Unsupervised Detection of Adversarial Collaboration in Data-Driven Networking

Matteo Sammarco, Miguel Elias Mitre Campista, Marcin Detyniecki, Tahiry Razafindralambo, Marcelo Dias de Amorim

Unsupervised Detection of Adversarial Collaboration in Data-Driven Networking

Matteo Sammarco*, Miguel Elias Mitre Campista†, Marcin Detyniecki*‡,
Tahiry Razafindralambo§, and Marcelo Dias de Amorim¶

*AXA, France †GTA/COPPE-PEE/Del-Poli – Universidade Federal do Rio de Janeiro ‡Polish Academy of Science
§LIM – Université de La Réunion ¶LIP6/CNRS – Sorbonne Université

Emails: {matteo.sammarco, marcin.detyniecki}@axa.com, miguel@gta.ufrj.br,
tahiry. @univ-reunion.fr, marcelo.amorim@lip6.fr

***Abstract*—Data-driven networking in combination with machine learning is a powerful way to design and manage networked systems. In this paper, we consider the case of participatory collection of wireless traffic, which is an inexpensive way to infer the wireless activity in a locality. Since such a type of measurement system leans on the goodwill of the end users, it opens a new venue for malicious actions. Possible consequences of attacks are changes in the underlying communication substrate or even the collapse of the network. We assess the influence of these adversaries by identifying possible hostile actions and propose a method to detect them based on unsupervised machine learning models. Through an experimental campaign in various scenarios, we show that attacks with critical impacts are systematically detected, while unidentified attacks produce only a negligible impact in the measurement system.**

***Index Terms*—Network data analytics, data-driven networking, wireless networks, collaborative measurements, and machine learning in network measurement.**

I. Introduction

Following the upward trend of wireless network traffic and the impact of machine learning in the resolution of complex problems, the data-driven networking (DDN) paradigm will be a central component in the future of networks [1], [2]. The idea is to improve the design and the management of networked systems through the analysis of network data and measurements. In the context of Wi-Fi wireless networks, the focus of this paper, several monitoring systems exist to record wireless networking activity [3]–[6]. These systems aim at capturing packet exchanges between nodes as accurately as possible. The resulting dataset reflects the network behavior and may serve multiple purposes, including the analysis of user mobility or the improvement of the network quality of service. For instance, an access point can be moved from an over-provisioned to an under-provisioned area to increase the overall capacity of the system.

Collecting wireless traffic and monitoring network activity in small environments is relatively easy: all the information collected must first be logged into data traces, which are merged together to produce a single file containing a coherent global view of the wireless activity. The problem becomes tricky in larger scenarios where multiple monitoring devices operate simultaneously for the sake of completeness. The first obstacle, in this case, is how to scale the monitoring system. A promising solution is to rely on a participatory approach where end-users collectively contribute with individual measurements [7]–[9] in exchange of benefits, such as better connectivity [10], [11].

One shortcoming of the participatory methodology is the possibility of malicious actions [12]. Users may feel motivated to insert fake traces in the monitoring system to maximize their benefits or to deteriorate the overall system performance. These results could be achieved by, respectively, attracting additional infrastructure towards the malicious user (*attractive attack*) or purging near infrastructure away (*repulsive attack*) [13].

In this paper, we identify attractive and repulsive attacks from malicious participants as well as two trace-forging strategies, namely *synthetic* and *camouflage*. In the synthetic strategy, the attacker introduces fake traces containing only hypothetical source-destination pairs, while in the camouflage strategy, the attacker also introduces fake traces of wireless exchanges between real nodes. Our system can identify a suspicious trace inserted by a malicious user for any combination of attack and strategy, relying on wireless trace manipulation, dimensionality reduction, and unsupervised machine learning techniques. Concerning supervised approaches, we do not need a labeled dataset of wireless traces which, in a mobile and dynamic context, could result biased.

In summary, the main contributions of this paper are:

- **We identify possible weaknesses of collaborative monitoring systems for DDN concerning participatory measurements.** We address two possible attacks based on the insertion of fake traces, i.e., *attractive* and *repulsive* attacks, and two trace-forging strategies: *synthetic* and *camouflage*.
- **We propose a procedure to infer the participation of a malicious user in the monitoring system.** We consider the captured traces as text where words are source-

This work was partially funded by the French National Research Agency (ANR) under the project ANR VERSO RESCUE (ANR-10-VERS-003). Miguel Elias M. Campista would also like to thank the Brazilian agencies CAPES (Coordenação de Aperfeiçoamento de Pessoal de Nível Superior) Finance Code 001, CNPq, FAPERJ, and São Paulo Research Foundation (FAPESP) grant #15/24494-8 and #15/24490-2.

destination pairs (SD). Like a bag-of-words approach, each trace is a point in an N-dimensional space, where N is the SD set cardinality. Upon identifying clusters of similar traces in such space, we can infer fake traces.

- **We conduct *in vivo* experiments using three different scenarios**, called *collocated*, *scattered*, and *outdoor*, having different density of monitoring devices and environmental conditions. Results from any combination of circumstance and attack show that it is possible to detect a malicious user as soon as it does not inject a paltry amount of fake flows in the trace.

In the rest of the paper, we refer to any device recording a wireless trace collected in the collaborative measurement system (i.e., infrastructure access points, sniffing devices provided by collaborative or malicious users, or any other monitoring device) as monitoring nodes or just monitors, while a wireless trace is the set of IEEE 802.11 frames made up of source-destination (SD) flows. A malicious user (or attacker, or adversary) produces one malicious trace embedding fake SD flows.

This paper is organized as follows. Section II details the addressed problem and introduces possible attacks and trace-forging strategies. The proposed clustering-based procedure to detect fake traces is presented in Section III. In Section IV, we describe the three testbeds used to validate our approach, while in Section V we show the results for each scenario and each attack. Finally, Section VI qualitatively compares our proposal to the state-of-the-art, while Section VII concludes this work and indicates future directions.

II. Problem Statement

In a collaborative measurement system, network users contribute to the monitoring system, acting as additional monitors. Such a "participatory" behavior is an increasing trend as smartphones are expected to measure the performance of mobile networks to improve users' experience [14], [15]. Additionally, since the only requirement is to run a packet sniffer for her benefit, we assume this should not be cumbersome for most users.

A. Participatory passive wireless monitoring

In the proposed system, monitoring devices (including users) produce traces to be analyzed. Figure 1 shows the normal operation of a collaborative wireless monitoring system. Supposing that each source-destination flow consists of four frames, collaborating nodes L_1 and L_2 can capture each fifty percent of the total wireless traffic composed of two flows: $A \leftrightarrow B$ and $C \leftrightarrow D$. Merging their traces, i.e., T_{L_1} and T_{L_2}, the wireless traffic captured increases to 75%.

B. Attacks against participatory monitoring

The completeness of the merged trace depends on several parameters, including geographic dispersion and the number of monitoring devices covering a given area. The trade-off of users' collaboration in wireless measurement systems, however, is the introduction of vulnerabilities. In this work,

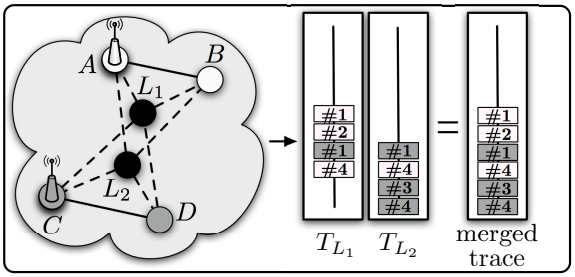

Fig. 1. The fraction of captured frames increases after merging two individual traces from different monitoring nodes. Each individual monitoring node, L_1 and L_2, has captured 50% of the total frames numbered in sequential order from 1 to 4. After merging, the fraction of captured frames increases to 75%.

we consider two different adversarial actions with opposite consequences as follows:

- **Attractive attack.** It consists of generating fake traces with *almost empty* sequences of frames, that is, forging SD flows containing only the first and the last frames indicated by the sequence number in the frame header.
- **Repulsive attack.** It consists of inserting fake traces with *complete* sequences of frames, that is, forging SD flows containing all frames.

Depending on the type of incentive, a malicious user could find benefit on triggering one of the two identified attacks. For instance, if users receive incentives, e.g., improved wireless connectivity to indicate under-provisioned areas, they could adopt the so-called *attractive behavior* to maximize such gains. We name it "attractive" because the monitoring system infers from traces containing an almost empty sequence of frames that, in the wireless network nearby, the users are under-provisioned, i.e., frames are lost quite often due to lack of wireless infrastructure. Therefore, the corresponding area needs additional resources to operate under normal conditions. If, on the other hand, users are willing to disrupt the network operation, they could perform the opposite action, indicating to the system the existence of over-provisioned areas. Such an act corresponds to the so-called *repulsive behavior*, in which users send traces with complete sequences of frames to the monitoring system. Without significant frame losses, the system can infer that the area nearby the monitoring user is over-provisioned and therefore can share part of its infrastructure with other areas if needed.

The abovementioned attacks can be triggered in any wireless network by a malicious user aiming to launch a counterproductive action. Independently of the attack, the outcome is always a network where malicious users influence the disposition of the network infrastructure. Hence, instead of only contributing to the system in exchange of, for instance, flat performance, malicious users are willing to either maximize their connectivity in detriment of all the others or even disrupt the network operation. Moreover, each attack could come with one of the following trace-forging strategies:

- **Synthetic strategy**. With this strategy, a malicious user

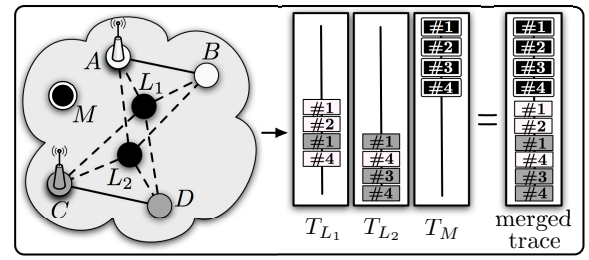

(a)

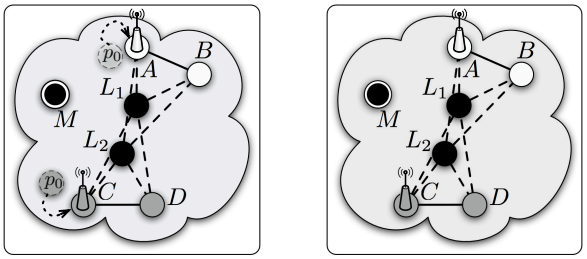

(b)

(c)

Fig. 2. Repulsive attack. a) Malicious node M forges a trace T_M containing a flow with a complete sequence of frames. As a consequence, the fraction of captured frames increases from 75% to 83% and the wireless infrastructure is repulsed to other areas. b) Infrastructure nodes in their initial position p_0. c) Network final position.

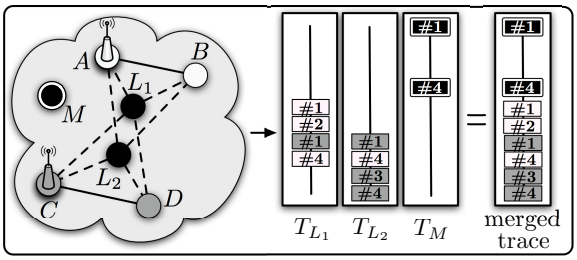

(a)

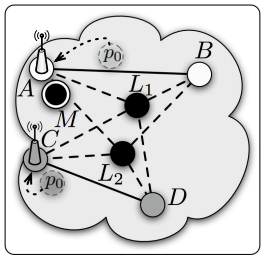

(b)

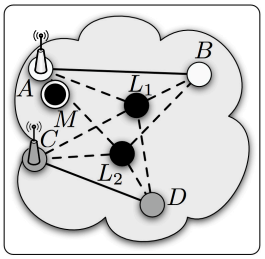

(c)

Fig. 3. Attractive attack. a) Malicious node M forges a trace T_M containing a flow with an empty sequence of frames. As a consequence, the fraction of captured frames reduces from 75% to 66% and the wireless infrastructure is attracted toward M. b) Infrastructure nodes in their initial position p_0. c) Network final position.

forges a wireless trace containing only fake SD pairs.

- **Camouflage strategy**. Attackers forge traces mixing fake and real SD pairs.

Figures 2 and 3 illustrate the repulsive and the attractive attack, respectively. We have a malicious node M, under control of a malicious user, inserting a fake trace T_M into the system. In a repulsive attack, as illustrated in Figure 2(a), the malicious node M forges a trace containing a flow with all four frames from a given SD pair. Note that if SD is fake (i.e., M creates from scratch a pair of communicating nodes), we say that the synthetic trace-forging strategy is applied. Otherwise, if SD is an existing pair, the strategy is the camouflage. Anyway, after merging, the fraction of captured frames over the total is approximately 83% ($3/4$ from $A \leftrightarrow B$, $3/4$ from $C \leftrightarrow D$, and $4/4$ from T_M), which is higher than the 75% captured in normal operation (Figure 1). As a possible attack outcome, the available infrastructure (nodes A and C) can be moved to an under-provisioned area as seen in Figures 2(b) and 2(c). On the other hand, in the attractive attack (Figure 3(a)), after the merge, 66% of the frames are considered captured ($3/4$ from $A \leftrightarrow B$, $3/4$ from $C \leftrightarrow D$, and $2/4$ from T_M). Thus, the fraction of the total frames becomes lower than the one obtained in the normal operation. The consequence could be the attraction of more infrastructure resources (nodes A and C) towards the malicious node (Figures 3(b) and 3(c)). Repulsive and attractive attacks are consequences of inserting fake traces into the system, considering infrastructure reallocation as user incentive. If other incentives were envisioned, new types of attacks could be possible as a consequence of the same malicious action.

C. Assumptions

In this paper, we assume that the monitoring system accepts only one trace from each user (could be a single trace or a merged one). This is an important assumption because the impact of a malicious user can be accentuated by the number of traces it adds to the system. In addition, monitoring nodes, legitimate or malicious, do only perform passive measurements. Hence, they do not add packets to the wireless network. At last, each trace has the same format, is obtained or forged during the same time frame, and contains a coherent amount of data. We consider that the merging procedure discards traces not matching these assumptions.

III. Proposed Clustering-based Procedure

The goal of the proposed system is to distinguish fake wireless traces produced by malicious users from legitimate ones.

The procedure, depicted in Figure 4, takes as input the recorded wireless traces from all monitoring nodes, including those produced by adversaries. Each trace is considered as a text composed of SD words, where the set of all unique SD pairs in the corpus of traces has cardinality N. Each trace is then described by a vector of size N, where each element is the completeness of frames recorded in that trace for one of the N SD pairs.

In our approach, completeness is important for the sake of accurate monitoring. It is defined as the percentage of frames

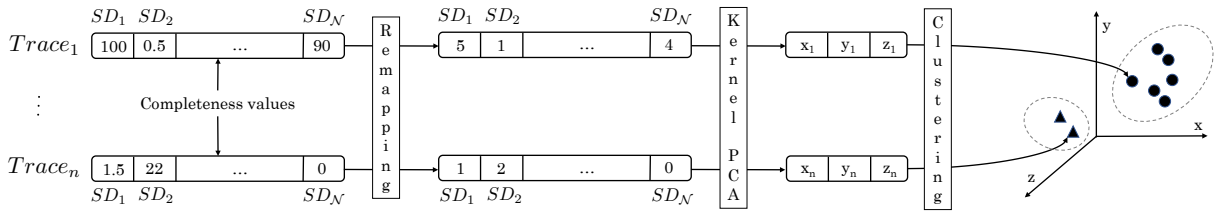

Fig. 4. Proposed clustering-based procedure.

captured by a monitor for each SD pair, checking the sequence number embedded in each frame header. Note, however, that for each SD frame exchange, the first and the last frame captured do not necessarily match the first and the last frame sent due to path loss and signal fading. Instead, they are the first and the last frame from SD communication that were captured and logged into a given trace. If the sequence of frames exchanged between SD is $\langle f_1, f_2, \ldots, f_{n-1}, f_n \rangle$, the completeness is then computed taking into account the first and the last frame captured, which could be f_2 and f_{n-1}, for instance. Hence, the completeness, computed in this example as $100 \cdot \frac{\texttt{\# frames captured}}{\texttt{\# frames in the sequence}}$ would account for frames from f_2 to f_{n-1}. Independent of the first and the last frame considered, the completeness always spans from 0 to 100. Also the sequence is obtained using the sequence number field in which the counter is zeroed after 2^{12} sequence numbers. Thus, we consider each trace as a point in a N-dimension space described by the completeness vector. Real traces share some SD and have some similarities in the distribution of completeness values, while malicious users are forced to forge novel SD and more extreme completeness values in order to get the desired effect.

Clustering points in the space puts together real traces and sets aside malicious ones. This approach, however, raises two problems:

1) **Distribution of completeness values is skewed.** Most of the values are equal to zero as many SD are unique to a single trace, especially when covering a wide area and when collaborating people are sparse. Some values are just above zero, nonetheless monitors which did not overhear anything from a SD can be far away from the source, whereas monitors which overhear just a little are still in proximity. Traces also present a considerable amount of SD exchanges complete at 100%, which usually correspond to control frames. Then, very few completeness values come from the interval $[0+\epsilon, 100-\epsilon]$.
2) **The number of SD pairs is usually very large.** N is usually a large number and, consequently, clustering becomes very hard to obtain. This problem, so-called "curse of dimensionality", is already well-known [16].

To tackle the first problem, we remap all completeness values from the continuous interval $[0, 100]$ to the discrete interval $[0, 1, 2, 3, 4, 5]$ in order to have a more uniform distribution of values. The second problem, on the other hand, is overcome through a non-linear Principal Component Analysis using a radial basis function kernel. In our scenarios, six principal components were enough to obtain 80% of variance.

Finally, a hierarchical agglomerative clustering is applied using Euclidean distance and with a merging strategy targeting the maximum distance between observations from pairs of clusters. The clustering algorithm is forced to find out two clusters which will take apart real traces from those forged by adversaries. With respect to density or centroid-based clustering, hierarchical algorithms produce a dendrogram which results helpful for further analysis and resource management.

In our implementation, we took advantage of `KernelPCA` and `AggloerativeClustering` functions from Python Scikit-Learn framework. We used, instead, WiPal as IEEE 802.11 traces manipulation software for fast frame decoding, trace merging, and wireless traffic sniffing [4].

IV. Testbed and Dataset

We have captured wireless traffic for 100 minutes in three scenarios, named *collocated*, *scattered*, and *outdoor*, to create a dataset of traces as much as close to reality.

In any scenario, we positioned eight Asus EEEPC-4G netbooks equipped with 3 USB Wi-Fi Netgear WG111v3 cards and running Xandros OS with a custom kernel and WiPal sniffing software. These laptops played the role of monitoring nodes and, as a consequence, they were just in charge of recording traces from the wireless activity. We underline that, in such experiments, the monitoring nodes capture any transmitted frames they overhear in the area (WiFi cards are in monitor mode). This means that the captured traffic can be from an access point within the scenario, but it can also be from a pedestrian with a WiFi mobile phone passing nearby. We assume that any incoming traffic must be captured by the measuring system, no matter how long it lasts or what kind of activity it is concerned with.

The collocated scenario was built inside a room within LIP6 computer science laboratory from UPMC Sorbonne Universités in Paris, France. All monitoring nodes were positioned side-by-side on the same table as illustrated in Figure 5(a). Individual traces have an average size of 253 MBytes, whereas the merged trace has a size of 450 MBytes. This scenario is our benchmark as the merged trace will have a very complete vision of the wireless activity in that room and individual traces will share large part of flows. The scattered scenario was built in the second floor of the IRCICA/LIFL computer science laboratory in Lille, France. In this case, monitoring nodes were positioned according to the availability of electrical outlets

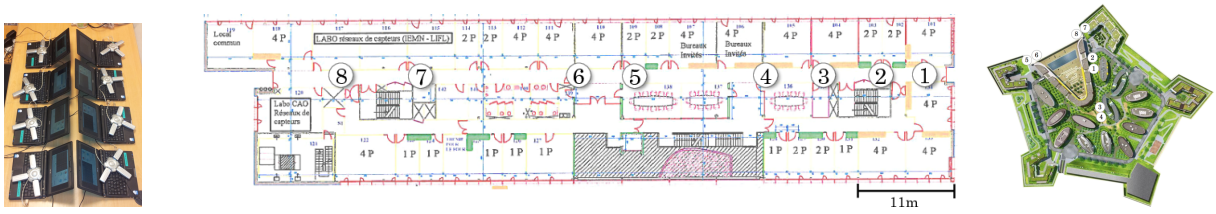

(a) (b) (c)

Fig. 5. Experimental scenarios and monitor deployment: (a) collocated, (b) scattered, and (c) outdoor.

along the corridor as illustrated in Figure 5(b). Individual traces have an average size of 205 MBytes, whereas the merged trace has a size of 1.2 GBytes. The larger size of the merged trace compared with the collocated scenario indicates that traces from more distant monitors are likely more complementary. For the outdoor scenario, monitors were placed outdoor in the Fort d'Issy of Issy Les Moulineaux, France, a 12 ha residential area and deployed as shown in Figure 5(c). It differs from the first two by the density of monitors in the area and mobility.

V. Experimental Results

This Section is divided in four experiment sets. First, we prove our claim regarding the advantages of collaborative measurements. Then, we show that it is possible to detect an adversary contributing with a fake trace using different strategies. Moreover, we apply the same procedure for an online detection and for multiple attackers. Since we are interested in finding malicious traces and, at the same time, restricting further investigations on just the selected ones, we mainly rely on the F1-score metric to evaluate the procedure.

A. Collaborative measurements

The scenarios collocated, scattered, and outdoor globally present 508, 612, and 580 unique SD, respectively. If we take single traces, we note that, on average, they include 401, 224, and 102 unique SD for the collocated, scattered, and outdoor, respectively. This result shows that merged traces bring more information from the target area than individual ones. The idea of collaborative monitoring is then important even in the collocated scenario where monitoring nodes are close. We can also presume that an attack must inject a number of SD pairs at the same scale to produce an effect.

Figure 6 shows the SD flow completeness for single traces and for the trace merging all of them, for all scenarios. In this figure, the X-axis represents all SD pairs in the network while the Y-axis represents the completeness level of each individual flow. The dots are the completeness level achieved by individual traces, whereas the red line is the completeness level achieve by the merged trace. Note that the merged trace shows always an upper-bound result, since it puts together all the individual traces collected.

B. Potential attackers detection

Last subsection validated the advantages of having collaborative measurements. This subsection shows how we could identify a malicious user infiltrated into the monitoring system. Before we start, it is important to give more details concerning the trace-forging strategies used in both attacks.

- **Synthetic strategy**. Wireless traces forged by malicious users contain only fake SD pairs, where the completeness of each communication pair is given by the attack type (100% and 0.05% for repulsive and attractive attacks, respectively).
- **Camouflage strategy**. Malicious users forge traces mixing together fake and really captured SD pairs communications. Completeness values for these latter are pulled from the same distributions as the legitimate traces. Such strategy better masks the attack, but can reveal the position of the attacker as real SD pairs come from monitors and access points in the same area.

Figure 7(a) shows clustering results for attractive and repulsive attacks in each scenario for the synthetic strategy. Clustering is considered to be correct if the malicious trace is circumscribed alone in a cluster while all the other traces compose a second cluster. For this experiment, we vary the number of fake SD embedded in a trace from 5 to 100. For the collocated scenario, real traces have so much in common that the clustering algorithm is immediately able to pick out the attacking trace. Nevertheless, in the scattered and outdoor scenarios, below 14 and 45 fake SD respectively, clustering is incorrect. Such amounts of flow correspond to just 2.3% and 7.7% w.r.t. N as reported in Table I. These low values will unlikely affect the network. Above these values, instead, fake traces become more and more distinguishable from the legitimate traces. This means that the more fake information an adversary adds, farther the forged trace becomes from all legitimate ones. In fact, they will introduce unique dimensions in the clustering space where they are placed. These considerations are valid for both attractive and repulsive attacks. Table I also shows the F1-score which, in the outdoor scenario, is just limited by the amount of fake flows injected during the experiment. Finally, with a synthetic strategy in action, the attacker should forge a trace with very few SD (maximum 7.7% of N in our testbeds) to circumvent the defense, but too few to generate an impact and to be taken into account.

Figure 7(b) shows the clustering results in all the scenario combinations for the camouflage strategy. In these experiments, malicious traces are always made up of 200 SD pairs. We, however, tune the amount of real SD pairs in the interval

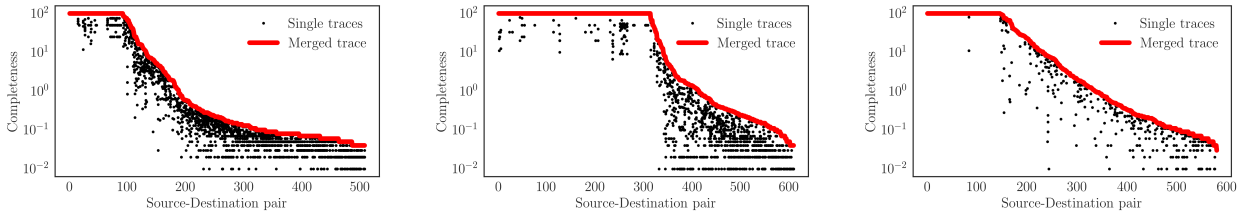

(a) (b) (c)

Fig. 6. SD flows completeness recorded in single traces and merged trace for experimental scenarios (a) collocated, (b) scattered, and (c) outdoor.

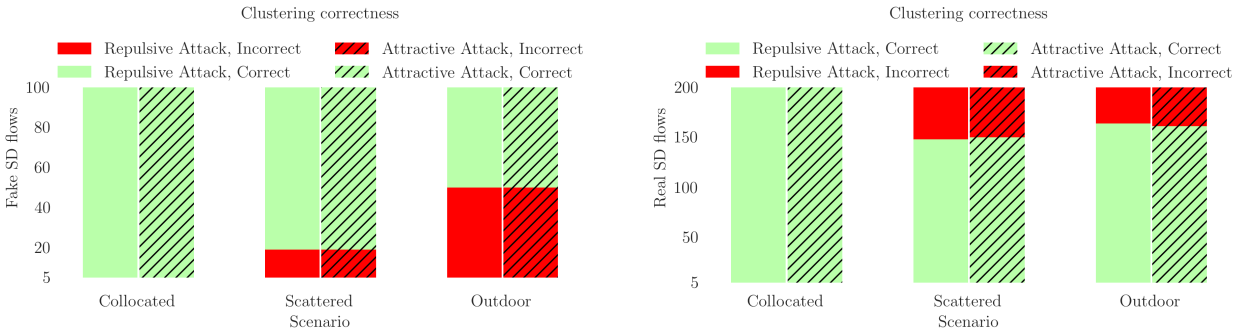

(a) (b)

Fig. 7. Clustering correctness results for repulsive and attractive attack, synthetic and camouflage trace-forging strategies and for collocated, scattered, and outdoor scenarios. (a) Synthetic attacking strategy. (b) Camouflage attacking strategy.

TABLE I
F1-SCORE AND MAXIMUM PERCENTAGE OF FAKE SD FLOWS TO INJECT IN MALICIOUS TRACES W.R.T. N TO AVOID CORRECT CLUSTERING.

		Synthetic strategy		**Camouflage strategy**	
Scenario	N	**Repulsive**	**Attractive**	**Repulsive**	**Attractive**
Collocated	508	1/—	1/—	1/—	1/—
Scattered	612	0.85/2.3%	0.85/2.3%	0.94/8.2%	0.95/8.5%
Outdoor	580	0.52/7.7%	0.52/7.7%	0.72/6.2%	0.70/6.7%

$[5, 200]$ and complement the malicious trace with fake SD pairs. The completeness of the fake SD pairs is according to the kind of attack. The collocated scenario remains difficult to hack, whereas in the scattered and outdoor scenarios, clustering is correct until real SD overcome 150 units (and fake SD flows become less than 50). In particular, regarding the scattered scenario, clustering is wrong when 52 and 50 fake flows remain in the trace for repulsive and attractive attack respectively, which, as shown in Table I, correspond to 8.2% and 8.5% w.r.t. N. This means that, an attacker must create a trace with less than 8.5% fake flows mixed with real SD flows to avoid further investigations. Regarding the outdoor scenario instead, the clustering becomes incorrect for repulsive and attractive attack with remaining 36 and 39 fake flows respectively, which correspond to 6.2% and 6.7% of N. Also in these conditions, attractive and repulsive attacks are detected almost at the same thresholds which are enough low to not negatively impact the network and the F1-score remains above 0.7.

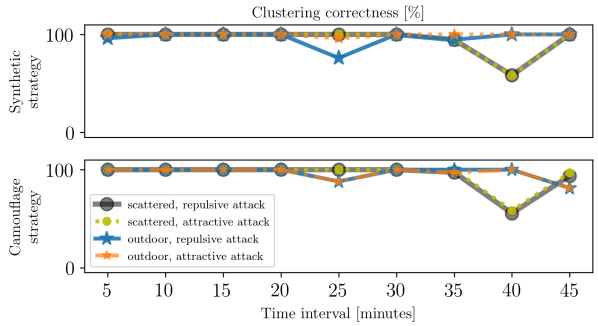

Fig. 8. Percentage of clustering correctness for online detection.

C. Online attackers detection

In order to be more proactive in attacking detection, we repeat the same workflow, clustering traces every five minutes instead of waiting for the entire traffic capture. We repeat experiments in scattered and outdoor scenarios, with all attack types and strategies, and with a malicious trace including at each repetition a raising number of SD flows from 5 up to the average recorded by legitimate traces in that time window. We compute the percentage of correctness as $100 \times \frac{\texttt{\# correct clustering}}{\texttt{\# repetitions}}$. As shown in Figure 8, clustering is 100% correct most of the time and always more than 76%, apart from minutes 40 to 45 in the scattered scenario. This happened because there was a legitimate monitor placed at the far end recording traffic from the outside. We conclude that online detection can be conducted considering periodical

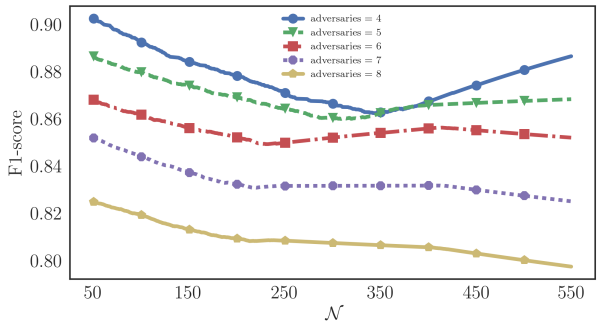

Fig. 9. F1-score for clustering with multiple adversaries.

interval of times in the order of a few minutes.

D. Collusion attack

Finally, we consider the case of more adversaries in collusion acting a repulsive, camouflage attack in the outdoor scenario as this is the most complex combination. In particular, up to 8 adversaries create traces embedding from 50 to 550 SD flows. To keep our initial assumption that malicious users are a minority, from every single original legitimate trace, we create nine extra traces, achieving 80 legitimate traces in total. These new traces have a random subset of the original trace SD set and completeness values picked from the same distribution of values as the original trace. Figure V-D shows the F1-score for every single clustering. We can notice that the score decays adding more and more adversaries, but remaining above 0.8 and it is just slightly dependent on the amount of flows N.

VI. Related Work

In this work, we are concerned with improving wireless networking measurements by increasing the number of monitoring nodes and measured traces in the same area of interest. There are solutions that advocate the use of radio monitors uniformly distributed over the target area, leading to higher complexity even in small scenarios [3], [17]. The main limitation of such approach is the lack of scalability since multiple monitoring devices have to be deployed. Another approach consists to optimize the number of monitors and their geographical position through trace similarity [5], [6].

Monitoring systems are also improved by relying on a combined approach, theoretical model plus experimental data. Finding a propagation model in an urban area is challenging because it has to deal with all sorts of obstacles such as streets, buildings, monuments, and so on. The trivial choice would be making an uncountable number of measurements until all possible regions are covered. This solution besides costly, can be unpractical depending on the area size. Hence, Robinson et al. proposed using digital maps to estimate the coverage of a certain number of access points based on theoretical models [18]. The goal is to characterize urban scenarios with only minimal measurements. First, they calculate the coverage borders of every access point on all directions, considering the attenuation introduced by the physical obstacles. Such coverage borders are found when a given evaluated metric reached its minimal value, e.g. SNR. After this first step, some measurements are conducted to refine the values found for the coverage borders. In opposition to our work, Chhetri et al. conducted a theoretical evaluation to compute the Quality of Monitoring (QoM), i.e., the amount of traffic collected by the monitoring system [19]. By using such a study, it is possible to improve the position of monitoring nodes.

As far as we know, works in the literature usually do not consider user participation during measurements. When they do, security issues regarding users' intervention are left aside. Ravindranath et al. enlarge the typical set of monitoring devices, normally limited to wireless interface cards, to encompass also other sensors available in mobile equipments, such as GPS, accelerometers, magnetic compasses, and gyroscopes [20]. By using such additional information, they can improve the network performance. For instance, a GPS could be used to obtain information about node mobility and, consequently, to calibrate routing protocol parameters. The work from Kanuparthy et al. is another which considers trustworthy user participation. They propose a tool for evaluating physical media conditions without needing any specific networking equipment [21]. All the measurements are conducted by users themselves and require the utilization of a probing tool and a server running in a PC connected to an access point. Based on information available in probes (One-Way Delay - OWD), it is possible to know delay components such as contention time, backoff, transmission delay, and certain constant interframe spaces as defined by IEEE 802.11 MAC protocol. Because the delay does not account for queuing at the sender, it is possible to estimate link layer properties and physical conditions.

Typically, malicious users are investigated by the measurement system. Paul et al., for instance, collect multiple traces to analyze interference among wireless nodes and, furthermore, to detect selfish behavior [22]. Authors claim that from trace analysis, it is possible to detect users who could maliciously gain access to the wireless medium by manipulating MAC protocol parameters, e.g., the backoff window size. Map [23] and DOMINO [24] are other monitoring systems specialized on capturing multiple traces from the network, merging them, and evaluating the presence of possible malicious users. All these systems have a different approach from ours, since in their work the malicious actions are external whereas in our work they can be part of the measurement system itself.

Even though also considering malicious users outside the monitoring system, Pedro Casas [25] compares some supervised machine learning models and techniques to detect attacks to networks. He observed that, based on his benchmarking, both neural networks and decision tree-based models provide the better results. In his work, the scenario is different from our own, he considered attacks in a wired network and anomalies in applications. In addition, the algorithms used for attack and anomaly detection are also different (supervised versus unsupervised) and the crowd is more related to multiple inputs than necessarily to the participation of multiple users. Instead

of providing local monitored traces, each crowd member provides the output of her machine learning technique.

VII. Conclusion

In this paper, we experimentally showed the impact of adversaries in collaborative wireless measurements for DDN. The geographical distribution as well as the number of collected traces can significantly improve the accuracy of the merged trace, allowing better informed decisions. To increase the number of monitoring nodes and traces, we rely on users' participation. On the flip side, malicious actions are possible.

In this direction, we identified two possible attacks, namely repulsive and attractive and two trace forging strategies: synthetic and camouflage. In addition, more adversaries might act in collusion. Thus, we proposed a detecting procedure, considering wireless traces as a set of SD flows. In according to the completeness of each SD flow, traces are represented as a vector indicating a point in a N dimensions space. Finally, applying an agglomerative clustering algorithm after a space dimensionality reduction, malicious traces are identified and separated from legitimate ones in a different cluster.

Results show that collaborative systems can indeed generate more complete information concerning the wireless monitoring of a target area. Since collaborative systems are worthy, we tested our attacker detection methodology in three different scenarios capturing real wireless traces, with two kinds of attacks and two trace-forging strategies. We show that in order to succeed, fake traces must include an important amount of fake flows which will be discovered by the detection system. Including very few fake flows will fool the detection but does not have the critical mass to achieve the desired effect. Our results, based on the F1-score and specific metrics, also show that the percentage of clustering correctness is maintained even if the attacking detection is performed online and in presence of more colluding adversaries.

As a future work, we plan to consider the case of mobile monitoring nodes, i.e., the effect of mobility on packet sniffing and trace merging. We also would like to define a trust metric to enhance the detection system and refine user behavior evaluation. The idea is to rely on trust to avoid misjudging monitoring nodes experiencing poor wireless connectivity.

References

[1] Cisco, "Cisco visual networking index: Forecast and trends, 2017 – 2022," Feb 2019.

[2] J. Jiang, V. Sekar, I. Stoica, and H. Zhang, "Unleashing the potential of data-driven networking," in *COMSNETS*, 2017.

[3] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage, "Jigsaw: solving the puzzle of enterprise 802.11 analysis," in *ACM SIGCOMM*, Sep.11-15 2006, pp. 39–50.

[4] T. Claveirole and M. D. de Amorim, "Manipulating Wi-Fi packet traces with WiPal: design and experience," *Software Practice & Experience*, vol. 42, no. 5, pp. 585–599, May 2012.

[5] M. Sammarco, M. E. M. Campista, and M. D. de Amorim, "Scalable wireless traffic capture through community detection and trace similarity," *IEEE Transactions on Mobile Computing*, vol. 15, no. 7, pp. 1757–1769, July 2016.

[6] M. Sammarco, M. E. M. Campista, and M. Dias de Amorim, "Trace selection for improved wlan monitoring," in *Proceedings of the 5th ACM Workshop on HotPlanet*, ser. HotPlanet '13. New York, NY, USA: ACM, 2013, pp. 9–14. [Online]. Available: http://doi.acm.org/10.1145/2491159.2491165

[7] P. Antoniadis, B. L. Grand, A. Satsiou, L. Tassiulas, R. Aguiar, J. P. Barraca, and S. Sargento, "Community building over neighborhood wireless mesh networks," *IEEE Technology and Society*, vol. 27, no. 1, pp. 48–56, Feb. 2008.

[8] C.-W. Yi, "A unified analytic framework based on minimum scan statistics for wireless ad hoc and sensor networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 20, no. 9, Sep. 2009.

[9] T. Razafindralambo, T. Begin, M. D. de Amorim, I. G. Lassous, N. Mitton, and D. Simplot-Ryl, "Promoting quality of service in substitution networks with controlled mobility," in *Ad Hoc Networks and Wireless (ADHOC-NOW)*, Jul.18–20 2011, pp. 248–261.

[10] I. Koutsopoulos, "Optimal incentive-driven design of participatory sensing systems," in *IEEE Infocom*, Apr. 2013, pp. 1402–1410.

[11] A. Boudries, M. Aliouat, and P. Siarry, "Detection and replacement of a failing node in the wireless sensors networks," *Computers & Electrical Engineering*, vol. 40, no. 2, pp. 421–432, Feb. 2014.

[12] B. B. Zarpelão, R. S. Miani, C. T. Kawakani, and S. C. de Alvarenga, "A survey of intrusion detection in internet of things," *Journal of Network and Computer Applications*, vol. 84, no. Supplement C, Apr. 2017.

[13] N. Bartolini, G. Bongiovanni, T. L. Porta, S. Silvestri, and F. Vincenti, "Voronoi-based deployment of mobile sensors in the face of adversaries," in *2014 IEEE International Conference on Communications (ICC)*, June 2014, pp. 532–537.

[14] D. Wu, R. K. C. Chang, W. Li, E. K. T. Cheng, and D. Gao, "Mopeye: Opportunistic monitoring of per-app mobile network performance," in *USENIX ATC*, Jul. 2017, pp. 445–457.

[15] L. Xue, X. Ma, X. Luo, L. Yu, S. Wang, and T. Chen, "Is what you measure what you expect? factors affecting smartphone-based mobile network measurement," in *IEEE INFOCOM*, May 2017, pp. 1–9.

[16] R. Bellman, *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press, 1957.

[17] J. Yeo, M. Youssef, and A. Agrawala, "A framework for wireless lan monitoring and its applications," 2004.

[18] J. Robinson, R. Swaminathan, and E. W. Knightly, "Assessment of urban-scale wireless networks with a small number of measurements," in *ACM MobiCom*, Sep.14–19 2008.

[19] H. N. Arun Chhetri, G. Scalosub, and R. Zheng, "On quality of monitoring for multi-channel wireless infrastructure networks," in *ACM MobiHoc*, Sep.20-24 2010, pp. 111–120.

[20] L. Ravindranath, C. Newport, H. Balakrishnan, and S. Madden, "Improving wireless network performance using sensor hints," in *USENIX conference on Networked Systems Design and Implementation (NSDI)*, Mar.30–1 2011, pp. 1–14.

[21] P. Kanuparthy, C. Dovrolis, K. Papagiannaki, S. Seshan, and P. Steenkiste, "Can user-level probing detect and diagnose common home-WLAN pathologies?" *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 1, pp. 7–15, Jan. 2012.

[22] U. K. Paul, M. M. Buddhikot, and S. R. Das, "Passive measurement of interference in WiFi networks with application in misbehavior detection," *IEEE Transactions on Mobile Computing*, vol. 12, no. 3, 2013.

[23] Y. Sheng, G. Chen, H. Yin, K. Tan, U. Deshpande, B. Vance, D. Kotz, A. Campbell, C. Mcdonald, T. Henderson, and J. Wright, "Map: a scalable monitoring system for dependable 802.11 wireless networks," *IEEE Wireless Communications*, vol. 15, no. 5, pp. 10–18, Oct. 2008.

[24] M. Raya, I. Aad, J.-P. Hubaux, and A. E. Fawal, "DOMINO: detecting MAC layer greedy behavior in IEEE 802.11 hotspots," *IEEE Transactions on Mobile Computing*, vol. 5, no. 12, pp. 1691–1705, 2006.

[25] P. Casas, "On the analysis of network measurements through machine learning: The power of the crowd," in *IFIP Network Traffic Measurement and Analysis Conference (TMA)*, Jun. 2018, pp. 1–8.