

Whitespace Networks Relying on Dynamic Control Channels

Jacek Kibiłda*, Justin Tallon*, Keith Nolan* and Luiz A. DaSilva*[†]

*CTVR, Trinity College Dublin, Ireland

[†]Virginia Tech, USA

Abstract—The emergence of whitespace networks, and whitespace communications in general, provides an opportunity to, at least partially, meet the ever-growing demand for mobile data communication. Most of the whitespace network solutions proposed so far realize coordination and rendezvous over licensed or unlicensed spectrum. In this paper we propose a protocol for networks that rely solely on whitespace spectrum. The proposed protocol allows both communication to the broader network (via the access point) and direct device-to-device links over whitespaces. To showcase the capabilities of the proposed solution we have implemented a proof-of-concept software defined radio experiment. Using the experimental platform we have evaluated the overheads of whitespace operation, which come in the form of an extra delay in association and a throughput loss of around 15% of that achievable with licensed spectrum. Our goal is to provide the groundwork for future radio systems that will operate in whitespace spectrum.

I. INTRODUCTION

Studies show that 100% and 58.6% of Internet traffic generated by smartphones and PCs, respectively, is carried over wireless interfaces. 69% and 57% accounts for WiFi, which operates in unlicensed spectrum, and 31% and 1.6%, respectively, accounts for cellular interfaces operating in licensed spectrum [1]. These numbers show that licence-exempt (or unlicensed) spectrum already plays a vital role in meeting the capacity challenge related to the mobile data crunch. The amount of transmitted mobile data will continue to grow, at an estimated compound annual growth rate (CAGR) of 78% from 2011 to 2016 [2]. To support this demand we need even more pervasive WiFi deployments, which are, however, limited by interference stemming from unlicensed operation. Another possible solution to meet this demand is to increase the cellular network's density, which comes in the form of small cells (e.g. femtocells) that operate in a licensed spectrum underlay to macrocells. However, there is an alternative at hand: whitespace spectrum and cognitive radio technologies.

Whitespaces are defined by the Internet Engineering Task Force (IETF) as portions of the frequency spectrum that are assigned to a particular use but are unoccupied at specific locations and times [3]. This definition implies the existence of incumbent services, which have prioritized access to the spectrum and whose signals should be protected from harmful interference stemming from other whitespace-operating services. An example of whitespaces are TV whitespaces (TVWS), which are portions of the frequency spectrum made available after the digital TV switchover in the UHF/VHF spectrum in certain geographical locations. To protect incumbent services of the UHF/VHF spectrum, such as Digital Video Broadcasting

- Terrestrial (DVB-T), communications regulators in the US and Europe selected a geo-location database technique as the most feasible, and, thus, the only mandatory solution [4], [5]. Hence, devices that desire to operate in the TVWS will have to interact with geo-location databases to obtain complete information about spectrum availability.

One of the objectives of TVWS regulation in Europe was to allow high efficiency and flexibility in spectrum usage with the widest possible ranges of uses and technologies [4]. Cognitive radios are ideally crafted for this purpose, as they are wireless communication systems aware of their environment, which learn from this environment and adapt to any statistical variations in it, to achieve, for example, higher reliability or spectral efficiency [6]. There is a number of scenarios envisaged for cognitive radios operating in the TVWS, for example, indoor/outdoor local area networks or ad-hoc (direct) communication between portable devices [4]. Realization of these scenarios will require a certain level of control and coordination between cognitive radio devices; in other words, the formation of a network. In [7] networks over TVWS are formed based on an enriched WiFi protocol and spectrum availability information determined based on local spectrum sensing. The latter, however, does not conform to the subsequent decisions made by the regulators to mandate geo-location databases as a means for protection of incumbent services. A more conservative approach to formation of networks operating over TVWS is to rely on out-of-band control messages using existing radio access technologies in licensed or unlicensed spectrum, e.g. [8], [9]. Drawbacks of this approach include the need for additional channels in some licensed band, or reliance on the congested ISM band. Having in mind these problems and the recent decisions of the major communications regulators, our goal is to design a network that relies solely on whitespace spectrum.

The purpose of this paper is to design, develop and evaluate a spontaneously created *whitespace network*, i.e. a network which relies solely on whitespace spectrum. In Figure 1 we depict an example instance of a whitespace network where control channels are deployed dynamically whenever and wherever possible to enable coordination and rendezvous between devices operating in whitespaces. Some of these devices, which have the capability to directly query the geo-location databases, may self-select to become whitespace access points, to arbitrate and control whitespace communications of other devices (subordinate devices). The subordinate devices, which could be, for example, sensors that belong to a home automation system, would typically have no means of communication

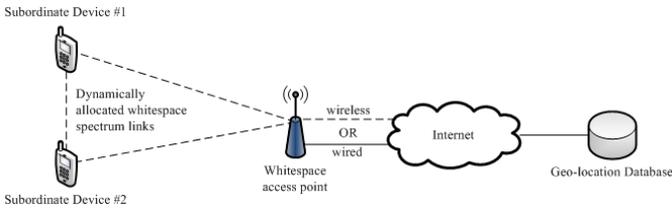


Fig. 1: An example whitespace network.

with the geo-location databases. Moreover, these subordinate devices would use whitespace spectrum intermittently to connect to the Internet (via the access point’s backhaul), or to perform direct device-to-device communications.

In this paper we propose a protocol that enables operation of whitespace networks with the use of geo-location databases, dynamic control channels deployed depending on the whitespace spectrum availability, cyclostationary signatures used for control channel identification, and performance monitoring to improve the whitespace allocation. The proposed protocol allows both communication to the broader network (via the access point) and direct device-to-device links. As part of our work, we have implemented a proof-of-concept software defined radio experiment that showcases the capabilities of the proposed solution. Using the experimental platform we have evaluated the tradeoffs related to operating exclusively in whitespaces, without relying on licensed spectrum for control channels. These tradeoffs come in the form of an extra delay in the order of hundreds of milliseconds and a throughput of up to 85% of that achievable with licensed spectrum links.

The rest of the paper is structured as follows. In Section II, we describe some core concepts important to the proposed solution. Then, in Section III we present the proposed protocol, which we have implemented as part of our proof-of-concept software defined radio experiment, described in Section IV. Then, in Section V we analyse the performance of the proposed solution based on the results collected from our experiment. Finally, we summarize our main conclusions in Section VI.

II. KEY ENABLERS OF DYNAMICALLY CREATED WHITESPACE NETWORKS

In order to build a network that solely operates in whitespace spectrum one needs to overcome a challenge related to the protection of incumbent services and to ensure coordination and rendezvous among the whitespace devices. In our work we overcome these challenges by relying on: geo-location databases, dynamic control channels, and cyclostationary signatures. In the following we give a brief introduction to each of the above mentioned concepts.

A. Geo-location databases

In principle, a geo-location database is a database that contains up-to-date information on the spectrum available at any given location and time instance, enriched with other types of related information, such as the maximum transmitted power or adjacent channel leakage ratio [3]. Geo-location databases are populated with information created by modelling the propagation of known incumbent transmitters (for example

as in [10]), where the model’s parameters and algorithms are selected by the authority operating the database. Such whitespace information is provided to the devices on a temporal basis, and whitespace devices need to periodically request the information, where the period is set according to the requirements of the local regulator. Whitespace devices are not allowed to transmit until they have successfully received up-to-date information on the available channels. When a device has no possibility to directly (without the use of whitespaces) connect to the database, another whitespace device may act as a proxy for the device’s queries [3]. In recent years, communication regulators world-wide have mandated geo-location databases as the only required solution to protect the incumbent services in the TV whitespaces, e.g. [4], [5]. Hence, in our work we rely solely on geo-location databases to protect incumbent services and to provide information on the whitespace spectrum opportunities.

B. Dynamic control channels

In general, control channels are deployed to organize mobile devices and convey network control information, for example, identification, synchronization, channel allocations (restrictions) or network policies. In order to facilitate the distribution of control channels for cognitive radios the European Telecommunications Standards Institute Reconfigurable Radio Systems Technical Committee (ETSI RRS TC) has recommended two ways forward: 1) out-of-band, where the control channels are distributed over a globally dedicated physical channel, 2) in-band, where the control channels are transported over a specific radio access technology using a separate or an existing control channel. The former has the disadvantage of requiring additional spectral resources and global harmonization. The latter is a viable solution for systems operating in licensed bands with fixed operational frequency and high level of coordination, which use whitespaces only temporally to extend network capacity. However, for systems that intend to rely solely on whitespaces it poses some difficulties, as the allocated operational frequency may change depending on the incumbent user behaviour. Herein, we propose a reliable solution for an in-band control channel in whitespaces, which is dynamically deployed depending on the whitespace availability. The centre frequency of this control channel is allocated based on the geo-location database information. Subordinate whitespace devices will acquire this centre frequency through the detection of a physical layer signature inserted in the transmitted waveform of the control channel, as described in the following subsection. Coordination of control channel allocations among whitespace-operating devices is achieved through the system of intermittent channel reservations added to the geo-location database information.

C. Cyclostationary signatures

Communication signals of contemporary radio systems have many inherent periodicities which come as a consequence of coupling stationary signals with, for example, periodical waveforms or training sequences. These periodicities may also arise as a consequence of typical communications procedures, such as sampling or multiplexing. One way to observe them is to perform first order and second order cyclostationary analysis to discover specific correlation patterns in time or

in the spectral domain of the signal, respectively. However, these periodicities may also be intentionally embedded into the physical signal as so called cyclostationary signatures [11]. A cyclostationary signature can be inserted in an OFDM signal $x(t)$ by mirroring over the center frequency a selected subcarrier subset [11], [12]. The arising periodicities can be observed through the spectral correlation function (SCF), at a cyclic frequency that corresponds to the ratio of the spectral distance between the mirrored subcarrier set and the useful symbol duration [11]:

$$S_x^\alpha(f) = \lim_{\Delta f \rightarrow \infty} \lim_{\Delta t \rightarrow \infty} \frac{1}{\Delta t} \int_{-\frac{\Delta t}{2}}^{\frac{\Delta t}{2}} \Delta f X_{\frac{1}{\Delta f}}(t, f + \frac{\alpha}{2}) \cdot X_{\frac{1}{\Delta f}}^*(t, f - \frac{\alpha}{2}) dt$$

where α is the cyclic frequency and Δt is the measurement interval, and

$$X_{\frac{1}{\Delta f}}(t, \nu) = \int_{t-\frac{1}{2\Delta f}}^{t+\frac{1}{2\Delta f}} x(u) \exp^{-i2\pi\nu u} du$$

is the complex envelope of the narrow-band-pass component of $x(t)$ with centre frequency ν and bandwidth Δf .

In a practical system the cyclostationary signature can be detected by sweeping across the bands of interest and performing circular correlations on the received signal samples. When the signature is present in the received signal, a spike in the SCF is observed and the receiver can start decoding the received signal. In case the signature can no longer be detected, the receiver will start sweeping the band until the signature is found again and a new centre frequency is determined. The cyclostationary signatures can be used to identify specific radio systems, specific access networks in coalitions of access networks [13], or to enable rendezvous in dynamic spectrum access networks [12].

III. THE PROPOSED DYNAMIC SPECTRUM ACCESS AND ALLOCATION PROTOCOL

In order to meet the challenges discussed above, with the use of the presented concepts, we have designed a dynamic spectrum access and allocation protocol (DSAAP) that allows for coordinated operation of a dynamic spectrum access network deployed in whitespaces for OFDM-based systems. In general, the DSAAP operations are performed as follows (the subsequent steps are also depicted in Figure 2): 1) When a new whitespace device that is able to connect to the Internet is switched on, it checks with the geo-location database for any available frequency channels. 2) If there are no other whitespace devices operating in the coverage region of the device and a channel is found, the device reserves this particular channel for secondary spectrum operation and becomes a whitespace access point. In the case that a whitespace access point is already operating in the coverage region, the device has to make a decision on whether or not to take the role of a whitespace access point. In this work we assume that every device that is able to connect to the Internet will become a whitespace access point, and we leave other decision schemes for future studies. A whitespace access point periodically transmits a broadcast signal that announces the availability of the whitespace access point in the specific frequency channel to

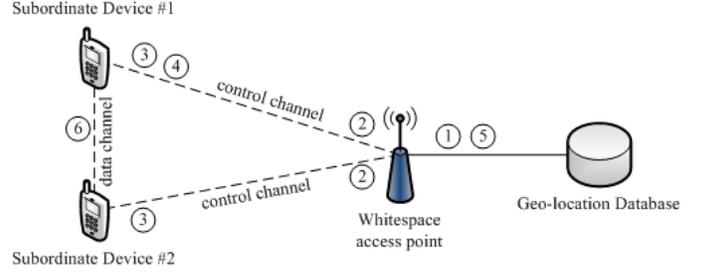


Fig. 2: The steps of DSAAP operation in a whitespace network.

any other whitespace devices. The transmitted broadcast signal has an embedded unique cyclostationary signature, which can be assigned as in [13] and detected with a cyclostationary feature detector described earlier. The broadcast signal carries information required to coordinate the cell's operation, such as the rendezvous channel or temporal spectrum allocations for whitespace devices. 3) When another whitespace device with no internet connectivity arrives in the coverage region of the whitespace access point, it sweeps the whitespace bands to detect the broadcast signal. If the broadcast signal is detected, the device decodes it and reads the cell's information. Then, using the rendezvous channel it associates with the access point and stays on the detected channel listening to the broadcast signal, becoming a subordinate device. If another whitespace device arrives, a similar procedure follows. 4) Whenever one of the subordinate devices requires transmission to another local device (or to the Internet), it requests (using the rendezvous channel) whitespace operation. 5) The access point queries the database and allocates a whitespace channel that meets the demands of the requested transmission, indicating to both the whitespace devices the centre frequency, assigned bandwidth, spectrum availability determination period, and the peer device's MAC address for direct device-to-device transmission. 6) The information is embedded to the control channel and both devices receiving the information reconfigure their radio front-ends to operate on the specific centre frequency and start the data transmission.

During the data transmission, the subordinate devices constantly monitor the connection quality. If the connection quality is sufficient and the spectrum availability determination time elapses, both devices leave the transmission channel and repeat the whole procedure. However, if during the transmission one of the devices observes a significant drop in the connection quality (by means of, for example, an increase in the frame error rate), then that device reconfigures to the rendezvous channel and sends a report to the access point. The access point will use the measurement information conveyed in the report to improve any subsequent data channel allocations.

IV. PRACTICAL REALIZATION

The proof-of-concept implementation of our proposed solution for whitespace networks leverages the open-source software defined radio platform Iris (formerly Implementing Radio In Software) developed at Trinity College Dublin [14]. Iris is a software architecture that has been implemented in C++ for General Purpose Processors (GPP) to provide a framework for building reconfigurable software radios. These radios are

constructed from a combination of a number of virtualized functional blocks for signal operations, in Iris nomenclature called *components*, and a number of *controllers*, which manage reconfigurability of the radio. Iris radios can be reconfigured at run-time in a number of ways. First, selected component parameters can be exposed to the controllers, which dynamically change the values of the parameters whenever certain predefined conditions are fulfilled. For example, in an OFDM modulator component a cyclic prefix length may be exposed to enable multi-path channel delay adaptation. Another way in which a reconfiguration can occur is through dataflow switching. A switch component is used to enable branching of signal samples to different components. This approach may be used, for example, in a cyclostationary signature detector, where one receiver branch is responsible for signal demodulation, and the other for the cyclostationary analysis. These reconfiguration mechanisms are supported with implementations of trigger and event handling mechanisms that allow for synchronous and asynchronous calls for reconfiguration.

Typically, we use Iris together with an off-the-shelf highly flexible radio frequency front-end known as the Universal Software Radio Peripheral (USRP), which depending on the selected daughterboard can operate in transmitter/receiver/transceiver mode on a multitude of radio frequencies (see www.ettus.com). In such a configuration, baseband in-phase and quadrature (IQ) signal samples are passed between the USRP and the host PC, where Iris is running as an application, using either Ethernet or USB connection. An example radio constructed for the purpose of the experiment is depicted in Figure 3. The radio consists of a single USRP device connected over Ethernet to a host PC where Iris is running. The constructed Iris radio is divided into two signal paths: receiver and transmitter, which are responsible for interactions with the USRP driver on the host PC and baseband modulation/demodulation processes. Demodulated OFDM signal samples are passed to the cyclostationary signature detector and to the Whitespace MAC (WS-MAC) layer. The former informs the Dynamic Spectrum Access (DSA) Controller about the detection of a signature, and the latter reads the MAC header of the received data to check whether the received packet data unit (PDU) belongs to the control or the data plane. All data packets are passed to an application, which translates them to user readable format, for example a text file or a video stream. All the components of the radio are monitored by the Dynamic Spectrum Access Controller, which collects reconfiguration events (e.g. signature detected) and issues reconfiguration commands to specific components, e.g. to Read/Write USRP - change the carrier frequency to N GHz. For the purpose of the experiment we have implemented two new Iris modules: the WS-MAC protocol and the DSA Controller.

Using Iris radios, implemented as in Figure 3, we have built a scenario that consists of three radio nodes that realize DSA operation. Figure 4 shows our demonstration scenario with the three transceiver nodes (a Whitespace Access Point (WSAP), a Whitespace Transmitter (WSTx) and a Whitespace Receiver (WSRx)) and the Whitespace Server (WSS). Each of the transceiver nodes includes the Iris physical layer, which consists of a number of signal processing components as described previously, and a role-dedicated WS-MAC component. In principle, the WS-MAC component handles frame compo-

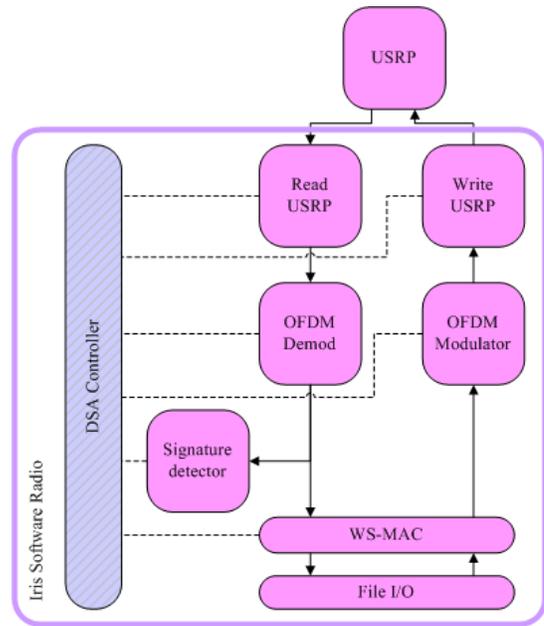


Fig. 3: Block diagram of an IRIS radio model implemented for the purpose of the experiment (solid lines denote signal paths, while dashed lines denote paths for reconfiguration events and triggers).

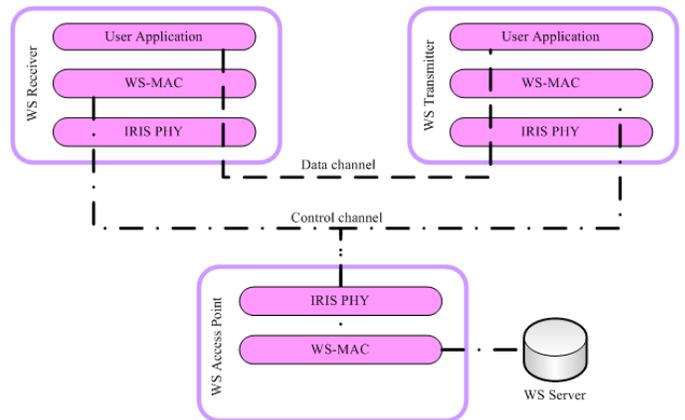


Fig. 4: Demonstration scenario.

sition/decomposition, schedules frame transmission times and handles device addressing. However, depending on the device's role each MAC component utilizes additional functions. For example, the WSAP-MAC component manages signature injection, which happens in the OFDM modulator, control frame generation, dynamic spectrum allocation, and device-to-device link setup. The dynamic spectrum allocation, performed by the WSAP-MAC, is based on the information retrieved from the WSS, which in our case is implemented as a database service on a host PC. The whitespace device MAC (WSD-MAC) component handles decoding of the control channel, dynamic spectrum access requesting, device-to-device association, and link quality performance monitoring. Eventually, both the WSTx and the WSRx implement a user application, for example a file transfer client or video streamer.

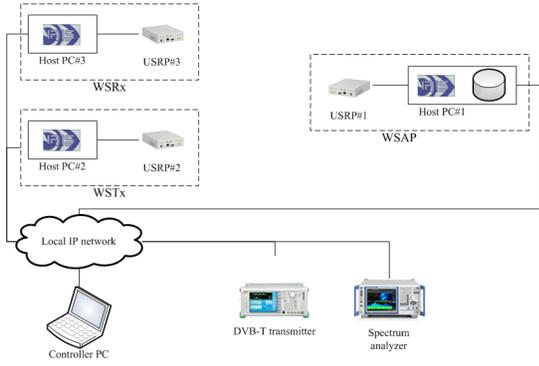


Fig. 5: Experimental setup.

V. PERFORMANCE ANALYSIS

In order to examine the performance of our system we set up an experiment, as shown in Figure 5. The experimental scenario consists of three whitespace devices, each of which consists of a USRP radio front-end, a host PC running Iris and the geo-location database (only WSAP), a vector signal generator (Anritsu MG3700A) transmitting a DVB-T signal, and a spectrum analyzer (Rohde & Schwarz FVSR) to observe the spectrum occupancy during the experiment. All the devices were connected to the local wired network and remotely controlled from a Controller PC. As radio front-ends for the whitespace devices we have used USRP N210 boards equipped with XCVR2450 dual-band (2.4-2.5 and 4.9-5.9 GHz) transmitter/receiver daughterboards. To reduce the impact of potential outside interference the experiment was carried out in a lab environment in the unlicensed part of the 5 GHz band. The devices were separated by the distance of approximately one meter, which allowed to establish (using USRPs at the given frequency range) reliable wireless communication links. The operational frequency spectrum was arbitrarily fixed at 20 MHz and divided into 20 channels (each 1 MHz wide). The control channel was randomly allocated to one of the currently available channels, while the data channel was dynamically allocated to the first available channel after the control channel. All the signals were QPSK-modulated samples, mapped onto 512-subcarrier OFDM symbols. In addition, the control signal carried a cyclostationary signature that was 8-subcarriers long with cyclic frequency of $\alpha = 62500$ Hz. The cyclostationary detector was configured to perform signature detection and carrier frequency estimation at the cyclic frequency, using a 272-bin Fourier transform. The data signal transmitted was a binary representation of a text file located on the Host PC#2 of the WSTx. An experimental cycle consisted of 20 (time to rendezvous) or 5 (relative throughput) minutes of continuous whitespace network operation, during which the WSTx was transmitting the data signal whenever a data channel was available. We have averaged our results over the results of fifteen of such experimental cycles.

We used two metrics to evaluate the performance of our system and quantify the overheads of whitespace operation: time to rendezvous (TTR) and relative throughput. We define the TTR as the time required to find and correctly detect the control channel frame for a device that has no a priori control channel information. We define relative throughput

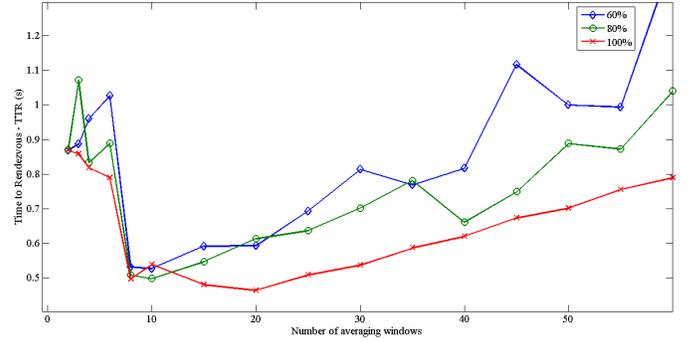


Fig. 6: Control channel discovery performance in practical realization - TTR vs. average number of window size samples for three different settings of the control channel duty cycle.

as the ratio between the achievable throughput of a system operating in whitespaces with our proposed dynamic spectrum access protocol to that of a system operating in non-whitespace spectrum with fixed allocations, which we have measured for the same experimental setup. We determined the TTR for varying averaging window sizes of the SCF and three different settings of the control channel duty cycle, where the duty cycle represents the percentage of time that the control channel is active as a fraction of the total observation time. We determined the relative throughput for three different settings of the spectrum availability determination period: 90, 60 and 30 seconds, where the 60 second spectrum availability determination period is mandated for mode II personal/portable devices by the US regulator [5].

Figure 6 illustrates the change in the TTR with increasing size of the averaging window for three different duty cycles. Based on the results presented in Figure 6 we note that the stable operating point of the cyclostationary feature detector (for any duty cycle of the control channel) is at an averaging window size of 8 samples. At that point the detection reliability and computational complexity are at balance, leading to a minimum TTR of roughly 0.5 seconds. Decreasing the averaging window size provides better computational performance but bears higher probability of false alarm [12], which leads to false detections that effectively increase TTR. Increasing the duty cycle of the control channel slightly reduces the TTR. The performance decrease observed for the operating point was in the range of 7%, when the control channel duty cycle was reduced from 100% to 60%. The benefits of the reduced duty cycle are: the increase in spectrum efficiency (as we can allow time division multiple access operation in the control channel) and reduced energy consumption of the access point node, related to shorter transmission times. Additionally, in our experiment switching an incumbent DVB-T signal on/off had a negligible impact on the TTR, as the presence of the DVB-T signal did not increase the false alarm rate, and, hence, the TTR was dependent mainly on the quality of the received control channel signal.

Figure 7 illustrates the relative system throughput, which quantifies the overheads of dynamic spectrum access utilization in comparison to exclusive spectrum access. The results presented herein show that these overheads account for roughly 19% of the throughput potentially achievable with exclusive spectrum usage in the case of spectrum availability determi-

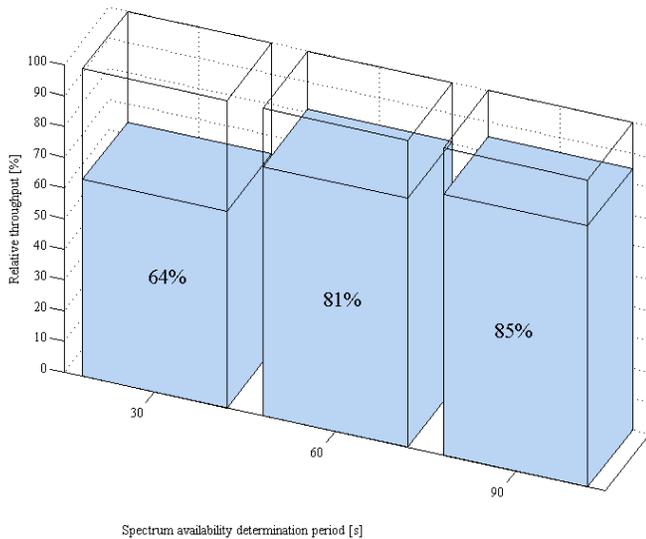


Fig. 7: Relative throughput performance in practical realization of whitespace data offloading.

nation performed every 60 seconds. In the case of spectrum availability determination performed every 90 and 30 seconds, the whitespaces overheads are around 15% and 36%. From the perspective of coexistence between secondary systems, this overhead could be seen as the price to pay for interference-free transmission. Co-channel interference may dramatically degrade throughput; for example, in WiFi systems a 0.01 mW interference signal may cause throughput reduction by a factor of four [15].

VI. CONCLUSION

The article presents a wireless network, where all the control and data communications happen over whitespaces, for example spatio-temporally unused TV channels. We have designed a channel access method and spectrum allocation protocol which enables reliable control channel deployment and efficient data communication over whitespaces. Moreover, our proposed solution enables realization of direct device-to-device links over whitespaces. In this article, we discussed some concepts that underlay our proposed solution and described the proposed protocol along with a practical realization, which we have used to evaluate the performance of whitespace operation in a proof-of-concept experiment involving software defined radios. What we have done with this experiment is to show that operation relying solely on whitespaces is possible between two peer devices without the need for a control channel over licensed spectrum. This result comes with certain tradeoffs, namely the delay in the discovery of the dynamic control channel, which in practical cases may be, as we show, in the range of hundreds of milliseconds, and the decrease in the achievable throughput, which depending on the spectrum availability determination period may vary between 15% and 36%. The goal of our work was to showcase and experiment with cognitive radio networks that rely solely on whitespaces, to build the groundwork for any future radio systems that will operate in whitespace spectrum. The new Iris components and controllers, developed along the way, will serve as the basis

for further practical experimentation with whitespace resource allocation algorithms.

ACKNOWLEDGMENT

The material is based upon works supported by the Science Foundation of Ireland under grants no. 10/CE/I1853 and 10/IN.1/3007, Framework Programme 7 (FP7) CREW project under grant no. ICT-258301, FP7 CogEU project under grant no. ICT-248560. This work was also supported in part by COST Action IC0905 TERRA.

REFERENCES

- [1] R. Thanki, "The Economic Significance of Licence-Exempt Spectrum to the Future of the Internet," Tech. Rep., Jun. 2012.
- [2] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011-2016," 2011. [Online]. Available: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf
- [3] A. Mancuso, S. Probasco, and B. Patil, (draft) *Protocol to Access White Space (PAWS) Database: Use Cases and Requirements*, Internet Engineering Task Force (IETF) Std., Rev. 09, Dec. 2012. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-paws-problem-stmt-usecases-rqmts-09>
- [4] "ECC REPORT 159: Technical and operational requirements for the possible operation of cognitive radio systems in the "white spaces" of the frequency band 470-790 MHz." Electronic Communications Committee (ECC), Tech. Rep., 2011. [Online]. Available: <http://www.erodocdb.dk/Docs/doc98/official/pdf/ECCREP159.PDF>
- [5] "'Unlicensed Operation in the TV Broadcast Bands, ET Docket No. 04-186 and 02-380; FCC 10-174," Federal Register Vol. 75, No. 233," Federal Communications Commission (FCC), Tech. Rep., Dec. 2010.
- [6] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, Sep. 2006. [Online]. Available: <http://dx.doi.org/10.1109/JSAC.2004.839380>
- [7] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White Space Networking with Wi-Fi like Connectivity," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 4, pp. 27–38, Aug. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1594977.1592573>
- [8] G. Gur, S. Bayhan, and F. Alagoz, "Cognitive femtocell networks: an overlay architecture for localized dynamic spectrum access [Dynamic Spectrum Management]," *IEEE Wireless Communications Magazine*, vol. 17, no. 4, pp. 62–70, Aug. 2010.
- [9] S. Sen, T. Zhang, M. M. Buddhikot, S. Banerjee, D. Samardzija, and S. Walker, "Dual Technology Femto Cell Architecture for Robust Communication using Whitespaces," *IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, Oct. 2012.
- [10] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl, "SenseLess: A Database-Driven White Spaces Network," *IEEE Transactions on Mobile Computing*, vol. 11, no. 2, pp. 189–203, Feb. 2012.
- [11] P. Sutton, K. Nolan, and L. Doyle, "Cyclostationary Signatures for Rendezvous in OFDM-Based Dynamic Spectrum Access Networks," in *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN 2007)*, Apr. 2007, pp. 220–231.
- [12] —, "Cyclostationary Signatures in Practical Cognitive Radio Applications," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 13–24, Jan. 2008.
- [13] J. Tallon, T. Forde, and L. Doyle, "Dynamic Spectrum Access Networks: Independent Coalition Formation," *IEEE Vehicular Technology Magazine*, vol. 7, no. 2, pp. 69–76, Jun. 2012.
- [14] P. Sutton, J. Lotze, H. Lahlou, S. Fahmy, K. Nolan, B. Ozgul, T. Rondeau, J. Noguera, and L. Doyle, "Iris: an architecture for cognitive radio networking testbeds," *IEEE Communications Magazine*, vol. 48, no. 9, pp. 114–122, Sep. 2010.

- [15] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of RF interference on 802.11 networks," *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 385–396, Aug. 2007. [Online]. Available: <http://doi.acm.org/10.1145/1282427.1282424>