Secure Interference Reporting for Dense Wi-Fi Deployments

Pantelis A. Frangoudis, Dimitrios I. Zografos, and George C. Polyzos
Department of Informatics, Athens University of Economics and Business
{pfrag, zwgrafos, polyzos}@aueb.gr

ABSTRACT

We study the problem of interference detection in dense Wi-Fi deployments, which are a reality in most modern metropolitan areas. Interference among neighbor Wi-Fi cells stems from the anarchic deployment of Wi-Fi access points (APs) and the fact that only few APs can operate at the same location on non-overlapping frequencies. Detection of interference conditions is the first step towards its mitigation. We follow a client-centric approach, where wireless clients monitor and report Wi-Fi AP presence. We focus on the security aspects of such a scheme, assuming that clients may attack the reporting mechanism by submitting fake information. Our early evaluation shows that for certain attacker strategies, simple mechanisms can effectively filter invalid reports with minimal loss of information.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.2.0 [**Computer-Communication Networks**]: General—*Security and Protection*

General Terms

Design, Performance, Security

Keywords

Wi-Fi, Interference

1. INTRODUCTION

With the proliferation of IEEE 802.11-based WLAN equipment, Wi-Fi pervades modern metropolitan areas. The low cost and ease of installation of Wi-Fi equipment, as well as its operation in unlicensed spectrum are the main reasons for its popularity. While in densely populated urban areas wireless coverage is no more an issue, unplanned and anarchic deployment of Wi-Fi networks comes with the cost of interference. For IEEE 802.11b/g there are only 3 non-overlapping frequency bands (channels) on which an AP can operate. In the scenarios we study, the probability of coexistence of more than 3 WLANs at the same spot is high.

Combating interference in chaotic WLAN deployments necessitates sophisticated interference mitigation strategies involving, among others, power control or channel selection. Information on spectrum usage is vital input to them. AP-centric interference detection schemes fail to capture interference at client locations (*hidden interference* [3]). Thus, we chose to follow a client-driven approach, where clients scan for Wi-Fi presence and report their results to a collecting entity.

Clients cannot always be assumed trustworthy, though. There may be disincentives to contribute truthful information, such as the overhead of spectrum monitoring [1]. Also, in an environment where clients subscribed with Wi-Fi providers visit one another's hotspots, one may be tempted to submit fake reports to manipulate the spectrum sharing mechanism to his affiliated provider's advantage, in return for better service or other benefits. Fake reports pollute the system's view of interference conditions and, consequently, affect interference mitigation schemes. In the *Cognitive Radio* context, similar problems are faced by Distributed Spectrum Sensing mechanisms, where the main goal is to detect the presence of primary (licensed) users [1]. Here, we study some classes of fake reporting attacks and propose suitable mechanisms for effectively combating them.

2. SYSTEM MODEL

We model our system as a weighted undirected graph. Vertices represent APs and edges represent user-perceived interference between neighbor APs. Clients situated in areas where neighbor cells overlap report this fact and the weights of the respective edges of the *Interference Graph (IG)* are updated accordingly. Notice that no edges are added in the case of overlapping cells which are not reported (e.g. due to lack of clients located there or to their refusal to report).

A client submits his reports to the AP he is attached to[2], and the latter forwards it to a collecting entity, which maintains the IG. For simplicity, we assume that all APs operate on the same channel. Since clients may misbehave, the IG does not necessarily encode the actual interference conditions. Fake reports contribute fake IG edges. Our mechanisms aim at eliminating these edges and revealing true interference.

[1] An IEEE 802.11 active scan may take more than 250msec, during which time the client station cannot transmit/receive application data. More advanced spectrum usage measurements may be more time consuming.

[2] The recently standardized IEEE 802.11k protocol could be used by clients and APs in the reporting process.

CoNEXT Student Workshop'09, December 1, 2009, Rome, Italy.
ACM 978-1-60558-751-6/09/12.

3. ATTACKS AND COUNTERMEASURES

In the first scenario that we study, clients are assumed to act independently and their reports are considered of equal weight. For each interfering AP pair reported, a unit is added to the respective edge's weight. We assume no cooperation among attackers, each of whom submits reports containing a number of random AP identifiers. Each of these fake reports contributes unit-weight edges to the IG. Thus, pruning all unit-weight edges from the IG practically eliminates the probability that a fake edge appears in it.

In the second scenario, APs belong to a number of wireless service providers and each user is affiliated with one of them. There are two classes of users: *roamers*, i.e. those attached to APs belonging to a different provider and *non-roamers*. Non-roamers always submit truthful reports, while roamers are not always honest and may form *colluding groups* as follows: all roamers affiliated with provider A who are currently attached to an AP of provider B, agree to report the same fake set of random APs. The filtering scheme of the first scenario is useless here; these fake reports would contribute edges with higher weight to the IG, which the filtering mechanism would fail to detect.

To counter this attack, each AP values more reports that originate from *trusted* clients, i.e., clients affiliated with the same provider. Roamer reports are *discounted* so that their cumulative weight per AP does not exceed that of a single trusted report. Depending on the number of roamers associated with it, each AP independently calculates the weight assigned to the reports of each of these n roamers so that

$$\sum_{i=1}^{n} w_i = 1 - e, \quad 0 < e \ll 1 \tag{1}$$

After discounting, a roamer's report is forwarded to the collecting entity. Now, for each AP pair contained in it, the weight of the respective IG edge is incremented by w_i. The filtering mechanism prunes all IG edges which have weight less than 1. The following conditions are sufficient for an edge to be reported: (1) At least 1 non-roamer reports the edge, or (2) there are sufficiently many roamers reporting it, so that the sum of the weights of their reports is at least 1.

The above process strictly bounds the weight of an edge reported by a colluding group below 1, which leads us to the observation that in the second scenario, pruning all unit-weight edges also eliminates fake edges.

In both scenarios, the filtering mechanism's efficiency is only limited by potential *false negatives*, namely, real edges which fail to reach the unit-weight threshold.

4. EVALUATION

We estimate the detection accuracy of our scheme via simulations. Our evaluation metric is the percentage of IG edges that are detected, after filtering has taken place. We simulated the topology of the Dartmouth Campus Wi-Fi network based on data available from the CRAWDAD archives. The simulated network included 521 APs whose locations were taken from [2]. For each AP, we set its transmission range to 80m and randomly placed 6 clients within its range. Clients scan for APs and report their (potentially fake) results, before the filtering mechanism is executed. After filtering, we compare the resulting IG with the real one. Our simulations were carried out in *ns2*[3].

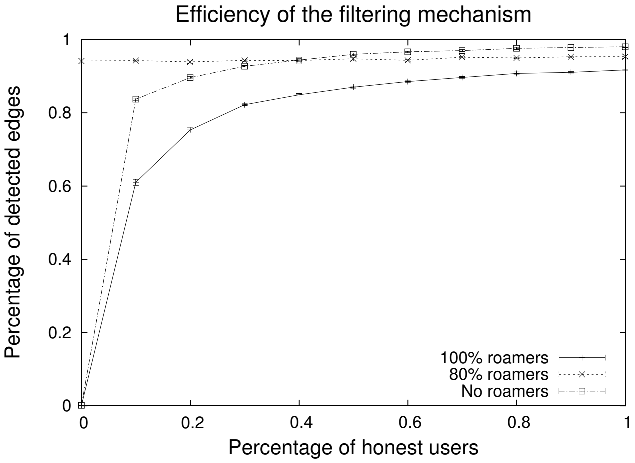

Figure 1: Efficiency of the filtering mechanism. The "No roaming" curve represents the first attack scenario

For the second attack scenario, all non-roamers are assumed truthful and all roamers associated with an AP form a single colluding group and submit the same fake report. Fig. 1 shows the percentage of real IG edges detected. When 80% of the clients are roamers, even if none of them is truthful, the rest of the (trustworthy) non-roaming users account for the high detection accuracy. In such a dense deployment (both wrt. to clients and APs), our results indicate that the simple mechanisms that we propose can effectively combat the described attacks without significant loss of information, even for large numbers of dishonest nodes.

5. CONCLUSION

We used a client-driven approach for detecting interference in dense Wi-Fi deployments, discussed potential attacks and proposed simple mechanisms to combat them. Our early evaluation has shown that under certain circumstances, these mechanisms can prove effective. Still, important issues need to be studied, with the effects of user mobility on the proposed mechanisms' performance and the potential of applying a reputation scheme to evaluate user reports topping the list. Also, more sophisticated attacks need to be studied.

6. REFERENCES

[1] R. Chen, J.-M. Park, Y. T. Hou, and J. H. Reed. Toward secure distributed spectrum sensing in cognitive radio networks. *IEEE Communications*, April 2008.

[2] D. Kotz, T. Henderson, I. Abyzov, and J. Yeo. CRAWDAD trace dartmouth/campus/movement/aplocations (v. 2004-11-09). Downloaded from `http://crawdad.cs.dartmouth.edu/dartmouth/campus/movement/aplocations`, August 2009.

[3] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. A. Arbaugh. A Client-Driven Approach for Channel Management in Wireless LANs. In *Proc. IEEE INFOCOM 2006*, Barcelona, Spain, April 2006.

[3] `http://www.isi.edu/nsnam/ns/`