



HAL
open science

Unity is strength: Improving Wi-Fi passive measurements through sniffer redundancy

Mohammad Imran Syed, Anne Fladenmuller, Marcelo Dias de Amorim

► To cite this version:

Mohammad Imran Syed, Anne Fladenmuller, Marcelo Dias de Amorim. Unity is strength: Improving Wi-Fi passive measurements through sniffer redundancy. *Ad Hoc Networks*, 2023, 151, pp.103287. 10.1016/j.adhoc.2023.103287 . hal-04207233

HAL Id: hal-04207233

<https://hal.science/hal-04207233v1>

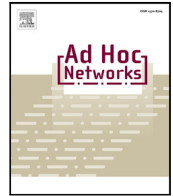
Submitted on 14 Sep 2023

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.



Distributed under a Creative Commons Attribution 4.0 International License



Unity is strength: Improving Wi-Fi passive measurements through sniffer redundancy

Mohammad Imran Syed^{*}, Anne Fladenmuller, Marcelo Dias de Amorim

Sorbonne Université, CNRS, LIP6, 4 Place Jussieu, Paris, 75005, France

ARTICLE INFO

Keywords:

Wireless
Passive measurements
Completeness
Redundancy
Wi-Fi
Experimentation
Traces

ABSTRACT

The passive capture of Wi-Fi traces using sniffers is a promising technique for characterizing the wireless activity of a target area without disturbing users with intrusive measurement tools. The main problem with this technique is that individual sniffers miss packets, which may lead to an inconsistent representation of the scenario. In this paper, we advocate that redundancy (i.e., collocating two or more individual sniffers) is necessary to achieve a reasonable picture of the wireless traffic at the time of the measurement. We formulate the notion of traffic completeness and investigate it experimentally by analyzing Wi-Fi traces obtained by up to 14 co-located sniffers (Raspberry Pi 4) in an office area. We make several observations, including the fact that all individual sniffers capture packets that none of the other sniffers capture. Our results confirm that, depending on the level of completeness that experimentation needs, redundancy is necessary. Moreover, we highlight the importance of characterization of the wireless environment. We define and analyze a few metrics for the characterization.

1. Introduction¹

Countless improvements in wireless network algorithms and protocols come from observations in real deployments [1–4]. Measuring wireless traffic is, however, extremely challenging because of the inherent volatility of wireless links [5]. Active traffic capture is potentially exact, but it requires the deployment of probes at as many nodes as possible. When the scenario involves nodes that do not belong to the same administrative entity, such as an open environment in a city, capturing a significant part of the traffic becomes difficult or even impossible. For example, to gather traffic from smartphones, one can either deploy probes at all access points associated with the user or create a measurement application and request that users install it on their devices. Volunteered users may represent an insufficient population sample, resulting in inaccurate, misleading, and biased results.

An efficient alternative is to perform passive measurements by deploying multiple *sniffers* throughout the target area [6–8].² Sniffers are devices operating in monitor mode that collect wireless packets regardless of the nature of the packets. It is a cost-effective and scalable measurement strategy that does not require users to be bothered by intrusive services. However, due to inherent characteristics of the wireless medium, such as multi-path, fading effects, or collisions, there are

no guarantees that a single sniffer can capture all packets it is exposed to, resulting in incomplete traces. In Fig. 1, we illustrate a typical scenario where four sniffers (s_1, \dots, s_4) do not have the same “view” of the wireless traffic due to detection impairments. A consequence of such an uneven sniffer behavior leads to discrepancies in the measurements, and further analyzes relying on such incomplete traces are likely to be biased or erroneous.

We advocate the use of *super-sniffers* to get around the problem of individual trace incompleteness. It involves adding redundancy to the system by co-locating two or more sniffers to increase the likelihood that at least one of the sniffers captures a packet. The super-sniffer $s_1 - s_2 - s_3 - s_4$ obtains a combined trace in Fig. 1 that presents a better view of the medium. The main question we address in this paper is *how the level of redundancy helps improve the measurement quality*. Assessing the value-add of increasing the redundancy is necessary to help designers balance the quality of the measure against the extra money required for the additional node. To answer this question, we propose a definition of a trace’s *relative completeness*, which gives the fraction of packets that a super-sniffer of a given size captures in comparison with the super-sniffer of maximum size (see Section 4 for the formal definition of “relative completeness”). As we will see in this paper, the

^{*} Corresponding author.

E-mail address: mohammad-imran.syed@lip6.fr (M.I. Syed).

¹ This paper is an extension of our work originally presented in the 2022 International Wireless Communications and Mobile Computing Conference (IWCMC 2022) [18] and in the 2022 IEEE 95th Vehicular Technology Conference: (VTC2022-Spring) [19].

² However, it is important to know which data to sniff while maintaining user privacy.




Fig. 1. Trace completeness. Because of the nature of the wireless medium, sniffers may miss several packets. We need to combine individual traces to get as close as possible to the complete trace.

relative measure of the environment fits well open environment where we do not know exactly what each source transmits.

The classification and characterization of the wireless network and environment is an important factor that can be achieved through passive measurements. It leads to interesting insights because of the unpredictability of the nature of the wireless networks [9,10]. However, the characterization is done with only one sniffer. We define the following metrics for the characterization of the wireless environment in Section 3: (i) access point completeness, (ii) detection of nodes, and (iii) presence of nodes for a shorter duration. We assess these metrics as an application of relative completeness.

We adopt a fully experimental approach. We focus on Wi-Fi traffic, although the methodology is general and can be applied to other technologies. To this end, we have designed and implemented a capture tool to circumvent issues with the individual sniffers, such as clock synchronization and drifts (see Section 5). As sniffers, we use Raspberry Pi 4 devices running on battery as the main module with a specific external Wi-Fi module. We assess the capture quality of individual sniffers and super-sniffers with up to fourteen-redundancy (i.e., a super-sniffer composed of fourteen co-located sniffers). We collected Wi-Fi MAC header traces in a private working environment for 24 h, corresponding to approximately 84 Gbytes of traces.

We make several observations related to completeness that we will discuss in detail in Sections 6–8. Firstly, we observe that the individual sniffers achieve relatively little completeness, which asserts the need for redundancy. We observe that three sniffers perform comparatively poorly based on our analysis of all fourteen individual sniffers, so we clean the dataset to remove those three sniffers from our analysis. Secondly, the symmetric difference of the traces from two individual sniffers is never null (around 20% difference), meaning that combining any two sniffers always brings new information. Even more surprisingly, each individual sniffer captured packets that none of the other sniffers captured. Thirdly, commercial off-the-shelf devices like Raspberry Pi are powerful enough to play the role of a sniffer, even in situations of high wireless activity. Fourthly, sniffers regularly miss packets, highlighting that the environment largely determines the quality of the sniffing process.

We similarly make observations for the characterization of the wireless environment that we explain in Section 9. Since we know the locations of access points in our experimentation system, we observe different levels of completeness for access points as a function of redundancy. The knowledge of the location of access points helps us to detect the presence of nodes in the vicinity of access points. Lastly, we observe that redundancy helps in the detection of more number nodes

that are present for a shorter duration which is critical for measuring mobility.

In summary, the contributions of this paper are:

- **Metric for relative completeness.** We propose a formal definition of completeness that incorporates the notion of redundancy. We formulate the completeness of the super-sniffer, which we evaluate experimentally.
- **Experimental evaluation.** We follow an experimental approach to assess the behavior of the temporal variation in completeness per 5 min over 24 h. We analyze how completeness varies over time.
- **Access point completeness.** We know the locations of access points in our experimentation area and we investigate the completeness they achieve as a function of the size of the super-sniffer.
- **Traffic load.** We investigate the effect of varying traffic loads on the completeness of traces. As the traffic capture duration is 24 h, the amount of traffic varies at different times of the day.
- **Role of sniffers.** We analyze the performance of individual sniffers to highlight the fact that all sniffers do not perform in the same manner and to help us identify the sniffers that perform consistently poorly despite the conditions in the wireless medium changing over 24 h.
- **Detection of nodes.** We detect the presence of nodes in the vicinity of access points based on the knowledge of the location of access points and RSSI values seen by the sniffers.
- **Pairwise completeness.** The sniffers complement each other in improving the quality of passive measurements when considering them in pairs. We call this pairwise completeness and evaluate it for eleven sniffers after cleaning the dataset.
- **Presence of nodes for a shorter duration.** The duration of the presence of a node at a certain location is vital for measuring mobility. We analyze the number of nodes present for a shorter duration and link it with redundancy in the number of sniffers.

The rest of the paper is organized as follows. Section 2 explains the state of the art and how our work stands out. In Section 3, we define the metrics for characterizing a wireless environment. We define and formulate relative completeness in Section 4. Section 5 explains the experimental methodology. In Section 6, we discuss the data collection and examine the characteristics of the data that we collect. We evaluate the performance of individual sniffers and thus, provide practical evidence of the need for redundancy in Section 7. We measure the improvement in the quality of sniffing by the introduction

of redundancy in Section 8. Section 9 provides the analysis for the metrics of wireless environment characterization that we define in Section 3. We discuss our findings in Section 10. We conclude the paper and list some open issues in Section 11.

2. Related work

Xu et al. merge the individual traces into a single and then run an inference procedure to reconstruct the missing [11]. It needs at least one packet of a conversation in a trace to infer the missing packets and its accuracy also depends on the capture percentage. The evaluation is dependent on a simulation where the process removes packets from the trace randomly whereas, we keep the packet with the best RSSI value. Wit is a tool to merge multiple traces and then reconstruct the missing packets by inferring if they were received by the destination by making use of the frames like Association Request and Response [7]. PMSW is a passive monitoring system that relies on sequence numbers to infer the missing packets in a wireless sensor network. However, it only captures data and acknowledgment packets, leading to a complex synchronization solution [12]. There are no conversation, data, or association frames as we rely on probe requests for contact traces. Sammarco et al. rank the traces collected from multiple sniffers based on similarity to determine which traces should be merged to achieve maximum completeness. It decreases the number of merge operations [13].

Schulman et al. estimate the number of missed packets using sequence numbers and re-transmission bit [14] but they do not capture traffic of their own and rely on datasets available on CRAWDDAD (now part of IEEE DataPort) [15], whereas, we collect our own traces. The dependence on the re-transmission bit would create some bias because it is hard to infer how many packets are re-transmitted because they have the same sequence number.

Mahanti et al. examine the beacon and acknowledgment frames, MAC-layer sequence numbers, and placement of the sniffer to address the incomplete traces [16]. They use the results from one sniffer to create a layout of four sniffers on three floors. The amount of packets captured in 24 h is nearly the same as that captured by our sniffers in 10 min. Garcia et al. develop a passive monitoring system called EPMOST which focuses on election to choose the nodes of the target area for their packets to be captured by the sniffers but more in terms of energy consumption which reduces the number of packets captured by 0.62% [6].

LiveNet provides a platform for monitoring and processing passive traces but the transfer of packets to the serial port seems to result in packet loss and the validation is also based on the data measured in a controlled environment [17].

Our work stands distinctive as we focus on redundancy for trace completeness based on real-world experiments in an uncontrolled environment and do an exhaustive temporal analysis over the period of 24 h in an office environment with a varying traffic load in the wireless medium. We remove the poorly performing sniffers from the analysis. Moreover, our solution is more oriented toward contact traces and mobility reconstruction.

3. Wireless passive measurements: How to characterize the wireless environment

Passive measurements have been around for a number of years and have found use in various domains particularly wireless as well as Internet measurements. In the context of wireless measurements, it is interesting to explore the wireless environment. However, wireless passive measurements are always done with a single sniffer placed at a specific location. As we mentioned earlier, a single sniffer is not enough to capture traffic representative enough of the wireless medium as it misses packets because of the inherent characteristics of the wireless

networks. It is, therefore, not possible to explore and understand the characterization of the wireless environment without redundancy.

The characterization of the wireless environment is an important aspect of passive measurements and it is measured in several ways by analyzing the traces and finding trends for the wireless traffic captured. The characterization helps to gain insights into the performance and behavior of the sniffers as well as the environment making use of the traces we capture by passive sniffing. To this end, propose and define several metrics in the following section.

3.1. Metrics for characterization

We perform real-world passive measurements and use the traces to help us characterize the wireless environment. We use three metrics to characterize the wireless environment of our capture: (i) access point completeness (ii) node detection (iii) duration of the presence of nodes. We explain each of these in the subsequent text.

Percentage of packets captured from access points. The access points are present at different distances and even on different floors with respect to the location of the sniffers. The percentage of packets that the sniffers receive from a certain access point leads us to represent this value as a function of redundancy in the number of sniffers.

Detection of nodes. The detection of the presence of a node at a certain distance is an essential factor for localization. It is possible to detect the vicinity of a node if some information about the location of the access points with respect to the sniffer is known.

Presence of nodes for a shorter duration. The duration of the presence of the nodes is a critical factor for measuring mobility as well as localization. The nodes that remain present in the vicinity of a sniffer for a short amount of time, need to be captured with precision to be later able to extract the correct trajectory of the users.

3.2. Fetching the metrics

We fetch the information to evaluate the above-mentioned metrics as follows:

1. We know all the MAC addresses and the locations of the access points present in our experimentation set-up. This information leads to the identification of the presence of packets sent by the access points in our traces. Once we have the traces for each access point, we perform the merging operation to remove the duplicate packets which new information each level of redundancy brings to the percentage of packets captured.
2. We assess the RSSI values of the packets we receive from the access points. Since we know the location of the access points, we categorize the presence of nodes at a certain distance from the sniffers. We utilize the ground truth of the distance of access points and their RSSI values to analyze the RSSI values of other nodes and hence, detect the presence of those nodes at a certain distance.
3. We have the MAC address and time-of-arrival of each packet in our traces. We organize the data to find how many sniffers captured each packet. This helps us to analyze the presence of nodes for a duration of up to 30 s as a function of redundancy.

4. Relative completeness

As we introduce redundancy in the number of sniffers, we coin the term *super-sniffer*. A super-sniffer of redundancy m is composed of m individual sniffers. In Fig. 2, we illustrate a super-sniffer of size three. We have 14 sniffers in our experimental set-up (see Section 5), so the maximum size of the super-sniffer in our experiments is 14.

Before proceeding, let us first define the notion of “relative completeness”:




Fig. 2. A super-sniffer. 3 individual co-located sniffers grouped together form a super-sniffer of size 3.

Definition 4.1 (Relative Completeness). The *relative completeness* is the proportion of packets that an individual sniffer or a super-sniffer capture “relative” to the number of packets captured by the super-sniffer of maximum size (i.e., $m = 14$ in our case).

We formalize it as follows. Let $S = \{s_1, s_2, \dots, s_M\}$ be the set of M sniffers that we have at our disposal to compose a super-sniffer, T_{s_i} be the trace (i.e., set of packets) captured by sniffer $s_i \in S$, and $\mathcal{T} = \{T_{s_1}, T_{s_2}, \dots, T_{s_M}\}$.

We define π^m as a subset of m elements of \mathcal{T} and Π^m be the set of all instances of different combinations of π^m :

$$\Pi^m = \{\pi_1^m, \pi_2^m, \dots, \pi_{\binom{M}{m}}^m\} = \{X = \{x_1, x_2, \dots, x_m\}, x_1, x_2, \dots, x_m \in \mathcal{T}, x_1 \neq x_2 \neq \dots \neq x_m\} \quad (1)$$

where $\binom{M}{m}$ is the number of combinations of super-sniffers of size m that can be built out of M sniffers.

The outcome trace of a super-sniffer is a single trace resulting from the combination of the individual traces of the sniffers composing the super-sniffer. We refer to such a trace as $A^{\pi_i^m}$, i.e., as the union of the traces $\pi_i^m \in \Pi^m, i = 1, 2, \dots, \binom{M}{m}$:

$$A^{\pi_i^m} = T_a \cup T_b \cup \dots \cup T_m, \quad T_a, T_b, \dots, T_m \in \pi_i^m, \quad (2)$$

and

$$T_a \neq T_b \neq \dots \neq T_m. \quad (3)$$

As underlined earlier, the *maximum reachable quality* is obtained when the super-sniffer is M -fold redundant (i.e., it is composed of all M individual sniffers):

$$A_{\max} = A^{\pi^M} = T_{s_1} \cup T_{s_2} \cup \dots \cup T_{s_M}. \quad (4)$$

We need to make two observations now. Firstly, note from Eq. (1) that Π^M has a single element, which is π^M . Therefore, the quality of a capture is denoted by $A^{\pi_i^m}$. The value of this measure quality is obtained by taking its ratio with the result of maximum value when all M sniffers are considered. Secondly, A_{\max} is the best result that we can obtain. That is why we consider it as the reference number to define the “relative” completeness:

$$C(A^{\pi_i^m}) = \frac{|A^{\pi_i^m}|}{|A_{\max}|}. \quad (5)$$

There are multiple super-sniffers of size m , each one resulting from a different combination of m out of M sniffers. Each of the $\binom{M}{m}$ super-sniffers leads to a different value of completeness. We can then define two special cases, which come respectively, from the super-sniffer that leads to the largest completeness and the super-sniffer that leads to the smallest completeness:

$$C_{\max}^m = \max_{i=1,2,\dots,\binom{M}{m}} C(A^{\pi_i^m}) \quad (6)$$

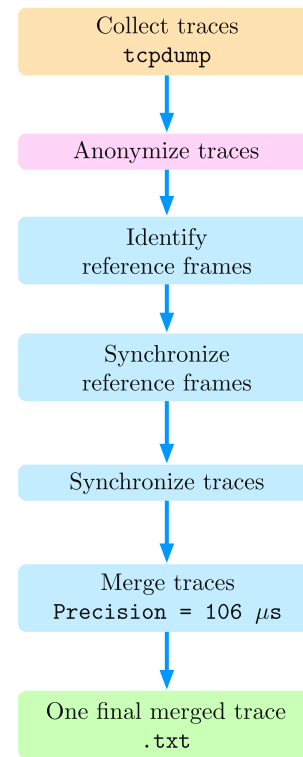


Fig. 3. Experimental methodology. The steps involved in the complete process of collection, privacy protection, processing, and analysis and outcome of the traces.

and

$$C_{\min}^m = \min_{i=1,2,\dots,\binom{M}{m}} C(A^{\pi_i^m}). \quad (7)$$

Number of traces per size of super-sniffer. We need to build traces π_i^m for all combinations of sniffers of different sizes. If we consider $m = 4$, then Π^m in Eq. (1) is equivalent to $\{\pi_1^4, \pi_2^4, \dots, \pi_{\binom{M}{4}}^4\}$ which means that we need to build traces for all combinations of $m = 4$ sniffers out of the total M sniffers. It represents all combinations of sniffer s_1 with combinations of three sniffers other than s_1 , similarly combinations of sniffer s_2 with three sniffers other than s_2 itself, and so on.

5. Experimental methodology

In this section, we explain our experimental methodology. The methodology follows the structure depicted in Fig. 3 (we will explain each step below). In previous companion papers [18,19], we collected traces using fewer sniffers and only for a very limited time (less than one hour). Here, we go much further as we use more sniffers and run the traffic capture for 24 h – we are able, then, to investigate, among others, the effect of the temporal evolution of the wireless traffic on the quality of the capture.

The measurement process follows the flow shown in Fig. 3. We describe each of the steps in the following.

Trace collection. Sniffers run `tcpdump` to collect traces [20]. We configure some filters to gather only the header fields we need for this work (for example, to avoid capturing personal data as discussed below). The outcome of the capture process is one pcap file per individual sniffer.

Trace anonymization. The privacy of the users is a top priority for us. We anonymize the traces by running several protection techniques

on the packets. Firstly, we do not disclose the geographic locations of our measurements. Secondly, we configure the sniffers to capture only the headers of the packets. In our work, we need the header as it brings the necessary information to combine traces from different sniffers. But, since headers contain the MAC addresses of the devices, which are considered personal information, we need to provide extra privacy guarantees. To this end, we hash and then truncate the MAC identifiers of the headers.

Identification of reference frames. We use a combination of different header fields to identify reference frames that are present in both traces: (i) sender's hashed and truncated MAC address, (ii) sequence number, (iii) frame sub-type, (iv) Frame Checksum Sequence (FCS), (v) fragment number, and (vi) timestamp as per the IEEE 802.11 standard [21].

Synchronization of reference frames and traces. The *beacon* and *probe response* frames are the closest representatives of real-time clocks. These frames lay the foundation for the synchronization process. The tool can only synchronize two traces at a time. Therefore, a reference trace and as well as the trace which has to be synchronized is the input to the tool. The first step is to extract the beacon and non-re-transmitted probe response frames from both traces independently. These frames are called *unique frames*. The next step is to extract the unique frames that are common in both traces. The coverage areas of the sniffers capturing these traces must overlap to execute this step. The common frames are referred to as *reference frames*. Next, the timestamps of reference frames are synchronized using *linear regression* over a sliding window of 3 frames. The synchronized reference frames are then used to synchronize the complete trace. The tool provides an additional option of concatenating or merging the synchronized traces.

Trace merging. The principle behind a super-sniffer is its ability to merge traces collected by its individual sniffers. The merging process requires that input traces be synchronized so that a packet that appears in multiple individual traces is identified unambiguously. We developed a Python tool called PyPa1 that performs such an operation [22].

6. Dataset: Collection and characteristics

6.1. Data collection

We have fourteen sniffers in our measurement set-up, whose main component is a Raspberry Pi model 4B (RPi4 hereafter) [23]. We use an external Wi-Fi module, Alfa AWUS051NH, one per sniffer [24]. The advantage of this specific external Wi-Fi module is that it can be easily set to monitor mode. The monitor mode is a radio mode that makes it possible for the Wi-Fi card to passively listen to all Wi-Fi traffic in the wireless medium. We choose the 2.4 GHz band and channel 1 for our measurements.³ Fig. 4 shows the components of our sniffer.

We deploy the sniffers in the form of super-sniffer as we depict it in Fig. 5. We capture the traces indoors in an office scenario where the traffic load is high [18], which allows us to study the variation in the amount of traffic over time. We co-locate the sniffers and they remain stationary for the whole duration of the capture. We perform one test for 24 h. We collected the traces from 17:30 on 12th July 2022 to 17:30 on 13th July 2022.

We position fourteen collocated sniffers atop two boxes, each measuring 24.1 cm in height. This arrangement effectively accommodates all sniffers within a room simultaneously. In earlier experiments detailed in companion previous publications [18,19], we have shown that the particular height placement of sniffers has no impact on the study's outcomes, with these experiments being conducted indoors and outdoors on surfaces such as floors and tables.




Fig. 4. A single sniffer. It is composed of a Raspberry Pi 4B node, an Alfa AWUS 051NH antenna, and a powerful external power supply to allow us to carry out the measurements for a longer period of time.

Due to the ongoing nature of the experiments, which run overnight, outdoor measurements are unfeasible as continuous equipment supervision is required. Therefore, we opt for an indoor setting. The chosen room holds a central location, allowing us to capture packets from many devices. Additionally, we carried out preparatory tests of smaller durations to validate the selection of this location, the results showed a similar pattern of results that we achieve with the 24-h collection, irrespective of the height or exact placement of the sniffers. Moreover, this room provides a secure environment for our testing equipment, mitigating tampering or unauthorized access concerns.

Trace organization. As we collect all the traffic over an extended period of time, the size of the traces that we collect is 84 GB. The traces themselves but the number of merge operations we do for measuring redundancy makes the processing complicated and overloaded. We analyze all possible combinations of sniffers for a given redundancy, which gives a total of $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ combinations. We have 16,369 possible combinations with 14 sniffers. As the first step after time synchronization, we select a time granularity of 5 min for ease of processing and speeding up the calculation and analysis. We split each trace into 288 5-min sub-traces to cover the whole 24-h period. We merge the traces for each 5-min slot independently for all combinations of super-sniffers of all sizes.

6.2. Data characteristics

In this section, we detail the characteristics of the data we collect. We discuss (i) the traffic load in the medium (ii) what kind of traffic we capture (iii) why we capture all kinds of traffic (iv) the number of different source MAC addresses detected over 24 h.

Types of 802.11 frames. There are three types of 802.11 frames namely management, control, and data. The “type” field in the 802.11 frames indicates if it is a management (0), control (1), or data (2) frame. The frames such as association request, probe request, beacon, authentication, probe response, and disassociation fall under the category of management frames. Each of these subtypes has a number associated with it. We can use 0 in the “type” filter to get the management frames. We can pair the “type” field with the “subtype” field to filter out a certain kind of management frame. Similarly, the Ready to Send (RTS), Clear to Send (CTS), acknowledgment (ACK), block ACK, and beamforming report poll are types of control frames. We discuss how the traffic for these frames varies over time in the following text.

Traffic in the medium. Fig. 6 shows the number of packets per 5 min for the duration of 24 h. We only capture headers so these curves do

³ Channel 1 in the 2.4 GHz band was the most active in our setup.



Fig. 5. Testing location and composition of the super-sniffer. The order in which we place our co-located sniffers to form a super-sniffer of maximum size 14.

not fully characterize the load in terms of the utilization of the medium i.e. the data packets are less in number but they may occupy a larger portion of the medium because of a bigger size. Note also that our capture is based on the sniffers that only support the IEEE 802.11n, so we do not capture IEEE 802.11ac and 802.11ax traffic, and data packets sent with the highest Modulation and Coding Scheme (MCS) values [25]. Management frames (probe request/response, beacons, etc. [21]) are often sent at a lower data rate, which is why they make up a large portion of the load and at a consistent rate.

The orange, light gray, and dark gray curves in the figure depict the data, management, and control frames respectively. The blue line represents all the traffic. The x -axis and y -axis represent time and the packets captured per 5 min respectively. We see in the figure that the traffic load is higher in the late afternoon and morning because more people are present during office hours. At the same time, there is comparatively less traffic in the evening and at night. The traffic remains consistent in the non-peak hours; in fact, it is mostly the management frames that make up the traffic during that time. The traffic varies more during the peak (or office) hours when we also see control and data frames. The data frames cause small peaks but the control frames contribute the most to the load variation over time. There are a few sharp spikes when the control frames shoot and they result in high fluctuation of the traffic load in the medium. This is the kind of variation that will help us understand the impact of traffic load on the level of completeness that the sniffers can achieve.

If there are no devices at all in the coverage area of an access point, the access point only sends broadcast beacons i.e. management frames. So when there are some devices present, there is communication between the devices and the access point starting with a probe request. There are a few management frames to initiate the connection




Fig. 6. Traffic load for the whole duration of 24 h. The traffic is split into the data, management, and control frames, per 5 min each.


with the access point but after that, it is actual data transfer that takes place. The number of data frames, therefore, increases. However, there are more control frames than data frames for one complete communication. As the medium access protocols like CSMA dictate, the communication starts with control frames when a node that wants to start a communication, sends a Ready to Send (RTS) frame. Similarly, there are other control messages like CTS, ACK, and block ACK that constitute a single communication instance. That is why the number of control frames per MAC address rises during peak hours when there are non-access-point devices present in our experimentation area.

Reason for capturing all traffic. We see that management frames have negligible variation over the course of 24 h and can, hence, not be used to study the impact of load variation on completeness over time. Similarly, the data frames have no variation either. Moreover, they are almost non-existent during the nighttime when there is no one in the office space. We see the most variation in the case of control frames but we cannot solely rely on them because they fall to a very low level during the night, indicating the lack of traffic for a meaningful analysis. Therefore, we consider all these frames in our experiments because they truly help us to understand the impact of the evolution of the traffic load in the wireless medium on the quality of trace capture.


Number of sources detected over 24 h.⁴ Fig. 7(a) presents the number of unique MAC addresses per 5 min on the y -axis as a function of time on the x -axis. We see that the number of unique MAC addresses per 5 min corresponds to the traffic load in Fig. 6. The traffic in the medium increases when the number of unique addresses increases i.e. there are more devices present in the office area. The value also decreases during the night because fewer people are present in the vicinity. We observe a surge in the number from 13:30 onwards, as there are more devices in the office area during that period.

Fig. 7(b) displays the average number of packets per MAC address per 5 min (on the y -axis) over the course of 24 h. This figure helps us understand the density of the traffic present in the medium over time. With more devices during the daytime, the average number of control frames per MAC address rises. The average number of management frames per source also drops during this time because a higher number

⁴ A source MAC address does not equate to a device because of the MAC address randomization implemented by the devices [26]. But our study spans 5 min and the MAC randomization happens either after 15 min or never for older devices.



(a) Number of unique MAC addresses per 5 minutes.



(b) Average number of packets per MAC address per 5 minutes.

Fig. 7. The distribution of the number of unique MAC addresses and the average number of data, management, control, and all frames per MAC address. The results are per 5 min.

of devices results in higher activity. Similarly, during the nighttime, the average number of data and control frames per MAC address falls to a negligible level. We have packets from 19 access points in our traces but the amount of packets from 13 of those is negligible. We capture packets from only 6 access points for the whole duration of our experimentation. The packets during the night are there from at least 6 or at most 19 access points. It implies that a smaller number of devices are transmitting packets during the night. Therefore, the ratio number of packets per MAC address increases during the night. It is coherent with the fact that the number of unique sources drops during the night because of the presence of no people. Since we perform the testing in the office area, there are no people in the area during the night (except the security personnel) which means there are fewer laptops, smartphones, tablets, and other Wi-Fi-compatible smart devices. As a result, there is negligible communication between the devices and the APs, consequently, there would be an insignificant number of data and control packets.

However, the APs keep sending beacons (a type of management frame) periodically. Therefore, the average number of management frames per MAC address increases during the night even though the number of MAC addresses decreases. This average is low during the




Fig. 8. Average completeness of all single sniffers. The red lines represent the standard deviation of the completeness of all 14 sniffers individually. The higher standard deviation values show that all the sniffers do not perform in the same manner.

day because it is calculated over a more significant number of devices present in the vicinity. As there are only a few APs on the premises, that is the reason we see a spike in the average number of packets per MAC during the night despite having a low number of unique MAC addresses as no users connect to the Wi-Fi network.

All these data characteristics help us do the analysis of the performance of single sniffers as well as super-sniffers in the subsequent sections.

7. Individual completeness

In this section, we analyze the performance of each individual sniffer to identify if there are any faulty sniffers which could have a negative impact on the performance.

7.1. Completeness as a function of load

Fig. 8 shows the average completeness of all 14 sniffers per 5 min. The blue line represents the average completeness whereas the red bars are representative of the standard deviations of the completeness values of all 14 sniffers. We have the completeness values on the y-axis and time on the x-axis.

We observe that the completeness values decrease as the traffic load (Fig. 6) in the medium increases. The completeness crosses 60% when the load is low during the night, but it falls as low as 40% when the load is higher. In either case, the average completeness appears low. Each single sniffer misses around 40% to 60% of the packets depending on the time of the day.

When we look at the red bars, we notice that the values of the standard deviation differ, and at some points, they are larger comparatively. It means that there is a significant difference in the completeness values of the individual sniffers. A couple of questions arise: (1) are some of the sniffers faulty? (2) is it the same sniffer(s) that consistently performs poorly and leads to bad results? We answer these questions in Section 7.2.

7.2. Not all sniffers are good

Fig. 9 shows the completeness of all 14 individual sniffers over 24 h. We see that there are a few sniffers that perform consistently poorly. When we look at the zoomed parts of the figure, we identify that 3 sniffers, namely s_5 , s_{11} , and s_{14} , achieve low completeness values. We recall from Fig. 7(a) in Section 6.2 that the number of devices keeps changing over the period of 24 h. It means that the medium and the conditions change over time. This implies that the problem comes from the device itself, not the environment (multi-path, collisions, etc.).

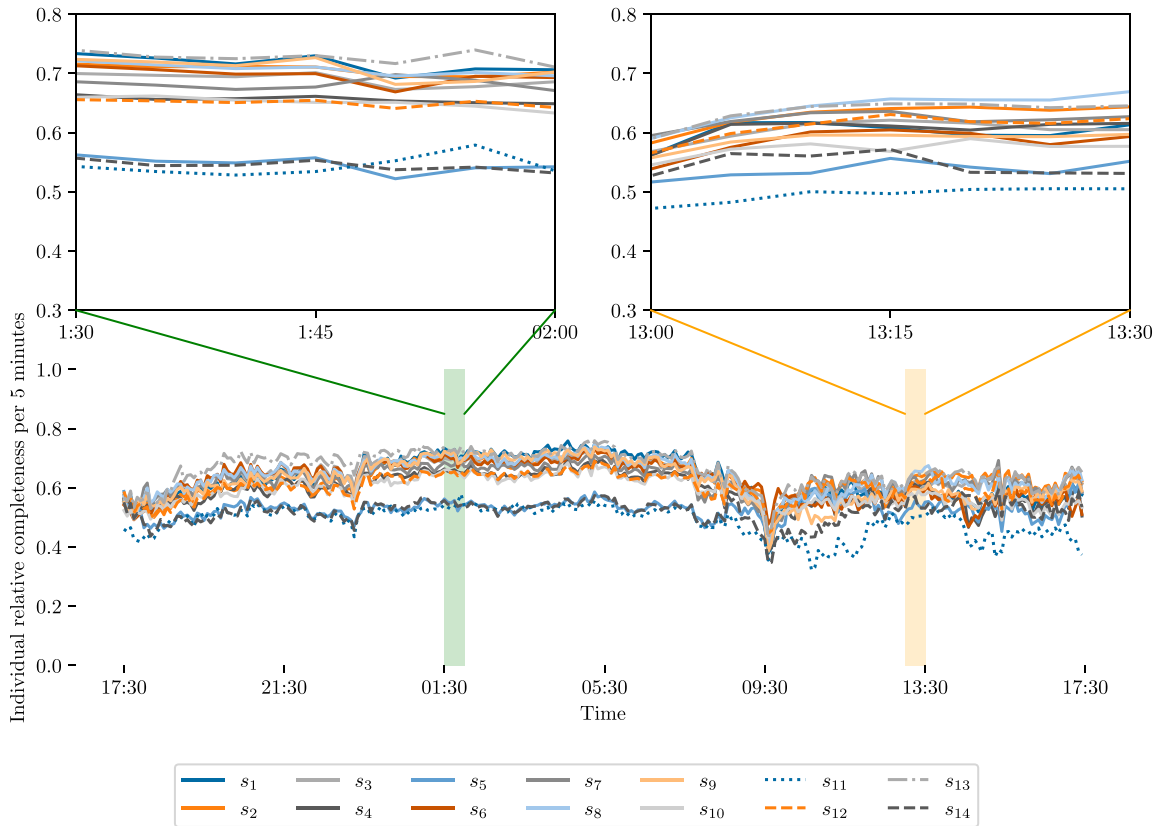


Fig. 9. Completeness of each sniffer. The completeness level of each individual sniffer. The zoomed area highlights that 3 sniffers perform consistently poorly.

7.3. RSSI

We explore the RSSI values of packets captured by each sniffer to further highlight the fact the aforementioned three sniffers are faulty. While the RSSI value is not an absolute measure of the quality of a link [4], quality of reception is often considered acceptable above -70 dBm and poor under -70 dBm [27]. Fig. 10 presents the percentage of packets captured with acceptable and poor RSSI, and packets missed, for each of the 14 sniffers. We see that the sniffers s_5 , s_{11} , and s_{14} capture the least percentage of packets with poor RSSI, proportionally. Sniffer s_{11} captures a negligible amount of packets with poor RSSI in comparison with other sniffers. It reiterates our finding that these sniffers are faulty and can lead to a biased analysis.

7.4. Cleaning the dataset

From this point on, we use a subset of the dataset for our analysis. We remove the worse performing sniffers s_5 , s_{11} , and s_{14} from the analysis part, we are, hence, left with 11 sniffers (i.e., a super-sniffer of maximum size 11).

We understand that the decision of pursuing with 11 sniffers can introduce some bias in our analysis but we want to continue with consistent sniffers. So, 21% of the sniffers lead to a poor dataset in our experiments. We need to investigate more in the future what could be the exact reason for this malfunction and whether would there be a possibility of fixing it. Along with that, we need to devise a strategy to select how to choose the best-performing sniffers for the experiments.

7.5. Are single sniffers enough?

Single sniffers may not be enough in several situations such as (i) Trajectory reconstruction: we see more than 40% packet loss, if a single sniffer misses the one important packet that is needed for measuring

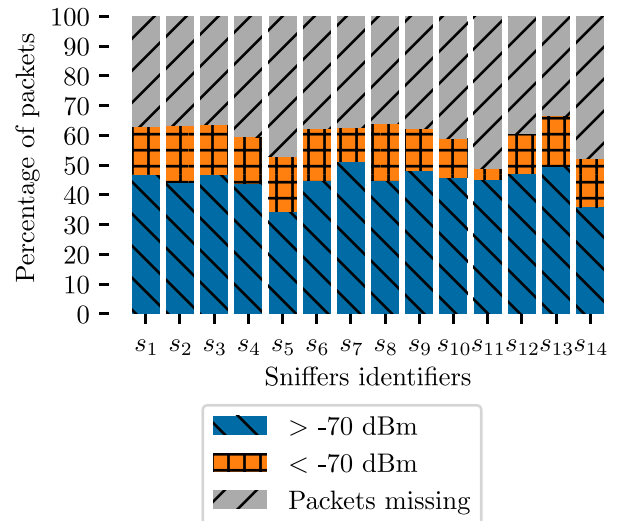


Fig. 10. RSSI. The percentage of packets captured with good (>-70 dBm) and bad (<-70 dBm) RSSI, as well as the percentage of packets missed by each individual sniffer.

the mobility, it can lead to flawed analysis (ii) Localization: such a significant packet loss can have dire effects on the search, rescue, and safety activities, especially in the case of emergency and disaster management.

These observations strengthen the case for the use of a super-sniffer. In the next section, we present our analysis to highlight the improvements that a super-sniffer brings.




Fig. 11. Pairwise completeness. The improvement in completeness 2 sniffers bring to each other when considered as a pair. The values inside the circles depict the completeness level of each sniffer and the edges' weights highlight the improvement 2 sniffers bring as a pair.

8. Super-sniffer completeness

To recall, we use a clean dataset for the analysis that we present from here on. We remove the three sniffers that performed poorly in comparison with the other 11 sniffers.

8.1. Pairwise completeness: Combination of 2 single sniffers

We define the *pairwise completeness* as a metric to rank the individual contribution of a sniffer towards completeness when paired in all combinations with other sniffers. In other words, when two traces are merged, how much information comes exclusively from the first trace, and how much comes from the second. Fig. 11 represents the pairwise completeness. We compare the pairwise completeness of each sniffer with s_{13} as it has the highest individual completeness.

The values inside the circles represent the values of completeness of individual sniffers over 24 h. The labels outside the circles identify the sniffer. The value on each edge of this star indicates the improvement in completeness the 2 sniffers bring as a pair.

The node s_3 brings an improvement of 11% when it is considered as a pair with node s_{13} , and s_{13} brings an improvement of 14% to s_3 . We see a 13% improvement when the node s_1 is paired with s_{13} , and the improvement is 17% when we pair s_{13} with s_1 . The minimum and maximum values of improvement are 9% and 18% respectively.

There is a trend that with increasing individual completeness of sniffers, their average pairwise gain also increases. We believe this metric can enhance the quality of measurements if only two sniffers are to be used. This graph can help us in the following ways, alternatively:

- It enables us to carry out an experiment of a small duration with all m sniffers co-located and then find out the best two sniffers for further experiments.
- We do the experimentation as planned and then we create this star as an initial analysis to select the two best nodes.

In this way, we are sure of getting the best traces for different analyses.

8.2. Completeness gain

We define the completeness gain as the improvement a super-sniffer of each size introduces. The completeness that we consider here is the average of all combinations of the super-sniffer of a specific size. Hake's method of finding the normalized gain is widely used for determining




Fig. 12. Normalized gain of average completeness over time for super-sniffers of all sizes.




Fig. 13. Super-sniffer completeness. Minimum, maximum, average, and reference completeness as a function of the size of super-sniffer for super-sniffers of all sizes.

the quality a new method brings in comparison with the existing methods/results [28]. The formula is as follows:

$$\langle g \rangle = \frac{\langle \text{post} \rangle - \langle \text{pre} \rangle}{100 - \langle \text{pre} \rangle}, \quad (8)$$

where $\langle \text{pre} \rangle$ and $\langle \text{post} \rangle$ refer to the results obtained before and after the improvements, respectively. The normalized gain is also known as the g -factor.

Fig. 12 represents the normalized gain of average completeness for super-sniffers of all sizes. We see that adding a single sniffer to compose a super-sniffer of size 2 results in a significant gain in completeness. There is still significant gain when we add another sniffer to the super-sniffer, i.e., $m = 3$. We get some gain as we keep adding sniffers until we get the super-sniffer of maximum size $m = 11$; however, the value of gain keeps reducing.

8.3. Overall completeness depending on the super-sniffer's size

In this section, we present the results of completeness for super-sniffers of all sizes as well as our reference super-sniffer. We build our super-sniffer as follows:




Fig. 14. Completeness of reference super-sniffer. The completeness of our reference super-sniffers of all sizes over the course of 24 h. This shows the variation in the level of completeness over time.

- Single sniffer: $m = 1$. The reference in this case is the single sniffer that gives the best completeness.
- Reference super-sniffer of size $m = 2$; which is the combination giving the best completeness among all super-sniffers of size 2 containing the reference single sniffer (i.e., $m = 1$).
- Reference super-sniffer of size $m = 3$ is the one giving the best completeness among all super-sniffers of size 3 containing the reference super-sniffer of size $m = 2$.
- We proceed the same way for the remaining super-sniffers up to $m = 11$.

Fig. 13 shows the minimum, maximum, and average completeness for all combinations of m sniffers, and reference completeness of our reference super-sniffer, represented by blue, yellow, green, and red lines, respectively. The numbers above the lines represent the improvement that our reference super-sniffer of each size brings to the table. The x -axis represents the time, while the y -axis gives the completeness for a combination of up to 11 sniffers. These are the results over 24 h.

We observe that the completeness improves by 13% by adding only one sniffer to make a reference super-sniffer of size $m = 2$. Adding one more sniffer brings a further improvement of 6% in the value of completeness. We keep seeing some improvement as we keep increasing the size of our reference super-sniffer. The rate of improvement keeps decreasing with every addition of a sniffer to the super-sniffer, but each sniffer brings new information to the super-sniffer. We observe that when we go from the super-sniffer of size $m = 9$ to $m = 10$ there is no improvement in reference/maximum completeness in our case; there is still an improvement of 2% in the case of minimum and average completeness though.

We also notice that the maximum and reference completeness are identical for super-sniffers of all sizes. It means that the super-sniffer of a certain size that achieves maximum completeness is part of the best-performing super-sniffer of the succeeding size. The minimum and maximum completences are also not too far apart.

Fig. 14 shows the completences of our reference super-sniffer as a function of time. We see that the completeness improves significantly by adding just one sniffer to make the super-sniffer of size $m = 2$. The improvement is around 10% for the whole duration of 24 h and it also varies concerning the traffic load. The rate of improvement keeps decreasing with every addition of a sniffer to the super-sniffer, but each sniffer brings new information to the super-sniffer. We note improvement in the quality of the trace capture with redundancy irrespective of the traffic load and time of the day.

The value of completeness decreases when there is more traffic in the medium during office hours, most notably around 10:00 in the morning. The use of a super-sniffer, however, helps improve the quality of capture even during the high load. It indicates that, comparatively, a higher number of sniffers are needed when the traffic in the medium is really high. In either case, our concept of super-sniffer increases the value of completeness.

8.4. RSSI

Fig. 15 depicts the percentage of packets captured by combinations of each number of sniffers with good or bad RSSI values as well as the percentage of the packets missed. We see that around 80% of the packets are missed if we use single sniffers because 20% of the packets are captured by only individual sniffers. 10% packets are captured by strictly 2 sniffers. Similarly, around 15% and 30% packets are captured by 10 and all 11 sniffers respectively. There are redundant packets that are removed in the process of merging, but there is a percentage of packets missed by a fewer number of sniffers. This finding further strengthens the use of a super-sniffer to improve the quality of the capture.

8.5. Discussion

We capture whatever is present in the medium at the time of experimentation irrespective of the environment and the surroundings. We realize that our experimental setup has a limitation as there is no benchmark set in a controlled environment. However, we plan to perform experiments in an anechoic chamber in the near future where we will not only have our sniffers but also a certain number of access points as well as users with a deterministic load. It will help us measure and define *absolute completeness* in the true sense. The individual completeness of single sniffers will then be measured based on this absolute completeness. We believe this will improve the performance of the super-sniffers. Moreover, we plan to study the results of redundancy by using multiple Wi-Fi adapters with a single Raspberry Pi node. We believe that this redundancy coupled with the benchmark set by an anechoic chamber will reduce the financial cost and improve energy efficiency.

9. Application of relative completeness: wireless environment characterization

In this section, we evaluate the metrics for the characterization of the wireless environment, which we defined in Section 3, through passive measurements.




Fig. 15. Super-sniffer wise RSSI distribution. The percentage of packets captured with good (>-70 dBm) and bad (<-70 dBm) RSSI, as well as the percentage of packets missed by the combination of each number of sniffers.

Table 1
Access point locations.

Access points	Floor	Corridor
AP1	1	Perpendicular
AP2	1	Same
AP3	2	Same
AP4	2	Same
AP5	4	Same
AP6	4	Same

9.1. Access point completeness

Our traces contain packets from a total of 19 access points (APs). However, we constantly capture packets over the whole duration of 24 h from only 6 APs. Therefore, we can compute completeness for these 6 APs for those we receive packets in all 288 5-min sub-traces. We have only 6416 packets from the rest of the 13 APs combined with several of them having no packets for several 5-min sub-traces. In fact, we capture only 1 packet from one of the APs that is installed at the 5th floor, while we place our experimentation set-up at the 1st floor. The 6 APs are placed as we show in Table 1. The value “same” in the corridor column means the AP is installed on the same side of the corridor as our sniffers whereas “perpendicular” means the AP is in the corridor that is perpendicular to the corridor where the sniffers are present. We see that 2 APs are located on the 4th floor and we capture their traffic for the complete 24 h. Whereas, there are APs in the same corridor as our sniffers but we see a negligible amount of packets for completeness in our traces. It is possible that these APs switch channels as we configure our sniffers to channel 1.

Fig. 16 highlights the completeness achieved by our reference super-sniffer for the 6 APs that we mention in Table 1. We see that AP5 achieves the lowest completeness which is not surprising because it is located on the 4th floor. However, AP6 from the same floor achieves higher completeness. AP1 is located in the perpendicular corridor and it has comparatively lower completeness. AP2, AP3, and AP4 achieve similar completeness. We see that completeness improves for all the APs with the increase in the size of the super-sniffer. The difference in the levels of completeness is very narrow from super-sniffer of size onwards.

To study the completeness of a specific node over a longer duration, we need to capture its traffic for longer periods. Moreover, the sniffers are able to capture the packets in the indoor environment even from far away nodes.




Fig. 16. Access point completeness. The completeness of our reference super-sniffer for the 6 access points that are present in our traces for the whole 24-h duration.

Table 2
Detection of nodes.

Floor	Corridor	% of nodes detected
1	Perpendicular	18.56
1	Same	14.43
2	Same	23.71
3	Same	4.12
4	Same	26.81
5	Same	12.37

9.2. Detection of nodes

Since we know the location of the APs, we analyze the RSSI values of the packets captured from each of those. As we mentioned in the previous section we captured only 1 packet from one of the APs, we, therefore, ignore this AP for the analysis we present in this section. We select one 5-min slot where we capture packets from all 18 APs. We calculate the mean, RSSI values for each AP; firstly we calculate the mean of the RSSI values for each packet that is seen by multiple sniffers, and then we take the mean of RSSI values of all packets received from a certain AP. The averaging of the values helps us to get the value that our sniffers see from a specific AP.

We use this information as a training set to compare the RSSI values of APs with the RSSI values of other devices and detect the presence of a node at a distance similar to the distance between the sniffers and the AP. We extract the RSSI values of all packets that the sniffers capture from each unique device present in our traces in a similar manner. We then calculate the average of the RSSI values for each unique device.⁵ We map the average RSSI values of these devices with those of the APs. Since we have multiple APs in a single corridor, we see a similar RSSI value at the sniffers with a similar standard deviation. We group these APs as a single location.

We present the results in Table 2. We see that by making use of the RSSI values, we are able to detect the presence of nodes at the same distance as that of a group of APs even on the 5th floor. The results are consistent for all 5-min sub-traces.

9.3. Presence of nodes for a shorter duration

The duration of the presence of nodes is a key factor for measuring pedestrian mobility for trajectory reconstruction. A mobile node is present at a certain location for a short time before moving on to the next point of reference. We present the analysis of the presence of nodes

⁵ Note: We consider unique addresses in our traces as unique devices. However, multiple MAC addresses might belong to a single device because of MAC address randomization.




Fig. 17. Presence of node for 30 s or less over a 5-min period. The number of nodes present for 30 s for non-peak-hour (night) and peak-hour time (afternoon) slots. The number on the top of the lines shows the number of new devices detected.

for a duration of up to 30 s in the context of redundancy in the number of sniffers. Note that we consider the packets received from non-access-point devices only because the presence of access points is of no use in this context as they stay there all time. We choose one peak-hour and one non-peak-hour 5-min slot for the clarity of analysis, the results are similar for all 5-min slots.

Fig. 17 represents the number of nodes present for up to 30 s as a function of the number of sniffers. The orange line represents the peak-hour slot whereas the blue line is for the non-peak-hour slot. The numbers at the top of the lines represent the number of new devices detected by a redundancy of each size. The number of nodes increases by a smaller proportion with the increasing number of sniffers in the non-peak time. The non-peak time slot is during the night when there are very less non-access-point devices present in the vicinity of the sniffers as discussed earlier in Section 6.2. However, we still detect new devices with the increase in the number of sniffers.

On the other hand, we detect a large number of devices even with a single sniffer in the peak-hour slot. The number of devices detected keeps increasing with the size of the redundancy. The number of new devices detected is also higher in this case e.g. we detect 68 new devices by adding just one sniffer to get a redundancy of size 2. Similarly, for 3 sniffers, we detect 116 additional devices as compared to a single sniffer. The detection number keeps decreasing with the increase in the number of sniffers.

The important point to note is that single sniffers miss nearly 56% of the devices. This is really critical in the case of trajectory reconstruction when the users are mobile and stay at one given location for a smaller duration. We end up missing a lot of devices if we use only one sniffer and we see from the figure that redundancy improves the results, thus, making it possible to detect more packets. The redundancy in the number of sniffers, therefore, improves the results of measuring mobility.

Reason for 30-s limit selection. If an access point is placed correctly it should offer a coverage of up to 150 ft or 54.72 m indoors for the 2.4 GHz band [29]. For contextualization, the average walking speed of an adult is between 1.2 m/s and 1.4 m/s [30,31]. An adult mobile user will take between 32.14 and 37.5 s at maximum, depending on the distance from the access point, to go out of the coverage range of a 2.4 GHz access point. We believe 30 s is a good limit for the duration of presence for measuring the detection of mobile nodes.

10. Discussion

We know that a single sniffer performs poorly irrespective of the conditions of the wireless medium. There is a need to introduce redundancy in the number of sniffers to improve the quality of the traces but

it comes at a financial as well as management cost. There is a trade-off between the cost of the sniffers and the level of performance in the quality of trace capture. We advise choosing a higher number of sniffers when the traffic load in the medium is high.

At the same time, one needs to be careful about the choice of sniffers. We notice that all the sniffers do not behave the same way despite the conditions of the medium changing over the course of 24 h. We show that the faulty sniffers can be removed from the detailed analysis after some initial diagnosis but it is not easy to have a consistent platform from the very beginning. We plan to explore, firstly, the exact cause of poorly performing sniffers, and secondly, address the need of setting up a good sniffing platform from the beginning of the experiments.

11. Conclusion

In this paper, we elaborate on the notion of trace completeness. We present the analysis for traces captured simultaneously by eleven co-located sniffers over 24 h. We highlight the importance of grouping sniffers into super-sniffers to improve completeness significantly. At the same time, we highlight that completeness varies over time depending on the traffic load in the wireless medium. We also evaluate the metrics access point completeness, detection of nodes, and presence of nodes for a shorter duration to characterize the wireless environment. Using passive measurements with only one sniffer, we miss a significant number of transient nodes that are present for a shorter duration. We plan to create a layout of the sniffing platform to avoid faulty sniffers from the very beginning, as well as identify the root cause of a few sniffers performing poorly. We also plan to evaluate MAC-address-based completeness. We finally intend to do measurements on different channels to study the impact of channel selection on completeness. Moreover, we plan to perform experiments in a controlled environment with an anechoic chamber to benchmark the absolute completeness.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

Acknowledgment

This work has been partially funded by the ANR MITIK project, French National Research Agency (ANR), France, PRC AAPG2019.

References

- [18] M.I. Syed, A. Fladenmuller, M. Dias de Amorim, Assessing the completeness of passive Wi-Fi traffic capture, in: 2022 International Wireless Communications and Mobile Computing (IWCMC), 2022, pp. 961–966, <http://dx.doi.org/10.1109/IWCMC55113.2022.9824970>.
- [19] M.I. Syed, A. Fladenmuller, M. Dias de Amorim, How much can sniffer redundancy improve Wi-Fi traffic? in: 2022 IEEE 95th Vehicular Technology Conference: (VTC2022-Spring), 2022, pp. 1–5, <http://dx.doi.org/10.1109/VTC2022-Spring54318.2022.9860874>.
- [1] A. Galanopoulos, V. Valls, G. Iosifidis, D.J. Leith, Measurement-driven analysis of an edge-assisted object recognition system, in: IEEE ICC, 2020, pp. 1–7, <http://dx.doi.org/10.1109/ICC40277.2020.9149069>.
- [2] W. Zhou, Z. Wang, W. Zhu, Mining urban WiFi QoS factors: A data driven approach, in: IEEE BigMM, 2017, pp. 9–16, <http://dx.doi.org/10.1109/BigMM.2017.12>.
- [3] P. De Vaere, T. Bühler, M. Kühlewind, B. Trammell, Three bits suffice: Explicit support for passive measurement of internet latency in QUIC and TCP, in: Proceedings of the Internet Measurement Conference 2018, IMC '18, New York, NY, USA, 2018, pp. 22–28, <http://dx.doi.org/10.1145/3278532.3278535>.

- [4] J. Wang, Y. Zheng, Y. Ni, C. Xu, F. Qian, W. Li, W. Jiang, Y. Cheng, Z. Cheng, Y. Li, X. Xie, Y. Sun, Z. Wang, An active-passive measurement study of TCP performance over LTE on high-speed rails, in: ACM Mobicom, New York, NY, USA, 2019, <http://dx.doi.org/10.1145/3300061.3300123>.
- [5] M.D. Corner, B.N. Levine, O. Ismail, A. Upreti, Advertising-based measurement: A platform of 7 billion mobile devices, in: ACM Mobicom, Snowbird, UT, USA, 2010, pp. 435–447.
- [6] F. Garcia, R. Andrade, C. De Oliveira, J. Souza, EPMOST: And energy-efficient passive monitoring system for wireless sensor networks, *Sensors (Basel Switzerland)* 14 (2014) 10804–10828, <http://dx.doi.org/10.3390/s140610804>.
- [7] R. Mahajan, M. Rodrig, D. Wetherall, J. Zahorjan, Analyzing the MAC-level behavior of wireless networks in the wild, in: Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Association for Computing Machinery, New York, NY, USA, 2006, pp. 75–86, <http://dx.doi.org/10.1145/1159913.1159923>.
- [8] Y.-C. Cheng, J. Bellardo, P. Benkő, A.C. Snoeren, G.M. Voelker, S. Savage, Jigsaw: Solving the puzzle of enterprise 802.11 analysis, in: Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Association for Computing Machinery, New York, NY, USA, 2006, pp. 39–50, <http://dx.doi.org/10.1145/1159913.1159920>.
- [9] K. Papagiannaki, M. Yarvis, W.S. Conner, Experimental characterization of home wireless networks and design implications, in: Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications, 2006, pp. 1–13, <http://dx.doi.org/10.1109/INFOCOM.2006.293>.
- [10] B. Pavkovic, F. Theoleyre, D. Barthel, A. Duda, Experimental analysis and characterization of a wireless sensor network environment, in: Proceedings of the 7th ACM Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks, in: PE-WASUN '10, Association for Computing Machinery, New York, NY, USA, 2010, pp. 25–32, <http://dx.doi.org/10.1145/1868589.1868595>.
- [11] X. Xu, C. Tong, J. Wan, Improve the completeness of passive monitoring trace in wireless sensor network, in: 2010 Asia-Pacific Services Computing Conference (APSCC 2010), IEEE Computer Society, Los Alamitos, CA, USA, 2010, pp. 560–566, <http://dx.doi.org/10.1109/APSCC.2010.112>.
- [12] X. Xu, J. Wan, W. Zhang, C. Tong, C. Wu, PMSW: A passive monitoring system in wireless sensor networks, *Int. J. Netw. Manage.* 21 (4) (2011) 300–325, <http://dx.doi.org/10.1002/nem.792>.
- [13] M. Sammarco, M.E.M. Campista, M. Dias de Amorim, Trace selection for improved WLAN monitoring, in: Proceedings of the 5th ACM Workshop on HotPlanet, Association for Computing Machinery, New York, NY, USA, 2013, pp. 9–14, <http://dx.doi.org/10.1145/2491159.2491165>.
- [14] A. Schulman, D. Levin, N. Spring, On the fidelity of 802.11 packet traces, in: *Proceedings of the 9th International Conference on Passive and Active Network Measurement*, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 132–141.
- [15] C. team, CRAWDAD, <https://ieee-dataport.org/collections/crawdada>.
- [16] A. Mahanti, M. Arlitt, C. Williamson, Assessing the completeness of wireless-side tracing mechanisms, in: IEEE WoWMoM, 2007, pp. 1–10, <http://dx.doi.org/10.1109/WOWMOM.2007.4351786>.
- [17] B.-r. Chen, G. Peterson, G. Mainland, M. Welsh, *LiveNet: Using passive monitoring to reconstruct sensor network dynamics*, in: *Distributed Computing in Sensor Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 79–98.
- [20] The Tcpdump Group, Tcpdump and libpcap, 2023, <https://tcpdump.org>.
- [21] M.S. Gast, *802.11 Wireless Networks: The Definitive Guide*, second ed., O'Reilly Media, Inc., 2005.
- [22] M.I. Syed, A. Fladenmuller, M. Dias de Amorim, PyPal: Wi-Fi Trace Synchronization and Merging Python Tool, Technical report, LIP6 UMR 7606, UPMC Sorbonne Université, France, 2022, URL <https://hal.archives-ouvertes.fr/hal-03618014>.
- [23] Raspberry Pi 4 model B, 2019, <https://tinyurl.com/2p89uund>.
- [24] AWUS051NH Wi-Fi adapter, <https://tinyurl.com/yk8vk3zv>.
- [25] IEEE Standard for Information Technology– Local and Metropolitan Area Networks– Specific Requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput, IEEE Std 802.11n-2009 (Amendment to IEEE Std 802.11-2007 as amended by IEEE Std 802.11k-2008, IEEE Std 802.11r-2008, IEEE Std 802.11y-2008, and IEEE Std 802.11w-2009), 2009, pp. 1–565, <http://dx.doi.org/10.1109/IEEESTD.2009.5307322>.
- [26] J. Martin, T. Mayberry, C. Donahue, L. Foppe, L. Brown, C. Riggins, E.C. Rye, D. Brown, A study of MAC address randomization in mobile devices and when it fails, 2017, <http://dx.doi.org/10.48550/ARXIV.1703.02874>.
- [27] Juniper, RSSI values for good/bad signal strength, 2019, <https://www.mist.com/documentation/rssi-values-good-bad-signal-strength/>.
- [28] R.R. Hake, Interactive-engagement versus traditional methods: A six-thousand-student survey of mechanics test data for introductory physics courses, *Amer. J. Phys.* 66 (1) (1998) 64–74, <http://dx.doi.org/10.1119/1.18809>.
- [29] Electric Power Board (EPB) of Chattanooga, How far will your Wi-Fi signal reach?, <https://tinyurl.com/4djad57b>.
- [30] A. Mansfield, E.L. Inness, W.E. Mcilroy, Chapter 13 - stroke, in: B.L. Day, S.R. Lord (Eds.), *Balance, Gait, and Falls*, in: *Handbook of Clinical Neurology*, vol. 159, Elsevier, 2018, pp. 205–228, <http://dx.doi.org/10.1016/B978-0-444-63916-5.00013-6>.
- [31] O. Mohamed, H. Appling, 5 - clinical assessment of gait, in: *Orthotics and Prosthetics in Rehabilitation (Fourth Edition)*, fourth ed., Elsevier, St. Louis (MO), 2020, pp. 102–143, <http://dx.doi.org/10.1016/B978-0-323-60913-5.00005-2>.



Mohammad Imran Syed received his B.Sc. degree in telecom engineering (silver medal) from Bahria University, Islamabad, Pakistan in 2010, and his M.Sc. degrees in Mobile and Satellite Communications (merit) from the University of South Wales (formerly University of Glamorgan), Treforest, UK in 2012, and in ICT Innovation from EIT Digital Master School (TU Berlin – sehr gut and Sorbonne Université – mention bien) in 2019. He has recently completed a successful defense of his Ph.D. in Computer Science at LIP6, Sorbonne Université, Paris, France where he also worked as a teaching assistant. Previously, he has worked at Alcatel-Lucent (UK), Mobilink (Pakistan), and INRIA (France). His research interests include wireless networks, wireless and Internet measurements, and opportunistic networks.



Anne Fladenmuller received her Ph.D. degree in 1996 from université Pierre et Marie Curie, France. She joined the University of Technology of Sydney, Australia for 2 years as a postdoc and then lecturer, before becoming a researcher at Alcatel Corporate Research Center, France in 1998. She got an assistant professor position at université Pierre et Marie Curie in 1999 and is now a full professor at Sorbonne Université in the Network and Performance Analysis team of the LIP6 research laboratory. Her major research interests are in wireless networks.



Marcelo Dias de Amorim is a research director with the French National Center for Scientific Computing (CNRS) and a member of the LIP6 Computer Science Laboratory at Sorbonne Université, Paris, France. His research interests focus on understanding, designing, and evaluating interactive dynamic networks. In the latest years, he has been investigating disruptive strategies to better handle the data tsunami problem in such networks. He is a former associate editor for IEEE Transactions on Mobile Computing and IEEE Communications & Tutorials, as well as a guest editor for IEEE Wireless Communications Magazine and Elsevier Computer Communications. He served on the technical program committee of several international conferences and was the general co-chair of IEEE Infocom 2019.