Remote Analysis of a Distributed WLAN using Passive Wireless-side Measurement

Aniket Mahanti
Carey Williamson
Martin Arlitt
University of Calgary

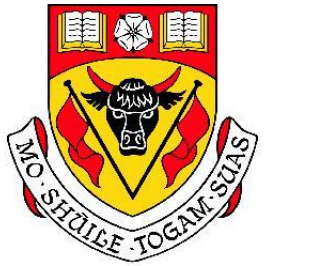

Introduction

- Wireless Local Area Networks (WLANs) are commonplace in many university campuses.
- Usage trends observed on a campus network often transcend many other WLAN environments, such as enterprises and public hotspots.

- As WLANs grow in size, scale, and complexity, the challenges for WLAN measurement also grow.
- The primary challenges for WLAN measurement include the geographic diversity of WLAN deployments, the physical proximity required for WLAN packet capture, and the need for a wireless-side view of the network.

Wireless Trace Collection Methods

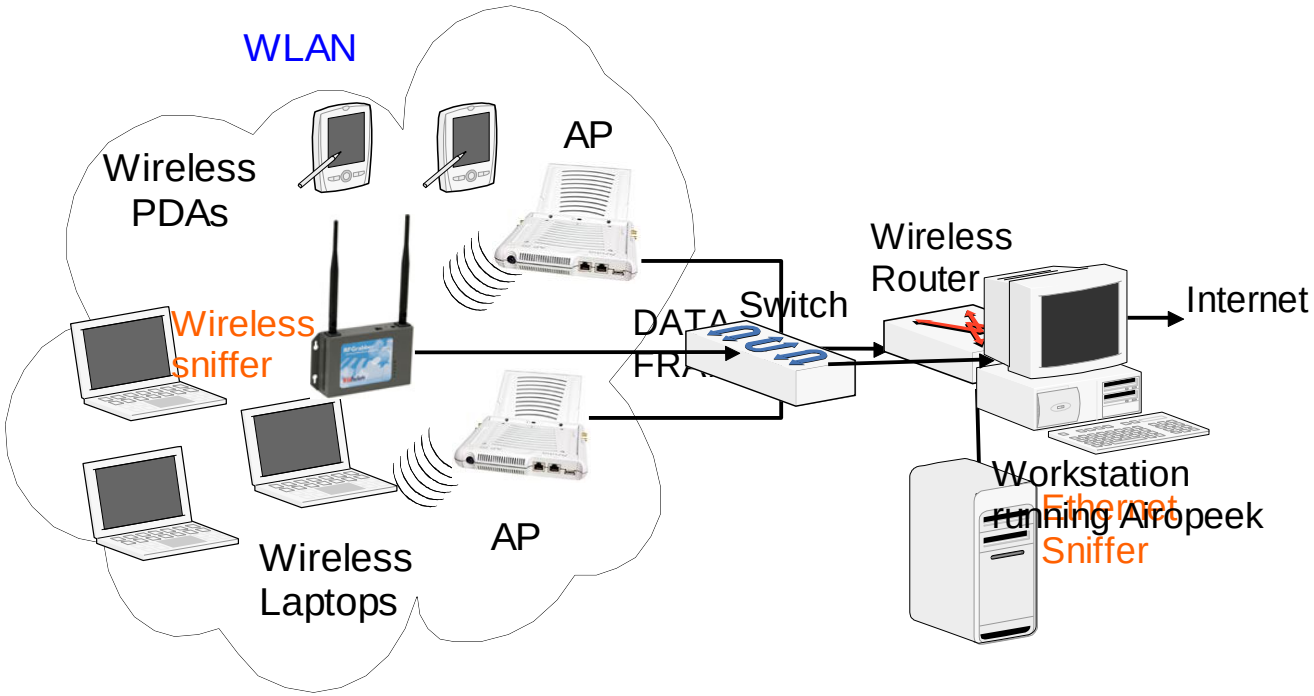

Wireless-side Measurement

Advantages of Wireless-side Measurement

- Wired-side Measurement
  - Does not capture Control or Management frames.
  - Wireless MAC header gets replaced by an Ethernet MAC header.
  - Obtaining MAC/PHY information is difficult.
  - Supplementary information required for complete WLAN analysis (e.g., SNMP polling, syslog).
- Wireless-side Measurement
  - RFGrabbers can capture all wireless frame types.
  - RFGrabbers capture the complete wireless MAC header.
  - Airopeek can provide MAC/PHY information such as data rate, frame directionality, signal strength, and retransmission flags.
  - No supplementary information required.

Objectives

- Demonstrate the feasibility of a practical and commercially-available solution for remote passive wireless-side measurement in a large distributed production WLAN.
- Present a comprehensive multi-layer analysis of our WLAN datasets, from the application layer to the wireless link layer.

Network Environment

AirUC is the wireless network available throughout the University of Calgary campus, provided by UCIT:

- Uses 802.11 a/b/g standard.
- Available to 28,000 students, and 5,000 faculty and staff.
- Non-encrypted infrastructure network consisting of 476 Aruba APs (2006).
- APs controlled by 6 central AP controllers.
- Uses three channel spectrum for ‘b/g’ mode (channels 1,6,11).

Aruba AP 70

Measurement Methodology

- We collected WLAN traces using a specialized trace capture program called Airopeek, which works in conjunction with network adapters to capture wireless frames.
- We used off-the-shelf adapters called RFGrabbers that can capture all 802.11 a/b/g frames at a remote location (i.e., “listen only” AP).
- The RFGrabber plugs into an Ethernet LAN and sends UDP-encapsulated copies of captured frames back to Airopeek running elsewhere on the network.

Wireless-side Trace Collection

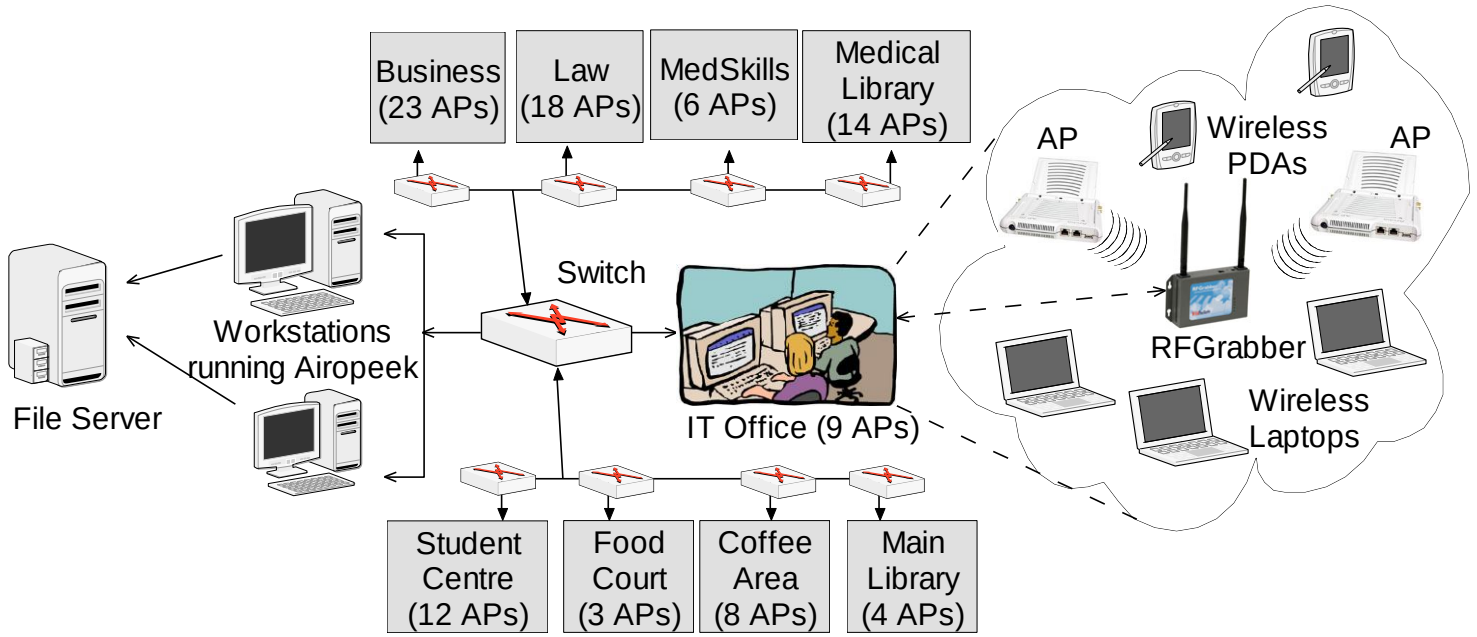

- RFGrabbers were configured to scan channels 1, 6, and 11 every 500 ms to capture WLAN traffic in the `b/g' mode.
- RFGrabbers captured packets from 97 APs at 9 locations, representing 20% of the WLAN.
- The RFGrabber probes see 95%–99% of the traffic transiting a nearby AP.

Trace Data Overview

Trace Duration	~6 weeks (Mar 3 – Apr 14, 2006)
Number of Frames	~ 1 billion 64% Management frames 36% Data frames
Number of Users	6,775 (based on MAC addresses)
IP Traffic Volume (Total = 102 GB)	Incoming = 58 GB Outgoing = 27 GB Local (Internal) = 17 GB
Avg. user sessions/day	1,481
User devices	50% of user devices had built-in wireless NICs (e.g., Intel, IBM, Mac)
Operating systems	60% Windows, 12% Mac OS

Multi-layer WLAN Analysis

- User view
  - WLAN usage
  - Usage regularity
- Application view
  - Application-layer protocols
  - Traffic directionality
- Mobility view
  - APs and locations visited
  - Mobility pattern
- User session view
  - Sessions per user
  - Session duration
  - Session activity
- Network view
  - AP load
- Wireless view
  - Channel usage
  - Error rates

User View

- Daily WLAN usage
- Hourly WLAN usage
- Usage regularity

Daily WLAN Usage

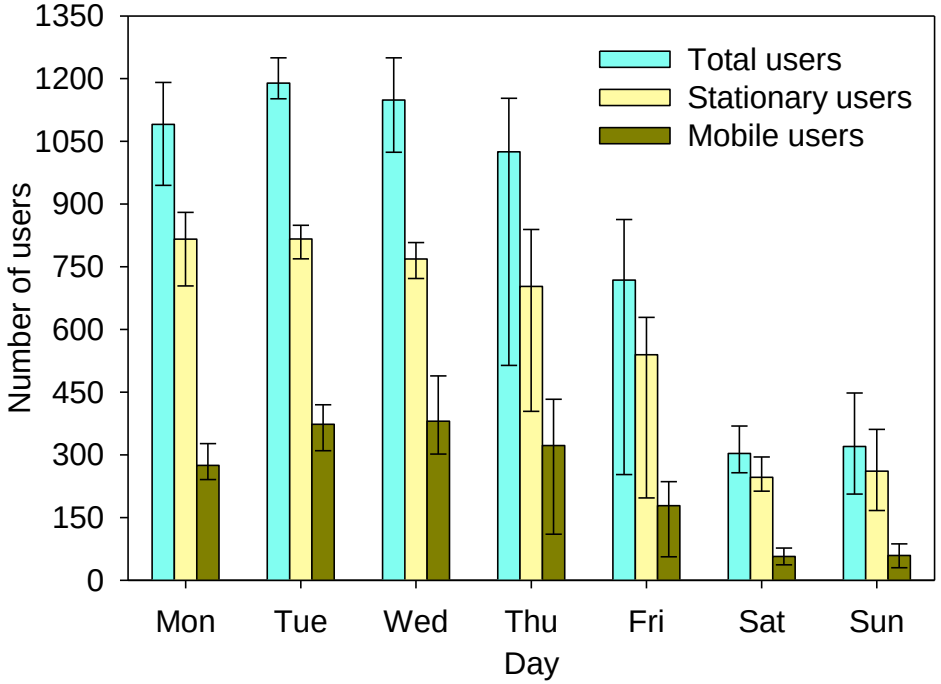

- More users used the WLAN during the early part of the week.
- On each day, about 25% of the observed users are mobile.

Hourly WLAN Usage

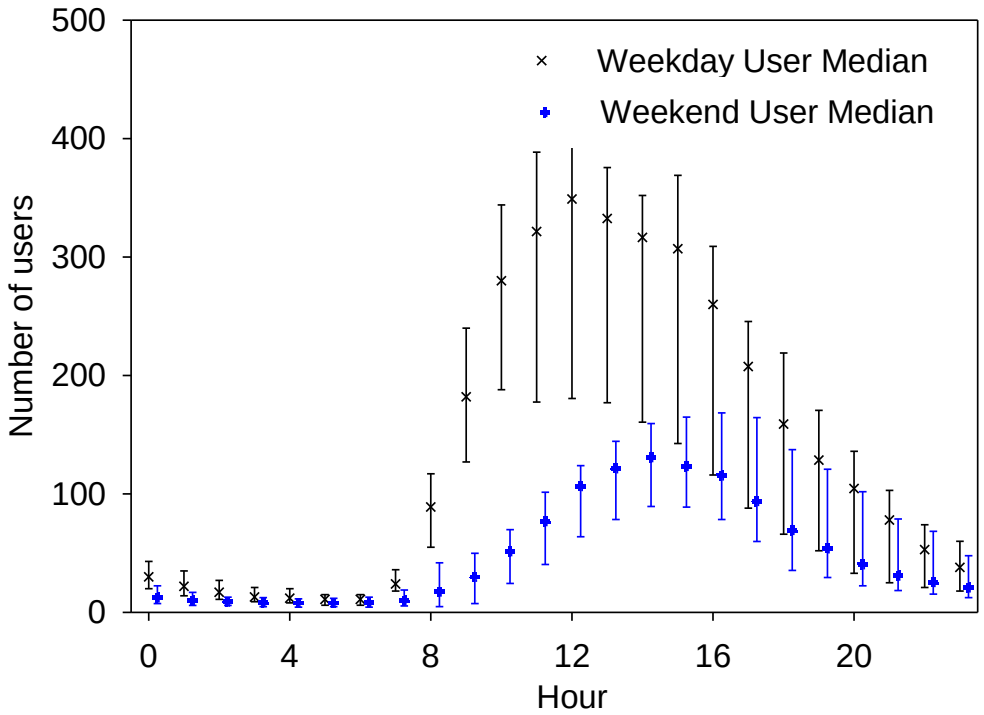

- Diurnal usage pattern is evident.
- The diurnal patterns observed were quite consistent across all of the 9 locations studied.
- The Main Library location differed slightly: activity persisted into the late evening, because of extended hours during the final exam period.

Usage Regularity

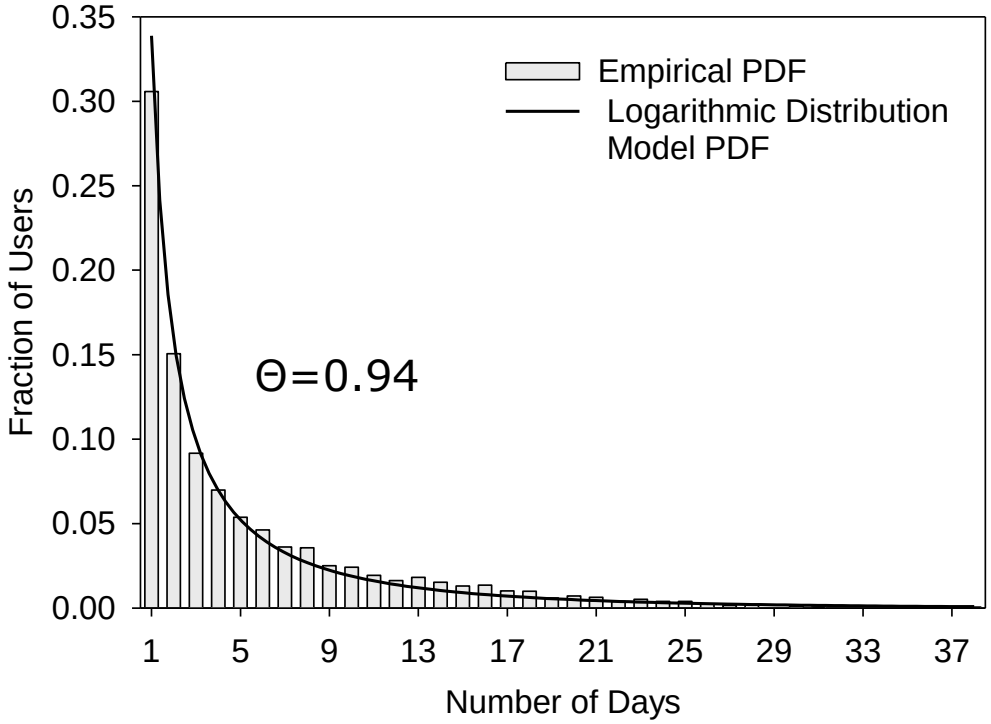

- Approx. 30% of users used the WLAN on only one day in trace.
- Only 3 users connected on all days during the trace period.

Application View

- Application-layer protocols
- Traffic directionality

Application-layer Protocols

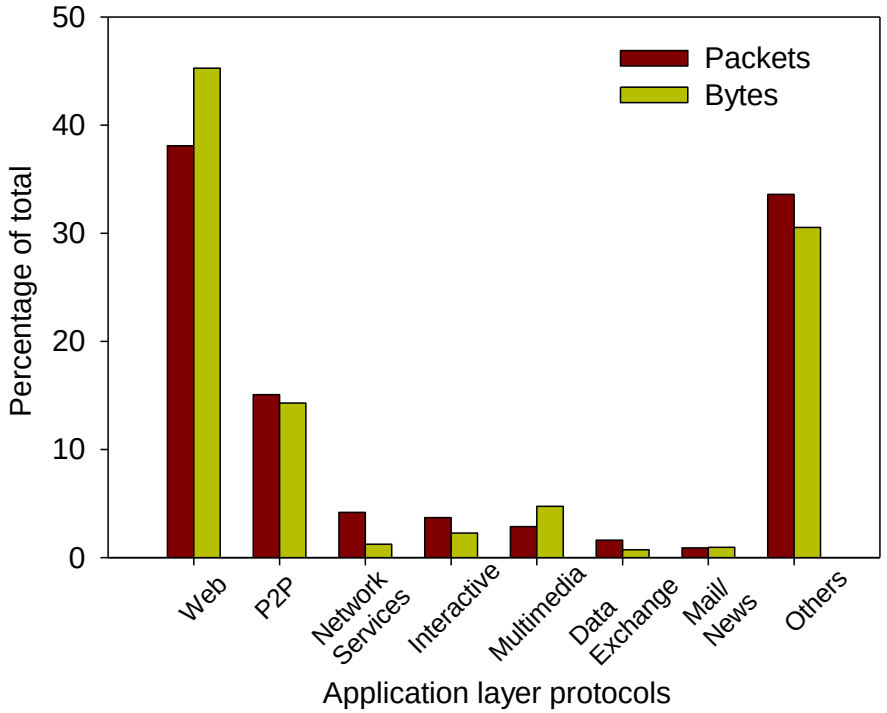

- We used a simple port number-based approach for traffic classification.
- About 46% of user traffic bytes was from Web surfing and 15% of user traffic was from known P2P applications.
- About 30% of traffic was “Others” (unknown).
- By applying payload-based signature classification on a separate 1-hour trace we found that a majority of the “Others” traffic was due to P2P.

Traffic Directionality

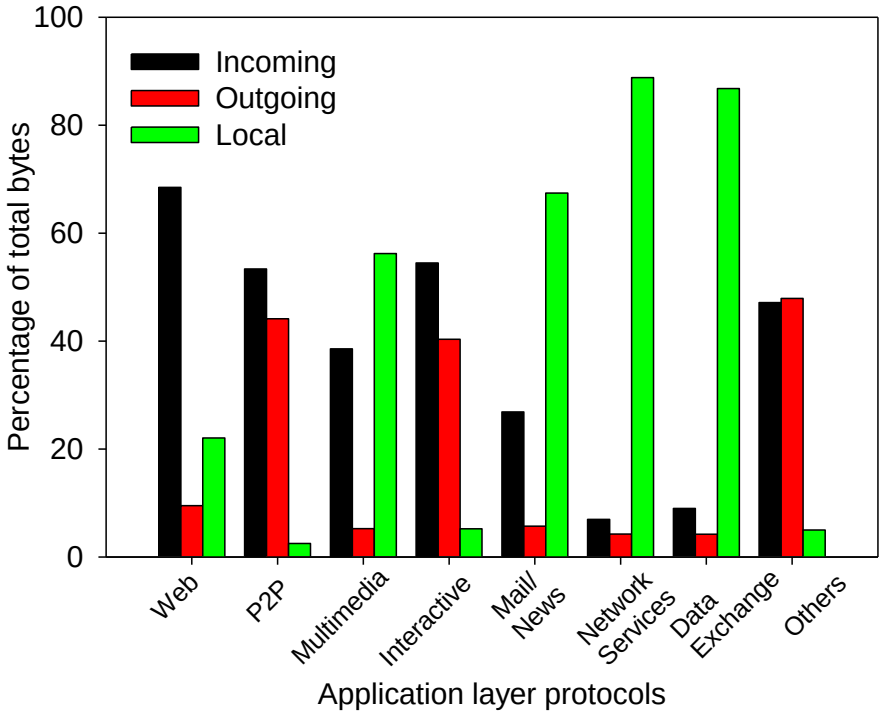

- Analysis reveals distinctive profiles for different network applications.
- Web: Users surfed off-campus Web sites more than local university sites.
- Data file system: Users are primarily accessing content from UofC file servers.
- P2P: Traffic balance between incoming and outgoing. Low internal P2P traffic suggest that these applications do not exploit local network topology well, or that users have such diverse interests that local file sharing is rare.

Mobility View

- APs and locations visited
- Mobility pattern

APs and Locations Visited

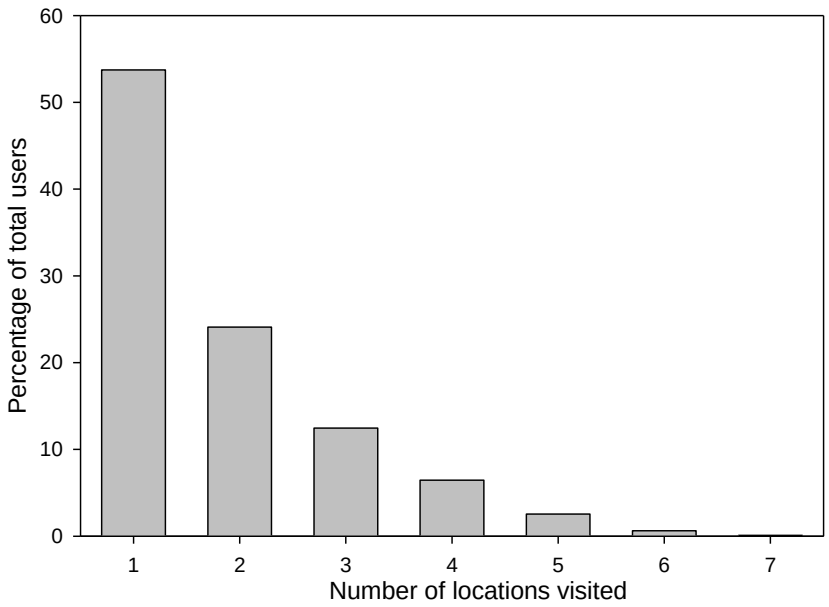

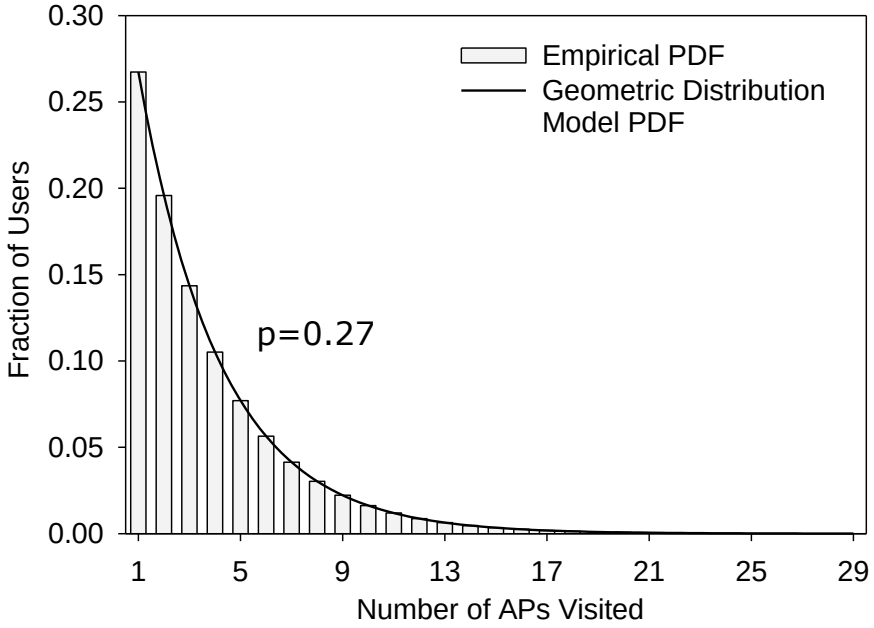

- About 54% of users were seen at only a single physical location.
- About 30% of the users were seen at only one AP.
- Visit behaviour differs slightly across locations, since it is influenced by the number of APs available.
- Few users were highly mobile; nonetheless, the distribution does have a pronounced tail.

Mobility Pattern

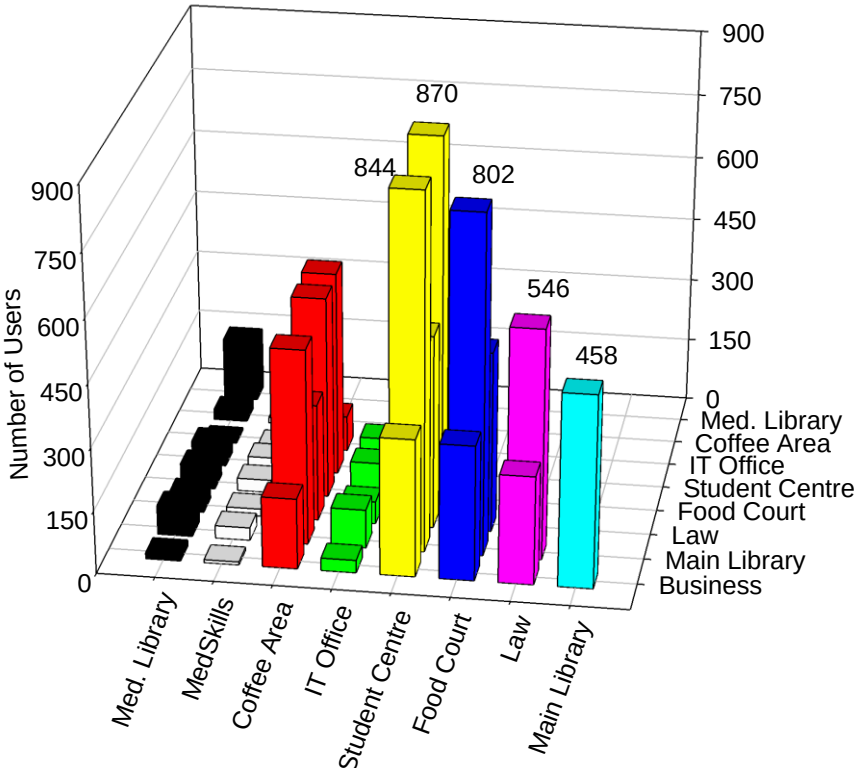

- The user mobility patterns observed are influenced by geographic proximity. For example, only 70 users from the two Medical Centre sites (2 kms away from the main campus) were observed using the WLAN at other campus locations.
- Many users are common between the Student Centre, Food Court, Law, and Main Library, considered pairwise. These results reflect the popularity of these locations with users.

User Session View

- Sessions per user
- Session duration
- Session activity

Session Duration

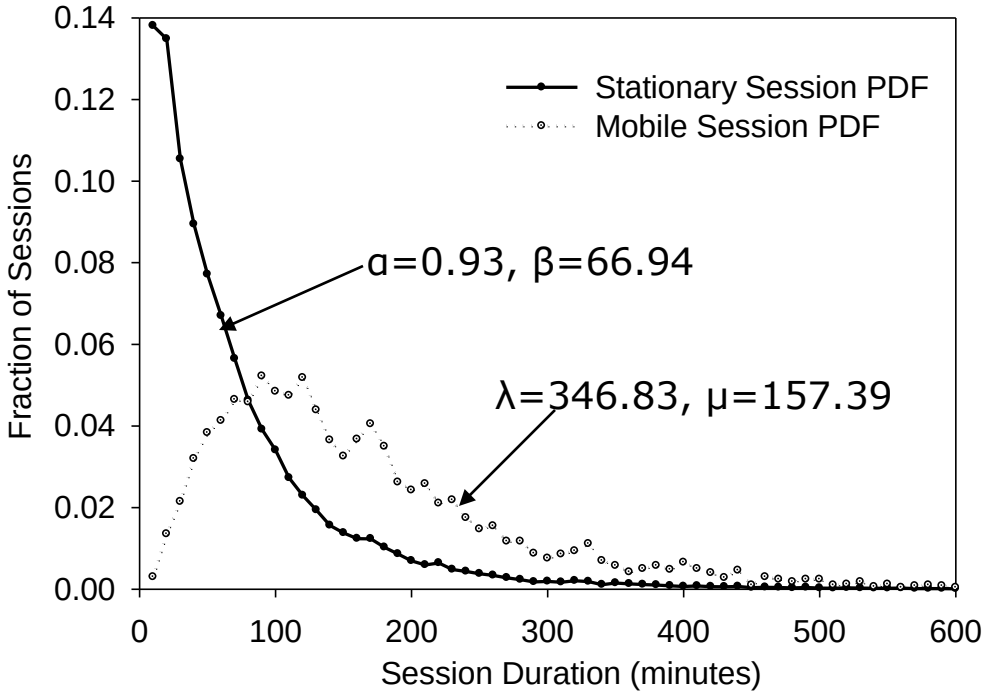

- Approx. 90% of all sessions ended within 3 hours.
- About 11% of all sessions are mobile sessions.
- Mobile sessions tend to last longer than stationary sessions. About 90% of all mobile sessions ended within 6 hours.
- The median duration for stationary sessions was 44 minutes, while the median for mobile sessions was 2 hours.
- Stationary session durations follow a Weibull distribution, while mobile session durations follow an Inverse Gaussian distribution.

Network View

- AP load

AP Load

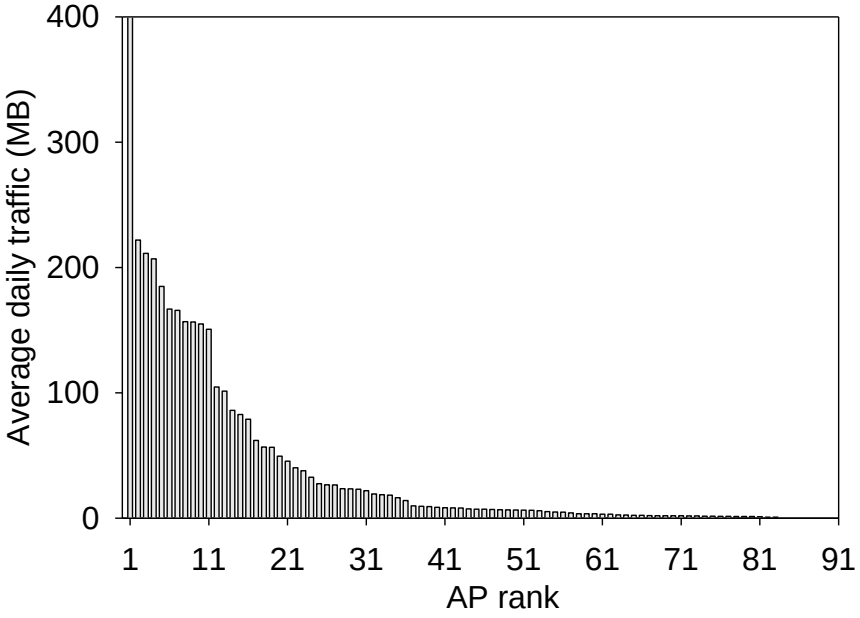

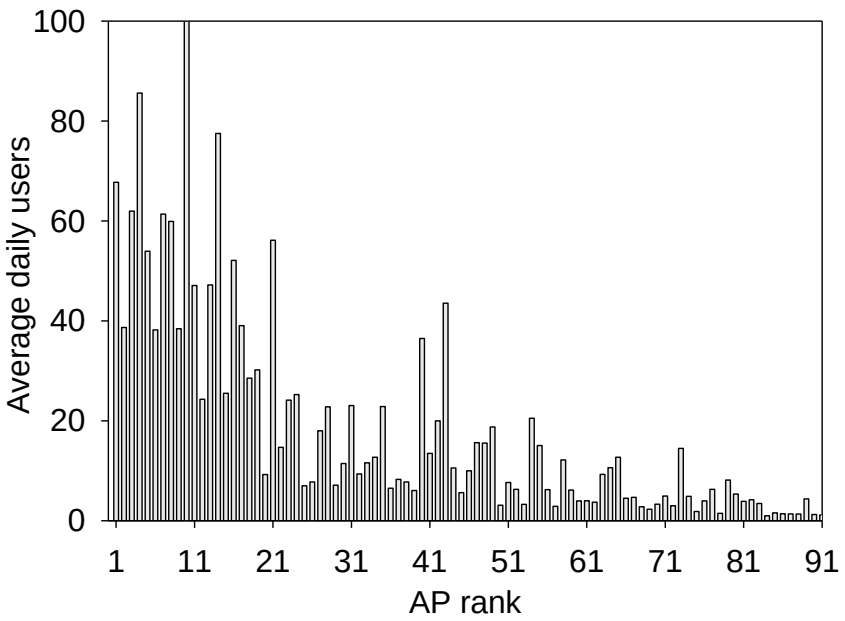

- Load is unevenly distributed across APs.
- Traffic load on APs is loosely related to number of users these APs (in the same rank order) handled.
- Non-uniform AP usage seems to be an inherent characteristic of WLANs.

Wireless View

- Channel usage
- CRC error rate
- Retransmission rate

Channel Usage

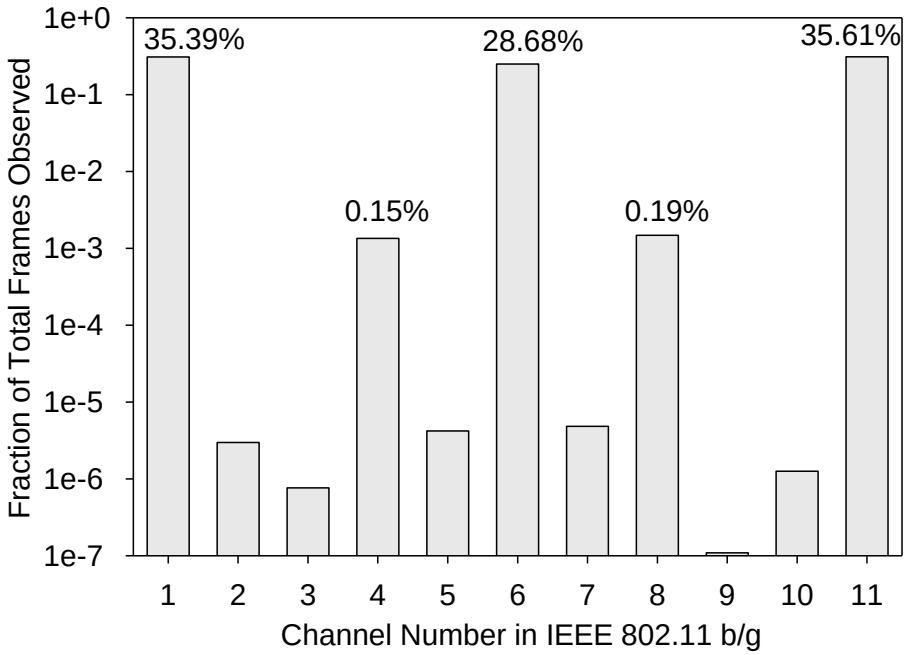

- Load is roughly balanced on channels 1, 6, 11.
- Frame transmissions are observed on all other channels too!
- Some APs in the Student Centre were configured to use overlapping channels (e.g., 1, 4, 8, 11). Such configurations have been found to be practically feasible.

CRC Error Rates

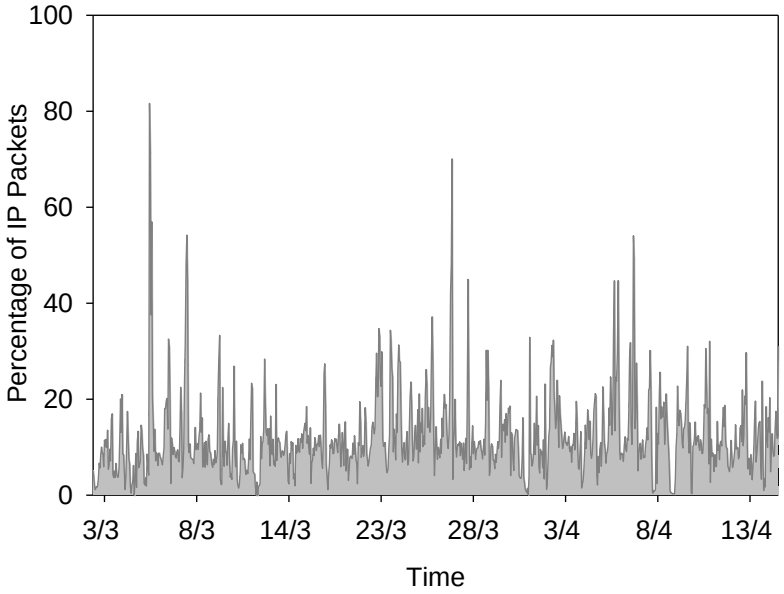

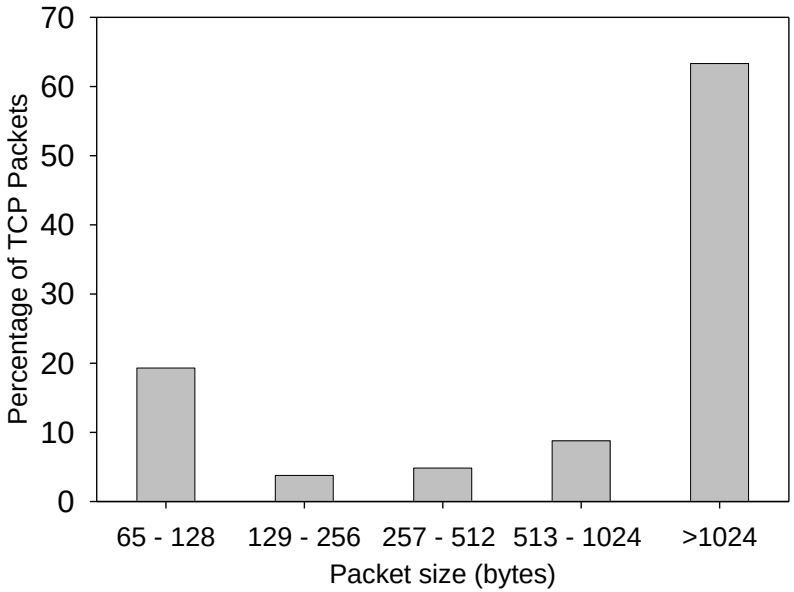

- CRC error rates were higher than expected, across all locations. CRC errors are caused by interference from nearby traffic on the channel, poor radio link, and channel noise.
- Errors are concentrated on the packet sizes that are dominant. Approx. 52% of TCP packets were of size 65-128 bytes and 31% of packets were bigger than 1 KB.
- Probability of packet corruption increases with packet size.

Retransmission Rates

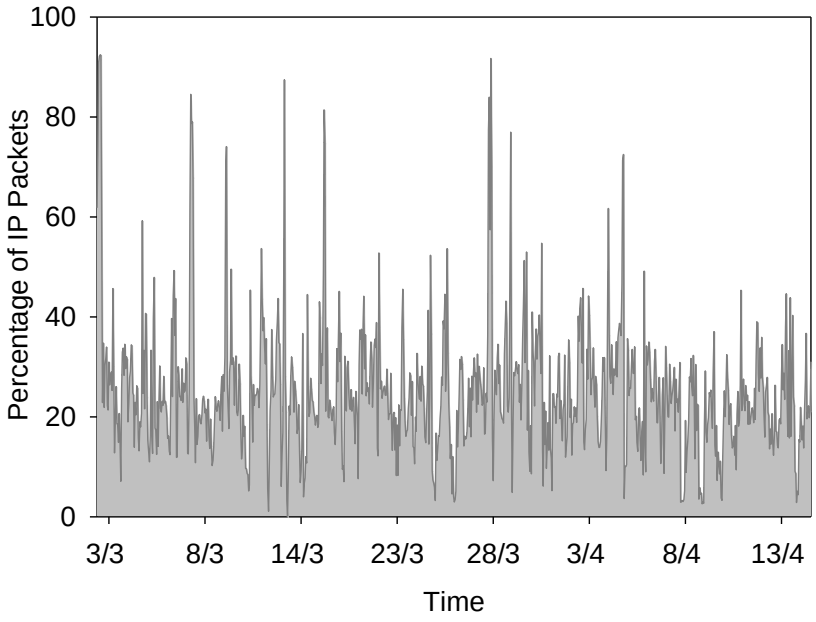

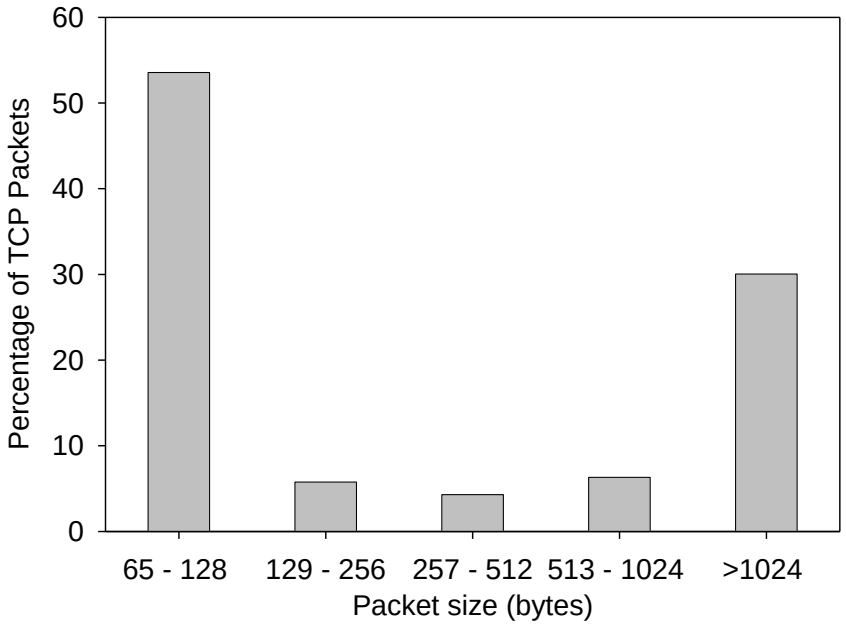

- Approx. 25% of Data frames observed were retransmissions.
- Approx. 50% of TCP retransmitted packets were small (<128 bytes).
- CRC errors are only one of the reasons for packet retransmission. Thus, there is no direct correlation between these results and the results in the previous slide.

Summary

- We presented a measurement study of a campus WLAN environment, with the data collected using remote passive wireless-side measurement.
- Our study demonstrated the feasibility and effectiveness of remote non-intrusive wireless-side measurement in a geographically-distributed campus WLAN environment.
- Analysis of our traces identified several trends consistent with prior campus WLAN measurement studies, including diurnal usage patterns, diverse network application usage, and limited user mobility, while offering new observations on session activity, mobility patterns, and wireless channel usage in our campus WLAN.
- Our analysis identified several emerging trends in application usage, user mobility behaviour, and WLAN deployment, as well as some performance-related issues at the wireless layer.