Coexistence of Dedicated Short Range Communications (DSRC) and Wi-Fi: Implications to Wi-Fi Performance

Gaurang Naik, Jinshan Liu, Jung-Min (Jerry) Park
Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA
{gaurang, jinshan, jungmin}@vt.edu

***Abstract*—The 5.9 GHz band is being actively explored for possible spectrum sharing opportunities between Dedicated Short Range Communications (DSRC) and IEEE 802.11ac networks in order to address the increasing demand for bandwidth-intensive Wi-Fi applications. In this paper, we study the implications of this spectrum sharing to the performance of Wi-Fi systems. Through experiments performed on our testbed, we first investigate band sharing options available for Wi-Fi devices. Using experimental results, we show the need for using conservative Wi-Fi transmission parameters to enable harmonious coexistence between DSRC and Wi-Fi. Moreover, we show that under the current 802.11ac standard, certain channelization options, particularly the high bandwidth ones, cannot be used by Wi-Fi devices without causing interference to the DSRC nodes. Under these constraints, we propose a Real-time Channelization Algorithm (RCA) for Wi-Fi Access Points (APs) operating in the shared spectrum. Evaluation of the proposed algorithm using a prototype implementation on commodity hardware as well as via simulations show that informed channelization decisions can significantly increase Wi-Fi throughput compared to static channelization schemes.**

I. Introduction

In 1999, the Federal Communications Commission (FCC) allocated 75 MHz of spectrum in the 5.9 GHz band for vehicle-to-vehicle (V2V) or vehicle-to-infrastructure (V2I) communications. The proposed vehicular communication technology, known as the Dedicated Short Range Communications (DSRC), is based on the IEEE 802.11p for its physical (PHY) and medium access control (MAC) layers. The spectrum band $5.850 - 5.925$ GHz has been allocated to DSRC applications.

Currently, the FCC is considering a proposal to allow Unlicensed National Information Infrastructure (U-NII) devices, in particular IEEE 802.11ac devices, to operate in the $5350-5470$ MHz and $5850 - 5925$ MHz bands [1]. The latter band is the one reserved for DSRC applications. DSRC users will remain primary users, and the 802.11ac users shall be unlicensed secondary users. The main objective of this plan is to assign more spectrum to 802.11ac in order to accommodate additional wide-bandwidth channels, and in turn, to enhance support for high data-rate applications. If the FCC decides to adopt the proposal, it would create a new set of rules for spectrum access in the 5.9 GHz band that would become U-NII-4.

In August 2013, the IEEE 802.11 Regulatory Standing Committee created a subcommittee called the *DSRC Coexistence Tiger Team* to explore possible band sharing techniques that will enable harmonious coexistence of DSRC and 802.11ac, and also help inform the regulatory process. In March 2015, the Tiger Team published their final report [2] that summarizes the issues surrounding the proposed band sharing ideas discussed in the group. Two key proposals are described in the report. In the first proposal, it is suggested that spectrum sharing is enabled by using the existing DSRC channelization and Clear Channel Assessment (CCA) in 10 MHz-wide channels [3]. This requires all 802.11ac (and other unlicensed) devices to be equipped with a new component to detect 802.11p preambles. This approach is similar to the Dynamic Frequency Selection (DFS) mechanism, and it requires that once an 802.11ac device detects a DSRC preamble, the frequency band from $5850 - 5925$ MHz must be declared busy for 10 seconds. The second proposal, on the other hand, requires modifying the DSRC channelization scheme [4]. The proposal suggests limiting the safety-critical applications to operate only in the upper DSRC channels, which are not shared with Wi-Fi, while the non-safety applications operate in the lower channels, which are shared with Wi-Fi. However, the report stressed the need for further analysis, simulations, and field testing to determine an appropriate coexistence approach.

In this paper, we present our findings from a comprehensive study on the coexistence of 802.11p and 802.11ac, with a particular focus on 802.11ac system performance. Unless stated otherwise, we use the terms 802.11ac and Wi-Fi, and 802.11p and DSRC interchangeably. The main contributions of this paper are summarized below.

(1) Based on our experimental findings, we show that under the current 802.11ac standard, not all of the Wi-Fi channelizations (i.e., channel and bandwidth combinations) in the U-NII-4 band—particularly the high bandwidth options (i.e., 40 and 80 MHz channels)—can be used by 802.11ac devices without causing significant interference to DSRC nodes.

(2) Under these channelization constraints, we propose a Real-time Channelization Algorithm (RCA) to maximize the throughput of Wi-Fi Access Points (APs) operating in the shared spectrum.

(3) We evaluate the effectiveness of the proposed RCA using a prototype implementation on commodity hardware. We also use simulations to validate our RCA in at-scale networks.

The rest of this paper is organized as follows. In Sec. II, we briefly describe previous work on Wi-Fi bandwidth adaptation and coexistence of DSRC and Wi-Fi. Some key 802.11p and Wi-Fi features related to our work are summarized in Sec. III.

This work was partially sponsored by NSF through grants 1314598, 1265886, 1431244, and 1547241, and by the industry affiliates of the Broadband Wireless Access & Applications Center and the Wireless@Virginia Tech group.

Details of our experimental testbed are provided in Sec. IV. In Sec. V, we present key findings regarding permissible Wi-Fi channelizations in the U-NII-4 band under the current IEEE 802.11ac standard. Taking into consideration these restrictions, we describe RCA in Sec. VI. We evaluate the performance of RCA on our testbed as well as using *ns-3* simulations. In Sec. VIII, we discuss the implications of our findings.

II. Related Work

Dynamic channel bandwidth adaptation for Wi-Fi APs has been previously studied. In [5], Chandra et. al. suggest dynamic adaptation of Wi-Fi channel bandwidth in order to maximize the throughput of a single AP-client pair. The IEEE 802.11n standard [6] introduced channel bonding that enables the use of wider (40 MHz) channels. However, the use of wider channels does not necessarily provide performance gains. Deek et. al. [7] and Arslan et. al. [8] show that naïve channel allocation and bandwidth selection schemes in heterogeneous 802.11 networks can lead to severe performance degradation. These studies consider 802.11 devices operating in enterprise networks and do not consider changing the channelization in response to varying traffic conditions.

In recent years, studies have characterized the performance of 802.11ac links in the presence of legacy 802.11 systems (802.11 a/b/g/n). In [9], Zeng et. al. investigate the performance of 80 MHz wide 802.11ac links in the presence of 20 MHz legacy 802.11 systems operating in the secondary channels of 802.11ac. The authors show that in such a coexistence scenario, the performance of the 802.11ac system suffers severe degradation. Similar results have been reported in [10].

As the above studies indicate, the characterization of 802.11ac (or 802.11n) performance in the presence of legacy 802.11 systems has attracted significant attention from the research community. In contrast, spectrum sharing between 802.11p and 802.11ac has not garnered much attention yet. Since 802.11p standard is derived from 802.11a, the coexistence of 802.11ac and 802.11p may seem similar to the above coexistence scenarios. However, there are key differences between the 802.11a/b/g/n-802.11ac coexistence and 802.11p-802.11ac coexistence. Firstly, 802.11p uses 10 MHz-wide channels. Wi-Fi transmitters have no support for preamble detection of 10 MHz-wide signals, while 802.11p cannot decode 20 MHz preambles. Thus, if 802.11p and 802.11ac transmitters were to operate in the same spectrum, both transmitters would detect each other's presence using only Energy Detection (ED). Since the detection threshold using ED is higher (by approx. 20 dB) than preamble detection [6], the sensing range of both transmitters would be much smaller as compared to that in traditional coexistence scenarios, resulting in increased probability of collisions in both networks.

Secondly, 802.11p is the primary user of the spectrum. Thus, an 802.11p transmitter must have higher priority in accessing the channel in the presence of Wi-Fi transmitters. Lansford et. al. [11] suggest increasing certain channel access parameters such as inter-frame spacing (IFS) to provide higher channel access priority to 802.11p nodes. In [12], Park et. al. show that increasing the Arbitration IFS (AIFS) value of 802.11ac transmitters indeed helps in mitigating the priority reversal problem. However, the authors ignore the 802.11ac performance degradation caused by the increased AIFS. Furthermore, in [12], 802.11a devices are used to emulate both 802.11p as well as 802.11ac devices. As a result, certain key features of 802.11ac—such as frame aggregation and channel bonding—that have severe impact on 802.11p performance were not taken into consideration.

III. Technical Background

A. Dedicated Short Range Communication

DSRC is a short-range to medium-range wireless communications technology that is designed to support vehicular applications. There are seven 10 MHz channels in the band as shown in Fig. 1. The spectrum includes one Control Channel (CCH), Ch. 178, and six Service Channels (SCHs), Ch. 172, 174, 176, 180, 182, and 184, and a 5 MHz guard band.

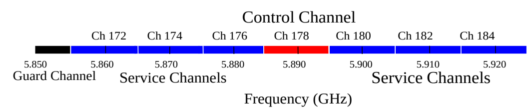

Fig. 1: DSRC spectrum map

The DSRC standard is an extension of IEEE 802.11a, designed for vehicles traveling at high speeds. Each DSRC node can operate on the CCH as well as the SCHs. Vehicles with a single radio use time division to operate between the CCH and SCHs. Time is synchronized among all DSRC devices, and the radio switches between CCH and SCH periodically.

There are two kinds of vehicular applications. The first kind—safety applications—are communicated by the DSRC nodes using periodic broadcasts of basic safety messages (BSMs). BSMs contain information such as vehicle position, velocity, acceleration, etc. The information contained in the BSMs is processed at the higher layers for applications such as lane change detection, impending collision alerts, etc. The second kind of applications are the non-safety applications such as downloading digital maps, automatic toll collection, etc. While these services are offered only on the SCHs, the SCHs can also be used to transmit safety application packets [13]. Thus, successful delivery of packets transmitted by DSRC nodes on all channels is of critical importance.

B. IEEE 802.11ac

IEEE 802.11ac is the latest standard in the IEEE 802.11 family. Unlike previous IEEE 802.11 standards, 802.11ac operates only in the 5 GHz band. 802.11ac achieves very high throughput by deploying the following features.

Channel bonding. Channel bonding uses two or more adjacent 20 MHz channels to form 40, 80 or 160 MHz channels in order to achieve higher system throughput. Among all the bonded channels, one 20 MHz channel, referred to as the *primary channel*, is used to transmit Wi-Fi beacons and other control information. All other bonded channels are referred to as *secondary channels*.

Frame aggregation. To reduce the overhead involved in the channel access mechanism at the MAC layer, 802.11n introduced frame aggregation. Frame aggregation involves sending multiple data frames in a single transmission opportunity. IEEE 802.11ac inherits this feature from 802.11n.

MIMO. The use of multiple antennas at the transmitter and receiver can be used to increase the reliability of transmissions (using *spatial diversity*) or to increase the rate of communication (using *spatial multiplexing*).

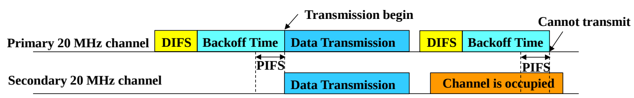

Fig. 2: Carrier sensing for primary and secondary channels.

We now discuss the channel access mechanism used in 802.11ac over the primary and secondary channels. An 802.11ac transmitter first senses the primary channel for Distributed IFS (DIFS) amount of time[1]. If the primary channel is sensed idle for a DIFS interval, the transmitter chooses a random backoff counter from its current contention window (CW), and continues to sense the channel. After each time slot, the transmitter decrements the CW by 1. During this backoff time, if the primary channel is sensed to be busy, the device freezes the backoff counter, and keeps sensing until the channel is idle again. When the channel is idle, it resumes the backoff counter. Once the backoff counter reaches 0, the transmitter sends the frame.

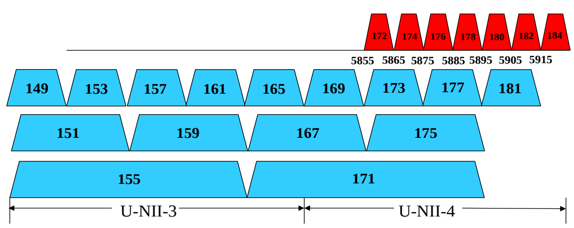

Fig. 3: Channelization of the U-NII-4 band.

On the other hand, the secondary channels are sensed for Point IFS (PIFS) interval of time immediately preceding the end of the backoff counter [9]. Thus, the secondary channel is sensed for a much smaller interval as compared to the primary channel, as seen in Fig. 2.

802.11ac supports both static and dynamic channel width operation. In static channel access, if any of the sub-channels is sensed busy, the 802.11ac transmitter will continue to sense the channels until all primary and secondary channels are idle. Consequently, under heavy traffic conditions, the probability of accessing a wide channel will be diminished [14]. In dynamic channel access, an 802.11ac transmitter first attempts to transmit over wide bandwidths. If any secondary channel is sensed busy, instead of waiting for all secondary channels to become idle, the transmitter transmits on the primary and adjacent idle secondary (if any) channels. Additionally, the enhanced-RTS and CTS mechanism can be used to combat the hidden node problem on the secondary channel [15].

[1]The actual value of IFS used depends on the type of Wi-Fi traffic. While DIFS is used for data frames in the Distributed Coordination Function (DCF) mode, 802.11ac mandates the use of Enhanced Distributed Channel Access (EDCA) mode which uses AIFS instead.

C. Channelization of the U-NII-4 Band

The channelization in the U-NII-4 band is shown in Fig. 3. The unlicensed use of the $5.850 - 5.925$ GHz band, if allowed by the FCC, will enable additional five 20 MHz, two 40 MHz, and one each of 80 MHz and 160 MHz channels. In the rest of this paper, we only consider channels in the U-NII-4 band, i.e. Ch. 169-181.

IV. Experiment Setup

We created and configured a testbed comprising of 4 802.11ac nodes and 6 802.11p nodes. In this section, we describe the testbed setup in terms of configuration of the hardware, the spectral scan and analysis procedure, and performance evaluation of DSRC and Wi-Fi devices.

A. Hardware Configuration

We use off-the-shelf hardware for 802.11ac and 802.11p. The RouterBoard RB911G-5HPacD is used as the 802.11ac AP as well as the client. The RB911G-5HPacD has two RF chains, thus enabling 2×2 MIMO. It uses Qualcomm Atheros QCA9882 chipset for 802.11ac. We use an embedded Linux distribution—OpenWrt (Chaos Calmer 15.05.1) as the operating system with ath10k drivers on the 802.11ac hardware.

The 802.11p hardware comprises of RouterBoard RB433AH as the processor, and RouterBoard RB52H as the RF card. The RF card is compatible with 802.11a/b/g and operates in the 2.4 and 5 GHz band. We use an OpenWrt distribution (Attitude Adjustment 12.09) with ath5k drivers provided by open-source vendor Componentiality [16] on the 802.11p devices. This distribution modifies the 802.11a hardware to operate as 802.11p with 10 MHz-wide channels.

As discussed in Sec. III, 802.11p operates in Ch. 172-184. The current regulations on 802.11ac, however, does not permit Wi-Fi operation above Ch. 165. Therefore, we configured the 802.11p nodes to operate on 10 MHz channels in the U-NII-3 band (Ch. 149-165). In this way, we created the spectrum sharing scenario of the U-NII-4 band in the U-NII-3 band. Although two 802.11p channels overlap with one 802.11ac channel in the U-NII-4 band, since partial overlap of 20 MHz channels is not allowed in the 5 GHz band, 802.11 hardware can operate only on certain fixed channels. Specifically, in our testbed, the 802.11p hardware can operate on Ch. 153, 157 and 161 with a channel width of 10 MHz. Although this channelization deviates from the one proposed for the U-NII-4 band (due to constraints imposed by the hardware and firmware), we believe that this does not negatively impact our study, because an 802.11ac transmitter must declare a 20 MHz channel to be busy if either of the two overlapping 10 MHz DSRC channels are occupied. In the rest of this paper, we refer to the channels by their original 10 MHz and 20 MHz channel numbers in the U-NII-4 band, i.e., Ch. 172-184 (for DSRC) and Ch. 169-181 (for Wi-Fi).

B. Spectrum Scan and Analysis

Each Wi-Fi radio senses the channel before its transmission. We believe that this spectral information can be leveraged in order to make informed channel and bandwidth allocations in Wi-Fi networks. The spectral information is required at the MAC layer, and is implemented at the kernel level. However, some Wi-Fi chipset manufacturers provide access to this spectral information at the user-space [17], [18].

In our testbed, we access the raw spectral information available from the 802.11ac QCA9882 chipset and analyze this data using a custom user-space application (based on [19]) to quantify spectral utilization. The 80 MHz spectrum (Ch. 169-181) is scanned approx. 1000 times in 1 sec. In each scan, 256 sample points are distributed across the 80 MHz spectrum. The scan and analysis is performed at the Wi-Fi AP itself; thus, permitting real-time usage of the spectral information.

For each 20 MHz channel, the spectral utilization is calculated as the fraction of time for which the signal level is above the energy detection threshold of the Wi-Fi AP (energy detection threshold is used because Wi-Fi AP cannot detect 802.11p preambles). For each 40 MHz channel, the spectral utilization is calculated as the fraction of time for which signal levels on any one of the two 20 MHz sub-channels is above the energy detection threshold. The spectral utilization for the 80 MHz channel is calculated in a similar manner.

By analyzing the spectral information, an AP can learn the usage of the spectrum, and adjust its channelization in an informed manner. Thus, the spectral scan and analysis enables the AP to take into account the interference from other devices.

C. Performance Evaluation

Different metrics such as Packet Error Rate (PER), Packet Delivery Ratio (PDR), link throughput etc. can be used to characterize a wireless link. We use the wireless link throughput as the metric for performance evaluation of 802.11ac links because 802.11ac has been primarily developed for high data-rate applications. On the other hand, we use the PER to analyze the performance of 802.11p systems, since 802.11p is expected to be used mainly for safety applications where low PER is of critical importance.

We use the packet generation tool Iperf for generating synthetic traffic and measuring the throughput and PER on the Wi-Fi and DSRC nodes. Only UDP traffic is generated because we wish to ignore the performance loss due to TCP overheads (handshake, congestion control, TCP retransmissions etc). In case of 802.11ac links, unless otherwise mentioned, we generate UDP traffic from the client to the AP.

The traffic patterns and topology of nodes for 802.11ac and DSRC are described alongside each test scenario. In all cases, we assume that a Wi-Fi AP and a client are both stationary. This is a reasonable assumption in most enterprise networks. The DSRC nodes are also stationary in our testbed. In practice, however, the DSRC nodes are mobile. We emulate this in our testbed using varying traffic patterns on DSRC channels.

Dynamic change of channelizations can lead to large switching delays due to hardware reconfigurations, dissociation and re-association of the client(s). To minimize such delays, we leverage the Channel Switch Announcement (CSA) feature of IEEE 802.11. In a Basic Service Set (BSS), CSA can be used by an AP to notify the clients of the exact time of switch and channel information of the new channel. By using CSA, the switching delay can be reduced to a few hundred milliseconds.

D. Custom Modifications

In Sec. V, we describe the need to adjust the AIFS value of the Wi-Fi transmitter based on the operating primary channel. We enable this change in the AIFS from the user-space by extending the *iw* command line utility.

The IEEE 802.11 standard specifies the procedure for selection of primary and secondary channels in case of channel bonding [6]. If multiple APs share the same 40 MHz (or 80 MHz) channels, all the APs must set the same primary and secondary channels. However, as explained in Sec. V-B, we need to set specific primary and secondary channels in our testbed. We achieve this by disabling the automatic shift of primary and secondary channels in the *hostapd* daemon.

V. DSRC PERFORMANCE IN THE PRESENCE OF 802.11AC

In the coexistence scenario between DSRC and Wi-Fi, DSRC are the primary users. Thus, before proposing RCA, we first analyze the impact of Wi-Fi transmissions in the shared spectrum on DSRC performance. We show that in order to provide adequate protection to DSRC transmissions, certain channelization options must be disabled at Wi-Fi transmitters.

A. DSRC in the Primary Channel of 802.11ac

One of the conclusions made in the previous studies [11], [12] is that increasing the AIFS of the Wi-Fi transmitter to a sufficiently large value can be used to mitigate the impact of Wi-Fi transmissions on DSRC performance. This is, however, valid only when the DSRC channel of operation overlaps with the primary channel of the Wi-Fi transmitters. We discuss DSRC's performance when it operates on a secondary Wi-Fi channel in the next subsection. In this subsection, using experimental results obtained from our testbed, we show that the conclusion made by investigators in the previous studies is indeed correct. However, we show that the value of AIFS must be increased to a value that is much higher than that suggested in [12], and justify this claim.

We consider a network comprising of 2 DSRC nodes and a Wi-Fi AP-client pair as shown in Fig. 4. The DSRC transmitter generates traffic on the CCH (Ch. 178), and the Wi-Fi AP operates on the 20 MHz channel that overlaps with the CCH (Ch. 177). The direction of the arrows on the solid lines indicates the direction of communication. The distance between the nodes are such that all nodes are within the transmission and sensing range of each other.

The source node sends UDP traffic to the sink node using Iperf such that the channel is fully saturated, but has a very small PER ($0 - 1\%$), i.e. any further increase in the packet transmission rate from the source node will lead to non-zero PER. The length of each DSRC packet is 200 bytes, and the packets are transmitted at a fixed PHY-layer rate of 6 Mbps.

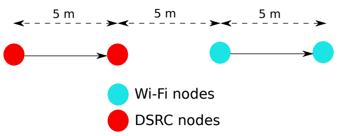

Fig. 4: Topology for testing DSRC system performance

We first observe the performance of DSRC nodes in the absence of Wi-Fi transmissions. The Wi-Fi transmitter then starts a UDP stream of 1500 byte packets to the receiver. The Wi-Fi transmitter uses its rate adaptation (RA) algorithm to identify the best Modulation and Coding Scheme (MCS) for its transmissions. Fig. 5a shows the PER at the DSRC receiver as a function of Wi-Fi AIFSN. AIFS is related to AIFSN by

$$\text{AIFS} = \text{SIFS} + \text{AIFSN} \times \text{slot\_time}, \qquad (1)$$

where SIFS is the short IFS used for acknowledgment packets and slot_time is the time duration of each Wi-Fi slot.

Clearly, DSRC's medium access priority can be increased by increasing the AIFS value of Wi-Fi. The optimal AIFS value that the Wi-Fi transmitters must use to adequately protect DSRC would depend on a number of factors, including vehicle density, network topology, etc. In our setup, we observe from Fig. 5a that beyond AIFSN = 100, the impact of Wi-Fi transmissions on DSRC's PER is negligible. Henceforth, we set the AIFSN value to 100 (AIFS = 916μs) for Wi-Fi transmitters when they share the spectrum with DSRC.

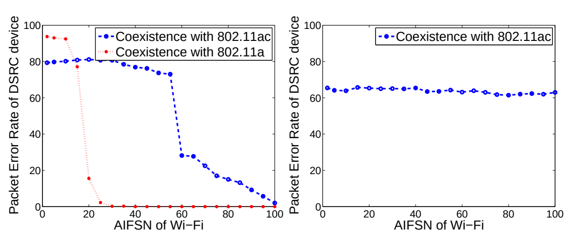

(a) DSRC on primary channel (b) DSRC on secondary channel

Fig. 5: PER of DSRC node on primary and secondary channel

We believe that the value of AIFSN required for prioritizing DSRC transmissions in our setup is much higher than that suggested in [12] due to frame aggregation in 802.11ac. To verify this, we let the interferer operate in 802.11a mode, where frame aggregation is disabled. The PER at the DSRC receiver in this case is also shown in Fig 5a. We infer that frame aggregation has severe impact on DSRC performance. This is because a single channel access of an aggregated frame is equivalent to access of multiple non-aggregated Wi-Fi frames as far as DSRC performance is concerned.

B. DSRC in a Secondary Channel of 802.11ac

We claim that even a sufficiently large increase in the AIFS value of Wi-Fi transmitters would fail to protect DSRC from priority reversal (i.e. when both, Wi-Fi and DSRC nodes, have a packet to transmit, Wi-Fi node transmits prior to the DSRC node), if the DSRC transmitters operate in the secondary channels of the Wi-Fi transmitter. We verify the legitimacy of our claim using experimental results. We use the same topology as shown in Fig. 4. The DSRC source node transmits packets to the sink node on the CCH such that the channel is saturated and the PER is very low. The Wi-Fi nodes operate on a 40 MHz link such that the CCH overlaps with the secondary channel of the Wi-Fi link. The curve in Fig. 5b represents the PER at the DSRC sink node operating in the secondary channel of the Wi-Fi link.

Results in Fig. 5b show that a sufficiently large increase in the Wi-Fi AIFS does not protect DSRC users (against performance degradation), if the shared channel is a secondary Wi-Fi channel. In such cases, reliable DSRC performance can be guaranteed if the *full backoff procedure* (AIFS+backoff) [10] is followed on the secondary channels along with the primary channel[2], although this is not supported by the current Wi-Fi standards. From these results, we can make an important conclusion. *Under the current Wi-Fi standards, certain channelization schemes cannot be used by Wi-Fi networks without causing significant degradation in the performance of DSRC nodes that share the secondary Wi-Fi channels.*

TABLE I: Permissible channelizations in the shared spectrum

Channel	169	173	177	181	173(P) + 169(S)
Channel Index	C_1	C_2	C_3	C_4	C_5
HT mode	HT20	HT20	HT20	HT20	HT40
AIFSN	2	100	100	100	100

Summary: Given that DSRC nodes are the primary users of the shared spectrum, and considering the importance of safety applications supported by DSRC networks, it is essential that Wi-Fi transmissions do not degrade DSRC performance. From the experimental findings presented in this section, we conclude that while it is possible for Wi-Fi nodes to share the spectrum with DSRC nodes, *a Wi-Fi AP should not be configured to operate its secondary channel(s) on the shared spectrum*. Consequently, Ch. 171 (80 MHz) and 175 (40 MHz) shown in Fig. 3 should not be used by Wi-Fi APs. The only channel bonding configuration available at the Wi-Fi AP is to bond Ch. 169 and 173 as the secondary and primary channels respectively. The possible channelizations are listed in Table I.

VI. Proposed Channelization Algorithm

The increased AIFS value for Wi-Fi transmitters in the shared spectrum can have severe implications to Wi-Fi system performance. In the presence of high DSRC traffic, Wi-Fi APs operating in the shared spectrum could experience substantial loss in throughput. In some cases, it might be beneficial for

[2]This could not be verified in our testbed as the secondary channel access procedure is implemented in the proprietary firmware of QCA9882 (as is the case with most other chipsets).

the Wi-Fi AP to operate on the 20 MHz non-shared spectrum. On the other hand, if DSRC nodes are absent, or if DSRC traffic is low, using the shared spectrum could yield satisfactory throughput despite the increased AIFS. Thus, efficient utilization of the spectrum relies on informed channelization allocation at the Wi-Fi AP. In this section, we describe our proposed RCA that makes use of the spectral information available at the Wi-Fi AP to determine the best channelization for an AP operating in the U-NII-4 band.

A. Problem Statement

We consider a shared spectrum network with 6 DSRC nodes and 2 Wi-Fi BSSs. Each BSS comprises of an AP and an associated client. Each AP can select a primary channel and bandwidth from $C = \{C_1, \cdots, C_k\}$ combinations, as shown in Table I. The objective of RCA is to pick the best channel combination (in terms of its achievable throughput) among the permitted combinations using the locally available real-time spectral information. We only consider static channel access mechanism in this paper since most channel bonding options are not available at the Wi-Fi AP in the U-NII-4 band. In dynamic channel access mechanism, the AP determines the channel bandwidth for each transmission opportunity based on their availability. If the use of secondary channel is permissible, RCA can be used to determine the best primary channel, and the secondary channels can be selected as available.

B. Approach

The performance of a Wi-Fi link depends on factors such as the surrounding environment, and particularly on the number of DSRC nodes or vehicles in its vicinity, DSRC traffic pattern, etc. Under such circumstances, the optimal channel combination at the AP cannot be determined without prior knowledge of the DSRC nodes. At best, the Wi-Fi AP can change its channelization in response to increased DSRC traffic on certain channels. RCA seeks to do just that — scan and analyze the spectrum to calculate the spectral utilization, estimate a bound on the achievable throughput on each channelization, and select the channelization that yields the highest throughput.

RCA keeps repeating the above process in varying traffic conditions. The main task of RCA is to look for a better channelization every time the link throughput drops. Additionally, RCA regularly tries to check for a better channelization even if the channel conditions do not fluctuate. The finer details of RCA are described next.

C. Proposed Algorithm

Algorithm 1 Real-time Channelization Algorithm (RCA)

1: Initialize: Hop across all C_i to determine $R_i, i = 1, \cdots, k$.
2: **while** *true* **do**
3: Scan the 80 MHz spectrum.
4: Compute U_i, and E^i_{Th}, for $i = 1, \cdots, k$.
5: Update $R_{ch} = R_{cur}$.
6: Switch to Channel $ch = \underset{i}{\operatorname{argmax}}([E^1_{Th}, \cdots, E^k_{Th}])$.
7: Initialize: $\tau_{ch} = 0$
8: **while** $\left(\Delta(Th_{ch}) \leq \Delta_{max}\right)$ and $\left(Th_{ch} \geq \beta E^{ch}_{Th}\right)$ and $\left(\tau_{ch} \leq T_{max}\right)$ **do**
9: Transmit on channel ch.
10: **end while**
11: **end while**

TABLE II: Algorithm Parameters

Param.	Description	Value
L_{data}	Length of data packets	1500 bytes
L_{ack}	Length of ACK packets	100 bytes
K	Number of aggregated frames	$2-32$
R_i	Data rate used in previous transmission	$6.5-400$ Mbps
PRR	Packet Reception Ratio	$0-1$
R_{cur}	PHY Layer Rate used in the current session	$6.5-400$ Mbps
U_i	Utilization ratio	$0-1$
SIFS	SIFS time of Wi-Fi transmitter	16μs
AIFS	AIFS time of Wi-Fi transmitter	Eq. (1)
E^i_{Th}	Expected throughput of the ith channel	Eq. (2)
β	Throughput reduction factor	0.4
$\Delta(Th_{ch})$	Two-window throughput difference	$variable$
Δ_{max}	Threshold for two-window throughput diff.	0.15
$T_{\max}$	Maximum time before next scan process	Eq. (3)
$t_{analyze}$	Time required to analyze sensing data	$3-4$ sec
τ_{ch}	Sojourn time on channel ch	$0-T_{\max}$
$E^{\max}_{Th}$	Expected throughput at highest MCS	400 Mbps
t_{data}	Time to transmit data frames at R_{cur}	$K{\cdot}L_{data}/R_{cur}$
t_{BO}	Expected time spent in backoff process	$\mathbb{E}$[CW]
t_{ack}	Time to transmit block acknowledgement	L_{ack}/R_{cur}

The detailed steps in RCA are shown in Algorithm 1. The variables used are listed in Table II. For each parameter, the presence of a subscript or superscript i indicates that the parameter is specific to channelization i, for $i = 1, \cdots, k$.

Line 1*:* In practical settings, Wi-Fi transmitters use some form of RA algorithm in order to minimize the PER or maximize the throughput. In RCA, we let the transmitter use its default RA algorithm to determine the best MCS scheme for its transmission. RCA starts by initializing the rates R_i for all $C_i, i = 1, \cdots, k$. The AP hops across all channelizations in order to determine the client PHY rate on each channelization. The AP notifies the client of the channel switch using RCA; thus, minimizing the impact of channel switching.

Lines $3-4$*:* The AP scans the 80 MHz spectrum. Using this scan data, spectral utilization of each channelization is computed as described in Sec. IV-B. The AP then calculates a bound on the expected throughput on each channelization using these utilization ratios. We use Eq. (2) to calculate the expected throughput for each channelization. In [7], the authors show that in the absence of external interference, except for the higher order MCS schemes, the actual throughput closely matches the expected throughput computed using Eq. (2). We introduce the term $(1 - U_i)$ in Eq. (2) to take external interference into consideration.

$$E^i_{Th} = \frac{K \cdot L_{data} \cdot (1 - U_i) \cdot PRR}{\text{AIFS} + t_{BO} + t_{data} + \text{SIFS} + t_{ack}} \qquad (2)$$

The value of K, the number of aggregated frames, depends on the transmitter MCS scheme. Using Wireshark traces, we determine K for each MCS scheme.

Lines $5-6$*:* The AP tunes its RF front-end to the channelization that is expected to achieve the highest throughput. Before the AP moves to the new channel, the AP (i) announces the channel switch decision to all the associated client(s) using CSA, and (ii) stores the value of PHY layer transmission rate used by the client on the current channel for subsequent computations of the expected throughput.

Lines $8-10$*:* The objective of RCA is to determine the best channelization for an AP in real-time. A scan process is, thus, initiated in the following two cases.

Case (1): The link throughput drops due to increase in the utilization of the operating channel. This sudden drop can be attributed to entry of additional nodes using that channel. To identify a better channel, a scan process is initiated. We use the following two mechanisms for detecting an increase in spectrum utilization of the operating channel.

(a) A sudden drop in the throughput can be detected by maintaining two windows for calculating the average throughput. The first window averages over a larger interval (say w_1), while the second window averages over a smaller interval (say w_2). A sudden drop in link throughput will cause the average over smaller window to drop faster than the average over larger window. If the difference in average throughput over the two windows, $\Delta(Th_{ch})$, is greater than some threshold $\Delta_{\max}$, we trigger a scan process. In our experiments, we use $w_1 = 10, w_2 = 3$. For these window sizes, we observe that $\Delta_{\max} = 0.15$ accurately detects a sudden drop in throughput.

(b) There is a time delay of $t_{analyze}$ between the scan process and reconfiguration of the hardware to the new channelization. There is a possibility that certain nodes start operating on the new channel within this time interval. If the entry of these nodes causes the actual throughput to be significantly lower than the expected throughput, we infer that the channel conditions have changed since the previous scan, and initiate a new spectrum scan. We quantify this using a parameter β. A scan process is triggered if $Th_{ch} \leq \beta \cdot E_{Th}^{ch}$. A high value of β can lead to frequent scans in the absence of any interference, specially at higher MCS schemes where the actual throughput can be less than the expected throughput. On the other hand, a small value of β can fail to detect the entry of other nodes in the operating channel. Using experiments, we observe that value of $\beta = 0.4$ serves as a good indicator of external interference on the operating channel.

Case (2): $T_{\max}$ amount of time has elapsed since the previous scan. Although the throughput does not drop, there might exist another channelization that yields a higher throughput. Thus, after every $T_{\max}$ interval, the RCA scans the spectrum for a better channelization despite no drop in throughput.

RCA calculates the value of $T_{\max}$ based on the expected throughput of the current channelization. If the expected throughput is equal to the maximum possible throughput, a spectral scan and analysis can lead to no increase in throughput. Ideally, if $Th_{ch} = Th_{\max}$, $T_{\max}$ should be infinite. On the other hand, the utilization of the entire 80 MHz spectrum can be high, resulting in a small throughput even on the best channel. When the expected throughput is close to zero, we let the device operate on the operating channel for at least $t_{analyze}$ amount of time. In such cases, frequent scans can detect a better channelization as soon as its utilization decreases. Keeping these conditions in mind, $T_{\max}$ is calculated as,

$$T_{\max} = \frac{t_{analyze}}{1 - \left(\frac{E_{Th}^{ch}}{Th_{\max}}\right)}, \tag{3}$$

where Th_{max} is the maximum possible throughput. Since only 20 and 40 MHz channels are allowed, we let $Th_{max} = 400$ Mbps (corresponding to the highest MCS scheme at 40 MHz).

D. RCA Implementation

We implement RCA completely at the AP, with no modifications required at the associated client(s). This is important because it is infeasible to change software packages and drivers at all clients; however, the AP can enable RCA using a simple patch. Since RCA uses real-time spectral information to adapt the channelization, it is implemented with low response time on commodity hardware using custom user-space scripts. In our implementation, given that the AP detects a significant drop in throughput in the operating channel, the spectrum is scanned and the best channel is selected within $4-5$ sec. The main sources of delay at the AP are the scan process (approx. 1 sec), and the processing time of spectral information (approx. $3-4$ sec). During the scan and analysis phases of RCA, communication between AP and client continues over the previous operating channel. If the AP switches to a new channel, it notifies the clients(s) of the switch using CSA.

VII. Performance Evaluation

A. Experimental Characterization

The topology used to evaluate the performance of RCA is shown in Fig. 6. One DSRC transmitter-receiver pair operates on each of the Ch. 172, 178 and 182.[3] We use two pairs of Wi-Fi nodes in our experiments — an interfering pair that operates with saturated traffic on Ch. 169, and an observation pair, which is used to evaluate the performance of RCA. All the nodes are within the sensing range of each other.

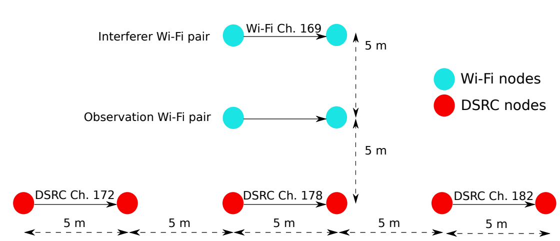

Fig. 6: Topology for evaluating RCA performance.

Traffic on DSRC channels largely depends on the number of DSRC nodes in the vicinity of the Wi-Fi AP. The larger the number of DSRC nodes, higher the utilization on the DSRC channels. While traffic on Ch. 172 (BSM traffic) and 178 (control traffic) are directly related to the number of

[3] Wi-Fi channels are odd numbered, while DSRC channels are even-numbered. DSRC Ch. 172 overlaps with Wi-Fi Ch. 173, Ch. 178 with Ch. 177, and Ch. 182 with Ch. 181.

vehicles, traffic on the SCH (Ch. 182) can be arbitrary. The traffic patterns of the four channels generated for our testbed experiments are shown in Fig. 7. The numbers inside each block indicate the spectrum utilization ratio at the observation AP as a fraction of saturated traffic. We change the traffic pattern every T seconds; by changing the value of T, we can emulate fast and slow varying traffic conditions. The duration of each experiment is $8T$.

The utilization of Ch. 173 and 177 first gradually increase and then decrease. On the other hand, the utilization of Ch. 181 is either 0% or 50%. We operate an interfering Wi-Fi BSS on Ch. 169 with saturated traffic. As a result, utilization of Ch. 169 varies in the range $30-60\%$. Using these traffic patterns, we emulate several types of varying traffic conditions.

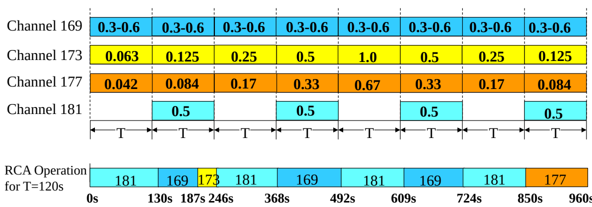

Fig. 7: Traffic pattern generated for RCA evaluation

First, we evaluate the performance of each static channelization allocation (shown in Table I) separately. We keep the channelization fixed throughout the duration of experiment ($8T$). Next, we run RCA on the observed AP and test its performance. The performance of RCA in comparison with all static allocations for different values of T is shown in Fig. 9a. Each bar represents the average throughput for the duration of the experiment averaged over 10 independent trials. For each value of T, the first five bars show the performance of static allocation to channelizations $1-5$. The sixth bar shows the average across the five static channelizations, and the sixth bar represents the performance of RCA.

Different values of time interval T have no impact on the performance of static channelization schemes because the same periodic pattern is generated, and only the duration of experiment run changes. The performance of the 40 MHz bonded channel is poor because the interfering Wi-Fi AP always operates on its secondary channel (Ch. 169). Fig. 9a shows that for $T=120$ sec, under the given traffic conditions, RCA outperforms the best static allocation by approx. 12%, and the average across all static allocations by approx. 50%. RCA efficiently uses the spectrum by using Ch. 169 when DSRC nodes occupy the shared spectrum while exploiting the idle time periods on Ch. 173, 177 and 181 whenever these channels can yield a higher throughout.

The operation of RCA on the observed AP for one instance of $T=120$ sec is shown at the bottom of Fig. 7. At the onset, spectral analysis reveals low utilization on Ch. 181; thus, RCA uses Ch. 181 with AIFSN = 100. Since the channel conditions do not change, the spectrum is scanned every $T_{\max}$ interval of time; however, Ch. 181 is picked every time due to its near-zero utilization. At $T=120$ sec, a DSRC node starts transmitting on Ch. 182 (overlapping with Ch. 181). The high AIFSN of the observation pair leads to a sudden drop in throughput, triggering the scan process. Next, Ch. 169 is picked as the best channel. Although the utilization of Ch. 169 is higher, RCA predicts that the low AIFSN (AIFSN = 2) when using this channel would result in a higher throughput in Ch. 169. In the next scan, RCA predicts a higher throughput in Ch. 173, and thus, switches to that channel. RCA continues operation in Ch. 173 until Ch. 183 becomes idle once again. Rest of the RCA operation is shown in Fig. 7.

Our experimental findings indicate that the Wi-Fi's opportunistic access of channels that are lightly utilized by DSRC can can provide a significant increase in throughput. As the value of T decreases, the performance gain of RCA over the best static allocation scheme decreases. However, even in the worst case (i.e., when $T=10$ sec), the throughput achieved by RCA is approx. 95% of that using the best static allocation scheme, and performs better than all other static allocation schemes. This drop in the relative gain of RCA over the best static assignment is because T becomes comparable to $t_{analyze}$. When an AP detects a throughput drop in the current channel, it takes about $4-5$ sec to tune to a new channel. Thus, the throughput gain in the new channel is just sufficient to compensate for the loss in throughput during the RCA analysis phase while the link continues to operate over the busy channel. In the next subsection, we evaluate RCA under more general settings via simulations.

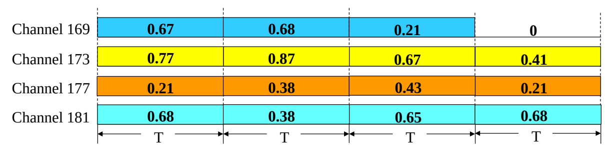

Fig. 8: Channel utilization ratio of the four channels

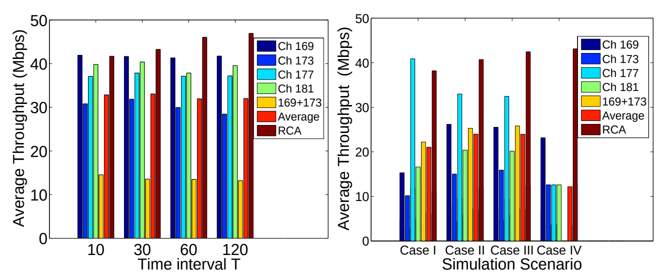

(a) Experimental results (b) Simulation results

Fig. 9: Evaluation of RCA

B. Simulation Study

We evaluate RCA at scale using *ns-3*. We observe the throughput of a particular Wi-Fi pair operating with saturated traffic. All DSRC nodes move along a straight road with a velocity of 20 m/s. The lane is 20m away from the Wi-Fi transmitter. We disable the RA algorithm and fix the MCS scheme to MCS-7. The 80 MHz spectrum is sampled $10,000$ times every second to compute the utilization ratio for each channel. 20 DSRC nodes operate in each of Ch. 173, 177 and 181, while two Wi-Fi APs operate on Ch. 169. All other parameters are similar to those used in the testbed experiments.

Results from our simulation studies are shown in Fig. 9b. While it is infeasible to validate RCA against all DSRC traffic patterns, we consider two special cases and two arbitrary traffic patterns. The four set of bars in Fig. 9b show the performance of RCA in comparison with each static channelization allocation in the following four cases. In Case I, the utilization ratio is constant across the entire simulation run. The utilization ratios are as shown in the first time interval $0-T$ in Fig. 8 (i.e. 0.67, 0.77, 0.21, 0.68 respectively). In this case, static allocation to Ch. 177 is arguably the best choice (due to its lowest utilization). If the Wi-Fi AP uses RCA, some performance loss results due to periodic scanning of the spectrum. Thus, in this extreme case, the throughput using RCA is smaller (approx. 7%) than the best static allocation scheme, but outperforms all other static allocation schemes.

Next, we generate arbitrary traffic patterns on the four channels (utilizations shown in Fig. 8). We set $T = 60$ sec in Case II, and $T = 90$ sec in Case III. In Case IV, we show the upper bound of the performance gain achieved by RCA. In the first time interval T (we use $T = 60$ sec), Ch. 169 is idle, while all other channels have full utilization. In the second time interval T, Ch. 173 is idle, and the other three channels have full utilization. Similar traffic patterns are generated for Ch. 177 and 181 in the third and fourth intervals respectively. In each of Case II, III and IV, we observe that the average throughput using RCA is higher than the best static allocation scheme, while the performance gains over other static allocation schemes are even higher.

Summary: We see that informed channelization allocation at Wi-Fi APs can yield substantial increase in the throughput over static channel allocation schemes, except in cases where traffic conditions remain constant forever or change over time intervals smaller than the analyzing delay. In practice, such frequent transitions in traffic patterns can occur at Wi-Fi APs located close to stop lights (depending on the interval of red and green lights). If this interval is small (order of $t_{analyze}$), certain static channel allocations can lead to higher throughput compared to RCA. However, this performance gap is very small as seen in Sec. VII-A.

VIII. Recommendations on the U-NII-4 Band Channelization

The Tiger Team, in its report [2], could not reach a consensus on band sharing options between DSRC and Wi-Fi. Based on the findings of this paper, we suggest the following recommendations for coexistence between DSRC and Wi-Fi in the U-NII-4 band.

In Sec. V, it was shown that DSRC nodes experience a high level of packet errors if the shared spectrum is configured as a secondary channel of some Wi-Fi AP. This renders the 80 MHz and one 40 MHz channelizations unusable by Wi-Fi APs. We propose a *combination of the two solutions* proposed in the Tiger Team report to enable harmonious coexistence between DSRC and Wi-Fi users; this approach is described below.

First, as proposed in [3], Wi-Fi transmitters must be able to detect 10 MHz wide 802.11p preambles. Enabling detection of 802.11p preambles at Wi-Fi transmitters will increase the sensing range of Wi-Fi with respect to DSRC nodes (due to lower threshold of preamble detection as compared to ED). This can also help in alleviating hidden node problems between Wi-Fi and DSRC nodes. Secondly, additional channelization options can be unlocked if the 802.11ac standard mandates full backoff on the secondary channels along with the primary channel. Finally, as suggested in [4], safety critical applications must have exclusive access to the highest three DSRC channels (i.e., channels 180, 182 and 184). On the other hand, the lower four channels, which can be shared with Wi-Fi, should be used for non-safety applications, since they can tolerate a higher level of packet loss.

Referring to Fig. 3, we note that the above recommendations exclude only one 20 MHz channel (Ch. 181) for Wi-Fi usage. On the other hand, all 40 MHz and 80 MHz channelizations can be used by Wi-Fi, thus enabling the use of wider bandwidth channels for supporting high data-rate applications.

References

[1] "Federal communications commission, in the matter of amendment of the commissions rules regarding dedicated short-range communication services in the 5.850-5.925 ghz band (5.9 ghz band)," tech. rep., FCC.

[2] "Final report of dsrc coexistence tiger team." https://mentor.ieee.org/802.11/dcn/15/11-15-0347-00-0reg-final-report-of-dsrc-coexistence-tiger-team-clean.pdf.

[3] "Proposal for u-nii-4 devices." https://mentor.ieee.org/802.11/dcn/13/11-13-0994-00-0reg-proposal-for-u-nii-4-devices.docx.

[4] "Proposal for dsrc band coexistence." https://mentor.ieee.org/802.11/dcn/13/11-13-1449-02-0reg-proposal-for-dsrc-band-coexistence.pptx.

[5] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *ACM SIGCOMM computer communication review*, vol. 38, pp. 135–146, ACM, 2008.

[6] "Part 11: Wireless lan medium access control (MAC) and physical layer (PHY) specifications," *IEEE Std*, 2012.

[7] L. Deek, E. Garcia-Villegas, E. Belding, S.-J. Lee, and K. Almeroth, "The impact of channel bonding on 802.11n network management," in *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, p. 11, ACM, 2011.

[8] M. Y. Arslan, K. Pelechrinis, I. Broustis, S. V. Krishnamurthy, S. Addepalli, and K. Papagiannaki, "Auto-configuration of 802.11n WLANs," in *Proceedings of the 6th International Conference*, p. 27, ACM, 2010.

[9] Y. Zeng, P. H. Pathak, and P. Mohapatra, "A first look at 802.11 ac in action: energy efficiency and interference characterization," in *Networking Conference, 2014 IFIP*, pp. 1–9, IEEE, 2014.

[10] M. Park, "IEEE 802.11ac: Dynamic bandwidth channel access," in *2011 IEEE International Conference on Communications (ICC)*, pp. 1–5, June 2011.

[11] J. Lansford, J. B. Kenney, and P. Ecclesine, "Coexistence of unlicensed devices with DSRC systems in the 5.9 ghz ITS band," in *Vehicular Networking Conference (VNC), 2013 IEEE*, pp. 9–16, IEEE, 2013.

[12] Y. Park and H. Kim, "On the coexistence of IEEE 802.11 ac and wave in the 5.9 GHz band," *IEEE Communications Magazine*, vol. 52, no. 6, pp. 162–168, 2014.

[13] "IEEE guide for wireless access in vehicular environments (WAVE) - architecture," *IEEE Std 1609.0-2013*, pp. 1–78, March 2014.

[14] O. Bejarano, E. W. Knightly, and M. Park, "IEEE 802.11ac: from channelization to multi-user MIMO," *IEEE Communications Magazine*, vol. 51, no. 10, pp. 84–90, 2013.

[15] M. S. Gast, *802.11ac: A survival guide*. " O'Reilly Media, Inc.", 2013.

[16] http://componentality.com/en/flexroad/.

[17] http://www.broadcom.com/blog/wireless-technology/a-breath-of-fresh-air-iq-for-enterprise-wi-fi-networks/.

[18] https://wireless.wiki.kernel.org/en/users/drivers/ath10k/spectral.

[19] https://github.com/simonwunderlich/FFT_eval.