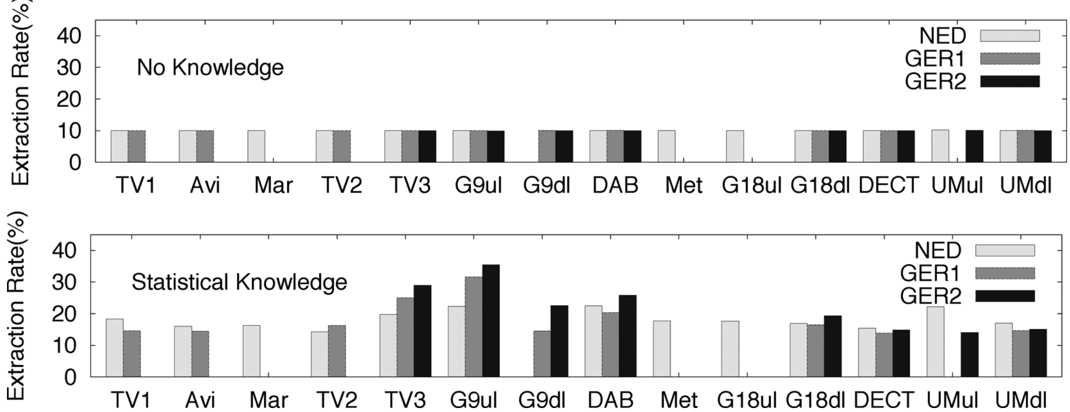

Fig. 4. Spectrum extraction rate with no knowledge (NKA) and statistical knowledge (SKA). The results are averaged over all segments for each service over a week. For GER1 and GER2, some services have no data because they do not have any partially available frequency channels. NKA only extracts 10% of available spectrum due to the 0.1 primary-user disruption limit. SKA increases the extraction rate to 15%–35%.

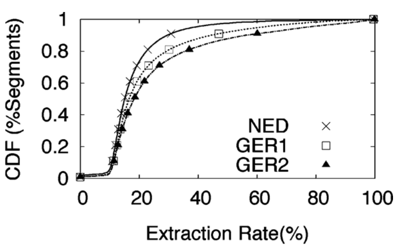

Fig. 5. Cumulative distribution function (CDF) of extraction rate of the SKA scheme for all segments across all services. For all locations, 80% of segments only get up to 37% extraction rate despite accurate statistical knowledge.

to being overly conservative. The first reason has been confirmed by the highly random distribution of primary-user idle time, shown in Fig. 3(b). A related study has also confirmed the difficulty in predicting primary-user access patterns [7]. The second reason, however, is that it is impossible to verify without the ground truth on primary-user spectrum usage patterns—the RWTH dataset is measured at the same 1.8-s intervals, preventing us from pinpointing the exact primary-user arrival and departure time that are required to evaluate the performance of systems using smaller slot sizes. In [16], we show that because original owners display highly random access patterns, reducing slot sizes helps but does not eliminate the need for conservative spectrum access. Thus, the problem of low extraction rate still remains.

C. Available Versus Extracted Spectrum

Our second question is whether the average spectrum availability is a reliable predictor of the amount of spectrum extracted. Answering this question is particularly important because many existing studies have been using the average spectrum availability to evaluate opportunistic access. Using the RWTH dataset, we reevaluate this claim by examining the relationship between the amount of spectrum extracted and the amount of spectrum available.

We first plot the extraction rate as a function of the average spectrum availability. Using the segments discussed in Section IV-B, Fig. 6(a) and (b) shows the spectrum extraction rate for all the GSM1800DL segments at NED, as a function of the average spectrum availability of each segment. As expected, NKA extracts about 10% available spectrum due to the 0.1 primary-user disruption limit. The results display some small variations, especially at low availability values. This is because some segments have fewer idle periods where the performance of a random access scheme like NKA does not converge to its expected value of 10%. Nevertheless, the extraction rate remains stable for all the availability values.

SKA's extraction rate, however, shows significant variance, especially at high spectrum availability regions. This is not triggered by the lack of idle instances, but the large variations in the distribution of primary-user idle time. While many segments display similar average availability, their primary-user idle time distributions and access strategies are significantly different, leading to notably large difference in their extraction rates. Overall, we observe a weak relationship between the extraction rate and the average availability.

Next, we compare the amount of spectrum extracted to the amount of spectrum available. Intuitively, a channel with larger availability will produce more usable spectrum using opportunistic access, which has been widely used to evaluate opportunistic access [4], [21]. Our results in Fig. 6(c) show that such a claim can be problematic. Again we observe significant variance in terms of the actual amount of spectrum extracted, particularly at high availability values. For example, for GSM1800DL at NED, the uncertainty (standard deviation/mean) of using the availability to predict the extracted spectrum is 36%. Therefore, an important conclusion from our analysis is that spectrum availability is no longer a sole metric to evaluate opportunistic spectrum access. One must also examine the access strategy as well as the primary-user idle time distribution when comparing two frequency channels.

D. Usability of Extracted Spectrum

We also wish to understand the feasibility of using extracted spectrum channels to serve traditional wireless applications. To do so, we examine the statistical patterns of the channel service and blocking time experienced by secondary devices. For each frequency channel, the service time defines the time a secondary user can continuously access the channel while the blocking

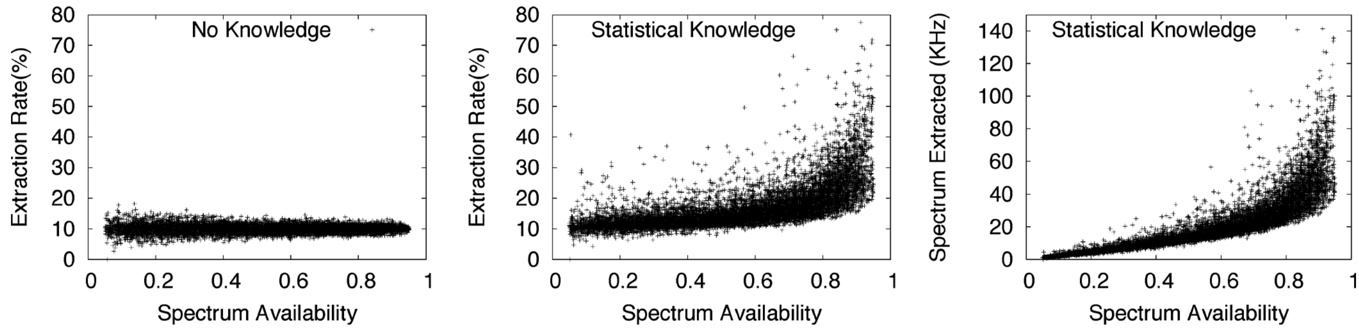

(a) (b) (c)

Fig. 6. Scatter plots of spectrum extraction rates of GSM1800DL at NED with 0.1 primary-user disruption limit. (a) NKA leads to roughly 10% extraction rate. (b) SKA becomes more effective when the spectrum availability increases, although there is significant variance at higher availability values. (c) However, the spectrum availability is no longer an accurate indicator of the spectrum extracted due to the large variance at high availability values.

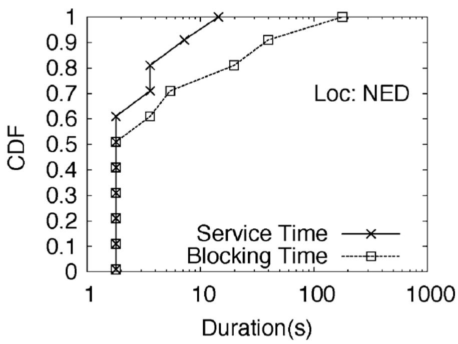

Fig. 7. Cumulative distributions of secondary user's blocking and service times, measured at NED location, using a 2-h segment of the channels used in Fig. 3(b), using SKA and 0.1 primary-user disruption limit.

time defines the amount of time a secondary user must wait before accessing the channel.

Fig. 7 shows the cumulative distribution of both metrics using the same set of channels in Fig. 3(b) and the SKA scheme. Comparing this result to that of Fig. 3(b) (the raw idle and busy duration of the channel), we see that the service time is one order of magnitude smaller than the primary-user idle time, while the blocking time is one order of magnitude larger than the primary-user busy time. While disappointing, this result is somewhat expected, given that the extraction rate of SKA is $<30\%$.

The absolute values are not promising. Secondary users experience prolonged blocking (2–200 s) and short service time (2–10 s). This means that secondary users have a very limited window for transmissions and face frequent interruptions. This type of access is unable to serve many of today's applications.

V. Frequency Bundling

The results in Section IV demonstrate that despite the abundant availability of partially used spectrum, the amount of spectrum actually accessible is much smaller than expected. More importantly, the extracted spectrum is heavily fragmented and scattered across time. Thus, the equivalent channels available to secondary devices are highly unreliable.

In this section, we examine the feasibility of building reliable transmission channels by combining together multiple unreliable frequencies, utilizing frequency diversity to compensate for the lack of reliability on individual channels. We refer to this method as *frequency bundling*.

Frequency bundling is both feasible in practice and attractive to primary and secondary users. Recent advances in radio hardware design make frequency bundling practical for secondary users. New frequency-agile radios can combine noncontiguous frequency channels to form a single transmission [22]. This bundling can be performed either before allocation by a primary user or spectrum regulator, or after allocation by the secondary users themselves. In the second case, care must be taken to avoid bundling contention between secondary users.

Challenges: Frequency bundling faces two key challenges. First, how should secondary users choose and group channels? To reduce blocking time, one should group channels that complement each other in time, i.e., negatively correlated in their spectrum usage patterns. This motivates us to examine the correlation across channels using our measurement dataset. Second, given a bundle of frequency channels, how should we design multichannel secondary access mechanisms that effectively utilize these channels? We address these questions in Section V-A and V-B, respectively, and examine the bundling performance in Section V-C.

A. Correlation Among Frequency Channels

In searching for bundling strategies, we start by examining the correlation among frequency channels in terms of their primary-user spectrum usage patterns. For this task, we again use the RWTH dataset because of its extensive coverage of frequency channels. We divide each channel trace into multiple 1-h segments and compute pairwise correlation among the channels by individual segments. We do not use our segmentation mechanism from Sections III and IV here because it produces variable-length segments among channels that cannot be used to calculate time-domain correlation. We study correlation between channels within the same service as well as across adjacent services, considering that frequency-agile radios are likely to combine channels in close proximity.

We use two metrics to quantify correlation: *Pearson's correlation coefficient* [23] and *mutual information* [24]. Pearson's coefficient ranges from -1 to 1, where -1 indicates strong negative correlation, 1 indicates strong positive correlation, and 0 indicates independency when X and Y are jointly normal [23].

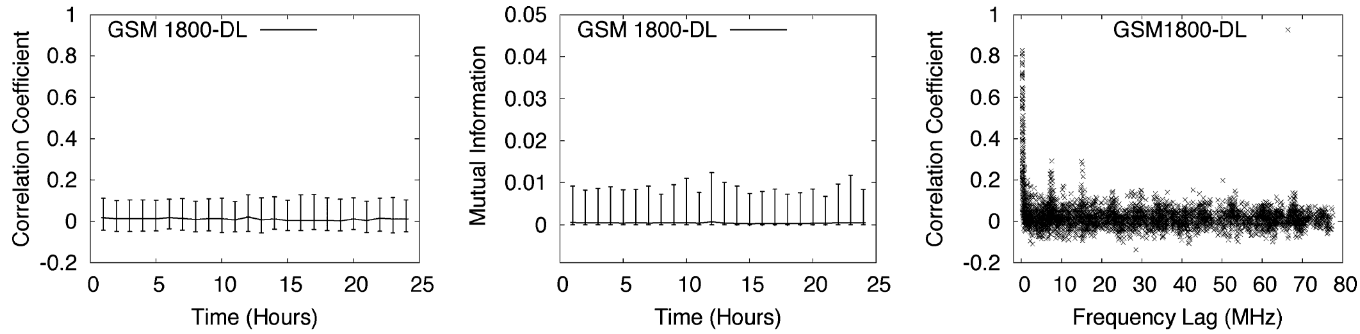

(a) (b) (c)

Fig. 8. (a), (b) Pairwise correlation of GSM1800DL channels at NED across different hours of the day. Both (a) correlation coefficient and (b) mutual information are close to 0. (c) Correlation coefficient as a function of frequency separation. Adjacent channels are highly correlated due to imperfect alignment between measurement and service channels.

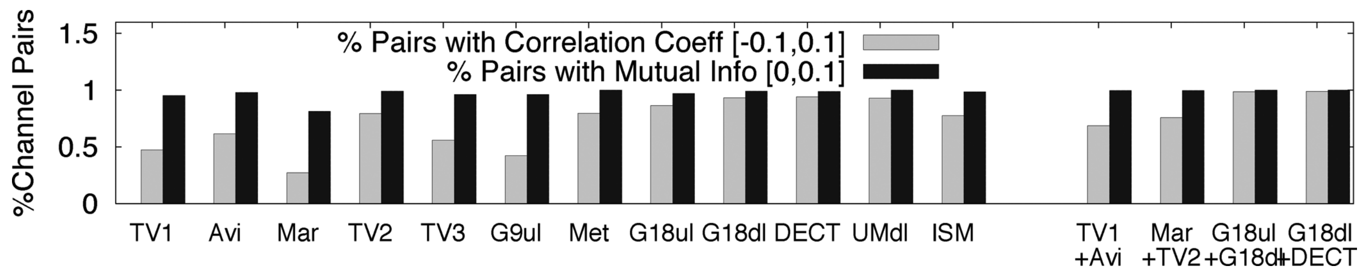

Fig. 9. Percentage of pairs with correlation coefficient between $[-0.1, 0.1]$ and mutual information between $[0, 0.1]$ at NED. A high percentage of channel pairs have very low correlation both within a service and across services.

While capturing both positive and negative correlation, this metric can only detect linear dependency. Mutual information ranges from 0 to 1, where it is 0 if and only if X and Y are independent. Unlike correlation coefficient, this metric detects general dependency.

Results: Our analysis on the RWTH dataset shows that channels display little dependency unless they are adjacent in frequency. As an illustrative example, Fig. 8(a) and (b) plots both correlation metrics over a day using all the GSM1800DL channels at NED. We segment the 24-h duration into 24 1-h segments, and for each hour calculate the pairwise correlation among all the channel pairs. We show our results by the median, 5%, and 95% values of the channel pairs. We see that all these values are close to 0, indicating minimum correlation between channels.

Fig. 8(c) shows a detailed trace of the correlation coefficient as a function of frequency separation. Again it shows that unless the two channels are adjacent to each other, there is no sign of strong correlation. The strong correlation among close pairs (those separated by less than 400 KHz) can be explained by two reasons. First, while the RWTH measurement channels are of the same width as the GSM1800 service channels (200 kHz), they are, however, not perfectly aligned with the GSM1800 service channels. Thus, adjacent measurement channels may map to the same service channel and hence appear heavily correlated. Second, adjacent channels can produce cross-band interference to each other, which makes them inherently correlated. The same was found from our UCSB GSM measurement results.

We have examined other services over different time periods, and the results show very similar trends. To illustrate the general trend across all the services, in Fig. 9 we show the portion of channel pairs with correlation coefficient between [−0.1, 0.1] and mutual information between [0, 0.1]. In addition to considering channel pairs within each service, we also include the result of channel pairs across adjacent services. We see that the majority of channel pairs, either within the same or adjacent services, display very little correlation. The correlation result is service-dependent because each service has different transmission properties and service channel width.

Summary of Findings: Our analysis on pairwise channel correlation leads to two key findings.

- Most of the channel pairs, either within a service or between adjacent services, display little correlation.
- Frequency channel pairs that are adjacent in frequency display relatively high correlation.

These results imply that opportunistic spectrum access across a frequency range will produce multiple channels with little correlation in their available spectrum patterns.

B. Bundling Frequency Channels

The availability independency across channels means that we can significantly improve overall reliability by simply bundling random channel pairs together. In the following, we first describe three candidate methods to access channels in a bundle, and then present our method for forming channel bundles.

Using Frequency Bundles: We propose three usage models, each mapping to a specific radio configuration and application type.

- *Channel Switching* (for simplified hardware): We consider secondary users with Wi-Fi-like radios that can only access a single channel, but can switch between channels on the

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

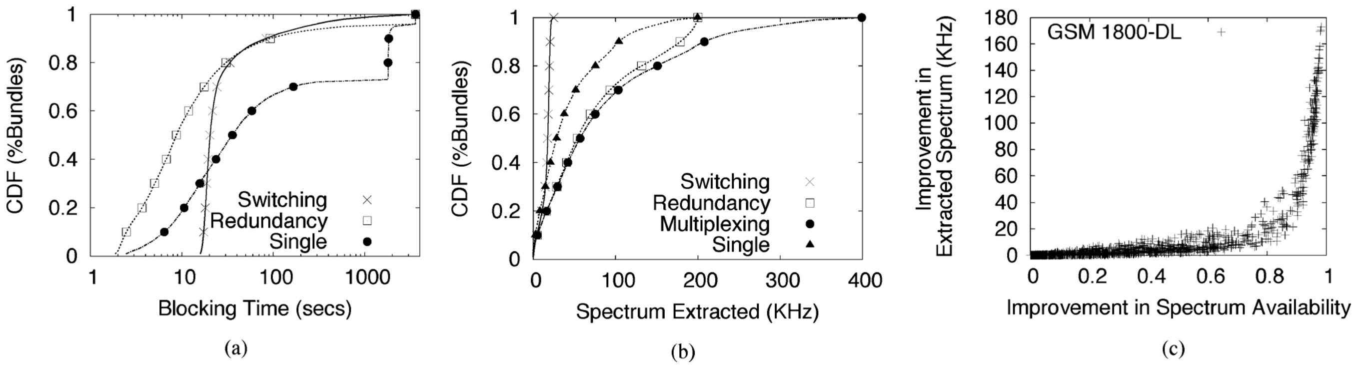

Fig. 10. Performance of 2-channel frequency bundling from all the 15 services at the NED location. Redundancy experiences the lowest blocking time, and Multiplexing enjoys the highest extracted spectrum. Yet for 70+% of bundles, Redundancy has similar extracted spectrum as Multiplexing. This is because of (c) the nonlinearity between the improvement in available spectrum and those in effective spectrum availability. (a) Blocking time. (b) Extracted spectrum. (c) Extracted spectrum versus availability.

fly. In this model, each user switches to another channel in the bundle when the current channel becomes busy or too risky to access. One artifact of this model is that because secondary users cannot monitor each channel continuously, they cannot use SKA, which requires the channel usage history. Instead, they can only use NKA and extract less spectrum.

- *Channel Redundancy* (for maximum reliability): In this model, secondary users can sense and communicate on multiple channels simultaneously. To maximize transmission reliability and minimize blocking time, this model sends the same data stream on all the idle channels in the bundle. When a channel becomes blocked, it skips the data stream. Because secondary users can sense and monitor each channel, they use SKA to access each channel independently. This model focuses on maximizing reliability—unless all the channels are inaccessible, secondary users can communicate continuously.
- *Channel Multiplexing* (for maximum bandwidth): This model also accesses multiple channels simultaneously using individual SKA, but multiplexes the data stream across current idle channels without any redundancy. Different from the Redundancy model, the effective transmission bandwidth varies over time.

Forming Frequency Bundles: We choose a random bundling method. It takes as input, k, the bundle size, and randomly selects k channels from the channel pool to form a bundle. We choose this method because of two reasons. First, the best strategy to minimize blocking time for all three models is to combine channels that complement each other, i.e., negatively correlated. Yet, because the majority of channel pairs show no sign of correlation, random bundling wins due to its simplicity. Second, we use random bundling to understand the performance trend of opportunistic access with different bundle sizes and to evaluate practical situations where secondary users have a small pool of channels for bundling. We only consider partially used channels for bundling since adding idle channels simply increases the bundle capacity by a fixed amount.

C. Bundling Performance

Using the RWTH data set, we evaluate the effectiveness of frequency bundling by combining channels from the same services. We divide a 1-day trace into 1-h segments, randomly bundle channels together, and simulate the three usage models on each segment. As usual, we only consider channels with daily average availability within [0.05, 0.95] and assume a primary-user disruption limit of $\eta = 0.1$.

We evaluate frequency bundling by the resulting channel's blocking time and bundling efficiency. In this case, the blocking time of a frequency bundle is the duration where all the channels are busy or too risky to access. Bundling efficiency is defined as the ratio of extracted spectrum when the channels are bundled together to when the channels are accessed independently. Note that bundling efficiency for Multiplexing is 1. Finally, we primarily included the results of SKA scheme in this section because the performance trends from using NKA scheme were similar to that of SKA.

2-Channel Bundling: Fig. 10(a)–(b) plots the cumulative distribution of the secondary user's blocking time and extracted spectrum using 2-channel bundles. We compare the performance of Single-channel, Switching, Redundancy, and Multiplexing. The performance of Single-channel is the mean of the two channels bundled together. Fig. 10(a) shows that Redundancy has the least blocking time by utilizing every available channel to avoid blockage. On the other hand, Switching experiences 16+ s blocking time. This is because Switching uses NKA due to lack of continuous channel monitoring. With a 0.1 primary-user disruption limit, on average its users will be blocked by 90% of time, or a blocking time of $9 * 1.8 = 16.2$ s. On the other hand, if we extend Switching to monitor each channel continuously, its performance will approach that of Redundancy for the 2-channel case.

Fig. 10(b) examines the actual spectrum extracted from these bundles. As expected, Multiplexing extracts the largest amount of spectrum by avoiding redundancy across channels. Yet surprisingly, Redundancy performs similarly to Multiplexing for 70% of the bundles. This is due to the nonlinear mapping between spectrum available and spectrum extracted [Section IV, Fig. 6(c)]. While Multiplexing improves the effective spectrum availability, its improvement in the spectrum extracted is limited. We confirm this hypothesis in Fig. 10(c), plotting the improvement in extracted spectrum as a function of the improvement in the effective spectrum availability. Even after adding 0.8 (or a raw 160 kHz) to the effective availability, the actual extraction improvement is only 20–30 kHz.

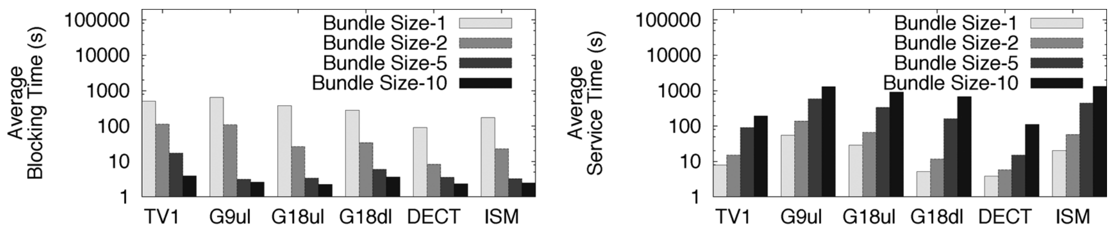

(a) (b)

Fig. 11. Impact of bundle size on (a) average blocking time and (b) service time, using random bundling and the Redundancy model, for six services over a period of 1 day. The improvement in both blocking time and service time scales exponentially with the bundle size k.

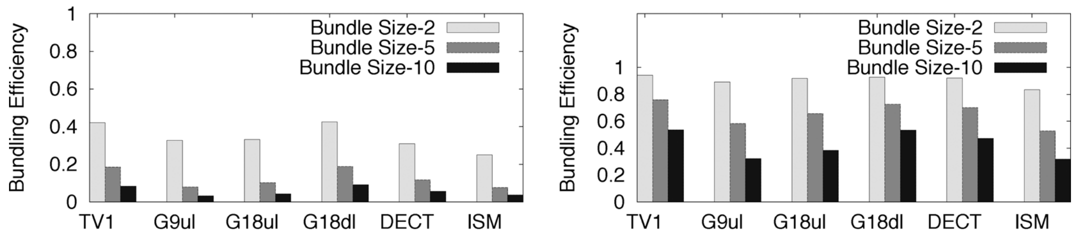

(a) (b)

Fig. 12. Impact of bundle size on average bundling efficiency of (a) Channel Switching and (b) Redundancy models. For both models, bundling efficiency decreases with bundle size because of the increase in the overlap of idle slots among the channels. Channel switching has relatively lower bundling efficiency because it uses a random exponential access scheme.

Impact of Bundle Size: Next, we investigate how the performance of frequency bundling scales with the size of the bundle. Using the same pool of channels, we vary the bundle size k between 2 and 10 and measure the resulting secondary user blocking time, service time, extracted spectrum, and bundling efficiency.

Results for the Redundancy model in Fig. 11 show that bundling can effectively reduce blocking time and increase service time. In fact, a linear increase in the bundling size k leads to one order of magnitude reduction in blocking time and improvement in service time. As k increases beyond 5, the performance quickly converges because additional channels do not offer any new availability. These results clearly demonstrate the effectiveness of frequency bundling.

The absolute values of average blocking and service times look very promising. For the six services shown in this result, bundling $k = 10$ channels randomly creates a pseudo single channel that enjoys on average a prolonged service time of 120–1300 s and occasionally 2–4-s interruptions. These numbers are almost two orders of magnitude better than the single-channel performance.

Fig. 12 plots the bundling efficiency for various bundle sizes for Channel Switching and Redundancy models. For both usage models, bundling efficiency decreases with increase in bundle size. This is because as more channels are added to the bundle, the overlap of idle slots increases. However, overlapping idle slots does not contribute to extracted spectrum in Channel Switching and Redundancy models. Channel Switching has very low bundling efficiency [Fig. 12(a)] because it has to resort to NKA scheme when accessing any given channel, whereas SKA scheme can be used when accessing channels independently. Redundancy model has a high average bundling efficiency of 0.7 [Fig. 12(b)], even with five channels in the bundle. This shows that our simple random bundling strategy works well in practice.

D. *Summary of Findings*

Our analysis in this section leads to two key findings.

- In terms of their spectrum availability patterns, the majority of frequency channel pairs in our dataset (200 kHz in size) display little correlation, unless they are adjacent in frequency.
- Frequency bundling can effectively build reliable and high-performance frequency channels from multiple unreliable channels. Even with random bundling, the improvement in the secondary user's service and blocking time scales exponentially with the bundle size.

VI. Related Work

We classify the related work into spectrum measurement studies and opportunistic spectrum access.

Spectrum Measurements: Several measurement campaigns have studied spectrum occupancy across the globe [3]–[8]. All of them have discovered significant opportunities for opportunistic spectrum access. An extensive measurement on 30-MHz–3-GHz frequency bands at six US locations [3] identified a maximum 13% spectrum occupancy. Measurements on the 2006 Football World Cup at two Germany locations show significant variations in spectrum usage before, during, and after the match. Significant variance was also found on cellular network's spectrum usage, using call logs over three weeks [8]. A recent measurement study at four Chinese locations detects strong dependency across frequency channels and applies a pattern-matching algorithm to predict channel state from past observations [5]. Finally, the Mobnets group from

RWTH performed extensive measurements at three European locations [7].

Our work differs from existing works by examining the actual spectrum accessible to secondary users without violating the primary-user disruption limit. Even with accurate knowledge of primary-user access statistics, we show that the accessible spectrum is significantly less than the available spectrum. We then propose and evaluate frequency bundling that builds high-quality transmissions out of many scattered spectrum fragments. Different from [5], our analysis shows that channels are mostly independent in their spectrum occupancy patterns. These differences might be attributed to two factors: 1) differences in usage at different measurement sites; and 2) inclusion of completely busy and idle channels in [5] during correlation calculation.

Opportunistic Spectrum Access: Research efforts in this area have developed both analytical access strategies and models [9], [25]–[27] as well as practical algorithms and systems [28]–[30]. They have motivated us to consider practical opportunistic access systems and to quantify the actual accessible spectrum. While most of these works either assume analytical models on primary-user access patterns or focus on realizing sensing and accessing in real systems, our work offers a complementary study that uses real-world measurement traces to understand the feasibility and effectiveness of opportunistic spectrum access.

VII. Conclusion

Little is known about how well secondary devices in dynamic spectrum networks can make use of the partially utilized channels occupied by primary users. We present in this paper the first comprehensive study on the level of "usable" spectrum available to secondary devices while respecting hard limits on disruptions to primary users. Our analysis of extensive fine-grain spectrum usage traces shows that even with extensive statistical knowledge on primary-user access patterns, and while running optimal algorithms, secondary devices can only extract 20%–30% of available spectrum in a channel. While this means current access schemes cannot provide usable channels to support traditional applications, we can regain reasonable levels of reliability by bundling multiple unreliable channels together. Our analysis shows very little to no correlation in spectrum usage patterns across channels, which leads us to choose a simple random frequency-bundling scheme. We also show that performing fine-grain extensive spectrum measurement is critical to understanding the performance and limitations of opportunistic spectrum access, and that the granularity of current measurements is not enough to fully capture original owner's spectrum usage variations.

References

[1] M. McHenry, "NSF spectrum occupancy measurements project summary," Shared Spectrum Company, Vienna, VA, Rep., 2005.

[2] K. Doppler, M. Rinne, C. Wijting, C. B. Ribeiro, and K. Hugl, "Device-to-device communication as an underlay to LTE-advanced networks," *IEEE Commun. Mag.*, vol. 47, no. 12, pp. 42–49, Dec. 2009.

[3] Shared Spectrum Company, Vienna, VA, "Spectrum reports," 2004 [Online]. Available: http://sharedspectrum.com/measurements

[4] V. Blaschke, H. Jaekel, T. Renk, C. Kloeck, and F. K. Jondral, "Occupation measurements supporting dynamic spectrum allocation for cognitive radio design," in *Proc. CrownCom*, 2007, pp. 50–57.

[5] D. Chen, S. Yin, Q. Zhang, M. Liu, and S. Li, "Mining spectrum usage data: A large-scale spectrum measurement study," in *Proc. ACM MobiCom*, 2009, pp. 13–24.

[6] O. Holland, P. Cordier, M. Muck, L. Mazet, C. Klock, and T. Renk, "Spectrum power measurements in 2G and 3G cellular phone bands during the 2006 Football World Cup in Germany," in *Proc. IEEE DySPAN*, 2007, pp. 575–578.

[7] M. Wellens, J. Riihijärvi, and P. Mähönen, "Empirical time and frequency domain models of spectrum use," *Phys. Commun.*, vol. 2, pp. 10–32, 2009.

[8] D. Willkomm, S. Machiraju, J. Bolot, and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *Proc. IEEE DySPAN*, 2008, pp. 1–11.

[9] S. Huang, X. Liu, and Z. Ding, "Optimal transmission strategies for dynamic spectrum access in cognitive radio networks," *IEEE Trans. Mobile Comput.*, vol. 8, no. 12, pp. 1636–1648, Dec. 2009.

[10] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty, "NeXt generation/dynamic spectrum access/cognitive radio wireless networks: A survey," *Comput. Netw.*, vol. 50, pp. 2127–2159, 2006.

[11] Q. Zhao and A. Swami, "A survey of dynamic spectrum access: Signal processing and networking perspectives," in *Proc. IEEE ICASSP*, 2007, vol. 4, pp. IV-1349–IV-1352.

[12] F. F. Digham, M. S. Alouini, and M. K. Simon, "On the energy detection of unknown signals over fading channels," in *Proc. IEEE ICC*, 2003, vol. 5, pp. 3575–3579.

[13] RWTH Aachen University, Aachen, Germany, "RWTH Aachen University static spectrum occupancy measurement campaign," 2009 [Online]. Available: http://download.mobnets.rwth-aachen.de

[14] M. Wellens and P. Mähönen, "Lessons learned from an extensive spectrum occupancy measurement campaign and a stochastic duty cycle model," in *Proc. TridentCom*, 2009, pp. 1–9.

[15] S. Shellhammer and G. Chouinard, "Spectrum sensing requirements summary," IEEE 802.22-05/22-06-0089-05-0000, 2006.

[16] V. Kone, L. Yang, X. Yang, B. Y. Zhao, and H. Zheng, "On the feasibility of effective opportunistic spectrum access," in *Proc. ACM IMC*, 2010, pp. 151–164.

[17] L. Cao and H. Zheng, "Distributed spectrum allocation in ad hoc networks via local bargaining," in *Proc. IEEE SECON*, 2005, pp. 475–486.

[18] C. Peng, H. Zheng, and B. Y. Zhao, "Utilization and fairness in spectrum assignment for opportunistic spectrum access," *Mobile Netw. Appl.*, vol. 11, pp. 555–576, 2006.

[19] W. Zhang, R. Mallik, and K. Letaief, "Cooperative spectrum sensing optimization in cognitive radio networks," in *Proc. IEEE ICC*, 2008, pp. 3411–3415.

[20] E. Keogh, S. Chu, D. Hart, and M. Pazzani, "Segmenting time series: A survey and novel approach," *Data Mining Time Series Databases*, pp. 1–21, 1993.

[21] D. Datla, R. Rajbanshi, A. M. Wyglinski, and G. J. Minden, "Parametric adaptive spectrum sensing framework for dynamic spectrum access networks," in *Proc. IEEE DySPAN*, 2007, pp. 482–485.

[22] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng, "Supporting demanding wireless applications with frequency-agile radios," in *Proc. NSDI*, 2010, p. 5.

[23] K. Pearson, "Contributions to the mathematical theory of evolution, II: Skew variation in homogeneous material," *Philos. Trans. Roy. Soc. London*, vol. 186, pp. 343–414, 1895.

[24] R. Fano, *Transmission of Information: A Statistical Theory of Communication*. Cambridge, MA: MIT Press, 1961.

[25] S. Srinivasa and S. Jafar, "The throughput potential of cognitive radio: A theoretical perspective," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, 2006, pp. 221–225.

[26] A. Sahai, N. Hoven, and R. Tandra, "Some fundamental limits on cognitive radio," in *Proc. Allerton Conf.*, 2004, pp. 1662–1671.

[27] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: A POMDP framework," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 589–600, Apr. 2007.

[28] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with Wi-Fi like connectivity," in *Proc. ACM SIGCOMM*, 2009, pp. 27–38.

[29] S. Deb, V. Srinivasan, and R. Maheshwari, "Dynamic spectrum access in DTV whitespaces: Design rules, architecture and algorithms," in *Proc. ACM MobiCom*, 2009, pp. 1–12.

[30] M. Wellens, A. de Baynast, and P. Mähönen, "Performance of dynamic spectrum access based on spectrum occupancy statistics," *IET Commun.*, vol. 2, pp. 772–782, 2008.

Vinod Kone received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology Guwahati, Guwahati, India, in 2006, and the M.S. and Ph.D. degrees in computer science from the University of California, Santa Barbara, in 2010 and 2011, respectively.

He is currently a Software Engineer with the Runtime Systems Group, Twitter, San Francisco, CA. His research interests are in network measurements, wireless networks, and distributed systems.

Lei Yang received the B.E. and M.S. degrees in electronic engineering from Tsinghua University, Beijing, China, in 2004 and 2006, respectively, and the Ph.D. degree in computer science from the University of California, Santa Barbara, in 2011. His Ph.D. thesis work focused on designing reliable, efficient, and enforceable dynamic spectrum-sharing systems using cognitive radios.

He is a Research Scientist with the Wireless Communications Lab, Intel Labs, Hillsboro, OR. His research interests are in wireless networking and systems, with a current focus on mobile systems, location sensing, and dynamic spectrum networks.

Xue Yang received the B.E. and M.S. degrees from the University of Electronic Science and Technology of China, Chengdu, China, in 1993 and 1996, respectively, and the Ph.D. degree from the University of Illinois at Urbana–Champaign, in 2005, all in electrical engineering.

She is currently a Senior Staff Research Scientist and echnical lead of several projects with Intel Labs, Hillsboro, OR, including multiradio coexistence, radio virtualization, and ubiquitous precision location. She is the author of over 40 journal/conference papers and has more than 60 pending or granted patents. Her current research interests are in the areas of wireless networking, mixed networks, mobile virtualization, and positioning.

Ben Y. Zhao received the B.S. degree from Yale University in 1997, and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 2000 and 2004, respectively, all in computer science.

He is an Associate Professor of computer science with the University of California, Santa Barbara. His research spans the areas of networking, security and privacy, distributed systems, simulation, and modeling. His recent interests are in online social networks, cognitive radios, and wireless data centers.

Dr. Zhao is a recipient of the National Science Foundation's CAREER Award, MIT *Tech Review*'s TR-35 Award, and ComputerWorld's Top 40 Technology Innovators.

Haitao Zheng received the B.S. degree in electrical engineering (with highest honor) from Xian Jiaotong University, Xi'an, China, in 1995, and the M.S.E.E. and Ph.D. degree in electrical and computer engineering from the University of Maryland, College Park, in 1998 and 1999, respectively.

She is currently an Associate Professor of computer science with the University of California, Santa Barbara. Previously, she worked with Bell Labs in New Jersey and Microsoft Research Asia, Beijing, China. Her recent research interests include wireless data center networking, dynamic spectrum access and cognitive radios, and social networks.

Dr. Zheng is a recipient of MIT *Tech Review*'s TR-35 Award, the Bell Labs President's Gold Award, and the George Harhalakis Outstanding Graduate Student Award.