MobiCom 2010 Poster: Passive Online Wireless LAN Health Monitoring from a Single Measurement Point *

Xian Chen[a] **Bing Wang**[a] **Kyoungwon Suh**[b] **Wei Wei**[c]
xian.chen@engr.uconn.edu *bing@engr.uconn.edu* *kwsuh@ilstu.edu* *weiwei@cs.umass.edu*
[a]CSE Department, University of Connecticut, Storrs, CT, USA
[b]School of Information Technology,llinois State University, Normal, IL, USA
[c]CS Department, University of Massachusetts, Amherst, MA, USA

Wireless LANs that are deployed in corporations and university campuses can contain a large number of access points and wireless hosts. In this paper, we develop a scalable approach that relies on measurements collected by a single monitor for WLAN health monitoring. We propose two metrics, local RTT and local loss rate, that indicate whether a wireless host experiences performance degradations inside a WLAN, and both can be obtained from measurements at the single monitor. Preliminary evaluation through data collected by a monitor at a university gateway router demonstrates the effectiveness of these two metrics.

I. Introduction

Wireless LANs (WLANs) have been widely used in corporations and university campuses. These WLANs can involve hundreds or even thousands of access points (APs), and many more wireless hosts. In a WLAN of such a scale, users may experience various performance problems on a daily basis [1]. Monitoring the health of a WLAN to quickly detect and predict performance degradations is the first step in taking remedy actions, and hence is crucial to maintain the health and normal operation of the network.

In this paper, we develop a scalable approach that relies on measurements collected by a *single* monitor for WLAN health monitoring. This monitor is placed at an aggregation point of a local network, capturing all traffic coming into and going out of the network. It has a global view of the network, and hence is at an ideal location to monitor the health of *all* hosts inside the network.

We propose two metrics that indicate whether a wireless host experiences performance degradations inside a WLAN. These metrics are *local RTT* and *local loss rate*, representing respectively the delay and the loss rate a TCP flow experiences inside the WLAN. Both metrics can be obtained based on measurements at the monitor. Preliminary evaluation through data collected by a monitor at the University of Connecticut gateway router indicates the effectiveness of these two metrics. We are in the process of developing online forecasting and detection techniques for WLAN health monitoring based on these two metrics.

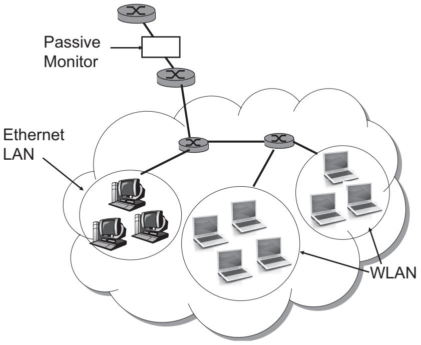

Figure 1: Measurement setup: a single monitor placed at an aggregation point of a local network captures incoming and outgoing traffic of the network.

Most existing studies on WLAN monitoring use distributed air sniffers that are spread out inside a WLAN (e.g., [1, 2, 3]), while our approach uses a single wired monitor at the edge of a network. A single wired monitor is also used in [4] for identifying WLAN traffics from Ethernet traffics, while our study focuses on WLAN health monitoring.

The rest of this paper is organized as follows. Section II describes the problem setting. Section III presents a high-level description of our approach. Section IV presents evaluation results using data collected at UConn gateway router. Finally, Section V concludes the paper and describes future work.

II. Problem Setting

We consider a local network (e.g. a university campus or an enterprise network), as illustrated in Fig. 1.

*This work was supported in part by NSF CAREER Award 0746841 and Caterpillar Scholar award.

A monitor is placed at an aggregation point (e.g. the gateway router) of the network, capturing all incoming and outgoing traffic. End hosts inside this network are connected to wired Ethernet or wireless 802.11 WLAN to access the Internet. Therefore, the monitor captures a mixture of wireless and wired traffic. Since we are interested in monitoring the health of the WLAN, we only keep track of wireless traffic at the monitor. This can be achieved by identifying WLAN traffic (e.g., using the passive online approach in [4]) or through IP addresses (when a network uses separate IP address pools for Ethernet and WLAN hosts).

Our goal is to detect WLAN hosts that are experiencing performance degradations or are about to experience performance degradations in realtime through measurements gathered at the single measurement point. For this purpose, we must answer the following two questions: (i) What metrics can we use to indicate the running status of a WLAN host? (ii) Based on these metrics, how to effectively predict and detect performance degradations in realtime?

III. Proposed Approach

We analyze the traffic headers collected passively by the monitor in realtime, and estimate delays and losses that a WLAN host experiences *inside* the WLAN. Based on the estimates, we can identify hosts that suffer from performance degradations. Furthermore, for each host, by analyzing the estimates over time, we can detect the change in the performance experienced by a host, and identify hosts that are about to experience performance degradations in the near future. Inspired by [5], we propose two metrics, *local RTT* and *local loss rate*, that measure respectively the delay and loss rate that a host experiences inside a WLAN.

Fig. 2(a) illustrates how to obtain a sample of local RTT. We consider a TCP flow that is from an outside server destined to a WLAN host (a typical scenario). Consider a data packet in this TCP flow. Let t_d and t_a denote respectively the time that this data packet and its corresponding ACK reaching the wired monitor. Then $(t_a - t_d)$ is a sample of the local RTT for the host, as shown in Fig. 2(a). When identifying a data packet and its corresponding ACK, we further impose the following additional restrictions to take account of several practical issues. First, we exclude all ACKs whose corresponding data packets have been retransmitted or reordered. Second, we infer whether delayed-ACK is implemented, and if so, we also exclude the ACKs released due to expiration of delayed-ACK timers.

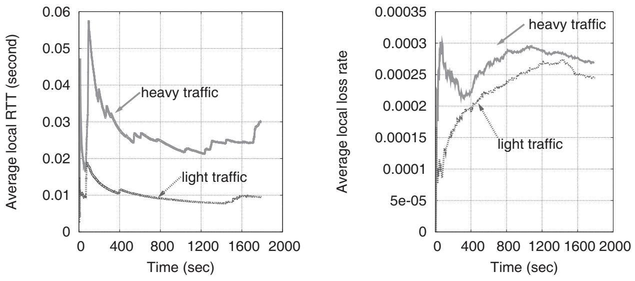

(a) Local RTT. (b) Local loss rate.

Figure 3: Cumulative average local RTT and local loss rate of a host under light and heavy traffics.

To obtain the local loss rate, we first divide time into intervals. Consider an arbitrary interval for a wireless host. Let n denote the number of data packets destined to this host. Of these n packets, a packet is a *local loss indicator* if and only if it is a retransmission and its corresponding ACK has not been observed by the monitor. For instance, in Fig. 2(b), the retransmissions of packets i and j are local loss indicators, while the retransmission of packet k is not. Let n_l be the number of local loss indicators in the time interval. Then, the local loss rate in the time interval is n_l/n.

IV. Performance Evaluation

We evaluate whether local RTT and local loss rate are effective metrics through experiments conducted at UConn. The monitor is a commodity PC, equipped with a DAG card that captures packet headers with accurate timestamps. This monitor is tapped into the UConn gateway router and captures all incoming and outgoing traffic. UConn has two types of WLANs, one is a secure 802.11i/WPA2-enterprise authenticated and encrypted network, and the other is an open/unsecure public network. We refer to these two WLANs as secure-WLAN and public-WLAN, respectively. The monitor only captures the traffic to the secure and public WLANs, and filters out other traffics. We next report evaluation results from a small-scale controlled experiment and a large-scale uncontrolled experiment.

In the small-scale controlled experiment, we monitor the local RTT and local loss rate of a secure-WLAN host, H, that downloads a large file (using TCP) from a remote server for 30 minutes under two settings. One setting has light traffic, where H is the only host accessing the Internet in the office. The other setting has heavier traffic, where three more secure-WLAN hosts in the office also access the Internet, using the same channel and the same AP as

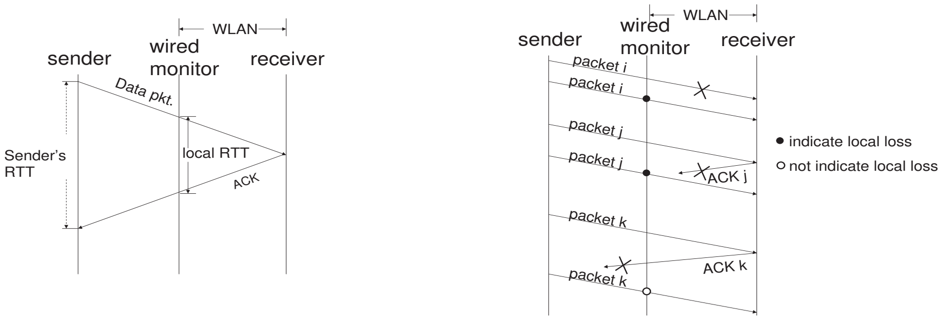

(a) Local RTT.

(b) Local loss.

Figure 2: Illustration of how to obtain local RTT and local loss based on measurement collected at the monitor.

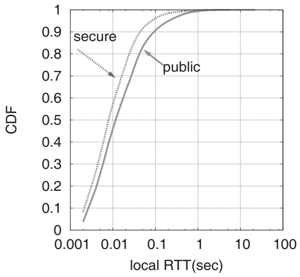

(a) Local RTT.

(b) Local loss rate.

Figure 4: CDF of local RTT and local loss rates of UConn secure and public WLANs.

H. Figures 3(a) and (b) plot the cumulative average local RTT and local loss rate of H under these two settings. We observe that indeed H experiences larger local RTTs and local loss rates under heavier traffic, indicating that these two metrics are good indicators of the performance experienced by the host.

Our large-scale uncontrolled experiment is conducted on May 5, 2009 during noon time for 30 minutes. We observe 2,110 and 1,261 hosts in the secure and public WLANs, respectively. For each host, we calculate its average local RTT and local loss rate over all of its TCP flows within that 30 minutes. Figures 4 (a) and (b) plot the cumulative distribution functions of local RTT and local loss rate. As expected, hosts in secure-WLAN tend to experience shorter average local RTTs and smaller local loss rates than those in public-WLAN since traffics in public-WLAN also need to be processed additionally at NAT translation servers and firewalls. We further investigate flows with high local loss rates and infer their TCP congestion windows using the techniques proposed in [5]. Our results confirm that local losses are often correlated with decrease in TCP congestion windows.

V. Conclusion and Future Work

In this paper, we proposed a scalable approach that relies on measurements collected at a single measurement point for WLAN health monitoring. We proposed two metrics to measure the delay and loss that a host experiences inside a WLAN. Preliminary evaluation using data collected at UConn gateway router indicates that these two metrics are effective.

As ongoing work, we are investigating how to use these two metrics to identify hosts with performance degradations. We are planning to conduct more controlled experiments to gain insights into the characteristics of faulty hosts and/or APs.

References

[1] A. Adya, V. Bahl, R. Chandra, and L. Qiu, "Architecture and techniques for diagnosing faults in IEEE 802.11 infrastructure networks," in *Proc. of ACM MobiCom*, September 2004.

[2] J. Yeo, M. Youssef, and A. Agrawala, "A framework for wireless LAN monitoring and its applications," in *Proc. of WiSe*, 2004.

[3] 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 *Proc. of ACM SIGCOMM*, (Pisa, Italy), September 2006.

[4] W. Wei, K. Suh, B. Wang, Y. Gu, J. Kurose, D. Towsley, and S. Jaiswal, "Passive online detection of 802.11 traffic using sequential hypothesis testing with TCP ACK-Pairs," *IEEE Transactions on Mobile Computing*, vol. 8, pp. 398–412, 2008.

[5] S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, and D. Towsley, "Inferring TCP connection characteristics through passive measurements," in *Proc. of IEEE INFOCOM*, 2004.