DEMO: Dynamic Adaptations of WiFi Channel Widths Without TX/RX Coordination

Alice Lo Valvo, Ilenia Tinnirello, Fabrizio Giuliano, Giuseppe Santaromita
University of Palermo, Italy
name.lastname@unipa.it

ABSTRACT

Most modern standards for wireless communications support physical layer adaptations, in terms of dynamic selection of channel central frequency, transmission power, modulation format, etc., in order to increase link robustness under time-varying propagation and interference conditions. In this demo, we demonstrate that another powerful solution for extending physical layer flexibility in OFDM-based technologies is the dynamic adaptation of the channel width. Although some standards already define the possibility of utilizing multiple channel widths (e.g. 20MHz, 10MHz, 5MHz for IEEE 802.11a standards), such an utilization is limited to a static configuration of a value defined during the network set-up. Conversely, we demonstrate that channel width adaptations can be performed in real-time during network operation, even on a per-packet basis. To this purpose, we propose an innovative and efficient receiver design, which allows the transmitter to take decisions about the channel width without explicitly informing the receiver.

1. INTRODUCTION

Spectrum agility has been traditionally considered as the capability to set up a wireless communication link over different spectrum blocks, by shifting from one central frequency to another. This capability is of primary importance for cognitive radio systems or for systems working in ISM bands, which are usually unplanned and characterized by significant spatial variations of spectrum availability. Available spectrum portions can differ in size. Therefore, another desirable capability of spectrum-agile technologies is the possibility of adapting the channel width to the available spectrum bandwidth and/or aggregating independent (non-contiguous) spectrum portions into a single logical link.

In this demo we deal with channel width adaptations for OFDM-based systems. As a reference technology, we consider wireless nodes based on the IEEE 802.11a OFDM PHY and MAC specifications. The PHY of legacy 802.11a nodes includes the possibility of working with 5MHz and 10MHz channel widths, in addition to the usual 20MHz configuration. Indeed, the possibility of adapting the channel width has been demonstrated to be beneficial for different performance figures, such as energy consumption, link reliability, throughput and fairness, [1], by generalizing the concept of *rate adaptation*. Differently from usual rate adaptation, where modulation formats are specified in the PHY fields of each frame, requests for channel width adaptation are transmitted in special control frames which require a confirmation by the intended receiver. A similar adaptation of modulation formats and channel widths is proposed in [2] for OFDM systems. Rather than changing the transmitter clock as in [1], in this work bandwidth adaptations are supported by selecting a set of subcarrier groups for each frame transmission. The subcarrier group ordering is piggybacked by the receiver to the sender in each acknowledgement. Randomized hopping between different channel widths have been proved to increase robustness against jamming attacks of fixed bandwidths [3]. In this case, the hopping sequence used by the transmitter is known to the receiver, which recovers the per-packet channel width by means of a synchronization mechanism with the transmitter hopping sequence. In [4] the benefits of heterogeneous channel widths are achieved by configuring multiple coexisting networks working with different (static) channel widths. To speed-up the scanning of networks working on multiple channel widths, the authors propose to passively perform a temporal analysis of typical channel timings that can be related to the transmission bandwidth employed in each network (such as the duration of acknowledgement transmissions or the DIFS interval).

In our demo we do not focus on the logic for selecting the channel width, but rather on the receiver architecture which enables the possibility of implementing different decision logics at the transmitter side. Although the logic could in principle benefits on context information signaled by the receiver, our architecture allows the transmitter to take decisions about channel widths, without explicitly signaling the decision to the intended receiver. As in [1], channel width adaptations are implemented by changing the clock of the OFDM transmitter. The receiver architecture has been implemented on the well-known WARP [5] research board, which is a FPGA-based SDR platform, for which it is available a reference implementation of legacy 802.11 PHY (including 802.11a/g OFDM modulations). We demonstrate that a spectral analysis of the preamble of each incoming frame can be performed on time for reconfiguring the internal clock of the receiver consistently to the transmitter one.

MobiCom 2017, October 16-20, 2017, Snowbird, Utah, USA.

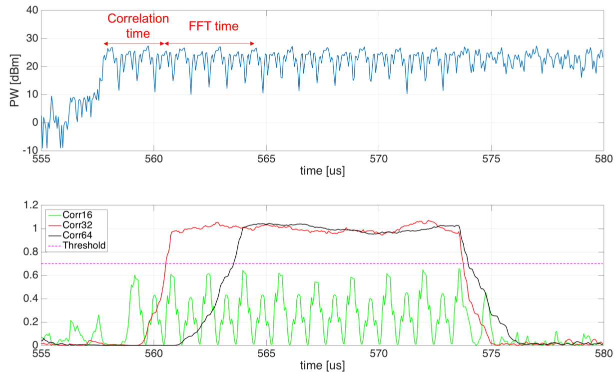

(a)

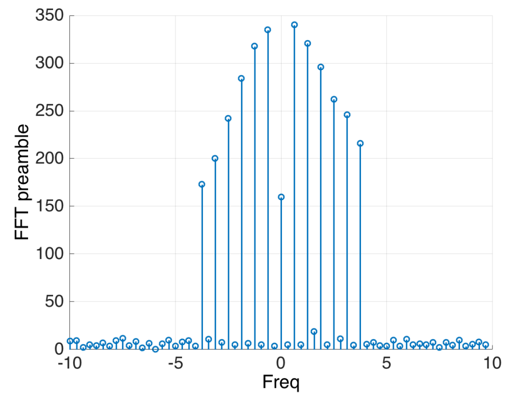

(b)

Figure 1: (a) Preamble time and correlations, (b) FFT shift of a preamble sent at 10MHz

2. TRANSCEIVER ARCHITECTURE

We designed a OFDM transceiver under the assumption that wireless nodes belonging to the same network work with a fixed central frequency, while the channel width used for transmissions can be arbitrarily selected by the senders in each transmission attempt, without explicitly signaling the choice to the receiver.

At the transmitter side, channel width adaptations are implemented by simply scaling the basic clock of the system (80MHz for the WARP board) to a desired output clock by means of a dynamic reconfiguration port. At the receiver side, we considered more general extensions of a typical OFDM receiver, devised to enable the correct identification of the channel width of an incoming frame within the reception of the short OFDM preamble. Specifically, we modified the peak detection block of a OFDM receiver working with fixed channel width, and we added an FFT block for performing a preliminary spectral analysis of the received signal when a preamble is detected.

Peak detection. The legacy 802.11 short preamble has a periodic structure with 10 identical symbols, each one lasting $0.8\mu s$, $1.6\mu s$ or $3.2\mu s$ in case the channel width is set to 20MHz, 10MHz or 5MHz. Usually, OFDM receivers detect the beginning of a short preamble by identifying this periodic structure, i.e. by correlating two windows of signal samples corresponding to two consecutive symbols. Working at a fixed sampling rate of 20Msample/s, a preamble symbol includes 16, 32 or 64 samples according to the channel width used in transmission. Therefore, rather than working with a correlator whose window value is statically configured, we replicated the correlator blocks, in order to perform three parallel correlations on windows of 16, 32 and 64 samples. The channel width of an incoming frame can be estimated by observing the minimum window size which gives an high correlation result. The total detection delay is 2 symbols, because two preamble symbols are enough for detecting the first correlation peak. For deciding if the correlation is positive or not, the correlation result is compared with a percentage of the sum of the modules of the first window W of samples used for correlation (i.e. with the theoretical correlation result in case of perfect periodic samples of the signal, $s(nT+WT) = s(nT)$ for $n = 0, \cdots W-1$ and $T = 1/20MHz$).

Note that a positive detection of a short preamble with a given channel width does not necessarily imply that the incoming frame belongs to the receiver network and is transmitted at the network (static) central frequency. Indeed, by sampling the signal at 20Msamples/s, we can recognize valid preambles sent with a central frequency shifted of +/-10MHz from the receiver one.

Spectral analysis. In case a valid short preamble is detected, a spectral analysis of the sub-sequent preamble symbols is performed by means of a 64-point FFT. The number of preamble symbols required for collecting 64 samples are obviously dependent on the channel width, and in particular are equal to 4 symbols transmitted at 20MHz, 2 symbols transmitted at 10MHz and 1 symbol transmitted at 5 MHz. In the worst case, such an analysis lasts 4 symbols of the preamble, thus leading to a total time of 6 symbols before taking a decision on the channel width. The results of the FFT allows to immediately reject preambles transmitted with a different central frequency. Moreover, they also provide an additional evidence of the channel width used for transmission. Indeed, because of the preamble structure, sub-carriers with non-null coefficient appear in contiguous positions, spaced of two sub-carriers or spaced of four sub-carriers when the channel width is, respectively, 5MHz, 10MHz, or 20MHz. In our implementation, the decision logic works by comparing the sub-carrier amplitude with a threshold calculated by averaging the amplitude of 12, 18 or 24 sub-carriers around the central frequency. We chose to limit the average operation to these sets of sub-carriers, rather than considering the whole set of 64 sub-carriers, in order to discard the power of interfering signals occupying adjacent bands. If six sub-carrier amplitudes are higher than the threshold in contiguous positions or spaced of a null sub-carrier or spaced of three null sub-carriers around the central frequency, the frame is recognized as a valid frame transmitted at 5MHz, 10MHz or 20MHz. This result is used for reconfiguring the clock of the system.

Clock Reset. After completing the reception of the frame, the receiver switches the clock to the basic value (i.e. switch to a reference bandwidth of 20MHz). In case the receiver is the destination of the frame, the clock reset is deferred until the completion of the ACK frame transmission.

Figure 1-a shows the temporal structure of a short preamble (top part of the figure) transmitted at 10MHz, from

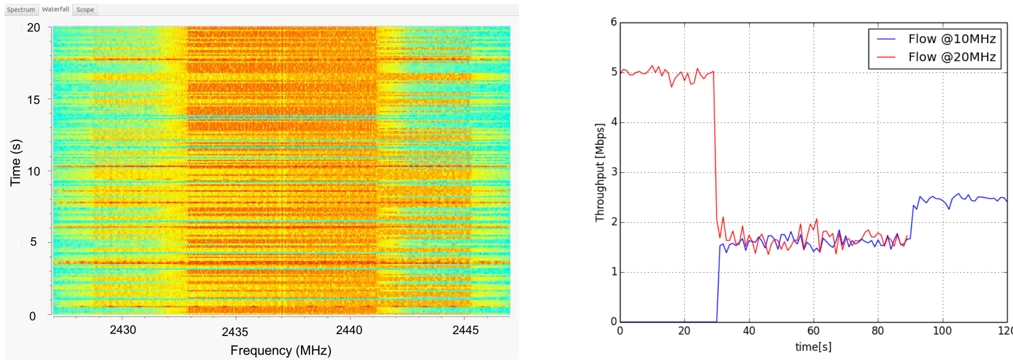

(a) (b)

Figure 2: Agile Receiver Performance: (a) waterfall and (b) throughput of two flows at 10 and 20MHz

which it is evident that the power ramp involves two symbols, and the parallel correlations of windows with 16, 32 and 64 samples (bottom part of the figure). Figure 1-b shows the results of the 64-point FFT. Both the correlation results and the FFT analysis allows to identify that the channel width of the incoming frame is set to 10MHz; moreover, the FFT results indicate the position of the central frequency used in transmission.

3. DEMONSTRATION DESCRIPTION

We consider a simple network with three nodes only, devised to showcase the functionalities of our bandwidth-agnostic receiver: two senders, A and B, able to dynamically tune the channel width used for transmissions, act in proximity of a receiver C. All nodes are implemented on top of the WARP board, for which it is also available a complete implementation of the 802.11 DCF access scheme. A further monitoring node, implemented by using a USRP [6], acts as a spectrum analyzer for showing in real-time the spectrum occupancy of the network in an output monitor (which also displays the throughput results obtained by the receiver). The transmitter A is statically configured for working at the same central frequency f_c of the receiver C, while the transmitter B can be tuned on the same central frequency f_c or on a central frequency in the range $[f_c - 10MHz, f_c + 10Mz]$. On each transmitter, it is possible to configure the rate of the traffic source (implemented with an iperf client transmitting UDP packets of 1470 bytes) and the channel width mode (static, at a given value, or dynamic with random patterns).

In case both the transmitters A and B are configured for working on the frequency f_c, we can demonstrate that the receiver is able to correctly demodulate the frames transmitted without any a-priori knowledge about the used channel width. Indeed, in this scenario the sequence of channel widths used in consecutive frames is unpredictable, due to the fact that transmitters A and B randomly win consecutive contentions. Therefore, even pseudo-random sequences of channel widths adopted by A and B cannot be mapped into a deterministic sequence of channel widths for the receiver. In other words, random contention prevents the adoption of any signaling mechanism based, for example, on the possibility of pre-announcing in each frame the channel width used for the next transmission. Figure 2 shows the results obtained in a real experiment based on this scenario, in which nodes A and B are statically configured for transmitting, respectively, at 20MHz and 10MHz, with a data rate set to 6Mbps. The figure shows the temporal sequence of random spectrum occupancy due to the contention mechanism (Figure 2-a), and the throughput results (2-b) when only node A is active, both nodes A and B are active, and when only node B is active. The receiver is able to simultaneously works with both the transmitters, which achieve a similar throughput when simultaneously active, because of the performance anomaly phenomenon [7].

In case transmitter B is configured for working on a different (opportunistically spaced) f_c frequency, we will demonstrate that the receiver is still able to detect node A transmissions despite of node B interference.

For showing properly this demo, we need a desk and a monitor (at least 24 inches), which displays the waterfall plot and the throughput of two traffic flows.

4. REFERENCES

[1] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. SIGCOMM'08, Seattle, WA, Oct. 2008.

[2] V. Pejovic and E. M. Belding, Whiterate: A context-aware approach to wireless rate adaptation, IEEE Transactions on Mobile Computing, vol. 13, no. 4, April 2014

[3] M. Liechti, V. Lenders, and D. Giustiniano. Jamming mitigation by randomized bandwidth hopping. In Proc. of ACM CoNEXT 2015. ACM, New York, NY, USA-

[4] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. In Proc. of ACM SIGCOMM 2009. ACM, New York, NY, USA, 27-38.

[5] A. Khattab, at al. WARP: A flexible platform for clean-slate wireless medium access protocol design. SIGMOBILE Mobile Computing Communication Review, vol. 12, no. 1, pp. 56-58, 2008.

[6] M. Ettus, "Universal Software Radio Peripheral", [online] Available: http://www.ettus.com.

[7] M. Heusse, et al. Performance Anomaly of 802.11b, INFOCOM 2003.