Enabling Wireless LAN Troubleshooting

Ilias Syrigos[1], Stratos Keranidis[1], Thanasis Korakis[2], Constantine Dovrolis[3]

[1] University of Thessaly, Volos, Greece
[2] NYU Polytechnic School of Engineering, New York, USA
[3] Georgia Institute of Technology, Atlanta, USA

Abstract. Particular WLAN pathologies experienced in realistic scenarios are hard to detect, due to the complex nature of the wireless medium. Prior work has employed sophisticated equipment, driver modifications, or even application-layer techniques, towards diagnosing such pathologies. The key novelty of our approach lies in the identification of metrics able to characterize the root causes of individual pathologies, while also being directly extractable from MAC-layer statistics available in today's wireless equipment. Through the development of the proposed framework as application-layer software on top of commercial hardware and its experimental evaluation, we validate the efficiency and applicability of our approach.

1 Introduction

With home WLANs becoming increasingly popular and the plethora of wireless devices operating in the limited unlicensed spectrum, the performance degradation experienced by end-users is almost inevitable. Common home WLAN pathologies are related with low-quality channel conditions. However, even high quality links may suffer from anomalies that are inherent to the operation of the 802.11 standard, such as contention for medium access. In addition, the well-known 802.11 impairments of *"Hidden-Terminal"* [1] and *"Capture-effect"* [2], which are identified in dense topologies, frequently appear in closely spaced WLAN environments.

As administrators/users of home WLANs are not aware of such pathologies, performance issues are usually interpreted incorrectly and the blame is attributed to ISPs. Troubleshooting WLAN performance is hard, due to the complex and dynamic nature of the wireless medium and requires collection of low-level information hardly interpreted even by experts.

Prior work in diagnosing wireless networks performance has considered a variety of approaches, ranging from in-depth studies [3–5] of specific pathologies through sophisticated equipment, to solutions [6] relying on vendor-specific drivers or modifications and application-layer frameworks [7] that are directly applicable to commercial WLAN devices. Considering the different categories of approaches, a tradeoff exists between the achievable detection accuracy and the applicability in common home WLAN setups. Towards bridging this gap, novel frameworks need to be developed that combine the advantages of both worlds.

In this work, we develop user-level detection mechanisms, which exploit low-level information that can be revealed by commercial Access Point (AP) devices. MAC-layer statistics are collected and updated as part of the Physical

layer (PHY) rate adaptation mechanism. These statistics include, but are not limited to, the number of transmission attempts as well as the number of which were successful. The key novelty of our approach lies in the identification of metrics based on the aforementioned statistics that are able to characterize the root causes of WLAN pathologies. Through extensive experimentation, we concluded in the identification of unique trends that performance experiences, in terms of the proposed metrics, when 802.11 links are affected by different WLAN pathologies. Our detailed findings have been incorporated in a combined detection methodology that has been implemented on commercial APs from different vendors. The main outcome of our research is an application-layer framework that is automatically activated upon the detection of degraded performance to accurately determine the underlying pathology and report it to the end user. Our work intends to highlight the importance of having MAC-layer statistics accessible from the application-layer through a standardised way and encourage all manufacturers of 802.11 equipment to adopt this approach.

2 Related Work

A great variety of research approaches has proposed mechanisms towards diagnosing common WLAN pathologies. Several works have focused on the detection of specific pathologies, such as distinguishing between frame losses resulting due to low signal or collisions in [3], or the identification of device types generating cross-technology interference in [5, 4, 8, 9]. Another class of approaches [10, 11] has proposed advanced anomaly detection frameworks that provide increased accuracy by combining measurements obtained from several nodes. In addition, [12, 13, 9] are based on the elaboration of multiple monitoring devices and require the application of synchronization protocols [13, 9], hence rendering them applicable only to centrally managed WLAN deployments. On the other hand, [14, 6] are based solely on observations derived from a single node, thus being applicable in independently owned home WLANs.

The various aforementioned approaches also differ on the specified implementation requirements. More specifically, in [5, 4] the use of sophisticated equipment is necessitated, while the approaches presented in [3, 14, 12, 6] require vendor-specific drivers or modifications. The main drawback of the aforementioned approaches is that they are not hardware agnostic. In [7], the first user-level approach able to infer the MAC-layer effects of common home WLAN anomalies is proposed. In our work, we take a step further by developing a systematic approach able to detect the root causes of an extended list of pathologies, by taking advantage of the detailed information offered by MAC layer statistics, while still being accessible by the application layer.

3 MAC-layer Statistics

Commercial 802.11 devices that are developed by major vendors of wireless products, such as Atheros and Intel can be controlled through well-known Open-Source drivers (ath9k, iwlwifi, Mad-WiFi and ath10k [15]). Such drivers constantly collect detailed MAC-layer statistics that are updated as part of the PHY rate adaptation procedure, including among others, information related to the total number of attempted frame transmissions and retransmissions.

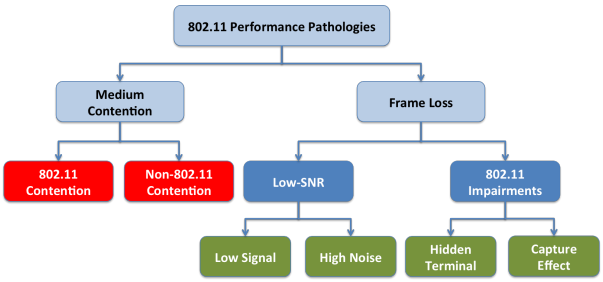

Fig. 1: Taxonomy of IEEE 802.11 Pathologies

In this work, based on this information we define and utilize two metrics. Firstly, the *Normalized Channel Accesses* : $NCA = CA/MCA$ where CA and MCA denote the attempted Channel Accesses and Model-Based Channel Accesses per second respectively, for a specific PHY rate and a specific frame length. MCA is calculated according to the 802.11 a/g performance model presented in [16]. We validated these values with experiments on various types of wireless chipsets under idle channel conditions at the fixed frame length of 1500 bytes. The NCA metric characterises the access (uninterrupted or not) to the wireless medium by a station willing to transmit data frames. Secondly, we define the *Frame Delivery Ratio* : $FDR = ST/CA$ where ST denotes the number of Successful Transmissions per second. The FDR metric is an indicator of the link quality which is responsible for the successful or not delivery of a frame.

In the following Sections, we present how these metrics can be exploited towards characterising the impact of commonly identified pathologies on WLAN performance.

4 IEEE 802.11 Related Pathologies

Perfomance of 802.11 stations first depends on the availability of channel access opportunities and second on the efficiency of frame delivery, whenever medium access is granted. We build our pathology identification mechanism on top of this initial observation and categorize pathologies into two classes. The first one considers pathologies occurring in cases that the transmitter identifies the medium as busy and thus defers from transmitting (**Medium Contention**). In the second category, we group pathologies occuring in cases that the medium is detected as idle, thus enabling the transceiver to proceed with frame transmissions that fail to be delivered at the receiver (**Frame Loss**). We present the taxonomy of the considered pathologies in Fig. 1.

4.1 Medium Contention

Contention-based pathologies frequently occur in dense WLAN deployments, where multiple 802.11 devices concurrently attempt to access the medium. However, as the unlicensed spectrum is also exploited by other wireless protocols (e.g. Bluetooth, Zigbee) and a large range of RF devices (e.g. cordless phones, security cameras), the medium is further congested due to non-802.11 emmissions. The resulting decrease in available channel access opportunities is directly dependent both on the channel airtime captured by 802.11 transmissions, as well as the transmission Duty Cycle (DC) of non-802.11 RF devices.

The crucial impact of medium contention is clearly highlighted in cases of contending stations that utilize diverse PHY rate configurations, which leads to performance anomaly [17]. The high bitrate stations observe a higher throughput degradation in comparison with the lower rate nodes. This degradation is a result of the low number of CA attempts due to the high channel airtime utilization by the low bitrate stations. Consequently, we expect the NCA metric to decrease across increasing PHY rate configurations of the concerned station. However, regarding the FDR metric, higher bitrates should result in higher number of collisions, due to simultaneously expiring back-off timers, and thus to a decrease in FDR, but not in that extent of considering it as a significant trend.

In case of non-802.11 contention, devices with fixed transmission DC, such as microwave ovens, can be interpreted as low bitrate stations which do not comply with the 802.11 standard and hence do not perform a *"Backoff"* procedure. As a result we expect a decrease in NCA metric across increasing PHY rates, as it happens in 802.11 contention. Another consequence of the absence of backoff mechanism in non-802.11 devices is that collisions can occur in the middle of a frame transmission and so higher PHY rates will result in lower probability of collisions. Taking that into consideration, we expect an increasing trend in the FDR metric.

4.2 Frame Loss

In this category, we group pathologies generated in scenarios that the 802.11 *"Channel Sensing"* mechanism constantly identifies the medium to be idle and grants uninterrupted medium access. However, conditions experienced at the receiver side may lead in reduced probability of successful Frame Delivery and subsequent doubling of the Contention Window (CW) parameter. As the reduced MAC-layer Frame Delivery Ratio (FDR) is the root cause of this phenomenon, we identify it as the key metric for characterizing the impact of Frame Loss related pathologies.

Fundamental causes of receiver side underperformance are usually related with the *low-SNR conditions* experienced as a consequence of the low Received Signal strength, resulting from channel fading and shadowing or due to high-Power non-802.11 emissions that result in Noise level increase. Considering that complex modulation schemes require higher link SNR to ensure reliable communication, in comparison with basic schemes, we expect the FDR performance to significantly decrease across increasing PHY rates, under low-SNR conditions. Furthermore, the decrease in FDR would also lead in an decrease in the NCA metric, as the doubling of the CW results in fewer CAs.

In addition, significant frame delivery inefficiencies may also be attributed to *802.11 impairments*, phenomena appearing in cases that concurrent channel access and subsequent frame collisions cannot be avoided through the 802.11 *Channel Sensing* mechanism. More specifically, the *"Hidden-Terminal"* anomaly occurs in cases that the receiver node lies within the transmission range of two active 802.11 nodes that are mutually hidden and cannot sense each other resulting in frame collisions. In cases that no remarkable difference is observed in the received signal strength of colliding frames at the intermediate node, the *"Hidden-Terminal"* phenomenon appears symmetrically for both flows. How-

ever, the most frequently observed case is the *"Capture-effect"* phenomenon, in which case a considerable difference in RSSI values is observed, resulting in a higher probability of successful decoding for the high-power frames. As a result, the link *"capturing"* the medium experiences lower collision probability accessing the medium more frequently and resulting in higher performance penalty for the affected links.

Longer duration transmissions experiencing higher probability of collision, so we expect to see an FDR increase across increasing PHY rate values of the affected link. However, hidden nodes suggest longer distances from the AP and consequently an underlying low-SNR pathology, so we also expect an FDR decrease in high PHY bitrates. In overall, we should identify a highly varying FDR metric accross PHY rates and additionally more notable variations under *"Capture-effect"* scenarios where the impact is more severe. As regards the NCA metric, although the underlying low-SNR conditions should impose a decreasing trend, the impact of FDR variation, which as mentioned before is higher under *"Capture-effect"* scenarios, would enforce NCA to not display a clear trend.

5 Detection Methodology

Having defined the key metrics of NCA and FDR, we next focus on developing a detection methodology able to identify unique trends on the way each individual pathology affects performance of both metrics. Before that, we need to decide upon the existence or not of a pathology. This is accomplished by a simple throughput test of fixed length frames at the maximum PHY rate, the result of which is compared to the analytical value calculated by the aforementioned model in [16]. In case that achieved throughput is lower than the 80% of the theoretical one, we initiate our proposed framework presented below. By taking advantage of the relation between the PHY rate of the affected link and the proposed metrics, we design an active probing mechanism that probes the WLAN channel with multiple packet trains, where each train consists of several packets that are transmitted at varying PHY rates. We call the proposed test as *Varying Bitrate Probing* and each train is transmitted in each one of the 802.11a/g compatible PHY rates, selected from the vector $R = (6, 9, 12, 18, 24, 36, 48, 54)$ Mbps. Each train provides a unique sample - we need multiple samples to make any statistical inference. In parallel with the probing procedure, the NCA and FDR metrics are calculated per each configured PHY rate.

Next, we apply the non-parametric *Theil-Sen Slope* estimator on the collected samples to identify trends in the relation between the two metrics and the PHY rate of the affected link. The output of the *Theil-Sen Slope* estimator consists of the slope estimation with 95% confidence interval, plus the p-value, where both aid in determining the existence of a trend and its characterization as increasing or decreasing. P-values are interpreted as follows: $p < 0.01$ indicates very high significance and $p < 0.05$ is considered significant and the null hypothesis (of the slope being equal to zero) is rejected in both cases. P values greater than 0.05 indicate failure to reject the null hypothesis and thus no trend is detected.

In such occasions, we have to distinguish between two further cases, where in the first case, data points present highly varying values and partially present both

significantly increasing and decreasing trends, while in the second the considered input data is roughly constant and result in approximately zero estimated slope. Although in the first case, no specific trend can be reliably detected, several scenarios might present high start-to-end variation, a trend that we also need to identify. To this aim, we enhance our test, by employing the *Pairwise Difference Test* metric $S_{PDT} = \frac{FDR^8 - FDR^1}{\sum_{k=2}^{8} |FDR^k - FDR^{k-1}|}$, where $k \in \{1, 2, ..., 8\}$ denotes the configured PHY rate. It is obvious that -1 $\leq S_{PDT} \leq$ 1. If there is a strong trend, either increasing or decreasing, S_{PDT} approaches 1 or -1. Identification of the second case is based on the evaluation of statistical dispersion through the measure of standard deviation. We consider specific standard deviation thresholds, as derived from our experimentation and described in the following section.

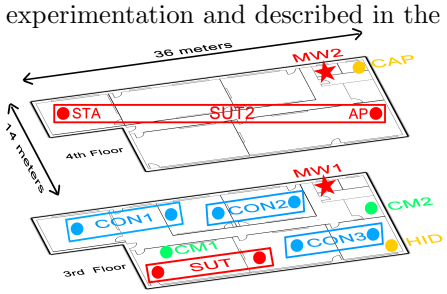

Fig. 2: Testbed Topology

6 Experimentation with Proposed Metrics

The experimental setup that is used as the basis of our experimentation consists of a single communicating pair of nodes that we refer to as System under Test (SUT). Both nodes feature the Intel 5300 chipset, implement the 802.11a/g protocol and operate in infrastructure mode, through the *iwlwifi* driver. In the following experiments, we reproduce each considered pathology and investigate how the performance of the SUT link is affected in terms of the NCA and FDR metrics, while it performs the *Varying Bitrate Probing* test. The devices participating in the following experiments are closely located within a double floor indoor office environment at the University of Thessaly premises, as depicted in Fig. 2. A representative subset of the various executed experiments that replicate each individual pathology is detailed in the following Sections.

6.1 Contention with 802.11 terminals

Through this first experiment, we aim at investigating the impact of medium contention with 802.11 compatible devices, and for this purpose we establish 3 contending links in close proximity and on the same channel with the SUT link. More specifically, we use Ch. 44 of the 5 GHz band that is totally free of transmissions in the testbed premises.

In the first 2 Scenarios, we activate only the CON1 link to transmit 5 Mbps of traffic load, at the PHY rates of 6 Mbps and 54 Mbps accordingly. Fig. 3(a) depicts the NCA performance and shows a significant decreasing trend across increasing PHY rates in Sc.1. due to the *"802.11 performance anomaly"*, while in Sc. 2 only minimal variation is detected across increasing PHY rates, as a result of the high PHY rate. In Sc. 3, we still activate only the CON1 link to

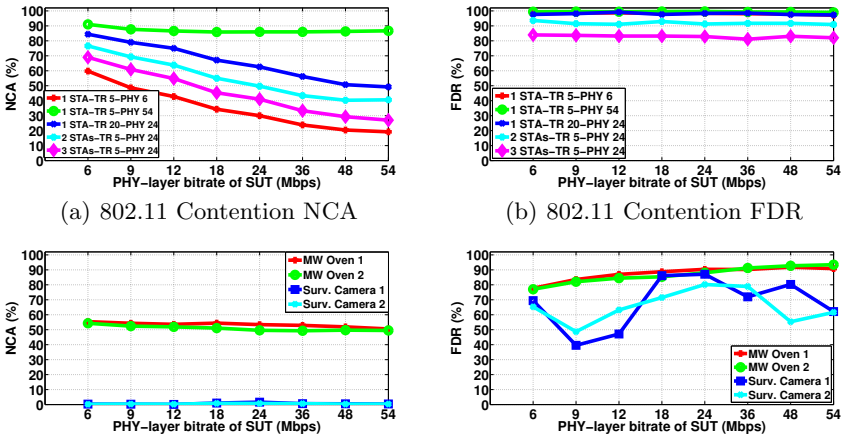

(a) 802.11 Contention NCA (b) 802.11 Contention FDR

(c) non-802.11 Contention NCA (d) non-802.11 Contention FDR

Fig. 3: NCA and FDR Performance of Medium Contention related Pathologies

transmit at the PHY rate of 24 Mbps with 20 Mbps of traffic load. We observe that in Sc. 3, the NCA values per PHY rate have decreased in comparison with Sc. 2, while a significant decreasing trend is clearly identified. Finally, in Sc. 4, we simultaneously activate links CON1 and CON2 to transmit 5 Mbps of traffic load, at the PHY rate of 24 Mbps, while in Sc. 5 we replicate the configurations of Sc. 4, but simultaneously activate the 3 links CON1, CON2 and CON3. In both cases, significant decreasing trends are identified by the Theil-Sen estimator. Summarising the above scenarios, a significant decreasing trend is detected with a p-value of 0.01 except for the Sc. 2 where the p-value of 0.4 is derived and thus no trend is detected. Fig. 3(b) plots the resulting FDR performance across all the considered Scenarios and presents only minimal variation across different PHY rates (standard deviation of 0.95). We notice that the increasing number of contending stations results in decreased FDR, as also observed in [18], a fact related with the increased probability of collisions when the back-off timers of multiple terminals simultaneously expire.

6.2 Contention with non-802.11 devices

In this second experiment, we aim at characterising the impact of different types of non-802.11 devices. More specifically, we consider a Microwave Oven (MW) that typically emits high RF energy in 2.44-2.47 GHz frequencies with DC of 0.5 and a Surveillance Camera that constantly (DC = 1) transmits with 10 dBm power, occupying 18 MHz of bandwidth on various frequencies of the 2.4 GHz band. The two devices are located at positions MW1 and CM1 of the 3rd floor accordingly. We set the SUT link to operate on the commonly configured Ch.6 (2437 MHz) and the camera on 2432 MHz.

In Fig. 3(c), we clearly observe that the continuously emitting Surv. Camera results in NCA values that are close to zero across PHY rates, as the SUT link constantly detects the medium to be busy. On the other hand, the MW that is activated with the DC of 0.5, provides a fixed amount of time available for medium utilisation per period. This phenomenon affects performance in terms

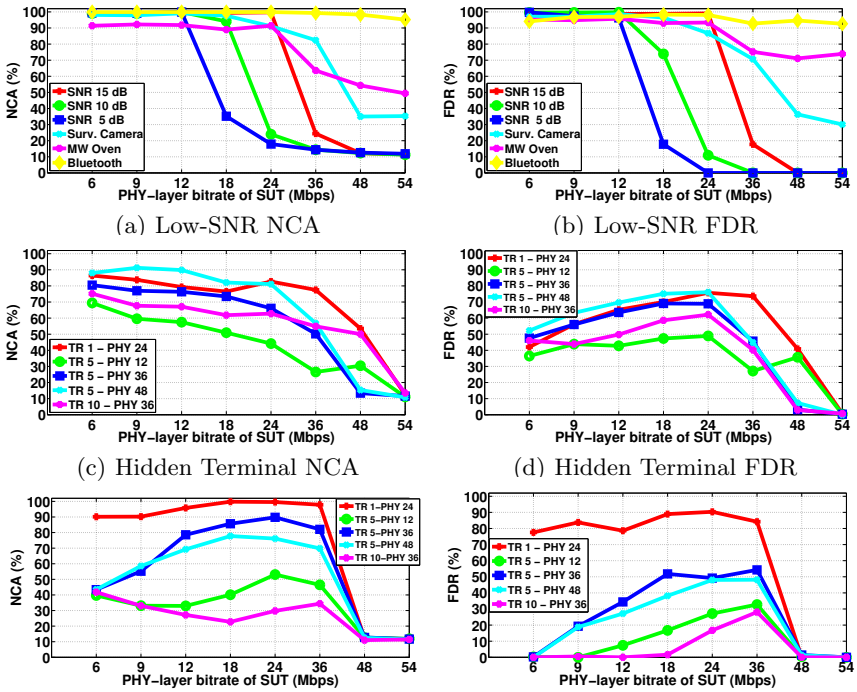

(a) Low-SNR NCA (b) Low-SNR FDR

(c) Hidden Terminal NCA (d) Hidden Terminal FDR

(e) Capture Effect NCA (f) Capture Effect FDR

Fig. 4: NCA and FDR Performance of Frame Loss related Pathologies

of the NCA metric, in a way similar to the *"802.11 performance anomaly"*, thus leading the NCA values to decrease across increasing PHY rates. The resulting FDR performance is depicted in Fig. 3(d), where in the case of the MW, an increasing trend is observed with the p-value of 0.01. Considering the FDR evaluation of the Surv. Camera, no specific trend is identified, as the FDR highly fluctuates due to the extremely low number of attempted transmissions.

6.3 Low SNR

In this experiment, we jointly investigate the impact of low-SNR conditions resulting either in *Low Signal* or *High Noise* scenarios. Considering the *Low Signal* case, we generate varying low-SNR topologies, by establishing a remote 802.11 link (SUT2) on the 4th floor and properly tuning the transmitter's Power, to result in links of 15 dB, 10 dB and 5 dB SNR. In Figures 4(a) and 4(b), we observe that performance regarding NCA and FDR is unaffected as long as the SNR provides for robust communication at the given PHY rate. However, in cases when the SNR requirement of the configured bitrate exceeds the SUT's link SNR, a remarkable drop in NCA and FDR is noticed resulting in a decreasing trend (p-value 0.01) for both metrics. Towards experimenting under High Noise conditions, we use the SUT link of the 3rd floor and activate the MW at position MW2 of the 4th floor, along with the Surv. Camera at position CM2 of the 3rd floor. In both scenarios, the high power emissions of the remotely located interfering devices are not detected to exceed the high Energy Detection threshold at the location of the SUT link. However, due to their high DC, the

Noise level is constantly increased, hence generating low-SNR conditions and approximating the performance obtained in *Low Signal* scenarios, in terms of both metrics.

6.4 Hidden Terminal

Towards experimenting with *"Hidden-Terminal"* scenarios, we establish a hidden to the SUT2 link, by activating the transmitter at position HID of the third floor and resulting in equally received signal strength at both link receivers. In the first scenario we measure the impact of 1 Mbps traffic load, transmitted at the PHY rate of 24 Mbps. As shown in Fig. 4(c), the NCA metric presents a decreasing trend, across increasing PHY rates, while no specific trend is identified for the FDR metric, as depicted in 4(d). More specifically, a small increase across the first PHY rates is followed by a sharp decrease due to the underlying low-SNR pathology. In the next three scenarios, we fix the traffic load at 5 Mbps and vary the PHY rate of the hidden link, between 12, 24 and 48 Mbps respectively. Across all the tested scenarios, identical trends as in Sc. 1 are detected. In Sc. 5, we fix the PHY rate of 36 Mbps and transmit at the traffic load of 10 Mbps, noticing similar performance as in previous scenarios. In all scenarios a decreasing trend is detected regarding the NCA metric with p-value of 0.01, while the null hypothesis regarding FDR is rejected (p-value > 0.05).

6.5 Capture Effect

In this last set of experiments, we examine the performance fluctuations under various scenarios experiencing the *"Capture-effect"* phenomenon. For this purpose, we use the node located at position CAP of the 4th floor, as the interferer of the SUT2 link. In Sc.1, we start by injecting the light traffic load of 1 Mbps in the wireless medium, while configuring the interfering link at the PHY rate of 24 Mbps. As presented in Fig. 4(e) and Fig. 4(f), similar trends are observed as in the considered *"Hidden-Terminal"* scenarios. In Scenarios 2, 3 and 4, we increase the traffic load of the interfering link to 5 Mbps and vary the PHY rate among 12 Mbps, 36 Mbps and 48 Mbps, while in Sc. 5 the traffic load is further increased to 10 Mbps and the PHY rate is fixed at 36 Mbps. Across all the considered scenarios, the NCA metric presents no significant trend, as the calculated p-values lie above 0.05. As regards to FDR, no trend is detected in all Scenarios (p-values above 0.23), although the high performance penalty in comparison to *"Hidden-Terminal"* is depicted with close to zero FDR values. Regarding the performance obtained in Sc. 1, we remark that as both 802.11 impairments pose similar impact on both metrics under low traffic load conditions, discrimination between the two phenomena will be challenging under such cases.

6.6 Framework Enhancement and Result Summary

In many cases of frame loss pathologies, Theil-Sen estimator falsely concludes that the existing pathology is the hidden terminal one. For that reason, the introduction of the aforementioned PDT metric enhances our test with further refinement of the identification of trends. More specifically, in cases where the Theil-Sen estimator detects decreasing trend in NCA attempts and no trend in FDR metric, we apply the PDT metric in FDR statistics. Through extensive experimentation, we concluded that $S_{PDT} <$ -0.8 denotes low-SNR pathology, -0.8 $\leq S_{PDT} \leq$ -0.32 denotes hidden terminal pathology and $S_{PDT} >$ -0.32 denotes

capture effect pathology. The outcome of our study is presented in Table 1, which lists the specific trends that can be detected through the proposed methodology. Not every combination of metrics' trends is mapped to a specific pathology and that cases may correspond to the existence of multiple simultaneous pathologies, which we do not consider in our current work. The derived findings have been incorporated in an application-layer framework that is automatically activated upon the detection of degraded performance to uniquely determine the underlying pathology.

		Frame Delivery Rate (FDR)			
		Constant	No Trend	Increasing	Decreasing
Normalized Channel Accesses (NCA)	Decreasing	802.11 Contention	Hidden Terminal	Non-802.11 Contention DC < 1	Low SNR
	No Trend		Capture Effect		
	Constant		Non-802.11 Contention DC = 1		

Table 1: Identified Trends per considered Pathology

7 Framework Evaluation

In this section, we extensively evaluate the detection performance of the developed framework, under two specifically designed sets of experiments. In each set, we replicate a specific anomaly under various scenarios and by measuring the number of true positives we quantify the perceived detection accuracy. In the first set, we generate scenarios of contention with 802.11 devices, while in the second we experiment with low-SNR scenarios and 802.11 impairments. For the sake of completeness we also examine cases, where our throughput test does not consider as pathologies.

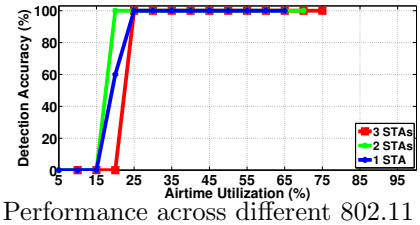

Fig. 5: Detection Performance across different 802.11 contention scenarios.

7.1 Contention with 802.11 terminals

We start by configuring 3 different topologies, consisting of 1, 2 and 3 contending stations that coexist within the 3rd floor of the testbed. We replicate 36 different scenarios in each different topology (108 in total), by varying both the configured PHY rate and traffic load parameters to generate diverse medium utilisation conditions. Through an extra wireless node, we monitor the percentage of Airtime that is captured by the contending links in each topology, towards highlighting the impact of Airtime Utilization on the resulting detection accuracy. In Fig. 5, we clearly observe that the detection performance improves across increasing medium utilisation conditions as the impact of contention is becoming more evident. More specifically, as medium utilisation increases above 25% the mechanism successfully detects the 802.11 contention pathology, across all the corresponding scenarios in the 3 different topologies. Low medium utilisation conditions (below 15%), however, are always detected as no pathology by the initial throughput test and thus are not considered as detection failures.

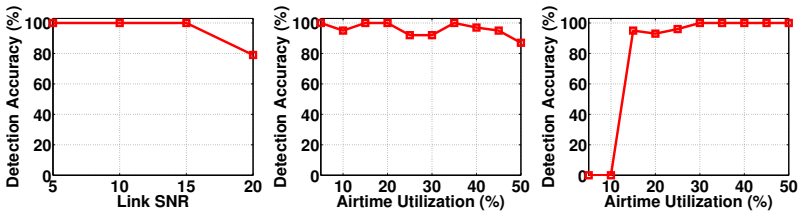

(a) Low-SNR (b) Hidden Terminal (c) Capture Effect

Fig. 6: Detection performance across Frame Loss related pathologies.

7.2 Frame Loss

In this set of experiments we evaluate the detection performance under low-SNR scenarios, by placing the transmitter of the SUT link at the 3rd floor and the receiver across 20 different locations at the 4th floor. We also vary the transmission Power of the SUT transmitter to further vary the SNR levels in each link and result in 80 different topologies. For each different topology, the SUT link executes our detection mechanism, in order to investigate whether the low-SNR conditions are detected. The resulting scenarios are grouped in 5 different SNR classes. We observe in Fig. 6(a) that the detection accuracy is 100% for all SNR classes, except for the 25 dB case, which poses no significant impact and is not detected as pathology from the initial throughput test.

Towards replicating the Hidden-Terminal and Capture Effect phenomena, we activate an interfering link at a fixed position in the 4th floor which is hidden to the transmitter of the SUT. By observing the RSSI values of the transmitted frames, we notice that 4 of the topologies lead to nearly equal (approximately 0-3 dB of difference) values between SUT's transmitter and interfering link's transmitter and consequently are vulnerable to the hidden terminal pathology. Furthermore, 9 links present a notable (>15 dB) difference in RSSI values and hence are vulnerable to the capture effect phenomenon. We evaluate our algorithm in the corresponding topologies that potentially suffer from 802.11 impairments, while inducing traffic of varying load and PHY bitrate at the interfering link and consider 36 different scenarios for each topology. In Fig. 6(b), we observe that the Hidden Terminal pathology is successfully detected across the various tested scenarios that are presented in order of airtime utilised by the Interfering link. In the case of the capture effect pathology, we notice in Fig. 6(c) that the obtained accuracy presents low performance for low airtime utilisation, due to the pathology causing similar impact upon the suffering nodes as the *"Hidden-Terminal"* one does.

8 Conclusions and Future Work

The proposed detection framework of WLAN pathologies causing performance degradation showed encouraging results by accurately detecting all the considered pathologies. Our approach of utilizing the MAC-layer statistics offered from some wireless devices' vendors pointed out the importance of making these accessible, as they are already implemented, to user-level. As we demonstrated this will be of great advantage to WLANs administrators in their effort of troubleshooting low performance. As future work, we seek to integrate our framework

with access points of volunteers, in order to further evaluate the existence of various pathologies in realistic environments. Finally, we aim at the improvement of our framework in terms of detecting all simultaneously existing pathologies, contrary to the current work, which reports just the prevailing one.

9 Acknowledgment

This work was funded by a Google Faculty Research Award.

References

1. L. Feilu, L. Jian, T. Zhifeng, T. Korakis, E. Erkip, and S. Panwar. The Hidden Cost of Hidden Terminals. In *Proc. of ICC*, 2010.
2. J. Lee, W. Kim, S. Lee, D. Jo, J. Ryu, T. Kwon, and Y Choi. An experimental study on the capture effect in 802.11a networks. In *Proc. of WinTECH*, 2007.
3. S. Rayanchu, A Mishra, D. Agrawal, S. Saha, and Suman Banerjee. Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal. In *Proc. of INFOCOM*, 2008.
4. S. Gollakota, F. Adib, D. Katabi, and S. Seshan. Clearing the RF Smog: Making 802.11N Robust to Cross-technology Interference. In *Proc. of SIGCOMM*, 2011.
5. K. Lakshminarayanan, S. Sapra, S. Seshan, and P. Steenkiste. RFDump: An Architecture for Monitoring the Wireless Ether. In *Proc. of CoNEXT*, 2009.
6. K. Lakshminarayanan, S. Seshan, and P. Steenkiste. Understanding 802.11 Performance in Heterogeneous Environments. In *Proc. of HomeNets*, 2011.
7. P. Kanuparthy, C. Dovrolis, K. Papagiannaki, S. Seshan, and P. Steenkiste. Can User-level Probing Detect and Diagnose Common home-WLAN Pathologies. *SIGCOMM CCR*, 42(1), January 2012.
8. S. Rayanchu, A. Patro, and S. Banerjee. Airshark: Detecting non-WiFi RF Devices Using Commodity WiFi Hardware. In *Proc. of IMC*, 2011.
9. S. Rayanchu, A. Patro, and S. Banerjee. Catching Whales and Minnows Using WiFiNet: Deconstructing Non-WiFi Interference Using WiFi Hardware. In *Proc. of NSDI*, 2012.
10. Y. Cheng, J. Bellardo, P. Benkö, A. Snoeren, G. Voelker, and S. Savage. Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis. In *Proc. of SIGCOMM*, 2006.
11. R. Chandra, V. Padmanabhan, and M. Zhang. WiFiProfiler: Cooperative Diagnosis in Wireless LANs. In *Proc. of MobiSys*, 2006.
12. A. Sheth, C. Doerr, D. Grunwald, R. Han, and D. Sicker. MOJO: A Distributed Physical Layer Anomaly Detection System for 802.11 WLANs. In *Proc. of MobiSys*, 2006.
13. V. Shrivastava, S. Rayanchu, S. Banerjee, and K. Papagiannaki. PIE in the Sky: Online Passive Interference Estimation for Enterprise WLANs. In *Proc. of NSDI*, 2011.
14. D. Giustiniano, D. Malone, D. Leith, and K. Papagiannaki. Measuring Transmission Opportunities in 802.11 Links. *IEEE/ACM Trans. Netw.*, 18(5), oct 2010.
15. *"Wireless Chipsets Drivers"*, http://wireless.kernel.org/en/users/Drivers.
16. J. Jangeun, P. Peddabachagari, and M. Sichitiu. Theoretical maximum throughput of IEEE 802.11 and its applications. In *Proc. of NCA*, 2003.
17. M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance anomaly of 802.11b. In *Proc. of INFOCOM*, 2003.
18. G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE JSAC*, 18(3):535–547, March 2000.