Less Transmissions, More Throughput: Bringing Carpool to Public WLANs

Wei Wang[⋆†], Yingjie Chen[‡], Qian Zhang[†‡], Kaishun Wu[†], Jin Zhang[§†]
[⋆]School of Electronic Information and Communications, Huazhong University of Science and Technology
[‡]Department of Computer Science and Engineering, Hong Kong University of Science and Technology
[†]Guangzhou HKUST Fok Ying Tung Research Institute
[§] Department of Electrical and Electronic Engineering, South University of Science and Technology of China
Email: {gswwang, cyjocg, qianzh, kwinson, jinzh}@cse.ust.hk

***Abstract*—The proliferation of WiFi hotspots in public places enables ubiquitous Internet access. These public WiFi hotspots usually serve scores of mobile devices and suffer from extremely poor performance in terms of low goodput and severe delay. In this paper, we first study the traffic characteristics in public WiFi networks, and demonstrate that the main causes of such poor performance are media access control (MAC) inefficiency and downlink-uplink traffic asymmetry. To cope with these issues, we call attention to transmission carpool, which facilitates an access point (AP) to send multiple frames for different mobile stations (STAs) in a single transmission. It reduces contention and conveys more frames in each channel access. As such, each downlink transmission carries more payload and thus improves efficiency and solves traffic asymmetry simultaneously.**

***Index Terms*—MAC efficiency, PHY/MAC Design, Frame Aggregation**

I. Introduction

With the growing popularity of WiFi-based wireless local area networks (WLANs), WiFi hotspots in public places have gained significant importance over the last decade. A recent industry report projected an annual growth rate of 350% for such public WiFi deployments in the near future. Those WiFi hotspots are usually deployed in crowded places such as large conventions, malls and cafeterias. These public places are referred to as *large audience environments*, in which each WiFi access point (AP) normally serves a crowd of mobile stations (STAs) simultaneously.

Unfortunately, the performance of WiFi hotspots in such crowded environments is usually extremely poor. A previous study [1] on the SIGCOMM 2008 trace show that as the number of active STAs increases, traffic demands grow accordingly; while the throughput of the whole network diminishes significantly. High traffic demands in large audience environments even incurs network failure such as the wireless network collapse during the Steve Jobs iPhone 4 keynote, where the wireless network recovered only after most of the audience turned off their wireless devices.

The fundamental causes of this problem are high media access control (MAC) overhead and downlink-uplink traffic asymmetry. In large audience environments, a significant number of STAs contend a limited number of channels within a carrier sensing range, which results in intensive contention. High contention in distributed coordination function (DCF) based WiFi networks can cost a large amount of overhead, including carrier sensing, backoff, and high collision probability. Moreover, several studies on SIGCOMM traces [2], [3] report that downlink traffic volume is about four times larger than uplink traffic volume. However, the DCF-based WiFi networks provide equal opportunity for APs and STAs to access channel, which conflicts with the downlink-uplink asymmetric traffic pattern. This conflict results in congestions in the downlink, thereby incurring severe downlink throughput degradation [1].

A major solution for such large audience environment scenario is from the view point of coordination and scheduling. Centralized coordination is advocated in many enterprise WLAN proposals to reduce unnecessary contention among APs. There are also several existing priority control schemes that prioritize AP's channel access over STAs to address the traffic asymmetric issue. However, these works focus on coordination and scheduling, which does not address the inefficiency issue caused by the MAC overhead. As the most recent standards – IEEE 802.11n and 802.11ac – have largely improved maximum PHY data rates over the previous IEEE 802.11a/g from 54 Mbit/s to 600 Mbit/s and over 1Gbit/s, the MAC efficiency of WiFi networks degrades rapidly in current high speed WiFi networks due to reduced time used for data transmission (while the time occupied by MAC overhead like contention overhead and preamble remain unchanged). Thus, it is urgent and necessary to cope with the MAC inefficiency issue.

To improve the MAC efficiency of each single link transmission, frame aggregation [4] is proposed in IEEE 802.11n to reduce contention by aggregating multiple frames for the same destination together at MAC layer. The applications of MAC frame aggregation are limited to bulk transmissions as the sender needs to wait to collect enough payload before actual transmission, and thus it is not applicable to real-time applications like VoIP and other short data flows such as short HTTP transactions. However, our measurements on campus library and SIGCOMM traces [2], [3] reveal that a large portion of the downlink frames is less than 300 bytes, which is not amendable to the MAC frame aggregation.

II. Traffic Characterization and Carpool Design

To further validate the commonly cited characteristics for large audience environments, we measured the downlink traffic in a typical campus library Wi-Fi network. This library measurement is complementary to the SIGCOMM public traces [2], [3] in the following three aspects: i) Library is another large audience environment besides conference; ii) The library measurement captures traces for all STAs in certain locations, while the commonly used SIGCOMM traces [2], [3] are collected from the measurement AP, which associates only with a subset of STAs in the conference; iii) The library WLAN adopts the IEEE 802.11n standard, while the traced SIGCOMM WLANs adopt the IEEE 802.11a/b/g [2], [3]. We have the following observations on both SIGCOMM and library traces:

1. Concurrent downlink requests. Similar measurements on SIGCOMM'08 trace are conducted in [1]. The results show that there are concurrent downlink requests for different STAs in large audience environments.

2. Downlink traffic dominance. The downlink traffic volume is about four times larger than uplink traffic volume. As distributed coordination function (DCF) based Wi-Fi provides equal opportunity for APs and STAs to access channel, this downlink-uplink asymmetric traffic pattern results in congested APs and severe downlink throughput degradation [1].

3. High ratio of short frames. More than 50% and 90% of the downlink frames are smaller than 300B in the SIGCOMM and library traces, respectively. High ratio of short frames combined with downlink dominance indicates intensive contention in the downlink transmission.

Based on the above features of WLANs traffic, we observe that, it is very inefficient to send one short frame to one STA in each channel access when there are a crowd of STAs waiting for the AP's responses. The contradiction between single-destination transmission and multiple STAs' requests renders downlink traffic backlogged at the AP, and thereby results in low throughput and high delays in the downlink [1].

This contradiction in public WLANs motivates the design of Carpool. Instead of restricting each downlink transmission to one STA, if we can "carpool" frames for multiple STAs into a single transmission, downlink contention would be significantly reduced and each downlink transmission could carry more traffic. One direct benefit of Carpool is the MAC efficiency improvement by reducing the contention overhead. Another benefit of Carpool is that it inherently relieves downlink traffic congestion at AP caused by downlink-uplink traffic asymmetry since each downlink transmission conveys more traffic for multiple receivers. Carpool is a PHY/MAC design that enables frame aggregation for multiple receivers in OFDM-based WLANs, especially for large audience environments such as conference rooms, airports, and coffee shops. Frames for multiple STAs queued at the Carpool AP are aggregated as a single large frame, where an aggregation header is inserted to indicate the destination of each subframe.

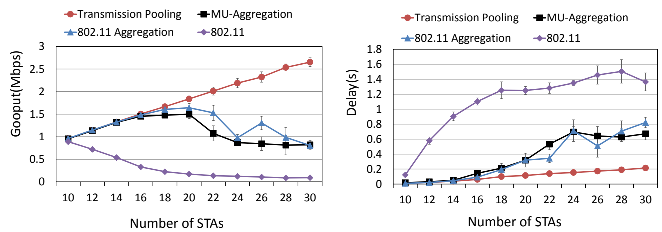

(a) Downlink goodput (b) Downlink latency

Fig. 1. Goodput and latency performance for VoIP with background traffic.

III. System Evaluation

We have implemented the basic mechanisms of the proposed transmission pooling design atop the OFDM implementation of GNURadio/USRP platform. We implement the entire PHY design directly in the USRP Hardware Drive (UHD). The frame synchronization and channel equalization algorithms are implemented according to IEEE 802.11a. Nodes in our experiments are equipped with RFX2450 daughterboards as RF front-end, which is configured to operate in the 2.4-2.5GHz range. In Fig. 1, we evaluate the performance with a challenging scenario – delay-sensitive VoIP traffic in large audience environments. To emulate real large audience environments, we inject background traffic according to the SIGCOMM 2008 trace [3], where the average times for TCP and UDP are 47ms and 88ms, respectively. The VoIP traffic is an ON/OFF UDP stream with a peak rate of 96Kbit/s and frame size of 120 bytes according to IEEE 802.11n requirements [5]. The results show that transmission pooling outperforms other schemes in terms of throughput and latency, especially when the number of STAs associated with one AP is large.

IV. Conclusion

In this paper, we investigated the characteristics of WiFi traffic in large audience environments and find that transmission carpool is a promising approach to scale the performance of WiFi in crowded public places. We observe that by enabling frame pooling for multiple STAs in the downlink transmission, the main causes of poor WiFi performance in large audience environments can be addressed simultaneously. We believe that the measurements conducted in this paper help abstract guidelines for future study, and the proposed transmission pooling design facilitates a new dimension to improve the efficiency of public WiFi networks.

References

[1] A. Gupta, J. Min, and I. Rhee, "WiFox: scaling WiFi performance for large audience environments," in *Proc. ACM CoNEXT*, 2012.

[2] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proc. ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005.

[3] A. Schulman, D. Levin, and N. Spring, "On the fidelity of 802.11 packet traces," in *PAM*, 2008.

[4] Y. Xiao, "IEEE 802.11n: enhancements for higher throughput in wireless LANs," *IEEE Wireless Commun. Mag.*, vol. 12, no. 6, 2005.

[5] A. Stephens and et al., "IEEE P802. 11 wireless LANs: usage models," *IEEE 802.11-03/802r23*, 2004.