



HAL
open science

Identifying Channel Saturation in Wi-Fi Networks via Passive Monitoring of IEEE 802.11 Beacon Jitter

Laudin Molina, Alberto Blanc, Nicolas Montavont, Ljiljana Simić

► **To cite this version:**

Laudin Molina, Alberto Blanc, Nicolas Montavont, Ljiljana Simić. Identifying Channel Saturation in Wi-Fi Networks via Passive Monitoring of IEEE 802.11 Beacon Jitter. *MobiWac 2017 - 15th ACM International Symposium on Mobility Management and Wireless Access*, Nov 2017, Miami Beach, United States. pp.63 - 70, 10.1145/3132062.3132069 . hal-01656170

HAL Id: hal-01656170

<https://hal.science/hal-01656170>

Submitted on 5 Dec 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Identifying Channel Saturation in Wi-Fi Networks via Passive Monitoring of IEEE 802.11 Beacon Jitter

Laudin Molina, Alberto Blanc,
Nicolas Montavont
Institute Mines Telecom
IMT-Atlantique
IRISA
firstname.lastname@imt-atlantique.fr

Ljiljana Simić
Institute for Networked Systems
RWTH Aachen University
lsi@inets.rwth-aachen.de

ABSTRACT

Every day large numbers of users connect to IEEE 802.11 networks in order to access the Internet and all sorts of services. However, due to their unplanned and unregulated nature, and the lack of admission control and Quality of Service Guarantees, these wireless networks can experience traffic demand that exceeds the network capacity. In this case, if a device tries to send more traffic, or if a new device joins the network, the aggregate throughput does not necessarily increase. In this paper we show that it is possible for IEEE 802.11 stations to detect a saturated channel by passively monitoring the beacon frames. Access points (APs) send beacon frames periodically and encode them using the strongest modulation and coding scheme, so that even stations far away from the sending APs can decode them correctly. When sending beacons, APs sense the channel first and, if it is busy, delay sending the frame, resulting in unevenly spaced beacon frames, whenever other transmitters are active. We present several experiments, under varying traffic loads, and analyze the distribution of the beacon jitter, whose variance increases as the offered load increases. We show that it is possible to determine, with an acceptable error rate, whether a channel is saturated by comparing the distribution of the beacon jitter with a reference distribution corresponding to a saturated channel.

KEYWORDS

IEEE 802.11; Beacons; Channel Saturation; Passive Measurements

1 INTRODUCTION

Today, IEEE 802.11 Wi-Fi enabled devices (e.g., laptops, smartphones, tablets) are ubiquitous and widely used by a large number of users accessing all sorts of applications and services. To meet the ever increasing demand for wireless connectivity, different actors have deployed IEEE 802.11 APs: Internet Service Provider (ISP) customers in their homes for their own use; businesses in their offices for their own employees, but also in public places for their customers (e.g., airports, shops, malls); public institutions serving larger areas (e.g., local administrations providing network coverage in a city). A large number of studies (see, for instance [1, 14, 17]) and projects (e.g., WiGLE, OpenSignal, Sensorly) have shown that, especially in urban areas, several APs can be detected at any given location. Due to the unregulated and unplanned nature of Wi-Fi networks, APs in close proximity of each other often operate on the same channel, especially on the frequently used non-overlapping channels (1, 6, and 11 for the 2.4GHz band). This can result in poor performance, in particular when the traffic demand exceeds

the channel capacity, resulting in a *saturated channel* (i.e., where any increase in the offered load does not result in an increase of the aggregate throughput). In this case, users are often better off joining a different network, for instance another AP operating on a different, and non-saturated channel, or a network using a different access technology (e.g., cellular network, wired network).

As it is common for users to have access to multiple APs, it is important to choose the “best” AP. Although there can be different definitions of the preferred AP, depending on the specific circumstances, an AP operating in a saturated channel is definitely a non-ideal candidate. Thus, one should take into account channel saturation when selecting an AP.

In this paper we propose a simple method for stations (STAs) to detect a saturated channel by passively monitoring beacon messages, which are available to all STAs as part of the IEEE 802.11 procedures. This enables STAs to passively collect information and to determine whether a channel is saturated or not. Using beacons to characterize the channel condition presents multiple advantages: beacons are always present, transmitted in broadcast mode, and use the strongest modulation and coding scheme (MCS). By analyzing experiments conducted with different traffic loads, we show that it is possible to identify whether a channel is saturated based on the distribution of the beacon jitter. Even though APs send beacon frames periodically, they do have to wait for the channel to be idle, resulting in an additional delay that depends on the traffic intensity. Empirical results show that the beacon jitter follows a similar distribution whenever the channel is saturated. Our solution exploits this by comparing the beacon jitter distribution with a reference distribution, corresponding to a saturated channel. While the literature on AP selection and Wi-Fi network characterization is vast (see, for instance, [2, 4, 6–9, 11–13, 15, 16, 18, 19]), to the best of our knowledge, no existing solution is both (1) implementable without changing the APs and (2) passive, i.e., it does not require exchanging any additional frames.

The remainder of the paper is organized as follows: Section 2 reviews the relevant literature; Section 3 gives an overview of IEEE 802.11; Section 4 defines *beacon jitter*; Section 5 presents our proposal for detecting saturated channels; Section 6 describes the experimental setup; Section 7 analyzes the results of the experiments and the performance of the proposed classification method. Section 8 concludes the paper and discusses possible extensions.

2 RELATED WORK

Over the last couple of decades, since the appearance of Wi-Fi networks, researchers have proposed different metrics to characterize these networks. In this section, we briefly review the literature most closely related to our work.

One metric is channel delay, which measures the time needed to send a frame. Haghani et al. [5] estimate the distribution of the access delay, from the moment the MAC frame is ready for transmission (i.e., the first in the transmission queue) until the corresponding acknowledgment is received. For this, Haghani et al. [5] developed a theoretical model that takes as input the busy/idle periods in the channel, as this factor affects the channel delay. Kajita et al. [8] estimate the channel delay by using a machine learning approach, where the input parameters come from traffic around the AP, implying that this approach requires information about the frame transmitted in the neighborhood. This information may not be always available, since APs might not be able to decode frames if the MCS is too high. Moreover, recent standard amendments (e.g., IEEE 802.11n, IEEE 802.11ac) allow APs to use beamforming, multiple-input-multiple-output (MIMO) and multi-user MIMO (MU-MIMO) to enhance the link quality with the receiver STA, reducing the chance that other STAs would be able to decode those frames.

Another widely studied metric is achievable throughput. Following Bianchi's seminal paper [2], several authors [6, 9] have proposed extensions and refinements based on this model to estimate the achievable throughput. They all take the collision probability as an input parameter. It is non-trivial to estimate this parameter, as it depends on several independent factors, such as the number of active transmitters, offered load, and collision domains (i.e., which nodes are within carrier sensing range) just to name a few. These factors cannot be precisely known by a single node. A common approach is to approximate the collision probability with the *retransmission ratio* [6, 9], sometimes called loss probability [15]. The STA sniffs the medium, counting the number of transmitted frames and the corresponding retransmissions. This approach is appropriate if the STA is able to capture all frames in the medium and all retransmissions are due to collisions. However, these assumptions do not always hold. As previously mentioned, the STA may not be able to decode all frames. In addition, some retransmissions may be caused by bad channel conditions and not by collisions. Hong et al. [7] use a collaborative approach to estimate the available throughput in which the STA requests information from the APs, including channel utilization and number of frames sent. Therefore, the estimation depends on the information collected by the AP and the willingness of the AP to share that information.

While these metrics are related to channel saturation, the relationship is not always straightforward. Our approach is to use the distribution of the beacon jitter to identify saturated channels. This has multiple advantages: (1) STAs can perform the measurements independently, without requiring support from other STAs; (2) beacons are transmitted periodically by all APs, i.e., they are always present; (3) beacons are broadcast, therefore APs do not use beamforming, MIMO or MU-MIMO; (4) beacons are encoded using the strongest MCS, so that they are more easily decoded.




Figure 1: Beacons transmitted by one AP in an idle (left) and busy (right) channel.

3 IEEE 802.11 OVERVIEW

3.1 Medium Access Control

CSMA/CA is the fundamental medium access control (MAC) mechanism of the IEEE 802.11 networks [3]. Whenever a device has a frame to send, it must first sense the channel to determine if there is an ongoing transmission. If the medium is busy, the device must defer the transmission and perform a random backoff. A device must decrement the backoff counter only while the medium is idle.

3.2 Beacon Transmission

In an infrastructure basic service set (BSS) the AP must periodically broadcast beacon frames. APs must schedule beacons for transmission every beacon interval (BI), which is a configurable parameter. Beacon frames include the BI. Like other frames, a beacon transmission follows the MAC procedure, meaning that the actual and nominal BI might differ. Moreover, the actual beacon interval may also vary over time according to the fluctuations of the channel conditions and of the offered load. APs must schedule subsequent beacons at the undelayed nominal interval [3].

Figure 1 illustrates two possible cases: in the first scenario the medium is idle and the nominal and actual BI are the same. In the second scenario, with a busy channel, beacons are delayed due to frames transmitted by other devices, resulting in a difference between the nominal and actual BI.

3.3 Channel Saturation

It is well known that in the random access CSMA/CA MAC protocol, as the offered load increases, the aggregate throughput (i.e., the sum of the throughput of all the transmitters in the same collision domain) increases until the offered load reaches its maximum stable value, often called the *saturation throughput* [2]. Further increases in the offered load result in the same (or lower) aggregate throughput. In this paper, we use the term *saturated channel* to indicate that the offered load is greater than or equal to the saturation throughput.

4 BEACON JITTER

Beacon jitter, δ , is defined as the difference between the nominal, n , and the actual, e , beacon interval:

$$\delta_i = e_i - n. \quad (1)$$

The actual beacon interval is defined as the time elapsed between two consecutive beacons, b_i and b_{i+1} , where t_i and t_{i+1} represent the timestamp corresponding to the beacons i and $i+1$, respectively:

$$e_i = t_{i+1} - t_i. \quad (2)$$




Figure 2: Three possible cases for beacon jitter: 1) b_i and b_{i+1} on time; 2) b_i on time and b_{i+1} late; 3) b_i late and b_{i+1} on time.

We study the distribution of the beacon jitter distribution as a proxy for detecting a saturated channel. Figure 2 shows the three possible cases for the beacon jitter:

- (1) $\delta_i = 0$ indicates no jitter. This occurs when the two beacons experienced the same delay during the channel access; this usually happens when the channel load is low.
- (2) $\delta_i > 0$ means that the delay of b_{i+1} is larger than the delay of b_i .
- (3) $\delta_i < 0$ means that the jitter of b_i is larger than the jitter of b_{i+1} . This happens because APs must transmit their beacons at the nominal time, regardless of previously delayed beacons.

5 METHOD FOR IDENTIFYING CHANNEL SATURATION

As we will demonstrate in our experiments in Section 7.1, the distribution of the beacon jitter is similar whenever the channel is saturated. Based on this empirical observation, we propose the following passive method for identifying whether a channel is saturated: we first compute the beacon jitter distribution, from a given AP, operating in a given channel, c , then we compare this distribution with a reference one, r , corresponding to a saturated channel. We use the Kolmogorov-Smirnov (KS) test [10] for the comparison. The KS test is based on the maximum distance between two distributions, so that similar distributions have a small KS value. Thus, if the KS value resulting from the comparison of c and r is below a given threshold α , we conclude that the channel is saturated as the two distributions are sufficiently close to each other. Algorithm 1 details all the steps of the proposed method.

6 EXPERIMENTS

We conducted two set of experiments to gather data to analyze the distribution of the beacon jitter in real IEEE 802.11 networks. We conducted the first set of experiments in a dedicated testbed, where we could control all the input parameters of the experiments. We refer to these experiments as *controlled condition experiments* in the rest of the paper. Using these experiments, we verify that the distribution of the beacon jitter does vary in a consistent way based on the channel load. We carried out the second set of experiments

Algorithm 1 Channel classification

```

1:  $\alpha \leftarrow$  threshold
2:  $r \leftarrow$  beacon jitter distribution for the reference experiment
3:  $\delta \leftarrow$  beacon jitter distribution for the target experiment
4:  $d \leftarrow$  KS( $\delta, r$ )
5: if  $d < \alpha$  then
6:   Saturated
7: else
8:   Non-saturated
9: end if

```

using a different testbed that shared a channel with a production IEEE 802.11 network. We refer to this second set of experiments as *uncontrolled condition experiments*. Using these experiments, we validate the method described in Section 5. Table 1 summarizes the configuration of all the experiments.

6.1 Controlled Condition Experiments

The controlled condition experiments consist of six APs and six STAs. STAs are Linux (Ubuntu 14.04) computers equipped with network interface cards (NICs) BCM4360. APs are Asus RT-AC87U. Two servers are connected via an Ethernet cable to the six APs, three APs per server. All STAs and servers are further connected to an Ethernet control network that is used to coordinate the experiments. Devices are located in two rooms, with three APs, and three STAs in each room (one STA next to its corresponding AP). The received signal strength indicator (RSSI) detected by each STA was above -70 dBm, ensuring that all STAs are within range of each other. All devices operate using IEEE 802.11ac on channel 40 of the 5 GHz band. During the execution of the experiments there were no other devices using this channel.

Table 1 summarizes the configurations for the experiments. APs transmit UDP traffic at a constant rate to the STAs. We vary the number of transmitter APs (between two and five) and the offered load per transmitter (between 0 Mbps and 4000 Mbps). In each experiment the offered load is the same for all the active transmitters. Each experiment lasts 60 s and is repeated ten times. In all configurations there is only one STA per AP. One of the six APs is idle to avoid biases due to the hardware load; beacons used in our analysis come from this AP. One of the six STAs operates in monitoring mode; this STA captures the frames used our study.

6.2 Uncontrolled Condition Experiments

The uncontrolled condition experiments consist of five APs and five STAs in a single room. APs are Alix-2d2, STAs are Linux (Debian 8.7) computers equipped with IEEE 802.11g NICs. One server connects all ten devices via an Ethernet network. All devices operate on channel 1 of the 2.4 GHz band using IEEE 802.11g. An unknown number of other networks and users also operate in the surroundings on the same and adjacent channels.

These experiments are used to evaluate networks under real and uncontrolled conditions, as networks in the vicinity serve regular users. Table 1 summarizes the configurations, with a different number of transmitters (two or four) and a different offered load

Table 1: Total offered load (e.g., the sum of over all transmitting APs) of the performed experiments.

Number of Transmitters	Controlled conditions		Uncontrolled conditions
	Non-Saturated (Mbps)	Saturated (Mbps)	Saturated (Mbps)
2	0, 0.2 0.5, 1, 1.5, 2, 4, 6, 8, 10, 20, 30, 40, 50, 60, 80, 100, 120, 160, 200	1600	40, 60, 80, 100, 120
3	0, 0.3, 0.75, 1.5, 2.25, 3, 6, 9, 12, 15, 30, 45, 60, 75, 90, 120, 150, 180, 240, 300	2400	
4	0, 0.4, 1, 2, 3, 4, 8, 12, 16, 20, 40, 60, 80, 100, 120, 160	320, 3200	30, 60, 90, 120, 150, 180
5	0, 0.5, 1.25, 2.5, 3.75, 5, 10, 15, 20, 25, 50, 75, 100, 125, 150, 200, 400, 4000	4500, 8000, 20000	

(between 30 Mbps and 60 Mbps). As with the experiments in controlled conditions, the transmitters use *iperf* to generate UDP traffic at a constant bit rate. Each experiment lasts 60 s and is repeated ten times. One of the five APs is idle to avoid biases due to the hardware load; beacons used in our analysis come from this AP. One of the five STAs operates in monitoring mode; this STA captures the frames used in our study.

In the uncontrolled condition experiments we only have access to our five APs and five STAs. Therefore we do not know the actual channel load and thus, we only consider experiments under saturated conditions. We have empirically verified this condition by increasing the total offered load until the aggregate throughput of the nodes under our control did not increase any further.

7 EXPERIMENTAL RESULTS

Figures 3 to 6 show the aggregate throughput and the beacon jitter, δ , for the controlled condition experiments, for two, three, four and five transmitting APs, respectively. The beacon jitter corresponds to the beacons transmitted by the idle AP and measured by the monitoring STA.

Figures 3(a), 4(a), 5(a), and 6(a) show the aggregate throughput vs. the total offered load (i.e., the sum over all transmitting APs). To improve the readability of the plots, the x-axis is broken in multiple segments due to the large range spanned by the values of the total offered load. In each case, at first, the aggregate throughput increases linearly as the total offered load increases. It then flattens as the total offered load approaches the saturation throughput. In the remainder of the paper we consider as saturated all the experiments with an offered load larger or equal to the one that gives the largest aggregate throughput. For example, in the case of five transmitters (Figure 6(a)) the maximum aggregate throughput observed is 161 Mbps, corresponding to a total offered load of 4500 Mbps. Therefore we consider as saturated all the experiments with a total offered load larger than this value. In the case of two and three transmitters the maximum observed aggregate throughput corresponds to the experiments with the largest total offered load (Figures 3(a), 4(a)); in each of these cases we considered as saturated only one experiment, namely the one with largest aggregate throughput (3200 Mbps, 2400 Mbps and 1600 Mbps).

7.1 Beacon Jitter and Channel Load

Figures 3(b), 4(b), 5(b), and 6(b) show box plots of the beacon jitter for the controlled condition experiments. The box plots indicate that the median beacon jitter remains close to zero regardless of the

number of transmitters and the offered load. This is explained by the fact that the jitter can be positive as well as negative. However, as the total offered load increases, the variability of the beacon jitter increases, as more and more beacons arrive late compared to the nominal beacon interval.

Beacon jitter variability increases until the channel reaches the saturation throughput, at which point the beacon jitter distribution remains constant. For example, Figure 6(b), corresponding to experiments with five transmitters, shows that in the non-saturated experiments the interquartile range (IQR) of the beacon jitter increases monotonically, from zero milliseconds when the total offered load is zero to 2.2 ms when the total offered load is 200 Mbps. In the saturated experiments the IQR of the beacon jitter remains around 2.4 ms, regardless of the increases in the total offered load.

In the case of saturated experiments, regardless of the number of transmitters, the beacon jitter distribution is similar. In particular, the IQR is roughly 2.4 ms in each case. This is evident when comparing the last box plots of Figures 3(b) and 4(b), the last two box plots of Figure 5(b) and the last three box plots of Figure 6(b).

Figures 3(c), 4(c), 5(c), and 6(c) show a different representation of the jitter distribution: each segment is the fraction of beacon jitter smaller than $7 \mu\text{s}$, $12 \mu\text{s}$, $140 \mu\text{s}$, $1500 \mu\text{s}$ and $1300 \mu\text{s}$. Regardless of the number of transmitters, the fraction of beacons whose jitter is less than $7 \mu\text{s}$ decreases as the total offered load increases (from around 80 % for an idle channel to a negligible value for saturated channels). Conversely, the fraction between $1500 \mu\text{s}$ and $1300 \mu\text{s}$ increases as the load increases, from a negligible value for an idle channel to roughly 60 % for a saturated one.

Figure 7 shows the empirical distribution of the absolute actual beacon intervals e , given by (2). The top pane shows the CDF in log-linear scale to highlight the differences in the values between $1 \mu\text{s}$ and 1 ms. The bottom pane shows the complimentary CDF, in log-log scale, to better show the tail of the distribution, i.e., values larger than 1 ms. Solid lines correspond to experiments in a saturated channel, dashed lines represent the remaining experiments. We note that non-saturated experiments are spread on the mid-to-left side of the figure, while saturated experiments are concentrated on the right. Figure 7 shows that, as the offered load in the experiments gets closer to the saturation condition, the beacon jitter distribution moves to the right, and confirms that, after the channel is saturated, the distribution of the beacon jitter does not change significantly even as the total offered load is increased. This can be explained by the fact that, when the channel is saturated, further increasing the offered load only increases the number of frames in the transmitter

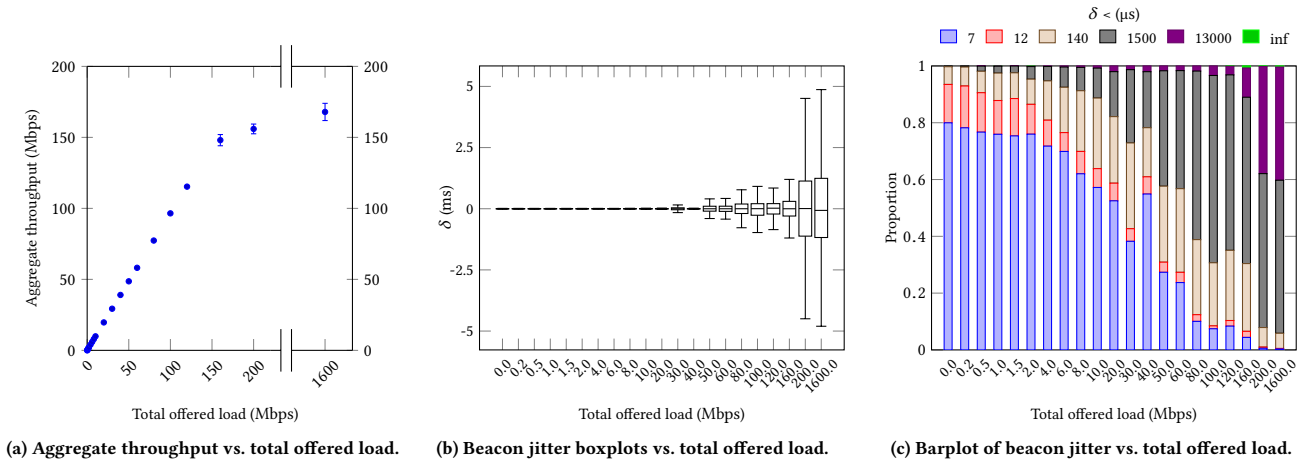


Figure 3: Throughput and beacon jitter results from controlled conditions experiments while using two transmitters.




Figure 4: Throughput and beacon jitter results from controlled conditions experiments while using three transmitters.

queue, without a direct impact on the traffic seen on the channel. Based on these observations, we can distinguish two cases:

- (1) The channel is saturated or close to saturation: beacon jitter distributions are similar to each other.
- (2) The channel is not saturated: beacon jitter distributions vary, but they are different from the distributions typical of saturated channels.

Figure 8 shows the actual beacon intervals for the saturated experiments with controlled conditions together with the saturated experiments with uncontrolled conditions. Note that the solid curves corresponding to the controlled condition experiments are the same in Figures 7 and 8. As previously mentioned, the uncontrolled condition experiments share the channel with a production IEEE 802.11 network, limiting our ability to identify the actual aggregate throughput. Thus, in uncontrolled condition experiments

we only consider saturated experiments. Compared to the unsaturated experiments represented in Figure 7, Figure 8 shows that beacon jitter distributions for saturated uncontrolled condition experiments are similar to each other and are concentrated on the right side of the figure. Additionally, Figure 8 shows that saturated controlled conditions experiments and saturated uncontrolled conditions experiments have similar beacon jitter distributions.

As mentioned in Section 5, we use the KS test to compare beacon jitter distributions. Figure 9 is the graphical representation of a matrix containing the KS value corresponding to all pairs of experiments. Experiments are ordered, in the same way for rows and columns, by increasing normalized aggregate throughput, i.e., the aggregate throughput of an experiment divided by the maximum aggregate throughput for the corresponding number of transmitters¹.

¹For instance, an experiment with four transmitters and an aggregate throughput of 100 Mbps has a normalized (aggregate) throughput of 100/156 Mbps.




Figure 5: Throughput and beacon jitter results from controlled conditions experiments while using four transmitters.




Figure 6: Throughput and beacon jitter results from controlled condition experiments while using five transmitters.

Experiments with low normalized throughput are in the upper-left corner, while experiments with high normalized throughput are in the lower-right corner. The diagonal corresponds to the experiments compared to themselves, therefore the KS value is zero. Areas close to the diagonal show low KS values (e.g., below 0.2), indicating that experiments with a similar normalized throughput have similar beacon jitter distribution. Note that experiments in the lower-right corner also have KS values below 0.2, which are the saturated experiments performed in both controlled and uncontrolled conditions. This confirms that the beacon jitter distribution is similar in all of our saturated experiments.

7.2 Empirical Validation

In this section we assess the performance of the method to classify a channel as saturated, as proposed in Section 5. First, we compute the

KS threshold value α , then, with the selected α value, we assess the sensitivity of our method with respect to the reference distribution.

To pick α we evaluate the performance of the method for values of α in the range $[0, 1]$. During the evaluation we use 10-fold cross-validation and the data obtained in the controlled and uncontrolled conditions experiments. To compare the performance of the different values of α we use the median Matthews Correlation Coefficient (MCC), which varies in the range $[-1, 1]$, where -1 means complete disagreement, 1 means perfect classification, and 0 is no better than random. Table 2 show the values of α that result in the best performance for a given reference distribution.

Figure 10 shows the MCC, precision, and recall scores for the different values of α . The figure corresponds the fist distribution in Table 2. $MCC = 0.73$ indicates that the algorithm performs better than a random classification. The method correctly classified all saturated experiments (recall = 1). From the experiments classified

Table 2: Values of α providing the best performance (highest MCC) with different reference distributions.

	1	2	3	4	5	6	7	8	9	10	11
α	0.22	0.21	0.21	0.23	0.23	0.22	0.22	0.18	0.19	0.17	0.19
MCC	0.73	0.73	0.73	0.73	0.73	0.73	0.73	0.70	0.71	0.73	0.74
Precision	0.66	0.66	0.66	0.66	0.66	0.66	0.66	0.64	0.64	0.67	0.68
Recall	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
	12	13	14	15	16	17	18	Mean	Std	Median	
α	0.17	0.21	0.17	0.19	0.23	0.20	0.20	0.21	0.03	0.21	
MCC	0.74	0.71	0.73	0.72	0.58	0.64	0.66	0.71	0.04	0.73	
Precision	0.68	0.64	0.67	0.66	0.53	0.60	0.62	0.65	0.04	0.66	
Recall	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.00	1.00	




Figure 7: Empirical distribution of the actual beacon interval in the controlled condition experiments.

Table 3: Performance of the proposed method for $\alpha = 0.21$.

	Mean	Std	Median	Min	Max
MCC	0.67	0.07	0.70	0.45	0.73
Precision	0.62	0.04	0.64	0.52	0.66
Recall	0.97	0.06	1.00	0.75	1.00

as saturated, 66 % are saturated experiments ($precision = 0.66$) and 34 % are non-saturated experiments. Note that, out of the false positives (i.e., experiments incorrectly classified as saturated), 38 % are within 90 % of the saturation throughput.

Table 3 gives the performance of the proposed method for $\alpha = 0.21$, which is the median of the α corresponding to the different reference distributions. We note that the median for the MCC, the precision and the recall are 0.70, 0.64 and 1, respectively. Moreover, the standard deviation for MCC, precision and recall are 0.07, 0.04 and 0.06, respectively. This shows that the method is robust with respect to the selection of reference distribution.

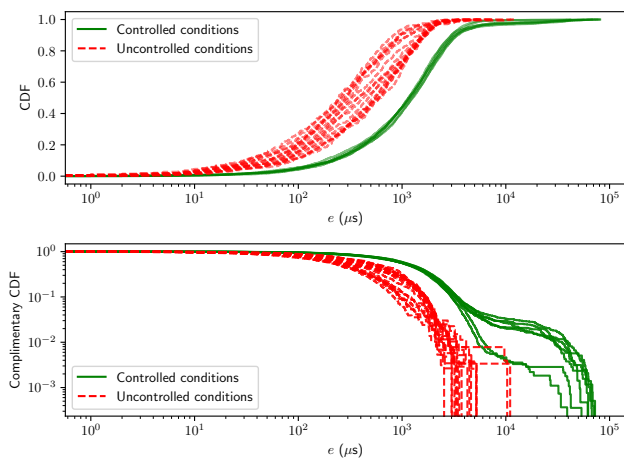


Figure 8: Empirical distribution of the actual beacon interval corresponding to the saturated experiments performed in controlled and uncontrolled conditions.

8 CONCLUSIONS

In this paper we studied the distribution of IEEE 802.11 beacon jitters as a proxy for detecting saturated channels. Using an extensive set of experiments we verified that the distribution of the beacon jitter varies consistently with the aggregate throughput. Specifically, we found that the beacon jitter distribution varies until the channel reaches the saturation throughput, at which point the distribution remains constant. Empirical results show that the beacon jitter distribution is similar and typical for saturated channels. On the contrary, non-saturated channels present a variable beacon jitter distribution that is different from the distribution typical of saturated channels.

Based on the insight gained from our experiments, we proposed a method for non-intrusive classification of Wi-Fi channels. The method takes as input the beacons transmitted by APs and captured by a monitoring STA. It correctly identifies 100 % of the saturated experiments, out of which 34 % are false positives. The experiments also demonstrated that regular Wi-Fi devices can passively collect the input needed, making our method a candidate strategy to be used for network and channel selection.




Figure 9: Graphical representation of the matrix containing the KS values for all pairs of experiments.

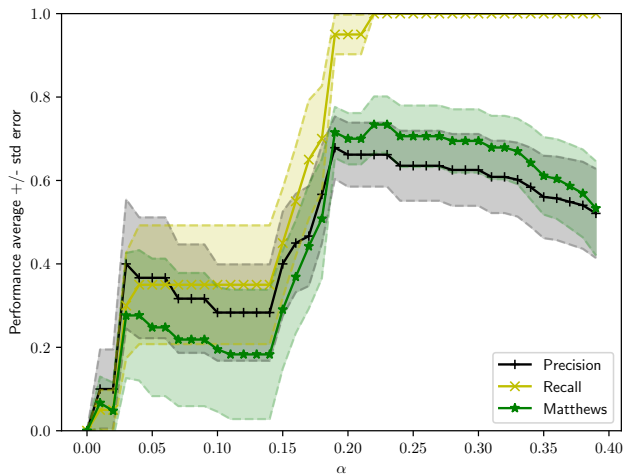


Figure 10: Performance of the channel classification method for different values of the threshold α .

Our current proposal is limited to a binary classification. We plan to extend it by using clustering techniques, to detect patterns in the distributions that could lead to a more fine-grained classification. We also plan to analyze for how long an STA would need to collect beacons in order to produce reliable classification results. This could also lead to a solution that can adapt to the time-varying nature of traffic loads in most Wi-Fi networks.

ACKNOWLEDGMENTS

This work has received a French government support granted to the CominLabs excellence laboratory and managed by the National Research Agency in the “Investing for the Future” program under

reference ANR-10-LABX-07-01. The EIT Doctoral School funded the stay of Laudin Molina at the Institute for Networked Systems (iNETS) at RWTH Aachen University, Germany.

REFERENCES

- [1] A. Achtzehn, L. Simic, M. Petrova, and P. Mahonen. 2015. IEEE 802.11 Wi-Fi access point density estimation with capture-recapture models. In *Int. Conf. Computing, Networking and Communications*. IEEE, 153–159.
- [2] G. Bianchi. 2000. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE J. Sel. Areas Commun.* 18, 3 (2000), 535–547.
- [3] IEEE Commitee. 2016. IEEE Standard for Information Technology-Telecommunications and Information Exchange between Systems Local and Metropolitan Area Networks-Specific Requirements-Part 11. (2016), 3534 pages.
- [4] A. Dhananjay and L. Ruan. 2008. PigWin: Meaningful Load Estimation in IEEE 802.11 Based Wireless LANs. In *IEEE Int. Conf. Comm.* IEEE, 2541–2546.
- [5] E. Haghani, M. N. Krishnan, and A. Zakhor. 2011. A Method for Estimating Access Delay Distribution in IEEE 802.11 Networks. In *2011 Global Telecommunications Conf.* IEEE, 1–6.
- [6] J.-H. Han and S.-J. Han. 2014. Nonintrusive Estimation of Expected Throughput for IEEE 802.11 Devices. In *Consumer Comm. and Networking Conf.* IEEE, 275–280.
- [7] K. Hong, J. P. Kim, M.-S. Kim, and S. Lee. 2016. Channel Measurement-Based Access Point Selection in IEEE 802.11 WLANs. *Pervasive and Mobile Computing* 30 (Aug. 2016), 58–70.
- [8] S. Kajita, H. Yamaguchi, T. Higashino, H. Urayama, M. Yamada, and M. Takai. 2015. Throughput and Delay Estimator for 2.4GHz WiFi APs: A Machine Learning-Based Approach. In *8th IFIP Wireless and Mobile Networking Conf.* IEEE, 223–226.
- [9] Heeyoung Lee, Seongkwan Kim, Okhwan Lee, Sunghyun Choi, and Sung-Ju Lee. 2008. Available Bandwidth-Based Association in IEEE 802.11 Wireless LANs. In *Proceedings of the 11th International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems*. ACM, New York, NY, USA, 132–139.
- [10] Frank J. Massey. 1951. The Kolmogorov-Smirnov Test for Goodness of Fit. *J. of the American Statistical Association* 46, 253 (1951), 68–78.
- [11] L. Mingzhe, M. Claypool, and R. Kinicki. 2008. WBest: A Bandwidth Estimation Tool for IEEE 802.11 Wireless Networks. In *33rd IEEE Conf. Local Comput. Networks*. IEEE, 374–381.
- [12] Q. Pang, S. Liew, and V. Leung. 2005. Design of an effective loss-distinguishable MAC protocol for 802.11 WLAN. *IEEE Comm. Letters* 9, 9 (Sep. 2005), 781–783.
- [13] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, and S. Banerjee. 2008. Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal. In *27th Conf. Computer Comm.* IEEE, 1409 – 1417.
- [14] P. Sapiiezynski, A. Stopczynski, R. Gatej, and S. Lehmann. 2015. Tracking Human Mobility Using WiFi Signals. *PLoS ONE* 10, 7 (July 2015), 1–11.
- [15] S. Vasudevan, K. Papagiannaki, C. Diot, J. Kurose, and D. Towsley. 2005. Facilitating Access Point Selection in IEEE 802.11 Wireless Networks. In *Proc. 5th ACM SIGCOMM Conf. Internet Measurement*. Berkeley, CA, USA, 26–26.
- [16] Hai L. Vu and T. Sakurai. 2006. Collision Probability in Saturated IEEE 802.11 Networks. In *Australian Telecomm. Networks and App. Conf.*
- [17] Y. Yiakoumis, M. Bansal, A. Covington, J. van Reijndam, S. Katti, and N. McKeown. 2015. BeHop: A Testbed for Dense WiFi Networks. *SIGMOBILE Mob. Comput. Commun. Rev.* 18, 3 (Jan. 2015), 71–80.
- [18] J.-H. Yun and S.-W. Seo. 2006. Collision detection based on transmission time information in IEEE 802.11 wireless LAN. In *4th Ann. IEEE Int. Conf. Pervasive Computing and Comm. Workshops*. IEEE, 414–419.
- [19] C. Zhao and C. Hua. 2014. Traffic-Load Aware User Association in Dense Unsaturated Wireless Networks. In *6th Int. Conf. Wireless Comm. and Signal Process.* IEEE, 1–6.