Poster: Detecting Client Mobility in WLANs Using PHY Layer Information

Li Sun
University at Buffalo, SUNY
Buffalo, NY 14260
lsun3@buffalo.edu

Souvik Sen
HP Labs
Palo Alto, CA 94304
souvik.sen@hp.com

Dimitrios Koutsonikolas
University at Buffalo, SUNY
Buffalo, NY 14260
dimitrio@buffalo.edu

ABSTRACT

The continuously increasing number of smartphones and tablets allow the users to access Wireless LANs (WLANs) while undergoing different types of mobility, posing new challenges to wireless protocols. Current history-based WLAN protocols do not work well in mobile settings where wireless conditions change rapidly. Thus, today's WLANs need to be able to determine the type of the client's mobility and employ appropriate strategies in order to sustain high performance. While previous work tried to detect mobility using hints from sensors available in mobile devices, in this work, we demonstrate how different mobility modes can be distinguished by using physical layer information – Channel State Information (CSI) and Time-of-Flight (ToF) – available at commodity APs, with no modifications on the client side. Our testbed experiments show that our mobility classification algorithm achieves more than 92% accuracy in a variety of scenarios. In addition, we demonstrate how fine-grained mobility determination can be exploited to greatly improve performance of client roaming and MIMO beamforming.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

Keywords

Wireless; WLAN; Mobility; PHY layer

1. INTRODUCTION

With the proliferation of smartphones, tablets, and the advent of the BYOD phenomenon, mobile devices are soon becoming the preferred medium of Internet access in Wireless LANs (WLANs). Due to their smaller form factor, these truly mobile devices allow the users to access the wireless network while undergoing different forms of mobility. However, client mobility poses difficult problems to the WLAN protocols. In static scenarios, the wireless channel remains stable and hence wireless protocols can refer to the past history, carefully adapting themselves to avoid failures and maximize performance. In contrast, during mobility, wireless conditions may change frequently requiring protocols to be agile and apply different strategies. Designing such strategies is not trivial because they depend on the actual type of mobility demonstrated by the client. E.g., the WLAN should try to roam the client to a better AP only when the client is moving away from her AP and should not consider roaming for the other scenarios because it is unlikely that a better AP will be discovered. Likewise, mechanisms to obtain higher throughput in modern 802.11n/ac WLANs, such as beamforming, rate adaptation, frame aggregation, multiuser-mimo (MU-MIMO), require different optimizations based on the intensity of client mobility – which determines the length of past history that a wireless protocol can refer to.

Today's WLAN protocols are unable to detect client mobility. They generally rely on the past history to generate optimal performance for stationary clients. However, such a general framework is slow to adapt to client mobility. Recent work [1] tried to detect mobility using on-board sensors in a smartphone. Although this technique can broadly classify between stationary and device mobility, it requires changes at both the client and the AP in order to benefit download traffic because client mobility state is unknown to the AP and it also requires the sensors to be always on consuming battery life. Rather than changing the client, which is a more difficult proposition, in this work we demonstrate how different mobility modes can be distinguished by using physical (PHY) layer information available at commodity APs.

2. MOBILITY CLASSIFICATION

We identify four broad categories of client mobility. If the client is stationary, it can be in the *static mobility mode* when there are no significant environmental changes affecting the channel between the AP and the client. A static client may also be in the *environmental mobility mode* when the channel changes due to external movements. Of course the client itself may be moving – a mobility mode that we call *device mobility*. The client may experience different speeds under device mobility. We identify two broad and dominant categories of device mobility in WLANs. First, the user may slowly move the device although she is stationary or her movement is confined within a small area, which we define as *micro mobility*. E.g., the user may be attending a VoIP call over WiFi and a little movement of her head may displace her smartphone. On the other hand, the client may change its location as its user walks from one location to another. In such scenarios, we classify the client to be under *macro mobility*.

We use HP MSM 460 APs with Atheros AR9390 chipset and Samsung Galaxy S5 smart phones. We tuned the AP at 5.805GHz us-

MobiCom'14, September 7-11, 2014, Maui, Hawaii, USA.
ACM 978-1-4503-2783-1/14/09.
http://dx.doi.org/10.1145/2639108.2642892.

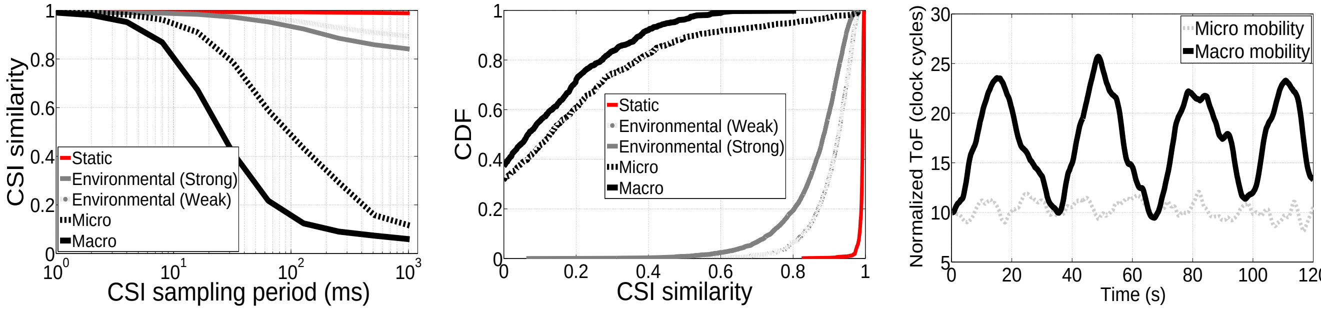

Figure 1: (a) Variation of CSI similarity values over time. (b) CDF of CSI similarity values for various mobility modes. (c) ToF values over time under device mobility. The user walks towards and away from the AP periodically in the macro-mobility scenario.

ing 40MHz channel and 802.11n protocol. We sent regular data packets from the AP and collected CSI, ToF, and Received Signal Strength Indicator (RSSI) information from the acknowledgment (ACK) sent by the client. We ran four different experiments to analyze different classes of mobility. Our goal is to develop simple classification schemes that can distinguish the above four scenarios.

2.1 Classifying Mobility using RSSI

We first explored the possibility of classifying the mobility mode of the client based on the RSSI of the client. In our experiments, we found that RSSI is quite stable in static scenarios in a quiet environment, but it is susceptible to any changes in the environment; often, the RSSI variation under environmental mobility is higher than the observed variation under device mobility. Therefore, we concluded that it is difficult to distinguish between environmental and device mobility solely based on RSSI. In the following, we will show how Channel State Information (CSI) can conclusively classify static, environmental, and device mobility, and how Time-of-Flight (ToF) can distinguish between the two types of device mobility.

2.2 Classifying Mobility using CSI

Any transmitted signal from the client undergoes reflections and arrives along multiple paths (multipath) at the AP. The wireless channel is expected to vary under environmental or device mobility because the fine-grained multipath structure may change. Those fine-grained variations are hard to be captured by RSSI because it aggregates all multipath components as one indicator. However, we found the same variations can be reliably detected by CSI because it captures the multipath characteristics from each subcarrier in the frequency domain. We characterize the *similarity* between two CSI samples by calculating the Pearson correlation coefficient with magnitudes on all subcarriers.

Figure 1(a) shows that the similarity stays close to 1 in the static case because the CSI matches across time on a stable channel. In contrast, the similarity drops sharply under environmental or device mobility. We further found that the similarity drops faster for device mobility than environmental mobility. This is because environmental mobility typically affects only a few multipath components, whereas if the client itself is moving, all multipath components will be affected. Figure 1(b) shows the distribution of similarity of consecutive CSI samples collected at every 0.5 seconds. Clearly, it is easy to find threshold values to distinguish between static, environmental, and device mobility. However, we found it is difficult to distinguish between micro mobility and macro mobility using CSI, even using larger sampling periods.

2.3 Classifying Device Mobility using ToF

To further classify device mobility, we utilize the intuition that the distance between the client and the AP under macro mobility changes more than that under micro mobility. Client's distance can be estimated based on RSSI [2] or CSI [3]. However, it was shown that RSSI and CSI are unreliable [4]. Instead, we utilize ToF [4, 5]. ToF is defined as the round trip propagation time of a signal transmitted between the AP and the client, which is proportional to the distance between them. The Atheros chipset can precisely compute the Time-of-Departure (ToD) of a data packet when it is sent out from the PHY layer. On correct reception of the packet, the client waits for a fixed SIFS duration and starts responding with an ACK. The chipset also reports an estimated Time-of-Arrival (ToA) of the ACK at the AP. The difference between the ToA and ToD contains the ToF between the AP and the client.

Figure 1(c) plots the ToF values over time for two different device mobility scenarios. For the micro mobility scenario, when the user naturally moves the device only within a small area, noisy ToF values can sometimes wrongly indicate changes in distance. However, the change in noisy ToF values in the micro-mobility case is quite random, while for the macro mobility scenario the ToF either steadily increases or decreases. This happens because within a reasonable time interval a walking user may either approach or move away from the AP, without changing her orientation. Therefore, we maintain a moving window of ToF values to detect macro mobility. Only if all the ToF values in the moving window suggest an increasing (moving away) or decreasing (moving towards) trend, we declare that the client is under macro mobility, otherwise the client is under micro mobility.

2.4 Evaluation of Mobility Classification

We implement our mobility detection scheme on the AP. We sample the CSI of the client every 500ms from existing data packet transmissions and maintain a moving average of the similarity between consecutive CSI values in 3 seconds. We compare the CSI similarity values with empirically chosen thresholds to distinguish between static, environmental and device mobility. If the CSI similarity indicates device mobility, we further consult ToF values to distinguish between micro and macro mobility (including moving directions). To deal with measurement noise [4], we sample the client's ToF readings every 200ms and aggregate them every second using a median filter.

We evaluate our scheme at more than 100 locations for over 24 hours. At each location we subject the client to different forms of mobility. Table 1 presents our overall performance results. We

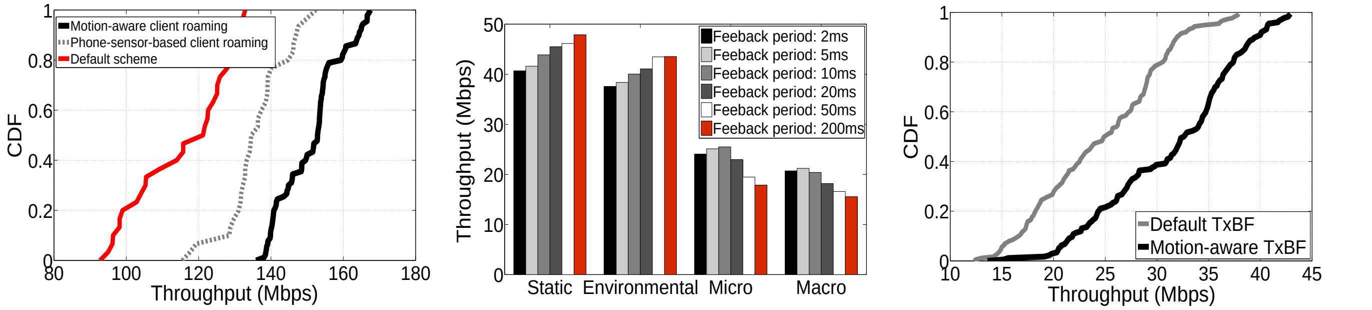

Figure 2: (a) Comparison of controller-based client roaming scheme with existing protocols. (b) Average throughput vs. CSI feedback period in different mobility scenarios with beamforming. (c) CDF of throughput gain for 300 different links subjected to a variety of mobility modes with beamforming.

find that the accuracy of our mobility detection scheme is more than 92% in all scenarios.

Table 1: Evaluation of mobility classification.

Ground truth	Detection result (%)			
	Static	Environmental	Micro	Macro
Static	99.7	0.3	0	0
Environmental	4.55	92.78	2.67	0
Micro-mobility	0	3.58	95.59	0.83
Macro-mobility	0	0	1.7	98.3

3. MOBILITY-AWARE PROTOCOL DESIGN

Client roaming. We first demonstrate how the knowledge of the client's moving direction at the AP can optimize client roaming. By default, a client associates with the AP with the strongest RSSI value and will trigger a handoff only when RSSI falls below a predefined threshold. Such a scheme is agnostic of the user's mobility; a moving client often remains wrongly connected to a far away AP, adversely affecting its own as well as the overall network performance. In [1], the authors proposed to use phone sensors to detect client mobility and request the phone to scan for better APs more aggressively. However, frequent scanning is time consuming and wastes energy and prevents the client from transmitting or receiving data (impacting throughput). To address this issue, rather than designing a scheme that requires changes at the client, we propose a controller-based protocol that roams the client to the appropriate AP whenever necessary.

In our scheme, the current AP of the client continuously determines the client mobility mode and shares it with the controller. If the current AP indicates that the client is moving away from it, the controller instructs the neighboring APs to periodically send NULL data frames to the client and compute the client's distance, RSSI, and heading information to themselves. If the client is moving towards another AP whose signal strength is similar or higher, we add that AP in the candidate set. If the controller finds at least one such better AP, it instructs the current AP to disassociate the client and asks only the APs in the candidate set to respond to the client's probe request. Consequently, the client roams to a better AP, ultimately improving performance. We evaluate our scheme on a 6-AP testbed in a 50x50 meter office by walking along its corridors. Figure 2(a) shows that our proposal performs better than the phone-sensor based scheme and improves median throughput by 31% over the default scheme.

Beamforming. Secondly, we studied how knowledge of client mobility can improve the performance of beamforming. Effective beamforming relies on timely CSI feedback from clients. Infrequent feedback will result in performance loss. However, too frequent feedback will be harmful because packets are transmitted at lower bit-rate in the feedback procedure which consumes more channel time. We use an AP as a client because we found none of the popular phones support explicit beamforming. Figure 2(b) shows that the optimal CSI feedback period decreases as the intensity of mobility increases and the default 200ms only benefits the static scenario. Therefore, we modified the WiFi driver at the AP to vary the feedback period based on the mobility mode of the client – 200ms for static, 50ms for environmental mobility, 10ms for micro-, and 5ms for macro-mobility. We conducted device-mobility experiments on 300 links and compared the performance of the adaptive feedback scheme with the default scheme which uses a statically configured 200ms feedback period. Figure 2(c) shows that our scheme outperforms the default scheme by 33% in the median case.

Client roaming and beamforming are not the only two wireless protocols that can benefit from client mobility. Currently, we are working on exploiting knowledge of client mobility to improve rate adaptation, frame aggregation, and MU-MIMO.

4. REFERENCES

[1] Hari Balakrishnan Lenin Ravindranath, Calvin Newport and Samuel Madden. Improving Wireless Network Performance Using Sensor Hints. In *USENIX NSDI*, 2011.

[2] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N Padmanabhan. Indoor localization without the pain. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, pages 173–184. ACM, 2010.

[3] Souvik Sen, Jeongkeun Lee, Kyu-Han Kim, and Paul Congdon. Avoiding multipath to revive inbuilding wifi localization. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*, pages 249–262. ACM, 2013.

[4] Alex Mariakakis, Souvik Sen, Jeongkeun Lee, and Kyu-Han Kim. Single access point based indoor localization. In *Proceeding of the 12th annual international conference on Mobile systems, applications, and services.*

[5] Domenico Giustiniano and Stefan Mangold. Caesar: carrier sense-based ranging in off-the-shelf 802.11 wireless lan. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, page 10. ACM, 2011.