

Assessing the Fidelity of COTS 802.11 Sniffers

Pablo Serrano

Departamento de Ingeniería de Telemática
Universidad Carlos III de Madrid
Leganés, Madrid, Spain
Email: pablo@it.uc3m.es

Michael Zink, Jim Kurose

Department of Computer Science
University of Massachusetts Amherst
Amherst, MA 01003, USA
Email: {zink,kurose}@cs.umass.edu

Abstract—Recent measurement studies have analyzed WLAN performance by means of wireless sniffers that passively capture transmitted frames. Also, for relatively large (enterprise) WLAN scenarios, previous work has investigated multi-sniffer deployments with devices placed far apart in order to capture all traffic in the network (even frames transmitted simultaneously by different nodes at non-interfering locations). However, for both these single- and multi-sniffer scenarios, little attention has been given to the fidelity of an individual device, i.e., the ability of a given sniffer to capture all frames that could have been captured by a more faithful device. We assess this fidelity (a term we make precise in this paper) by running controlled experiments inside an anechoic chamber and analyzing the similarities and differences between the trace file from the device under study and those of additional “shadow” devices placed in its close proximity. Our results show that fidelity varies significantly across sniffers, both quantitatively and qualitatively, and that performance may also depend on the nature of the experiment under study and on slight changes of the sniffer position.

I. INTRODUCTION

In recent years, a number of efforts have performed measurements studies in 802.11 networks with a variety of goals, ranging from wireless channel characterization in production wireless networks [1] to link interference measurement [2], to assessing the degree of standards conformance of 802.11 interface cards [3]. Recently, there has been an increasing interest in assessing the wireless measurement process itself. Authors have proposed the use of statistical tools to improve radio resource measurements from 802.11k reports [4] and have warned of the risks of proprietary chipset algorithms [5] that may introduce significant bias to the measurement results. However, little attention has been given to the performance of COTS (common off-the-shelf) passive monitoring devices, widely available and used in measurement-based work (such as Atheros chipset-based cards [5] or Aircap devices [6]). While existing work has shown that coarse-scale sniffer placement [1], [7] and data rate [8] can be crucial factors in the ability to sniff wireless traffic, an implicit assumption has been that a sniffer itself does not introduce significant measurement error.

In this paper, we experimentally investigate the fidelity (we will give a definition of our notion of fidelity in Section II) of four different passive wireless packet capture devices (sniffers) in an anechoic chamber, in which these devices capture frames being sent between an access point and a client. We characterize the extent to which a sniffer fails to capture (“misses”) a transmitted frame that is captured at another device and

investigate the variability of the missed frame rate within each receiver, the effect of frame transmission rate and frame type (RTS/CTS, DATA, ACK) on frame miss events at the sniffers, and the extent to which this rate is location-dependent. We find that there can be considerable variation in the average rate of missed frames among the diverse capture devices, and show (via hypothesis testing) that missed packet events among receivers are essentially independent. Together, these results provide valuable insights into the accuracy of frame capture and loss measurements reported by wireless packet sniffers.

The use of the anechoic chamber enables the assessment of fidelity in four ways: *i*) It provides an interference-free environment in which to run highly controlled experiments. In an anechoic chamber we can guarantee that only one station is transmitting at each time, preventing capture effects that could introduce significant bias in the measurements¹. *ii*) It also prevents interference from sources far away or in neighboring channels that could be received only by some devices (e.g., due to different antennae sensitivity) or that could affect some sniffer’s performance by reducing the SINR level. *iii*) It supports the repeatability of the measurements made, unlike many wireless measurement studies that, for example, run experiments during nighttime in the hope that interference will have a negligible impact. *iv*) Lastly, these “ideal conditions” provide quantitative values for fidelity that can be considered as “best case” scenarios. We expect that, in non-interference-free scenarios, the performance of the same device under the same experiments will be at most as good as the provided by our assessment.

The remainder of this paper is structured as follows. In Section II, we present our definition for fidelity and describe the measurement process, that is validated by means of extensive experimental measurements. We next analyze the statistical characteristics of the loss processes in Section III. In Section IV, we analyze the extent to which fidelity is affected by the physical placement of devices and the experiment under study. Related work in the area of wireless network measurements is presented in Section V. Finally, a summary of the results with the main conclusions from our study is given in Section VI.

¹Contrary to wired transmissions, on wireless channels when two or more frames are transmitted simultaneously it could be the case that, if the ratio between received powers is large enough, one of the frames is successfully received.

II. MEASUREMENT METHODOLOGY AND VALIDATION

In spite of the relatively large number of wireless network measurement studies to date (see Section V), the accuracy of the individual packet capture devices (sniffers) themselves has not been considered. If the measurements had been made by a different sniffing device, how different would the measurements have been? If the sniffer location had been moved only slightly (e.g., say 8 inches to the left or right) would the measurements have changed? Ideally, we would like to compare the performance of a given sniffer against that of a “perfect” reference sniffer; but of course no such “perfect” device exists. Moreover, because of the complex, time- and space-varying nature of the wireless channel, identifying packets that *should* be captured at the sniffer is itself difficult. If a packet is sent by one node and received at another but missed at the sniffer, would that packet have been captured by a “perfect” packet sniffer, or did physical channel characteristics prevent the packet from being receivable at that point in space, at that point in time? The fact that many wireless devices do not strictly adhere to the 802.11 standard [3] makes it even more difficult to infer which packets should have been received.

We address this challenge by approximating a “perfect” reference sniffer by a set of collocated, multiple packet sniffing devices within a very small region. Our intuition is that while one device may occasionally miss frames that were present in the area to be captured, the unioned set of received packets from two such collocated devices will contain even fewer missed frames, and three devices even less, and so on. We make this intuition precise, and discuss our measurement study for determining the number of collocated devices needed to well-approximate the “perfect” device shortly in Section II-B.

Suppose now that we determine that N collocated sniffers are enough to capture all frames and that we run a measurement experiment in which frames are captured at all N sniffers and stored in (individual) trace files. We define the *fidelity* of an individual packet sniffer to be the total number of frames captured in the trace file from that sniffer over the total number of frames that were captured by at least one of the N sniffers, i.e.,

$$F(i) = \frac{|T_i|}{\left| \bigcup_{j=1}^N T_j \right|} \quad (1)$$

where T_i is the trace file from device i , and all trace files start and stop at the same point in time.

Next, after we describe our testbed in Section II-A, we analyze how many devices are enough to capture all frames in the area under study. This is done in two steps:

- First, we measure the gain in terms of frames captured when adding a device from N to $N + 1$ sniffers, and assess when this gain is negligible (Section II-B).
- Then we analyze if sniffers tend to miss the same frames, or if instead missed packet events among receivers are essentially independent (Section II-C).

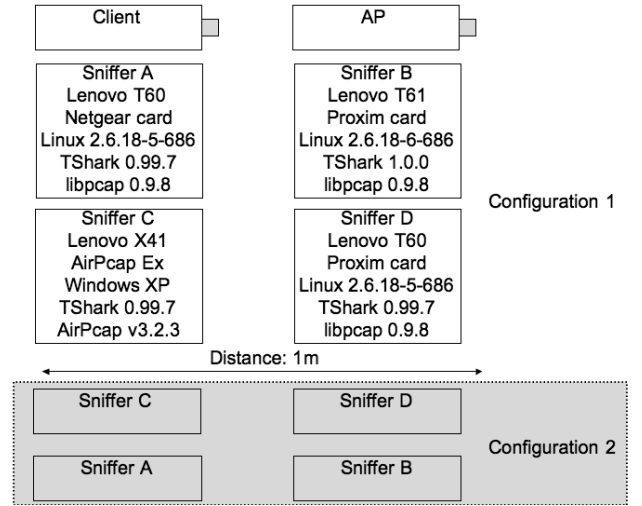


Fig. 1. Setup for the measurements

A. Testbed Description

We deployed our testbed in an anechoic chamber where no frame transmissions except those from our machines could be detected. Our testbed consists of 6 laptops. Two of these machines are used to set up the controlled experiment, one serving as an 802.11 Access Point, the other one as a client. Both of these laptops have Netgear WG511T PC cards as wireless adapters installed. The other four laptops serve as passive wireless sniffers (i.e., they do not send any frames), three running Linux (sniffers A, B, and D) and the fourth running Windows XP (sniffer C). All Linux-based sniffers use PCMCIA wireless cards. Sniffer A uses a Netgear WG511T PC card while sniffer B and D have a Proxim Orinoco 11b/g PC card installed. The Windows-based sniffer makes use of CACE’s AirPcap EX which is a USB 2.0 adapter. All 5 Linux-based laptops use madwifi 0.9.2 as driver for the wireless adapter, while for the Windows-based laptop the Airpcap v3.2.3 driver is used.

We make use of the Iperf² bandwidth measurement tool to send packets between the AP and the client. The packet capturing process is performed with tshark³ on the four sniffers, capturing the first 100 bytes of each frame, including all frame headers. One additional desktop machine, not shown in the Figure, serves as a controller for the entire experimental setup. All machines are connected via wired Ethernet. This allows for automated execution of the experiments and ensures that the control commands sent from the controller machine do not interfere with the wireless measurements.

Figure 1 also shows the relative physical placement of the sniffers in the measurement setup. In Section IV-A, we investigate the effect of sniffer placement on fidelity. We considered two different placement scenarios, which are depicted as Configuration 1 and 2 in Figure 1. In both scenarios we use the same set of physical locations, while changing the laptop

²<http://dast.nlanr.net/Projects/Iperf>

³<http://www.wireshark.org/docs/man-pages/tshark.html>

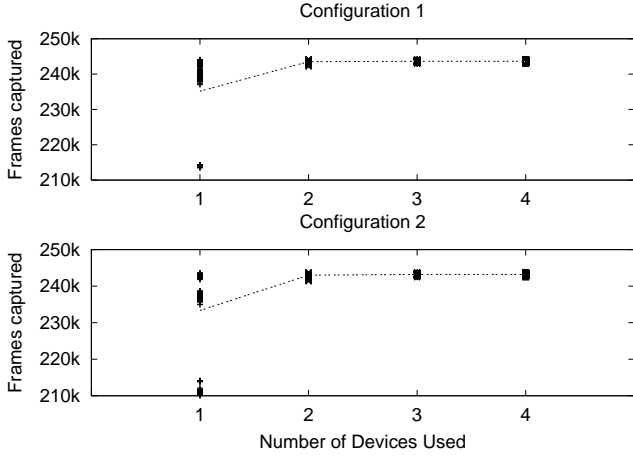


Fig. 2. Number of frames captured for the two configurations

placements in these locations.

B. Number of sniffers to use

The definition of fidelity (Eq. 1) assumes that the N devices collectively capture all the frames on a given location. To validate this assumption and obtain N , we count for the same experiment the total number of frames captured when using an increasing number of devices. When this number is such as there is a negligible gain in the number of captured frames when adding a new device, we claim we have *enough* devices.

More specifically, we proceed as follows. We use the following *base configuration* for our measurements: we set the stations to use the 802.11g at 54 Mbps and send 100 byte frames from the AP to the client at the maximum achievable rate using the RTS/CTS exchange. The experiment runs for 20 seconds and approximately 60k data frames are sent⁴. We repeat the experiment 100 times, computing for each run the total number of frames captured when using one device, two devices, etc. for every combination of devices⁵. We perform these experiments for the two configurations shown in Figure 1.

In Figure 2, we plot for each value of N the total number of frames captured by every combination of tracefiles (points) and its average (lines). We also provide in Table I numerical values for the average number of frames captured when using N devices, along with the improvement over the case of $N - 1$ devices (in parenthesis). The results can be summarized as follows:

- The total number of frames captured by a single device ($N = 1$) varies significantly, with a fidelity of 96–96.5% (when $N = 4$ values are used in the denominator to calculate the fidelity). Furthermore, in some cases the

⁴Because of the RTS/CTS exchange and the use of unicast mode –with acknowledgments–, the total number of frames should be around four times this value, i.e., 240k frames.

⁵To account for the total number of frames captured by more than one device we had to merge tracefiles, similarly to [1]. As the merging relies on accurate delays, first we had to assess the accuracy of timestamps, which we found to be μ s-accurate. For more details, see [9].

TABLE I
AVERAGE NUMBER OF FRAMES CAPTURED FOR DIFFERENT SCHEMES OF DEVICES AND CONFIGURATIONS

N	# Frames (Improvement)	
	Configuration 1	Configuration 2
1	235156.45 (n.a.)	233332.52 (n.a.)
2	243517.12 (3.56%)	243011.35 (4.15%)
3	243629.31 (0.05%)	243205.16 (0.08%)
4	243630.32 (0.0004%)	243207.76 (0.001%)

fidelity drops to 87.5%, this resulting in stand-alone tracefiles of reduced usefulness.

- The use of two sniffers results in a fidelity of approximately 99.9%, missing less than 0.10% of the frames. Note that still this is a relatively large number, as we are not considering a large scale deployment but rather a small, controlled and interference-free environment.
- Three shadow devices ($N = 4$) are sufficient for our purposes here, since the absolute and relative differences between the $N = 3$ and $N = 4$ cases are negligible.

Therefore, even though we have seen that COTS sniffers fail to capture all frames available on a location, we also conclude that there is no need to deploy more than three different devices if one wants a reasonably accurate estimate of the number of frames available at a scenario. However, because we are interested in the performance of the four devices, for the rest of the paper we always use $N = 4$.

C. Independence among sniffers

From the previous section we have seen that there is almost no gain in adding another device when using three sniffers to capture wireless traffic. It could be argued that this is because there is a strong correlation in the loss process across sniffers, so regardless the number of devices used, the same frames will be missed over and over. Our assumption is, on the contrary, that losses of frames that are available for capture in the area are independent between sniffers, so the only reason for the little gain is that almost all frames were already captured.

To reject the hypothesis that sniffers tend to miss the same frames we analyze the correlation of losses among sniffers. To this aim we introduce the *frame loss indicator sequence*⁶, this being a finite sequence that accounts for the frames missed by the device for a given experiment.

Frame loss indicator sequences: Given the union of the N individual tracefiles of size K frames ($K = |\cup_{j=1}^N T_j|$), the frame loss indication sequence (FLIS) for the tracefile T_i is the finite sequence b_i of binary numbers where, for each position k ($k = 1, \dots, K$), ‘0’ represents the k -th frame (f_k) is present on the tracefile, and ‘1’ represents the sniffer missed f_k (i.e., f_k was captured by at least one of the other devices):

$$b_i[k] = \begin{cases} 0 & f_k \in T_i \\ 1 & f_k \notin T_i \end{cases} \quad (2)$$

⁶Similarly to the packet loss indication sequence (PLIS) of [10].

TABLE II
CROSS-CORRELATION ANALYSIS

Sniffer pair	r	$Z(\mu = -0.3)$	$Z(\mu = 0.3)$
(A, B)	0.01046	-3.026	2.822
(A, C)	0.00061	-2.929	2.918
(A, D)	0.00043	-2.928	2.919
(B, C)	0.01176	-3.039	2.809
(B, D)	0.00220	-2.945	2.903
(C, D)	0.00137	-2.937	2.911

The relation of this sequence to the fidelity is straightforward. For the K frames transmitted, the FLIS marks a ‘1’ each time the sniffer misses a frame. Therefore the average of the sequence is the probability of missing a frame, i.e. the complementary of fidelity

$$1 - F(T_i) = \frac{1}{K} \sum_{k=1}^K b_i[k] \quad (3)$$

Using these sequences for each sniffer we are able to analyze the dependence between the loss processes across devices. For each experiment, we obtain the FLIS for every device and compute the (Pearson product-moment) correlation coefficient r for each pair of devices. This correlation coefficient r is defined for two sequences x and y of variables of length n as

$$r_{xy} = \frac{\sum x_i y_i - n \bar{x} \bar{y}}{(n-1) s_x s_y} \quad (4)$$

where \bar{x} (\bar{y}) and s_x (s_y) are sample estimators for the mean and standard deviation of x (y).

We plot in Figure 3 the value of the correlation coefficient for each pair of devices used for the 100 measurements of configuration 1 (we had very similar results for the other configuration). As absolute values for correlation are *small*, we could argue that variables are not strongly correlated. However, to have some statistical confidence in this statement, we apply Fisher’s r to z transform [11], defined as

$$z = \frac{1}{2} \log \frac{1+r}{1-r}, \sigma_z = \frac{1}{\sqrt{M-3}} \quad (5)$$

where M is the number of samples. With the aid of Fisher’s transform, we can run hypothesis tests on the correlation coefficient by means of the *standard score* (or *Z score*). This way, we can test if r is significantly different from $|\rho| = .3^7$, with the results shown on the third and fourth column of Table II. As all results fall outside the critical values of -2.58 and 2.58 , we can reject with 99% confidence the hypothesis that correlation is larger than $|\rho| = .3$ and therefore conclude loss rates are weakly or not correlated.

These results, together with the ones from the previous section, complete the validation of our methodology (the use of Eq. 1) to assess the fidelity of wireless sniffers, because *i*) the loss processes are independent among the diverse sniffers,

⁷This is a somehow arbitrary threshold to distinguish between *small* and *medium* correlated variables, proposed in Cohen, J. (1988), *Statistical power analysis for the behavioral sciences* (2nd ed.) Hillsdale, NJ: Lawrence Erlbaum Associates.

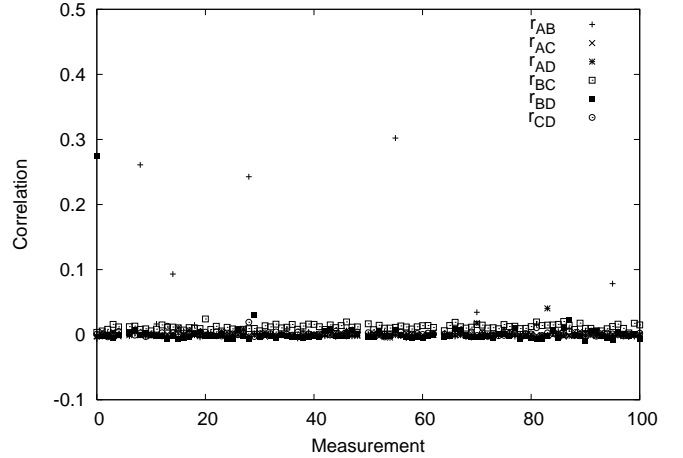


Fig. 3. Correlation between loss processes

which guarantees that adding a device always increases the number of captured frames, and *ii*) the gain from $N = 3$ to $N = 4$ is negligible, we thus use the union of the four tracefiles as in (1) to define the ‘‘ground truth’’ for the frames that are available for capture in the small measurement area. In the following section we analyze the losses seen by individual sniffers.

III. ANALYSIS OF THE LOSSES

Next we analyze the losses from each sniffer under study, with the 100 measurements from the base experiments run with the configuration 1 of Figure 1. In this section, we study the average values and statistical characteristics of the loss sequences from every sniffer. Then, in Section IV we study whether fidelity also depends on the nature of the experiment under study and on changes of the sniffer position.

A. Average Loss Rate

In Figure 4, we plot the average value of the FLIS (i.e., the loss rate) for the 100 measurements made. It is quite clear that there is significant variation in fidelity, both between sniffers and within the same sniffer. We have three easily distinguishable behaviors:

- Sniffer C (the Windows-Aircap device) provides practically the same average value for the loss rate for every measurement. This value is approximately 12.1%, corresponding to some of the values for the $N = 1$ case of Figure 2.
- Sniffer B (a Linux laptop) is around 10 times more faithful than sniffer C with a loss rate of 1.6%. However, the values from this loss process present more variability, ranging from 1% to 2.5% (this is easily seen in the empirical CDF for the loss rate of Figure 8).
- Linux-based sniffers A and D provide high-fidelity tracefiles, as the average losses are 0.08% and 0.15% respectively. On the other hand, the range of their losses is 0.02–0.3% and 0.01–0.6%, leading to much larger deviations than in the previous cases.

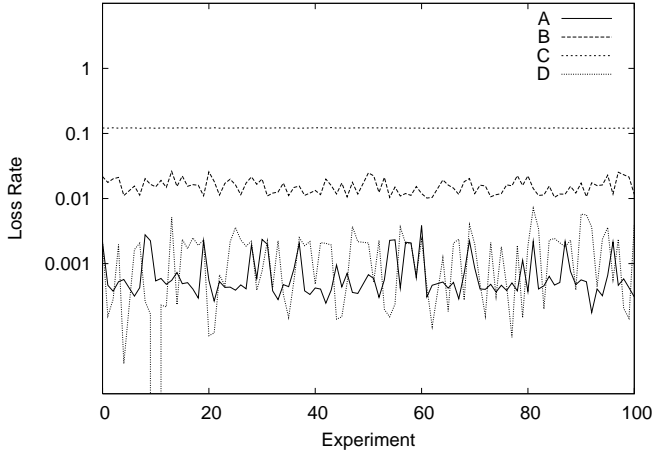


Fig. 4. Average loss rate

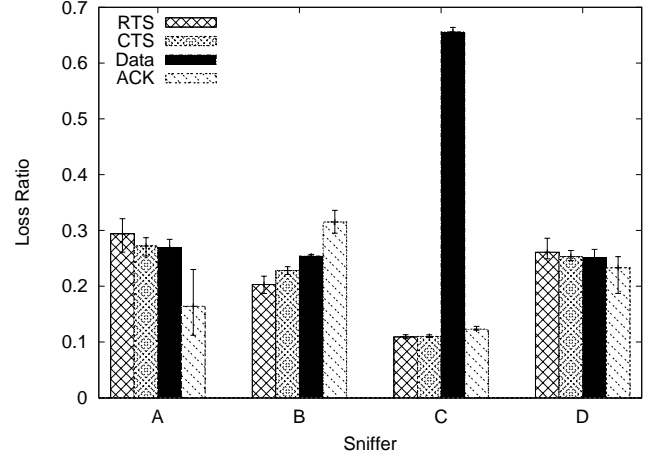


Fig. 5. Relative losses divided by frame type

We have thus seen that diversity in hardware and software leads to quite different performance figures across sniffers, even for the case of Linux-based devices. These differences are apparent not only for the average values, but also for the variability of the loss processes (as suggested by the relative distance between the maximum and minimum loss rate for the same device). In the following, we further analyze the characteristics and differences between the loss sequences.

B. Loss rate conditioned to frame type

In 802.11 WLANs, different types of frames can be sent not only with a different modulation scheme (e.g., a data frame is typically sent with a modulation rate higher than any other frame), but also after a different sequence of events (e.g., ACK frames are guaranteed by the MAC operation to be collision-free in most circumstances, while RTS frames are used to gain access to the medium and therefore can collide). This motivates us to analyze if all frames are equally likely to be lost, or if the loss rate might depend on the frame type.

Given the tracefile T_i from sniffer i , we compute the relative loss ratio for type t frames ($LR_i(t)$) by summing the FLIS for all frames of that type and dividing by the total number of missing frames, i.e.,

$$LR_i(t) = \frac{\sum_{k|type(f_k)=t} b_i[k]}{\sum_K b_i[k]} \quad (6)$$

Note that, in case all frames are equally likely to be lost, this ratio should be 0.25. In Figure 5, we plot the average values of all measurements for these relative ratios with bars, using lines for the .10 and .90 percentile values. The first apparent result is that all sniffers but D are not *fair* in terms of losses, as the chance of being missed significantly varies with the frame type. Given that the transmission order of the frames is typically RTS-CTS-Data-ACK, we can draw the following conclusions:

- Sniffer A is equally likely to miss any of the first three frame types, but acknowledgment frames are more likely to be captured.

- On the contrary, Sniffer B has an increasing bias in the capture process for every type of frame: the later the frame, the more likely is to be missed.
- Sniffer C has the largest bias in the loss process, with the chance of missing a data frame six times larger than for any other frame.

It is worth noting that, despite the fact that both CTS and ACK frames are sent in similar manner (with a lower modulation rate and immediately following a previous frame reception, with guaranteed medium availability), the behavior of sniffers A and B is exactly the opposite: the former tends to capture ACK frames, while the latter tends to miss them. These, along with the case for sniffer C, are quite unexpected results that could introduce a bias into the experimental evaluation of WLAN performance. For example, if the presence of an ACK is taken as an indication of a successful communication (see e.g., [12]), this uneven loss process between frame types could introduce significant bias in the measurement process, particularly in scenarios where most frames are successfully received.

C. Analysis of the loss sequences

Because not all types of frames are equally likely to be missing (as seen in Figure 5), we next analyze the statistical characteristics of the frame loss sequence. We first investigate the empirical cumulative distribution function (CDF) of the loss burst length, where the probability of a burst of length l is given by the number of times exactly l consecutive ‘1’s were found in the tracefile over the total number of bursts, for the 100 measurements. The resulting CDF is plot in Figure 6, showing again quite different behaviors for the sniffers under study:

- Sniffers C and D have similar CDFs, where more than 80% of losses happen in a single frame. However, as we will see next, this does not necessarily imply that the behavior of the loss processes is similar, as the burst length misses some temporal relationship among losses.

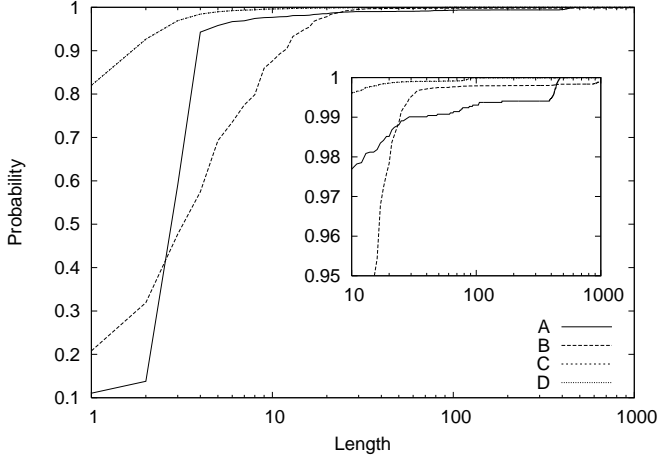


Fig. 6. CDF for the loss burst length

- Despite the similarities between sniffers A and D in Figures 4 and 5, their CDFs are quite different. While more than 80% of sniffer D losses just involve a single frame, for sniffer A around 80% of losses involve four consecutive frames, showing also a long-tail behavior (about 1% of the loss burst are larger than 50 frames).
- The length of the bursts for sniffer B seems to be uniformly distributed between 1 and 30, with 99% of the bursts involving 30 or less frames.

As in the previous section, the different and long-tailed behavior of the loss sequences could lead to wrong conclusions in measurement-based studies, as a typical explanation for long error bursts is path fading or an interference source. However, because we are running our experiments in a small and controlled environment, and due to the lack of correlation between any pair of loss sequences, the results strongly suggest that in this case this is due to the behavior of the measurement device itself.

To further analyze the temporal relationship among losses for a given sniffer we also compute the following estimator of the autocorrelation of the loss sequence

$$R(k) = \frac{1}{(n-k)\sigma^2} \sum_{t=1}^{n-k} (x[t] - \mu)(x[t+k] - \mu) \quad (7)$$

where μ and σ^2 are sample estimators for the mean and standard deviation of x , and plot the average of the 100 measurements in Figure 7. Apparently, there is no temporal correlation among losses at a device except for the Windows device (sniffer C), which presents a highly periodical pattern. Again, because we run our experiments inside an anechoic chamber and no other device presents a similar behavior, we argue that most of these losses should be happening in the path from the antenna to the application layer (through the USB port managed by the Windows XP operating system) and not in the path from the transmitter to the receiver.

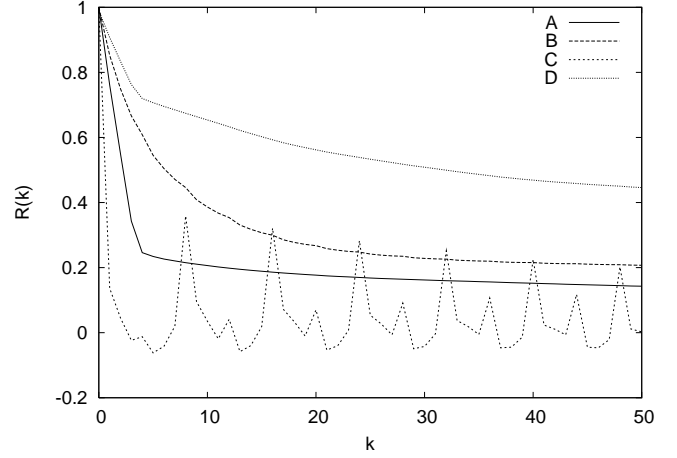


Fig. 7. Autocorrelation of the loss process

IV. CHANGES IN THE SCENARIO

We next consider the extent to which the parameters of the wireless communication under study and the physical deployment have an impact on the fidelity. That is, we are interested to assess if the performance of a wireless sniffer depends on minor changes of its position and on the characteristics of the wireless communication being monitored. We first compare the performance for two different deployments (Section IV-A), and then we analyze the impact of modifying the parameters of the 802.11 communication (Section IV-B).

A. Configuration of the sniffers

To analyze the impact of changes in sniffer placement, we change the position of sniffers to configuration 2 of Figure 1 and repeat the same 100 experiments of Section II-B⁸. We provide, in Table III, the average loss rate for each configuration (μ_1 and μ_2 , respectively). It seems that all sniffers are performing worse than for configuration 1, as all values are larger with the second configuration.

To investigate this quite unexpected result we compute the empirical CDFs of the average loss rate for both scenarios and plot them in Figure 8. Again, the behavior is different for each sniffer. It seems that the fidelity of sniffer D remains the same, while the CDF for sniffer A suffers from a small bias and the other two devices noticeably change their behavior. To gain more detailed statistical information we perform a one-tailed test on the difference between two means, the null hypothesis being the means are the same while the alternative hypothesis is that the second mean is larger. As for sniffers A, B and C the *Z score* is above the critical value of 2.326 (see Table III), we can reject the null hypothesis with 99% confidence and conclude that the mean has changed (this cannot be said for sniffer D).

Next we plot in Figure 9 the average loss rate of sniffers A, B and C for the two configurations (note that each set of 100 measurements was performed in order). We believe

⁸Please remind that in Section II we also did the validation of the measurement process for this second configuration.

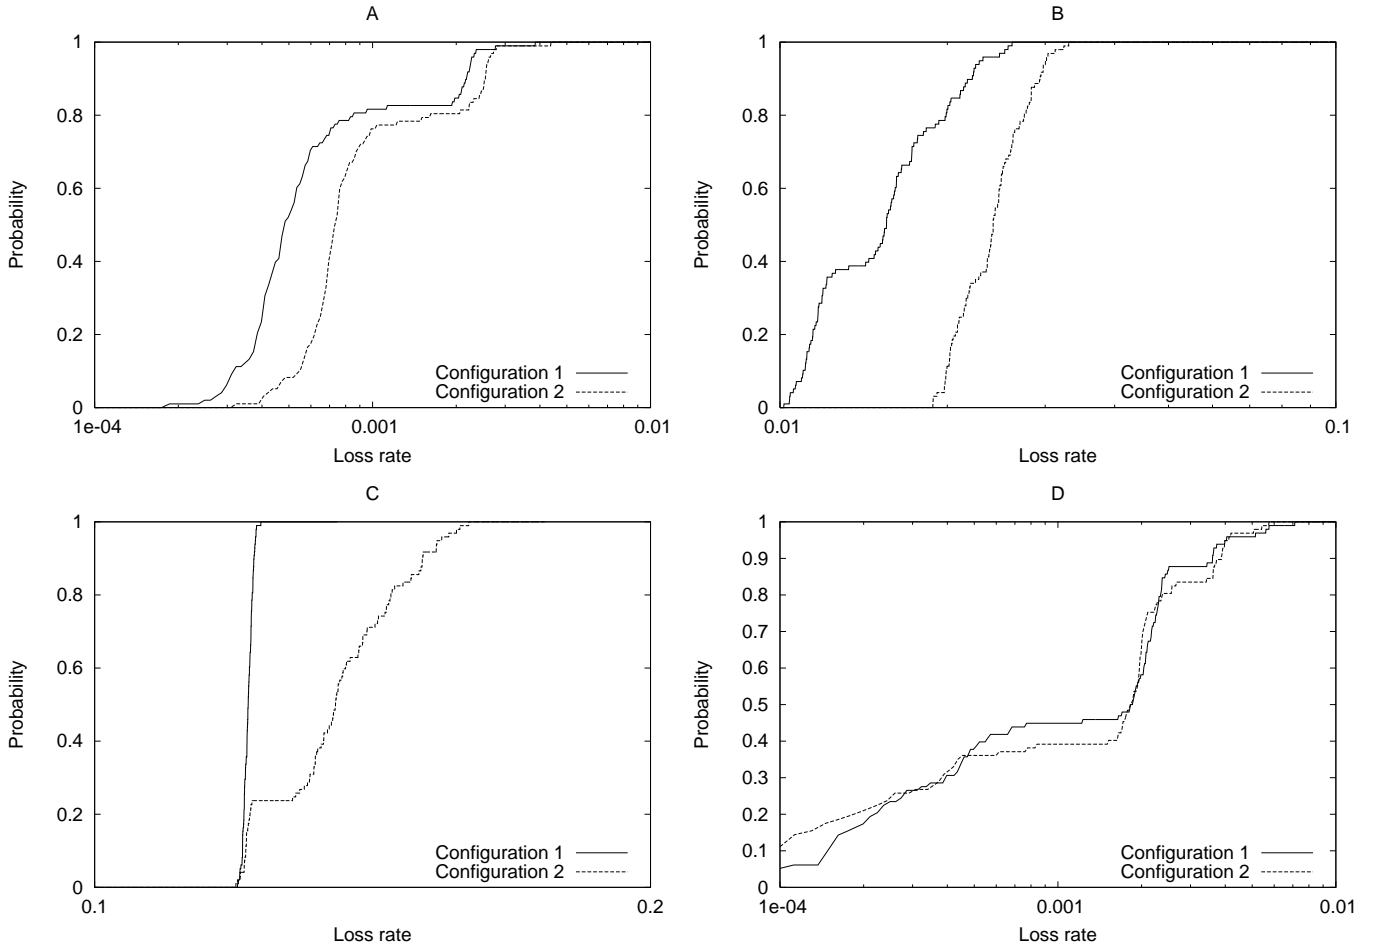


Fig. 8. CDF of the average loss rate for the two configurations

TABLE III
AVERAGE LOSS RATES FOR CONFIGURATIONS 1 AND 2

Device	μ_1	μ_2 (Increment)	$Z(\mu_2 - \mu_1 = 0)$
A	0.00079	0.00108 (36%)	4.02
B	0.01565	0.02428 (55%)	17.68
C	0.12110	0.13534 (12%)	14.72
D	0.00158	0.00166 (5%)	0.86

that the reason for the increase in the loss rate for sniffer C can be explained as follows. Before experiment # 20 this Windows device was providing the same fidelity, i.e., its behavior was not affected by the change of location (like sniffer D). However, from that experiment on, a resource consuming process might have started, degrading the device performance over time. Apart from this particular case, we observe that slight changes in position have an unpredictable impact on performance, as in our case some sniffers do not change their behavior (D) while others do (A and B).

B. Parameters of 802.11

In this section, we analyze if changing the parameters of 802.11 has an impact on sniffer performance. We have seen in Section III-B that, for some sniffers, some frame types are

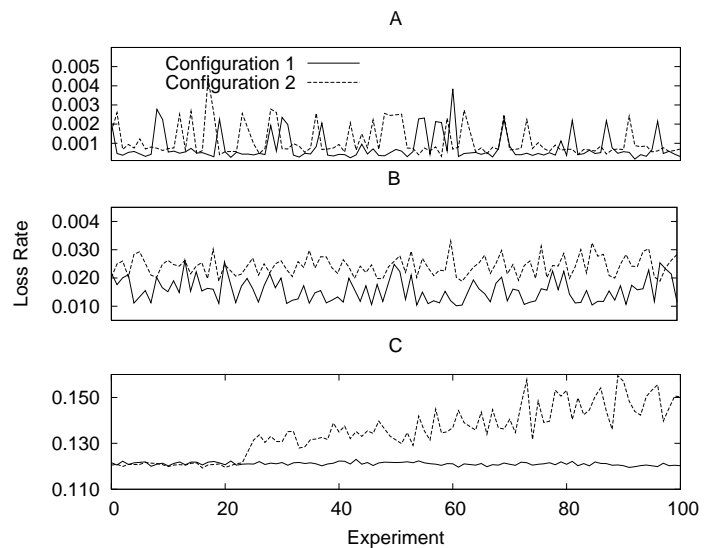


Fig. 9. Average loss rate for sniffers A, B and C in the two configurations

more likely to be missed. Now we want to assess the extent to which changes in the parameters impact sniffer performance. To this aim, we consider the following set of parameters

- RTS/CTS mechanism: on / off.
- Frame size: 100 / 1500 bytes.
- Modulation rate: 6, 9, 12, 18, 24, 36, 48, or 54 Mbps

that results in 32 different cases. We repeat each of these experiments 25 times, changing the duration of each experiment to generate approximately 60k data frames.

In Figure 10 we plot, for each sniffer, the evolution of the average value of the loss rate with the modulation used, for the four possible configurations of the {Frame size, RTS/CTS mechanism} parameters. Note that we use the same values in the y-axis for ease of comparison. Apart from sniffer A that remains oblivious to changes in the configuration parameters, it is clear that the modulation rate used (i.e., the load) has an impact on sniffer performance. This is quite noticeable for the {100 B, RTS/CTS} case, where devices B, C and D increase their loss rate by 10 when changing the modulation rate from 6 to 54 Mbps. On the other hand, for scenarios with low modulation rates, all Linux-based sniffers have loss rates below 0.1%.

The performance of the Windows device (sniffer C) is quite stable in the 6–18 Mbps range, where frame size has a noticeable impact on the loss rate. This result strongly suggests that losses happen in the path from the antenna to the tracefile. Also, for sniffers B and D, the loss rate is always higher for the {100 B, RTS/CTS} case than for any other configuration, which is again inconsistent with the assumption that the sole reason for missing frames is the wireless channel, as larger packets should lead to larger losses.

This relationship between the configuration of the 802.11 parameters and the fidelity of a sniffer can mislead performance evaluation measurements, as an increasing loss rate could be explained because of a low fidelity sniffer rather than a less robust modulation scheme. Furthermore, it is hard to determine the sniffer with the best performance as e.g., sniffer A performs quite stable regardless of the configuration of the 802.11 parameters, but its fidelity at light loads is worse than that of sniffer D.

V. RELATED WORK

Yeo et al. [7] were one of the first to report experiences and pitfalls from sniffing deployments, and advocated for the merging of tracefiles from sniffers placed far apart to obtain the most accurate picture of the behavior of a wireless LAN. However, because data is sent from both the AP and the client in a relatively large scenario, device performance is strongly dependent on relative positions between senders and sniffers. Our work uses a smaller and more careful deployment, with results taken from measurements inside an anechoic chamber, to analyze differences in performance between sniffers placed closed together, providing not only average values of performance but also insightful results about the characteristics of the loss processes.

Jigsaw [1] is a very sophisticated, distributed wireless monitoring infrastructure that is deployed in an office scenario to aid in networking diagnosis. Besides the fact that we also merge data from different sniffers, our goal is to investigate the limits of the measurement process (rather than to perform network diagnosis) using single and multiple sniffers.

In [8], the authors have analyzed single and multiple sniffer performance under extremely high frame rate scenarios, although their focus was on the hardware architecture limits rather than on the measurement process itself (actually, they used wired communication between devices). Schulman et al. [13] investigate the fidelity of 802.11 *packet traces*. Their goal is to measure the completeness of merged and independent wireless network traces and its relation with the load, by basically detecting missing sequence numbers in the tracefiles. However, this inference-based approach cannot identify the reason why a frame is missing (e.g., interference, collision, or lack of fidelity) and its quite limited to data frames.

VI. CONCLUSIONS

In this paper, we have thoroughly analyzed the ability of COTS sniffers to capture frames sent in a wireless communication, what we have defined as fidelity. We have observed that there is always an unpredictable loss at sniffers of frames that are received at other nearby sniffers, a performance limit that restricts the accuracy of measurements derived from tracefiles. Our analysis has shown that fidelity varies significantly across sniffers, both quantitatively and qualitatively, and that it may also depend on the nature of the experiment under study and on slight changes of the sniffer position.

To be able to perform this analysis we have carefully designed a methodology to assess the fidelity of sniffers - one of the main outcomes of this work. We believe that our methodology is the first to introduce a validation procedure, based on i) the completeness of the union of individual tracefiles and ii) the independence among sniffers' loss processes. By analyzing these tracefiles obtained through controlled experiments performed inside an anechoic chamber, we have seen that losses are independent across sniffers, and that three devices are enough to accurately capture a wireless communication (as the gain when adding a fourth sniffer is negligible).

The use of an anechoic chamber ensures the repeatability of the measurements performed. This way, the proposed methodology can be used to assess the performance of other sniffers, and can be extended to analyze the source of packet misses -both tasks constitute part of our future work.

ACKNOWLEDGMENTS

This work was funded in part by the National Science Foundation under grants, EEC-0313747 001, ANI-0325868, and EIA-0080119, and by the Ministry of Education and Science of Spain, under a José Castillejo grant, and POSEIDON project (TSI2006-12507-C03-01). We thank Ian Ricci for his support in setting up the testbed infrastructure and Gene Losik,

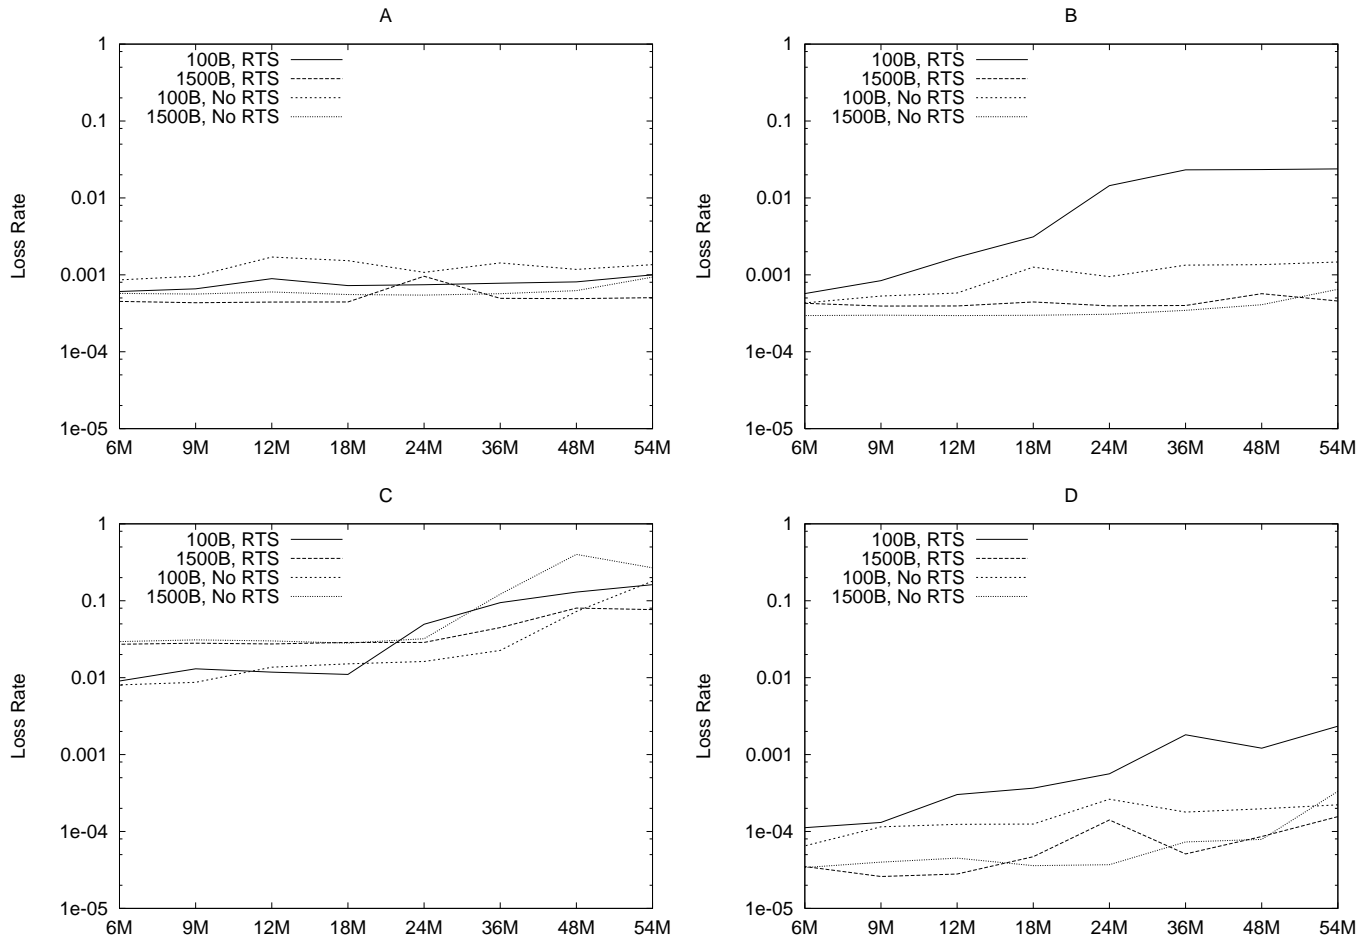


Fig. 10. Impact of the modulation scheme used on the sniffer loss rate for different configurations of the RTS/CTS mechanism and frame size

Chris Merola and Dan Schaubert for providing access to the anechoic chamber.

REFERENCES

- [1] Y.-C. Cheng, J. Bellardo, P. Benkő, A. C. Snoeren, G. M. Voelker, and S. Savage, "Jigsaw: solving the puzzle of enterprise 802.11 analysis," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 39–50, 2006.
- [2] J. Padhye, S. Agarwal, V. N. Padmanabhan, L. Qiu, A. Rao, and B. Zill, "Estimation of link interference in static multi-hop wireless networks," in *Proc. IMC*, 2005.
- [3] G. Bianchi, A. D. Stefano, C. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental assessment of the backoff behavior of commercial ieee 802.11b network cards," in *INFOCOM*. IEEE, 2007, pp. 1181–1189.
- [4] S. Mangold and L. Berlemann, "Ieee 802.11k: improving confidence in radio resource measurements," *Personal, Indoor and Mobile Radio Communications, 2005. PIMRC 2005. IEEE 16th International Symposium on*, vol. 2, pp. 1009–1013 Vol. 2, Sept. 2005.
- [5] D. Giustiniano, G. Bianchi, L. Scalia, and I. Tinnirello, "An explanation for unexpected 802.11 outdoor link-level measurement results," *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pp. 2432–2440, April 2008.
- [6] P. H. J. Perälä, M. Jurvansuu, and J. Prokkola, "Combined terminal and network measurement system for bottleneck localization," in *TridentCom '08: Proceedings of the 4th International Conference on Testbeds and research infrastructures for the development of networks & communities*. ICST, Brussels, Belgium, Belgium: ICST, 2008, pp. 1–7.
- [7] J. Yeo, M. Youssef, T. Henderson, and A. Agrawala, "An accurate technique for measuring the wireless side of wireless networks," in *WiT-MeMo '05: Workshop on Wireless traffic measurements and modeling*, Berkeley, CA, USA, 2005, pp. 13–18.
- [8] M. Portoles, M. Requena, J. Mangues, and M. Cardenete, "Monitoring wireless networks: performance assessment of sniffer architectures," in *Communications, 2006. ICC '06. IEEE International Conference on*, vol. 2, June 2006, pp. 646–651.
- [9] P. Serrano, M. Zink, and J. Kurose, "Assessing the quality of cots 802.11 b/g sniffers," University of Massachusetts, Tech. Rep. UM-CS-2008-13, Apr. 2008.
- [10] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz, "Measurements of a wireless link in an industrial environment using an ieee 802.11-compliant physical layer," *Industrial Electronics, IEEE Transactions on*, vol. 49, no. 6, pp. 1265–1282, Dec 2002.
- [11] P. R. Cohen, *Empirical Methods for Artificial Intelligence*. Cambridge, Massachusetts: MIT Press, 1995.
- [12] D. Giustiniano, D. Malone, D. Leith, and K. Papagiannaki, "Experimental assessment of 802.11 mac layer channel estimators," *Communications Letters, IEEE*, vol. 11, no. 12, pp. 961–963, December 2007.
- [13] A. Schulman, D. Levin, and N. Spring, "On the fidelity of 802.11 packet traces," in *PAM 2008, 9th Passive and Active Measurement conference*, Cleveland, Ohio, April 2008, pp. 132–141.