

A Walk on the Client Side: Monitoring Enterprise Wifi Networks Using Smartphone Channel Scans

Jinghao Shi*, Lei Meng[†], Aaron Striegel[†], Chunming Qiao*, Dimitrios Koutsonikolas* and Geoffrey Challen*

*Department of Computer Science and Engineering

University at Buffalo, Buffalo, New York 14260

Email: {jinghaos, qiao, dimitrio, challen}@buffalo.edu

[†]Department of Computer Science and Engineering

University of Notre Dame, Notre Dame, Indiana 46556

Email: {lmeng, striegel}@nd.edu

Abstract—During the one minute it takes to read this abstract, two billion smartphones worldwide will perform billions of Wifi channel scans recording the signal strength of nearby Wifi Access Points (APs). Yet despite this ongoing planetary-scale wireless network measurement, few systematic efforts are made today to recover this potentially valuable data.

In this paper we ask the question: “Are the smartphone channel scans useful in monitoring enterprise Wifi networks?” More specifically, can these client-side measurements provide new insights compared to the AP-side measurements that enterprise Wifi networks already perform? Beginning with two Wifi scan datasets collected on two large scale smartphone testbeds, we conduct case studies that show how smartphone channel scans can be used to (1) improve AP spectrum management, and (2) predict the impact of AP failure or overload. In each case, a walk on the client side yields valuable insights for network operators that are otherwise impossible to gain from AP-side measurements, and together our results demonstrate the value of smartphone channel scans.

I. INTRODUCTION

As mobile wireless devices proliferate and the demands they place on wireless networks continue to grow, monitoring the health and performance of large scale wireless networks is both essential and increasingly challenging. While site surveys and infrastructure-side data collection both play important roles in network monitoring, neither of these approaches can fully and continuously reflect actual network conditions experienced by the clients the network is intended to serve—particularly the mobile devices that generate most wireless traffic. Site surveys are snapshots that are neither spatially nor temporally representative, while AP-side data collection can only measure wireless conditions at AP fixed locations, not at the changing locations of mobile clients. As a result, multiple previous research studies have demonstrated the value of client-side measurements to improve network performance [13], [14]. Similarly the 802.11k amendment [1] provides a mechanism to allow APs to collect this information from clients to improve radio resource management.

Given the established interest in and emerging support for client-side wireless network measurement, it is surprising that today we are discarding a potentially valuable source of client-side information: the large number of channel scans generated continuously by mobile devices such as smartphones. These devices naturally perform frequent channel scans to cope with rapidly-fluctuating network environments caused by mobility. For example, Android smartphones scan every 15 seconds when unassociated to search for APs [25], and continue to

scan every 60 seconds even when already associated in order to identify better APs as wireless conditions change caused by client mobility, network interference, or other factors.

When considered as network measurements, channel scans share all the typical benefits of client-side measurements—such as capturing the real conditions experienced by end users, providing data that is impossible to gather using APs alone, and being simpler and more representative than site surveys. In addition, channel scans have several other unique benefits. First, note that clients spontaneously perform these channel scans in order to effectively use wireless networks. As a result, the client overhead required to collect them for monitoring purposes is limited to the costs of temporary storage (which is minimal) and telemetry (which can be minimized through delay tolerant and energy neutral data collection). Second, in many cases channel scans remain useful even after being stripped of information such as timestamps or device identifiers that could threaten client privacy and potentially limit the willingness of users to share this information.

While all Wifi clients perform channel scans, smartphones have several advantages that make them uniquely suited to collecting and providing measurements for monitoring enterprise wireless networks. First, unlike laptops, smartphones are always powered on, thus scan continuously both during and between periods of interactive use, providing better temporal resolution than devices that are regularly powered down or put to sleep. Second, being highly portable, users are likely to carry smartphones most of the time, giving these devices the ability to observe more of the network than would be seen by stationary or less-portable devices, making their measurements more representative of all locations where mobile users might utilize the wireless network. Finally, smartphone platforms already provide interfaces allowing apps (with appropriate permissions) to access channel scans, and app marketplaces provide an easy way to deploy the simple monitoring software required to collect them from billions of mobile devices. Compared to other mobile devices, smartphones make it easy to collect an enormous number of channel scans from everywhere that users might use the wireless networks.

But are these measurements actually useful? Can channel scans contribute valuable insights in monitoring enterprise wireless networks that would otherwise be impossible to gain from existing infrastructure-based approaches? That is the question that we set out to answer in this paper. Our goal

is to determine whether the billions of discarded smartphone channel scans represent a missed opportunity or redundant information that we can continue to safely ignore.

To do so, we take a walk on the client side. We utilize two Wifi channel scan datasets collected on two large scale smartphone testbeds (PHONELAB: 5,373,682 scans, 254 devices, 5 months; and NETSENSE: 32,564,809 scans, 125 devices, 32 months), described in more detail in Section II. To examine the case for *client-side* measurement, we ignore analyses that could be performed using measurements from APs, including APs that utilize extra radio hardware to continuously perform measurements in parallel with normal client traffic. We also focus our studies on aspects of network performance and behavior of interest to network operators, not just scientists, since we anticipate that large-scale collection of channel scans will only take place if the measurements are able to provide meaningful operational insights.

Our paper presents two case studies in Section III using one or both of our channel scan datasets:

- **Section III-A** looks at how the channel conflict graph differs from the client perspective, helping operators assess the effectiveness of infrastructure-only channel assignment algorithms at utilizing available spectrum.
- **Section III-B** investigates how load can be shifted among neighbor APs when a particular AP is not available, both for load balancing purpose and to evaluate network redundancy, allowing network operators to make better spatial planning decisions.

Section IV discusses related work before we present and discuss our conclusions in Section V. We believe that our case studies demonstrate the value of smartphone channel scans, and will publish the datasets so that others can extend our analyses or reevaluate our conclusion.

II. THE DATASETS

To examine the usefulness of Wifi scan results collected by smartphones, we analyze two large scan datasets¹ collected from two smartphone testbeds deployed in University at Buffalo (UB) and University of Notre Dame (ND): 5.3M scans from PHONELAB [15] at UB, and 32M scans from NETSENSE [23] at ND. Throughout the paper we refer to these datasets as **UB-Scan** and **ND-Scan**, respectively. Statistics summarizing both datasets are shown in Table I. Pending publication of this study, the UB-Scan and ND-Scan datasets will be made available to researchers for further study.

In addition, to compare the client- versus AP-side perspectives, we both (1) obtained the logs generated by the management software operating the Wifi network of UB and (2) performed additional data collection to address the limitations of that management tool. We refer to these two datasets as **UB-AP** and **UB-AP-Scans**, and describe them in more detail in Section II-D.

	PHONELAB	NETSENSE
Description	§II-A	§II-B
Identifier	UB-Scan	ND-Scan
Start	11/7/2014	5/1/2012
End	4/3/2015	3/31/2015
Duration (Days)	147	974
Participants	254	100–125
Device Type	Nexus 5	Mixed
Scans	5,374,406	32,564,809
Observed APs	30,604	72,001
Used APs	2742	2495
Wifi Sessions	160,886	149,863
Total Connection Time (Days)	23,322	50,969

TABLE I: **Dataset Summary.** Only Wifi scans and sessions observing the campus network are counted. Used APs refers to the subset of total APs that were used by the devices participating in the study. Total connection time includes only Wifi sessions with campus APs.

A. UB-Scan: PHONELAB Wifi Scan Dataset

PHONELAB is a large scale smartphone platform testbed at UB. Several hundred students, faculty, and staff carry instrumented LG Nexus 5 smartphones as their primary device, and receive discounted service in return for providing data to smartphone experiments. PHONELAB participants are distributed across university departments, making the results representative of the broader campus wireless network users.

The PHONELAB Android platform was instrumented to log the Wifi scan results and Wifi connection events naturally generated by the system. Note that the platform modifications are not necessarily required to collect such measurements: equivalent data can be collected by apps with the right permissions—as demonstrated by the ND-Scan dataset described next.

Each scan result contains multiple entries, one for each nearby Wifi AP observed by the smartphone. One entry includes the (1) scan timestamp, (2) AP SSID and BSSID, (3) RSSI and (4) AP channel. For this paper, we are only interested in scans that observe the UB campus network, therefore we remove scans that do not contain any UB campus APs.

B. ND-Scan: NETSENSE Wifi Scan Dataset

The ND-Scan dataset uses data from the NETSENSE study conducted at ND. NETSENSE participants were spatially concentrated in six undergraduate dormitories, with demographics (gender, major, and income) verified to be representative of the larger undergraduate population.

During the first two years of the study, NETSENSE participants were provided Nexus S devices flashed with the Cyanogenmod fork of the Android Open Source Project and running a user-level data collection app. In August 2013, participants were given the option to continue the study by purchasing their own replacement handset but continuing to receive free service, and fifty additional participants were recruited to replace those that chose to quit. From this point onward, NETSENSE relied only on the user-level data collection app.

The NETSENSE data collection app recorded scan results every three minutes including the (1) scan timestamp, (2) AP

¹Full IRB approval was obtained for the acquisition of both datasets.

SSID and BSSID, and (3) RSSI. Unlike UB-Scan, channel information was not recorded. Beginning in May 2012, Wifi connection events were also logged. For this paper, we utilize only the data collected from 5/1/2012 to 3/31/2015.

C. Differences Between the Scan Datasets

Compared to the UB-Scan dataset, NETSENSE devices recorded fewer sessions per participant day (1.5) than PHONELAB devices (4.3), despite logging similar numbers of session hours per participant day: 12.5 for ND-Scan v. 15.0 for UB-Scan. We believe that this is largely due to the difference between the Nexus S used by NETSENSE participants during the first two years of the ND-Scan dataset and the Nexus 5 used by PHONELAB participants during the entire UB-Scan dataset. In particular, the Nexus S is known to have poor Wifi sensitivity and is reluctant to connect to APs with low signal strengths. In addition, NETSENSE participants are all undergraduate students and spent various amounts of time on-campus during the three-year study period, leaving regularly for the summer or study-abroad programs. In contrast, PHONELAB participants are mostly faculty and staff and would have been mostly on campus during the six-month study.

D. UB Wifi Logs and AP Scans

To compare the client- and AP-side perspectives, we first obtained access to the system logs generated by the Cisco Prime system used to manage the campus Wifi network of UB. This dataset contains 8,041,604 Wifi sessions from 38,067 UB campus network users for 44 days from Mar 12 to Apr 25, 2015. Each record contains the following pieces of information: 1) the client's MAC address, 2) the AP's SSID and BSSID, 3) when the Wifi session began and ended, and 4) statistics such as bytes received and transmitted by the AP during the session. We also obtained an inventory of all UB campus APs, including their BSSIDs and course-grain location (campus, building and floor). Collectively we refer to this dataset as **UB-AP**.

Unfortunately, the Cisco Prime interface does not expose all information collected by the infrastructure network. For example, despite the fact that UB campus APs clearly perform periodic channel scans for purposes such as optimizing channel assignment and detecting rogue APs, we were unable to access the raw scan information—which is either not collected or hidden behind a proprietary database and not exposed to network administrators.

To address this limitation of the UB-AP dataset, we augment it with a more detailed dataset for the 14 APs on the 3rd floor of the CSE department building at UB. To reconstruct scan results from these APs, we colocated Nexus 5 smartphones on top of each AP and configured them to perform channel scans every second for 30 minutes, resulting in 1574 scans per AP on average. We configured a high scanning rate to try to compensate for the fact that smartphones typically have less sensitive radio hardware than commodity APs, but there is no way to perfectly account for these hardware differences, so our dataset should be seen as an approximation of the scans

that could have been collected by the colocated APs. We refer to this dataset as **UB-AP-Scans**.

III. CASE STUDIES

To see whether smartphone scans can generate new insights beyond what are already possible to gain from infrastructure side measurements, we conduct two extensive case studies to show how smartphone measurements can: (1) improve network spectrum management (§III-A), (2) help network operators make better spatial planning decisions (§III-B).

A. Spectrum Management

Channel assignment plays an important role in wireless network performance, and is typically modeled as a graph coloring problem on conflict graph $G = (V, E)$, where V is the set of APs, and $\langle AP_i, AP_j \rangle \in E$ if AP_i and AP_j interfere with each other when they are in the same channel. Previous works [13], [14] have shown that conflict graph constructed with only AP side measurement fails to capture all types of interference due to the hidden terminal problems. For instance, suppose AP_i and AP_j are beyond each other's communication range, with only AP's measurements, $\langle AP_i, AP_j \rangle \notin E$. However, a client that is associated with AP_i may still experience interference from AP_j .

In this section, we first show how the smartphone measurements can help build a more representative conflict graph that captures the interference experienced by clients. Then we demonstrate how to use this conflict graph to reduce client-perceived conflicts.

1) *Client-Assisted Conflict Graph Construction*: The conflict graph can be constructed with only infrastructure side measurements: each AP performs a Wifi scan and inserts an edge between itself and each of its neighbors. We refer to a graph constructed in this manner as the *Infrastructure-perceived* conflict graph, or $G_I = (V, E_I)$, where V is the set of APs, and $\langle AP_i, AP_j \rangle \in E_I$ if AP_i can overhear AP_j 's beacon frame, or vice versa.

To construct G_I , we need the AP-side scan results. Due to the limitation aforementioned in Section II-D, we are only able to construct G_I for the 14 APs in our department building, as shown in Figure 1. With the UB-AP-Scans dataset, for each AP_i , we add $\langle AP_j, AP_i \rangle$ to E_I if AP_j shows up in AP_i 's scan results. Note that in UB-AP-Scans dataset, we did not find any asymmetric AP pairs, effectively making G_I undirected.

However, from the Wifi client's perspective, any AP (other than its currently associated AP) within its carrier sensing range has potential conflict with itself, causing either extra backoff delay for uplink packets, or collisions for downlink packets. Therefore, a more representative conflict graph should also include edges between the client's associated AP and all other APs that appear in the client's scan results during the Wifi session. We refer to such graph as *client-perceived* conflict graph, or $G_C = (V, E_C)$, where V is the set of APs, and $\langle AP_i, AP_j \rangle \in E_C$ if any of AP_i 's clients can overhear beacon frames from AP_j , or vice versa.




Fig. 1: Floor Plan and AP position.




Fig. 2: Infrastructure-perceived Conflict Graph.




Fig. 3: Client-perceived Conflict Graph.

We use the UB-Scan dataset to construct G_C as follows. Given a scan result from a smartphone that was associated with AP_i , we add $\langle AP_j, AP_i \rangle$ to E_C if AP_j appears in the result. We also count the number of unique $\langle device, timestamp \rangle^2$ tuples for each conflict edge as its weight, to quantify the impact of the conflict: a large weight means the conflict is experienced by many devices and/or for a long period of time.

Figure 2 and Figure 3 show the constructed infrastructure-perceived and client-perceived conflict graphs respectively, where APs are positioned by their physical coordinates in Figure 1. First, we observe that G_C contains many edges that do not exist in G_I , which implies that clients experience more conflicts than the infrastructure can identify. Second, in this particular AP placement, we notice that E_I is a *proper subset* of E_C , which means our dataset not only reveals hidden conflicts, but also detects all conflicts that the APs already see.

We then take a closer look at each node's degree in both conflict graphs. The median node degree in G_I is 7, which is not surprising given the dense AP deployment in this floor. In G_C , the min degree is 10 and the median degree is 13. In reality, however, not all the conflict edges exist all the time: while G_I is stable, G_C may vary over time due to client mobility or association behavior. Therefore, such temporal fluctuations of G_C must be captured to effectively assign channels without wasting temporal channel reuse opportunities.

We seek ways to learn such temporal fluctuation patterns using our dataset. Specifically, we investigate *hourly* patterns. For each edge in G_C , we first obtain the set of unique $\langle device, timestamp \rangle$ tuples that we used to calculate edge weight, then we bin the tuples by the timestamp's *hour* field, and thus get the number of tuples in each hour of day.

Figure 4 shows the hourly tuple count distribution for all client-only conflict edges ($E_C - E_I$). As expected, all conflict edges are mostly seen during school hours (10 AM to 6 PM). Furthermore, different edges are mostly seen at various hours indicated by the darkest tile of each conflict edge. With such temporal fluctuation information, we can now construct the

time-variant $G_C(h) = (V, E_C(h))$, where h is a given hour of day, and $E(h)$ contains all the stable edges plus the edges that are reported at time h .

2) *Channel Assignment*: We then look at the effects of hidden conflict edges to the channel assignment of campus APs. According to UB IT staff, all campus APs are connected to central controllers, which collect the interference information from the APs and adjust each APs' channel to reduce conflicts. However, the measurement process and the channel assignment algorithm are all handled behind the scene, and little details were revealed or documented. Therefore, we again focus on the campus APs in our department building.

To monitor the APs' operating channels, we placed three smartphones in locations marked as triangles in Figure 1, and verified that the smartphones together can see all 14 APs in this floor. We configure the smartphones to perform a Wifi scan periodically. The experiment lasted for 6 days from Apr 23 to Apr 28, 2015. We fused the data from all devices and obtain each AP's channel history during the experiment period. In total, we detected no channel switch events in the 2.4 GHz band and 157 switch events in the 5 GHz band, which suggests the campus controllers only try to adapt channels in the 5 GHz band. Therefore, we hereby focus on the 5 GHz band only.

There are 9 orthogonal channels in the 5 GHz band in United States. We verified that the conflict graph G_I for the 14 APs is at least 7-colorable using the *largest first* strategy described in [10]. However, we still observed channel conflicts among the 14 APs during the experiment period. This is probably because those APs also have conflicts with campus APs in other floors, thus the actual conflict graph is larger than G_I we constructed.

We compare three channel assignment schemes: the observed channel assignment by campus AP controllers (*OBSERVED*), optimal channel assignment on G_I (*AP-OPT*) alone, and optimal channel assignment on $G_C(h)$ (*CLIENT-OPT*). For each scheme, we calculate the number of conflicts on $G_C(h)$ for each school hour (8 AM–6 PM) during our experiment period. This metric captures the actual number of conflicts experienced by the clients.

Figure 5a shows the CDF of number of conflicts of all 60 school-hours (10 hours/day $\times$ 6 days). We can see that when

²Timestamps are binned by hour—multiple scan results within an hour from the same device are only counted once towards the edge weight—so the weight is not biased by certain busy devices.






Fig. 4: $\langle \text{Device}, \text{timestamp} \rangle$ Count Distribution for Client-Only Conflict Edges. Hours before 8 AM and after 8 PM are omitted since their counts are all near 0. Edges are sorted by their node pair. Each tile are shaded to reflect the number of tuples in the hour—darker tile represents more conflicts.



(a) CDF of Conflict Edge Number of Different Channel Assignment.



(b) Number of Conflict Edges Over Different Window Size.

Fig. 5: Channel Assignment on Client- and Infrastructure-Perceived Conflict Graph.

considering only G_I , both the controller (*OBSERVED*) and even the optimal channel assignment (*AP-OPT*) cause many conflicts from client's perspective. With client's feedback, however, an optimal channel assignment on G_C (*CLIENT-OPT*) can be conflict free 50% of the time, and no more than 3 conflict edges at any time.

In practice, controllers may not be able to reassign channels as frequently as once per hour. Therefore, we extend the analysis to longer intervals. For a time window of $W (W > 1)$ hours, we obtain the conflict graph by combining the conflict edges from each hour in the time window. We then compare the number of conflicts of *AP-OPT* and *CLIENT-OPT* scheme. Figure 5b shows the results for $W = 2, 5$, and 10 respectively. As expected, both coloring schemes incur more conflicts than single hour case, since there are more edges in the graph to be colored. However, *CLIENT-OPT* can achieve significantly fewer conflicts than *AP-OPT*.

3) *Campus Network Conflict Graph*: We now examine the entire campus wireless network. The goal is to understand the complexity of large scale production wireless network deployments. We first construct campus wide G_C using the method described in Section III-A1. Figure 6a shows the CDF of edge weight for both UB and ND. Most conflict edges are seen rarely for both campuses, as both testbeds' participants only constitute a tiny portion of the entire campus population—0.5% for ND, and 0.6% for UB. To obtain more meaningful insights, we hereafter filter out conflict edges that are seen less than 10 times.

Figure 6b shows the CDF of node degree after filtering. The median degrees are 3 for ND and 5 for UB, thus the AP deployments at most of the campus areas seem not as dense as our department building, where the median conflict degree is 13. Again, these results may be biased by our small sample. We expect more conflict edges to be discovered with increased user coverage.

Even with such small user samples, however, there are certain APs with high conflict degrees in both campuses. For instance, at UB, 5% of APs (~ 90) have degrees larger than

20, posing challenges to spectrum management in such dense environments. Temporal patterns can be learned for such APs to discover channel reuse opportunities.

4) *Discussion*: Using the datasets, we demonstrate that the measurements from smartphones can help build a more representative conflict graph and do better channel assignment. Furthermore, patterns can be learned from longitudinal measurements to adapt for temporal fluctuations in order to discover channel reuse opportunities.


Note that it is also possible to collect client side conflict information in real time. For instance, 802.11k Radio Resource Measurement (RRM) amendment [1] defines a measurement type called *Frame Report*, in which the AP asks the client to report the number of packets the client receives from each unique transmitter on all channels. The AP can then use this information to infer other APs within the client's vicinity. However, such active measurements are disruptive since they require the client to switch channels.

B. Spatial Planning

In enterprise wireless deployments, APs are usually scattered evenly across a building to achieve good coverage, as in the case of our department building shown in Figure 1. However, there typically are certain public areas within the building, such as lounges and conference rooms, where users tend to use the wireless network more than other areas.

We first use the infrastructure side logs to study how the existence of public areas affects the load distribution of the campus APs. We use three metrics to quantify an AP's load during a period of time: number of Wifi sessions served, total duration of Wifi sessions, and total traffic of all sessions. For each of the 14 APs, we calculate these three load metrics during a 44 day period from 03/12/2015 to 04/25/2015.

Figure 7 shows the relative load distribution. We make several observations. First, as expected, the load of APs near public areas, such as AP 3, 4, 5, and 12, are much higher than others. Second, interestingly, certain APs near public areas, such as AP 7, are not as heavily loaded as other nearby APs.



(a) **CDF of Edge Seen Count.** Timestamps are binned by hour. 50% of edges are seen less than 2 (ND) or 3 (UB) times. (b) **CDF of Node Degree.** Edges seen less than 10 times are ignored.

Fig. 6: Characteristics of Campus-Wide Conflict Graph.

Moreover, loads are not evenly distributed among APs that are near the same public area, such as AP 3 and 5. Finally, we notice that the three load metrics may produce inconsistent AP load ranks: some APs may serve fewer Wifi sessions while still providing longer connection time or more traffic. In the following discussion, we use session count as the load metric, since the other two can be biased by factors such as applications or user behaviors.

From infrastructure logs, network operators can easily identify the hot spot APs that are more heavily loaded than their neighbors. Two actions can then be taken to deal with such load imbalance: *redirection* and *reposition*. With *redirection*, underutilized APs are left as is, so that when nearby hot spot APs are congested, clients can be redirected to them for load balancing. On the other hand, the underutilized APs can be *repositioned* to better locations to improve their utilization.

With only infrastructure measurements, however, it is not clear which approach should be taken for each of the underutilized APs. Furthermore, in both approaches, there are inherent challenges that are difficult to resolve. When redirecting the clients from a congested AP with a good signal to an idle AP with a possibly worse signal, it is not clear what the impact on clients' network performance will be. In addition, when there are multiple nearby underutilized APs as offloading candidate, it is difficult to determine each AP's signal-load tradeoff from clients' perspective. Finally when removing or repositioning underutilized APs, it is challenging to predict how their load would redistribute to nearby APs, or whether a coverage hole will be created.

In the rest of this section, we first show how smartphone measurements can be used to build an empirical load balancing graph to help the network operator make better load offloading decisions. Then we analyze how an AP's load would be redistributed upon removal, to further help network operators evaluate the impact of AP repositioning.

1) *Load Balance Graph:* In enterprise wireless environment, where all APs use the same set of SSIDs and authentication method, clients' association priority is solely based

on each AP's signal strength. However, it is well known that Wifi clients are usually reluctant to roam to new APs unless the associated AP's signal strength significantly drops below a certain threshold, which is as known as the "sticky client" problem. Therefore, the associated AP may not provide the best signal throughout the entire Wifi session.

From load balancing perspective, it is useful to identify the subset of an AP's neighbors which can potentially provide better signal to the AP's clients, such that when the AP is overloaded, its clients can be redirected without compromising their signal quality. With smartphone measurements, we can capture such relationships using an *empirical load balancing graph* $G_L = (V, E_L)$, where V is the set of APs, and $\langle AP_i \rightarrow AP_j \rangle \in E_L$ if a client associated with AP_i reports a better signal from AP_j than AP_i . We also assign a weight to each edge to quantify how many times such relationship is observed in the dataset.

Figure 8 shows the G_L constructed for the 14 APs in the 3rd floor of our department building. Such a graph is useful in two ways. First, it identifies backup APs for each hot spot AP for load balancing purposes. For example, a large portion of AP 5 and 11's load can be shifted to AP 2 and 9 respectively without degrading client perceived signal quality. Second, it helps to decide which action to take for each underutilized AP. For instance, edges with large weight, such as $\langle 5 \rightarrow 2 \rangle$, $\langle 4 \rightarrow 0 \rangle$ and $\langle 11 \rightarrow 9 \rangle$, indicate that those underutilized APs (2, 0 and 9) are better kept for load balancing purposes. On the other hand, APs such as 7, 10 and 13 seldom provide better signal than nearby hot spot APs, thus should be considered as reposition candidates. However, even some of these candidate APs can not provide good signal for load offloading purposes, they may exclusively serve certain users, thus removing them will probably create coverage holes. We further look into this in the next section.

2) *Load Redistribution Graph:* We now study how the load of a removed or broken AP would be redistributed among its neighbor APs. This information is useful in two ways. First, when repositioning underutilized APs, it is important to make




Fig. 7: Relative AP Load in Department Building. Each load metric is normalized by the maximum value in that category.




Fig. 8: **Load Balance Graph.** Nodes are shaded based on relative load. Edge width corresponds to $\langle \text{device}, \text{timestamp} \rangle$ count in our dataset.




Fig. 9: **Load Redistribution Graph.** Edge width corresponds to the number of Wifi sessions shifted from the originating node.




Fig. 10: **CDF of Campus AP Redundancy.**

sure that the removal of the AP neither increases the burden of nearby hot spot APs that are already heavily loaded, nor creates coverage holes for the clients that were only able to be connected to the removed APs. Second, when certain APs stop accepting connections either temporally (e.g., due to scheduled maintenance) or permanently (e.g., due to failure), network operators can use such information to evaluate the impact of such unusual events on clients' network connectivity.

To this end, we can build a *load redistribution graph* $G_R = (V, E_R)$ with smartphone measurements as follows. Given the Wifi scan result $R = (AP_1, AP_2, \dots, AP_n)$ that is reported *just before* a Wifi connection with $AP_c \in R$, we remove AP_c from R and find the AP with the best signal, denoted by AP_s . We draw an edge, $\langle AP_c \rightarrow AP_s \rangle$, and increase its weight by 1. Intuitively, this means if AP_c is removed, the client's next Wifi session will be shifted to AP_s . In a special case where AP_c is the only AP in R , we draw a loop edge $\langle AP_c \rightarrow AP_c \rangle$ and increase the edge weight by 1, indicating that this next session would be lost if AP_c is absent. To make R accurately reflect the device's network condition at the time of association, we require that R is reported no earlier than 10 seconds before the connection event, and Wifi sessions without such scan results are not counted towards G_R .

Figure 9 shows the load redistribution graph for the 14 APs. Among the underutilized APs, AP 7 has medium incoming weights, indicating it is capable to provide backup connections in case of neighbor AP failures. On the other hand, APs 10 and 13 have very small incoming weights, suggesting they may not be able to provide good signal for clients that need to be redirected when neighbor APs fail. Additionally, there are no loop edges for AP 10 and 13, meaning all the Wifi sessions they serve can be safely shifted to nearby neighbor APs when they are removed. Therefore, AP 10 and 13 can be repositioned to better places to increase their utilization. However, we must note that the above conclusion is biased by the Wifi sessions of PHONELAB participants, thus may not be representative of all campus AP users. Further investigations need to be conducted to reach a final decision.

We also observe from Figure 9 that, for those heavily loaded APs (e.g., AP 4 and 5), most of their loads can be shifted to

nearby APs, although certain number of connections will be dropped. To quantify this, we define the *redundancy* of an AP as follows:

$$\mathcal{R}(AP_i) = 1 - \frac{w_{AP_i \rightarrow AP_i}}{\sum_{j|AP_i \rightarrow AP_j \in E_R} w_{AP_i \rightarrow AP_j}} \quad (1)$$

where E_R is the edges of redundancy graph, $w_{AP_i \rightarrow AP_j}$ is the weight (i.e., number of sessions) of edge $\langle AP_i, AP_j \rangle$ in the load redistribution graph, and $w_{AP_i \rightarrow AP_i}$ is the weight of AP_i 's loop edge. A large redundancy means that most of AP_i 's connections can be safely shifted to its neighbors.

We are now ready to analyze the redundancy of campus APs. To make the analysis statistically meaningful, we only calculate the redundancy of APs whose total outgoing weight is larger than a certain threshold (100). In total, we found 396 such APs from UB-Scan dataset and 426 APs from ND-Scan dataset. Figure 10 shows the CDF of these APs' redundancy. The median redundancy is as high as 0.95 for UB and 0.9 for ND due to the dense AP deployment within campus. We do notice, however, that for about 5% of the APs in each campus, the redundancy is lower than 0.5. Such APs shall raise the attention of network operators and may require further investigation for better spatial planning.

3) *Discussion:* We studied how to use the smartphone measurements to help the network operator make load balancing decisions and evaluate the redundancy in AP spatial planning. Note that the accuracy and usefulness of the two graphs largely depend on the number of campus Wifi users who contribute measurements. In both graphs, the edge weight and even the edge existence are biased by the subset of users who report scan results. We discuss possible incentives to encourage participation in data collection in Section V-B.

In addition, similar to the conflict graph discussed in Section III-A, the load balancing graph could potentially be built in real time fashion. For instance, the congested AP could trigger a *Frame Report* on its clients and eagerly disassociate a client if it received better signal from another campus AP. However, such real time coordination may further burden the congested APs, thus an empirical graph built offline can still be useful.

IV. RELATED WORKS

Monitoring Wifi networks has received significant attention over the past decade. Early works [4], [7], [12], [20] used the existing AP infrastructure, SNMP logs, and traces collected on the wired side of the WLAN to analyze Wifi traffic. Yeo et al. [26], [27] introduced the idea of passive monitoring using a small number of wireless sniffers. The same approach was used in a number of follow-up works [8], [9].

The next step was the deployment of large-scale passive wireless monitoring systems. Jigsaw [6], [5] was the first deployed large-scale monitoring system. 192 sniffers were deployed to report all wireless events across location, channel, and time to a back-end server, which then merges different sniffer traces into one unified trace to identify the root cause of various performance artifacts. MAP [21] and its successor DIST [24] were security-focused wireless monitoring systems. Finally, as an alternative to dense sniffer deployment, wardriving was used to construct practical conflict graphs [30].

The drawback of these systems was the high cost and effort associated with the dense deployment of static wireless sniffers in order to achieve good coverage. A few works tried to mitigate the cost by exploiting existing infrastructure. DAIR [2], [3] used wireless USB dongles attached to employee desktop machines and used the collected traces for several applications including rogue AP detection, helping disconnected clients, and network performance monitoring. However, DAIR nodes were still static and suffer from some of the disadvantages of previous solutions.

The idea of using smartphones to monitor wireless networks and/or spectrum has been exploited in a few recent works [17], [16], [28], [22]. The work in [17] used smartphones to detect and map heterogeneous networks and devices in home networks. The smartphones periodically perform measurements and use them to detect new devices, determine the impact of one device to another, etc. The works in [16], [28] developed two crowdsourcing-based RF spectrum monitoring systems using smartphones. The smartphone was augmented with external hardware to perform spectrum measurements—a software defined radio in [16] and a frequency translator in [28]. Pazl [18] is a smartphone-based indoor Wifi monitoring system. The goal was to develop an indoor localization system to localize the measurements from different smartphones, but it was only evaluated via small scale experiments with five participants. The closest work to ours is MCNet [19], a system that uses smartphones to evaluate the user perceived performance in enterprise wireless networks. MCNet collects active performance measurements from smartphones, while we investigate the effectiveness of passive measurements (channel scans).

All these works share many of our ideas with respect to the advantages of a smartphone-based monitoring system over the previously developed systems using a collection of statically deployed sniffers and the site-surveys. Nonetheless, these works either instrumented the smartphones with external hardware to collect specific types of measurements, or

performed active measurements through deployed softwares. In contrast, in this paper, we tried to answer the question of whether the billions of Wifi channel scan results already being generated by smartphones could assist in wireless network monitoring, configuration, or troubleshooting.

V. CONCLUSIONS AND DISCUSSION

Through two extensive case studies, we demonstrated that smartphone measurements can contribute unique insights about large scale wireless network deployments which are difficult or impossible to glean from site surveys, statically deployed sniffers or infrastructure side logs. Quite simply, smartphones represent real network users, and their data is representative in the way that these other sources of measurements cannot be.

To conclude, we address two practical issues in collecting smartphone scan data: (1) Will smartphone scan data compromise user privacy? (2) What are the incentives for users to contribute data?

A. Privacy

One particular concern that arises when sharing channel scan data is privacy. To address this, we first point out the types of user data that the infrastructure can already collect. Then, we discuss the minimum requirements in terms of channel scan data for each case study.

In enterprise network environments, the Wifi session information, including user's identity, device MAC address and association time, is usually recorded by the management software. Therefore, a complete view of the user's Wifi connection activity is already available on the infrastructure side. Even when the user's device is not connected to any AP, the infrastructure can still detect the presence of the device by sniffing probe packets. Against this backdrop, providing channel scan information does not represent an additional loss of privacy compared to what infrastructure networks can already monitor about their users.

Next we discuss the minimum requirements in terms of scan data for each of the case studies in the previous sections.

- 1) For all case studies, device identification information is not utilized in any way; thus, it can be properly anonymized even before scan results data are uploaded.
- 2) Timestamp information is not needed either, as long as the multiple scan result entries can be correctly grouped together and be identified as from one single scan.
- 3) Since the channel scan data are mainly used to help monitor central managed networks with predetermined SSIDs, such information is not required from the scan results. Furthermore, the user can choose only to upload channel scan data when they contain a predefined set of SSIDs, to avoid revealing information such as home Wifi networks.
- 4) Signal strength information is necessary in analyzing and predicting the device's association behavior in the spatial planning case study, yet is optional in the other study. Additionally, this information can be replaced by

an ordering of the APs by signal strength, eliminating the need for absolute RSSI values.

- 5) Channel information is not required, since the infrastructure has the complete knowledge of the channel assignments for all APs at any instant.
- 6) Wifi connection information is only useful in identifying the association choice from the scan results in the spatial planning case study. If such information can be annotated in the scan results before data uploading, Wifi connection information is not required either. For example, the device can identify the scans happened just before a Wifi connection event, and mark the corresponding BSSID as chosen AP in the scan result.

To summarize, for the purpose of constructing the conflict graph or performing spatial planning as in our case studies, only the BSSID information is necessary for the infrastructure to identify the APs. Other potentially sensitive information can either be removed or anonymized.

B. Incentives

Next, We discuss incentives for users to participate in data collection. First, smartphones already perform Wifi scans aggressively, thus the natural scan rate already provides a stream of network measurements with sufficient temporal granularity for monitoring purposes. By only passively harvesting the scan results, the overhead incurred in the process is kept low. Further steps can be taken to reduce the energy consumption of data uploading as described in [11], [29]. Additionally, user privacy can be maintained using the anonymization techniques we discussed in Section V-A.

Finally, we look into incentives that encourage users to share the data. Note that in enterprise networks, such data collection consent can be incorporated as part of the authentication process. For instance, users may be required to install the data collection app when connecting to the enterprise network, and the installation can then allow the user's other devices to connect to the network. For public Wifi service providers such as Boingo, extra Quality of Service (QoS) or price discounts can be offered to incentivize participation.

In summary, the measurement overhead can be significantly reduced by passive data collection and asynchronous delay tolerant data uploading, and effective incentive mechanisms can encourage measurement participation.

REFERENCES

- [1] IEEE radio resource measurement amendment, 2008.
- [2] P. Bahl, R. Chandra, J. Padhye, L. Ravindranath, M. Singh, A. Wolman, and B. Zill. Enhancing the security of corporate wi-fi networks using dair. In *ACM MobiSys*, pages 1–14. ACM, 2006.
- [3] R. Chan, J. Padhye, A. Wolman, and B. Zill. A location-based management system for enterprise wireless LANs. In *USENIX NSDI*, 2006.
- [4] X. Chen, B. Wang, K. Suh, and W. Wei. Passive online wireless LAN health monitoring from a single measurement point. *ACM SIGMOBILE Mobile CCR*, 14(4):19–21, October 2010.
- [5] Y.-C. Cheng, M. Afanasyev, P. Verkaik, P. Benko, J. Chiang, A. C. Snoeren, S. Savage, , and G. M. Voelker. Automating cross-layer diagnosis of enterprise wireless networks. In *ACM SIGCOMM*, 2007.
- [6] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *ACM SIGCOMM*, 2006.
- [7] T. Henderson, D. Kotz, and I. Abyzov. The changing use of a mature campus-wide wireless network. In *ACM MobiCom*, 2004.
- [8] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding congestion in IEEE802.11b wireless networks. In *ACM IMC*, 2005.
- [9] A. P. Jardosh, K. N. Ramachandran, K. C. Almeroth, and E. M. Belding-Royer. Understanding link-layer behavior in highly congested IEEE 802.11b wireless networks. In *ACM E-WIND*, 2005.
- [10] A. Kosowski and K. Manuszewski. Classical coloring of graphs. *Graph Colorings*, 352:1–20, 2004.
- [11] F. Liu, P. Shu, and J. Lui. AppATP: An energy conserving adaptive mobile-cloud transmission protocol. 2015.
- [12] X. Meng, S. Wong, Y. Yuan, and S. Lu. Characterizing flows in large wireless data networks. In *ACM MobiCom*, 2004.
- [13] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted coloring based channel assignment for WLANs. *ACM SIGMOBILE Mobile CCR*, 9(3):19–31, 2005.
- [14] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. A. Arbaugh. A client-driven approach for channel management in wireless LANs. In *IEEE INFOCOM*, 2006.
- [15] A. Nandugudi, A. Maiti, T. Ki, F. Bulut, M. Demirbas, T. Kosar, C. Qiao, S. Y. Ko, and G. Challen. PhoneLab: A large programmable smartphone testbed. In *Proceedings of First International Workshop on Sensing and Big Data Mining*, pages 1–6. ACM, 2013.
- [16] A. Nika, Z. Zhang, X. Zhou, B. Y. Zhao, and H. Zheng. Towards commoditized real-time spectrum monitoring. In *ACM HotWireless*, 2014.
- [17] G. Nychis, S. Seshan, and P. Steenkiste. Using your smartphone to detect and map heterogeneous networks and devices in the home. In *ACM HotWireless*, 2014.
- [18] V. Radu, L. Kriara, and M. K. Marina. Pazl: A mobile crowdsensing-based indoor Wifi monitoring system. In *Network and Service Management (CNSM), 2013 9th International Conference on*, pages 75–83. IEEE, 2013.
- [19] S. Rosen, S.-J. Lee, J. Lee, P. Congdon, Z. M. Mao, and K. Burden. MCNet: Crowdsourcing wireless performance measurements through the eyes of mobile devices. *Communications Magazine, IEEE*, 52(10):86–91, 2014.
- [20] D. Schwab and R. Bunt. Characterizing the use of a campus wireless network. In *IEEE INFOCOM*, 2004.
- [21] 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*, 15(5):10 – 18, October 2008.
- [22] J. Shi, Z. Guan, C. Qiao, T. Melodia, D. Koutsonikolas, and G. Challen. Crowdsourcing access network spectrum allocation using smartphones. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*, page 17. ACM, 2014.
- [23] A. Striegel, S. Liu, L. Meng, C. Poellabauer, D. Hachen, and O. Lizardo. Lessons learned from the NetSense smartphone study. In *ACM SIGCOMM Computer Communication Review*, volume 43, pages 51–56. ACM, 2013.
- [24] K. Tan, C. McDonald, B. Vance, C. Arackaparambil, S. Bratus, and D. Kotz. From MAP to DIST: The evolution of a large-scale WLAN monitoring system. *IEEE Transactions on Mobile Computing (TMC)*, 13(1):216–229, January 2014.
- [25] Tianxing, X. Zhou, and A. Campbell. Cutting Wifi scan tax for smart devices. *Dartmouth Computer Science Technical Report TR2014-752*, 2014.
- [26] J. Yeo, M. Youssef, and A. Agrawala. A framework for wireless lan monitoring and its applications. In *ACM Workshop on Wireless Security (WiSe 2004)*, 2004.
- [27] J. Yeo, M. Youssef, T. Henderson, and A. Agrawala. An accurate technique for measuring the wireless side of wireless networks. In *USENIX/ACM WiTMeMo*, 2005.
- [28] T. Zhang, A. Patro, N. Leng, and S. Banerjee. A wireless spectrum analyzer in your pocket. In *ACM HotMobile*, 2015.
- [29] T. Zhang, X. Zhang, F. Liu, H. Leng, Q. Yu, and G. Liang. eTrain: Making wasted energy useful by utilizing heartbeats for mobile data transmissions. *Energy (Joule)*, 5:10.
- [30] X. Zhou, Z. Zhang, G. Wang, X. Yu, B. Y. Zhao, and H. Zheng. Practical conflict graphs for dynamic spectrum distribution. In *ACM SIGMETRICS*, 2013.