A Simple Queueing Network Model of Mobility in a Campus Wireless Network

Yung-Chih Chen
University of Massachusetts
Amherst MA, USA
yungchih@cs.umass.edu

Jim Kurose
University of Massachusetts
Amherst MA, USA
kurose@cs.umass.edu

Don Towsley
University of Massachusetts
Amherst MA, USA
towsley@cs.umass.edu

ABSTRACT

Although wireless networks have become ubiquitous, surprisingly few models of user-level mobility have been developed and validated against traces of measured user behavior. In this paper, we develop a simple, parameterized, open queueing network model of user mobility among access points in a campus network. Using CRAWDAD traces of user-access-point affiliation over time, we compare model-predicted performance with the performance actually observed in the traces, and find that a simple queueing model can indeed be used to accurately predict a number of performance measures of interest.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communications; C.4 [**Performance of Systems**]: Modeling techniques

General Terms

Performance, Theory

1. INTRODUCTION

Wireless networks, unlike their wired counterparts, have relatively few validated models of network user behavior and the traffic they generate. Yet, such models are crucial for studying wireless network protocols and architecture. Such models can also be used for network dimensioning, answering "what if" questions, such as how performance changes as the number of users or traffic scales up, or as the deployed network infrastructure evolves.

In this paper, we explore the use of simple *open queueing network models* of user mobility among access points (APs). The network model is "open" in that each mobile user enters the network, moves from AP to AP, and then leaves the network; each new arrival to the network is treated as a new, independent customer, considerably simplifying the computation of performance metrics. We model APs as $M/G/\infty$ queues; users are modeled as arriving to the network of wireless APs according to a Poisson process, and making independent probabilistic transitions among APs (or leaving the network). The model's inputs are the mean arrival rates, mean residency times at each AP, and the inter-AP transition probabilities, which are determined from empirical traces. Starting with AP-level CRAWDAD[1] traces of user-AP affiliation over time in a campus network, and comparing model-predicted performance versus the performance actually observed in the traces, our findings here are that such a simple model accurately predicts a number of performance measures of interest. However, in order to accurately predict certain performance measures, we find that it is necessary to partition users into two groups: those who only visit one AP after entering the network, and those who visit multiple APs.

The remainder of this paper is structured as follows. Section 2 describes the trace we use, and how we pre-process the trace. Section 3 presents our proposed queueing network model, which is validated in Section 4. Section 5 concludes this paper.

2. THE TRACE

There are several publicly available traces of long term user activity in a wireless LAN (WLAN) [6][3][2]. As we are interested in modeling user-level mobility among APs in larger-scale (e.g., campus-level) wireless networks, we seek traces that contain this information and where the network scale is large enough (both in terms of the number of APs and the user population) and the measurement period is long enough. We use the Dartmouth Simple Network Management Protocol (SNMP) trace [3], which was generated by the wireless LAN's mobility controller (i.e., a central server that coordinates all APs on campus) that polls each AP every five minutes[1]. At each such SNMP poll, an AP reports to the controller those clients that are currently associated with that AP. Although this information does not provide the precise time of a user's departure from an AP, we can infer the departure of a user by the absence of the user in a subsequent poll, as discussed below.

2.1 Trace Preprocessing

To circumvent the problem of diurnal user behavior (people go to sleep and do work different from their day time behaviors), we only consider mobile user activity during those periods of time when the university is most active. Hence, we extracted traces from 7 am to 7 pm of each day, and

S3'11, September 19, 2011, Las Vegas, Nevada, USA.

---

[1] The Dartmouth trace records wireless user activity for a 17-week period, from 11/2/2003 to 2/28/2004.

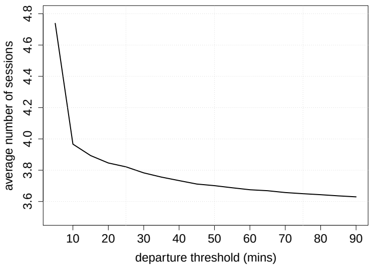

Figure 1: Average number of sessions for various departure length threshold.

removed all weekend, holiday, and inter-session periods as well[2]. The processed trace contains 547 APs across 6 different types of buildings (as listed in Table 2), with 6,452 distinct users.

We first define a *session* as the period of time during which a mobile user is continuously connected to the campus network; during a session the user may move from one AP to another. Thus, a session begins when the mobile user first associates with an AP (not having been previously associated with an AP) until the user disassociates from all network APs. As discussed previously, each AP periodically provides SNMP reports (at five-minute intervals) listing those mobile users that are currently associated with that AP. Occasionally, we find that a user disappears from the every-five-minute SNMP reports and then soon after reappears in the SNMP reports. There are three possible explanations:
1) The user actually left the network and returned;
2) The user was is motion, leaving one AP and then later associating with another AP;
3) An SNMP update was missing or lost.

Without explicit disassociations, it is difficult to infer which of these cases has indeed occurred. To distinguish true network departures from incorrectly inferred departures due to missing SNMP reports, we proceed as follows.

Let the departure length threshold, T_d, be the interval of time such that if the user does not appear in an SNMP report for T_d then the user is inferred to have left the network. Thus, periods of association by the same user that are separated by the amount of time $\tau > T_d$ (with no SNMP reports of that user during the intervening τ) are considered to be two separate sessions for that user.

Figure 1 plots the average number of sessions per-day per-user as a function of the departure length threshold. We note a sharp drop in the average number of sessions when the departure length threshold is less than 10 minutes (corresponding to an absence of that user in one or two back-to-back SNMP reports), and then a much slower decrease for larger threshold values. Thus, we chose a value of the departure length threshold of 10 minutes, and consider a user to have remained in the network if two intervals of activity (as reported by SNMP association reports) for that user are separated by ten minutes or less.

U	total number of users at steady state
N	total number of APs on campus
U_i	number of users associated with AP_i
$1/\mu_i$	the expected user stay time at AP_i
λ_i	arrival rate to AP_i
ρ_i	load of AP_i where $\rho_i = \lambda_i/\mu_i$
γ_i	exogenous arrival rate to AP_i
p_{ij}	empirical probability of moving from AP_i to AP_j

Table 1: Model Parameters

3. QUEUEING NETWORK MODEL

We model the campus wireless network of APs as a system of $M/G/\infty$ queues, where each AP is represented by an $M/G/\infty$ queue. We assume that the user arrival rate to each AP is a Poisson process, and each user's expected stay time at each AP is of general distribution. Each AP behaves as if there are infinite number of servers for each queue, and the AP can thus serve an infinite number of users[3]. We first introduce the key parameters in our model in Table 1.

Let the exogenous arrivals to AP_i be modeled as a Poisson process with rate γ_i. We then have the aggregate arrival rate to AP_i:

$$\lambda_i = \gamma_i + \sum_{j \neq i} \lambda_j p_{ji}, \ 1 \leq j \leq N \tag{1}$$

The probability that a user stays at AP_i and leaves the system is $p_{i0} = 1 - \sum_{j=1}^{N} p_{ij}$. Let $\pi(\vec{u})$ be the joint steady state population probability distribution s.t. U_i is the number of users at AP_i in steady state, and $\vec{u} = (U_1, \ldots, U_N)$. The corresponding marginal population probability distribution of AP_i is $P(U_i = u_i) = e^{-\rho_i} \frac{\rho_i^{u_i}}{u_i!}$. Hence, the joint steady state population probability distribution of those APs is of the following product form:

$$\begin{aligned} \pi(\vec{u}) &= P(U_1 = u_1, \ldots, U_N = u_N) \\ &= \prod_{i=1}^{N} \frac{\rho_i^{u_i} e^{-\rho_i}}{u_i!}, u_i \geq 0; 1 \leq i \leq N \end{aligned} \tag{2}$$

The predictions of user occupancy at each AP from our model are validated in the next section.

4. MODEL VALIDATION

We validate our model against the empirical trace data by considering the following metrics: user population distribution, user sojourn time distribution (i.e., session time of user's visit in the system), and the number of distinct APs visited by a user during a session.

4.1 User Population Distribution

Figure 2(a) shows the user population distribution of the most heavily loaded APs around campus. The red dashed line is the result predicted by the model (with the load, ρ_i, estimated by λ_i/μ_i of AP_i), while the solid line is the empirical population distribution. We note a good match between the model-predicted and empirically-observed values. To measure the closeness of the predicted results and the empirical ones, we use the Kolmogorov-Smirnov goodness-of-fit test (K-S test). The Kolmogorov-Smirnov test is used to determine whether a hypothesized distribution matches the empirical distribution, and is not sensitive to the bin-

---

[2] We also removed records from some of the APs that started to malfunction, repeatedly reporting the same number of users to the controller, from about Feb. 15, 2004 to the end of the measurement.

[3] In IEEE 802.11 specification, there is no user association limit for an AP. However, in practice, most AP manufactures have their recommendations for AP maximum capacity.

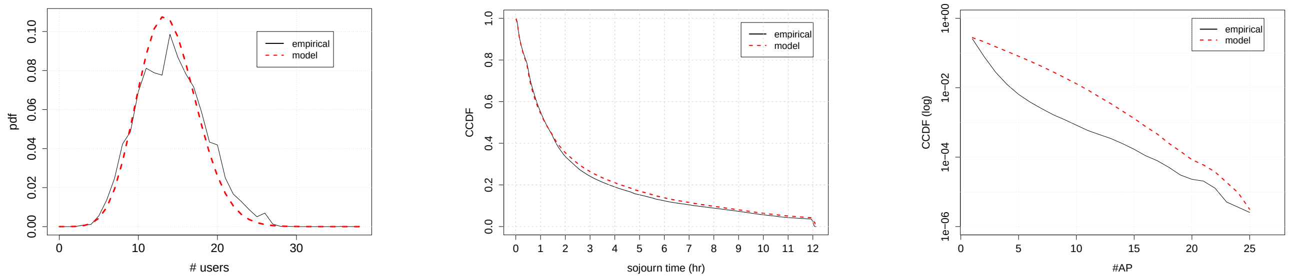

(a) (b) (c)

Figure 2: (a) user population distribution (b) sojourn time distribution (c) distribution of # visited distinct APs

AP Type	# passed	# in total	Ratio
Residential	204	212	96.26%
Academic	130	152	85.52%
Administrative	68	70	97.14 %
Social	45	45	100%
Library	40	49	81.63%
Athletic	19	19	100%

Table 2: Fraction of APs passed K-S test

ning of our data (in our case, the number of users), as it is in Chi-square test [4][5].

In this paper, we set the significance level of K-S tests to 0.05 (i.e., with 95% confidence level). Table 2 shows the acceptance ratio of K-S tests, that the predictions of our hypothesized model has a goodness-of-fit to the empirical distribution of user occupancy.

4.2 Sojourn Time Distribution

We validate the model by comparing the sojourn time (i.e., the duration of a user's session length) distribution generated by the model against to that of the empirical data. We ran simulations with exogenous arrival rate γ_i to AP_i determined from the trace. When a user associates to AP_i, he/she stays a time drawn from the stay time distribution of AP_i, and moves to AP_j according to the empirical probability p_{ij} computed from the trace. A user will leave the system when the sojourn time hits 12 hours or the next movement is to leave the system (i.e., $j = 0$). Figure 2(b) shows the sojourn time distribution of the empirical data and the one predicted from our model. we again note a good match between the model-predicted and empirically-observed values.

4.3 Number of Distinct Visited APs

In addition to looking at user-level sojourn time, we are also interested in a user's trajectory (the sequence of visited APs that a user moves in the system). We first look at the number of visited APs of a user as a metric. Due to the ping-pong effect[4] that frequently occurs, it is hard for us to show how precise our model is by comparing the *total number* of visited APs of a user, since one might alternatively associate with 2 or more APs every 5 minutes. Hence, instead of looking at the number of total visited APs of each visit, we focus on the *number of distinct visited APs*.

In the trace, we found that there is a great portion of mobile users who turn on their laptops, associate with an AP, use the Internet and then leave the network without making any transition. To model this detailed behavior, we found that mobile users should be modeled using 2 classes: those corresponding to exogenous arrivals making their first visit to the system and then leave, and those corresponding to arrivals from another AP, who have already visited at least one AP. For this first class of users, we associate a new parameter $P_s(i)$ (determined empirically from the traces), which is the probability that an exogenous arrival leaves the system after visiting only one AP. With probability $1-P_s(i)$, an exogenous arrival will proceed on to another AP upon leaving the first AP that is visited. With this additional distinction, Figure 2(c) plots the distribution of number of distinct visited APs from the trace and the model.

5. CONCLUSION

In this paper, we found that a simple, parameterized open queueing network of $M/G/\infty$ queues can capture user mobility and predict several network-level and user-level performance metrics of interest. Using the empirical dataset from the Dartmouth trace [3], the model-predicted performance was validated against the performance observed from the trace. Our parameterized model accurately predicts the user population distribution at APs and mean sojourn time. Our more detailed finding is: with two classes of users specified, those who only visit a single AP, and those with more mobility (with at least one transition among APs), we can more accurately predict the user's path length (the number of distinct visited APs during a session).

6. REFERENCES

[1] Community resource for archiving wireless data at dartmouth (crawdad) http://www.crawdad.org/.

[2] M. Balazinska and P. Castro. CRAWDAD data set ibm/watson (v. 2003-02-19). Downloaded from http://crawdad.cs.dartmouth.edu/ibm/watson.

[3] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of MobiCom*, pages 187–201. ACM, 2004.

[4] R. Hogg and E. Tanis. *Probability and Statistical Inference*. Prentice Hall, 6th edition, 2001.

[5] W.-j. Hsu and A. Helmy. On nodal encounter patterns in wireless lan traces. *IEEE Transactions on Mobile Computing*, 9:1563–1577, November 2010.

[6] M. McNett and G. M. Voelker. Access and mobility of wireless pda users. *SIGMOBILE Mob. Comput. Commun. Rev.*, 9:40–55, April 2005.

---

[4] The ping-pong effect is a phenomenon that a wireless device does not associate with just one AP and stay with it for a while but instead associates with a small, fixed set of APs[3].