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

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

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

Don Towsley
Computer Science Deptartment
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 such a simple queueing model can indeed be used to accurately predict a number of performance measures of interest.**

I. 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. Our work differs from much previous theoretical work in that we characterize user mobility using trace data, and use selected measured values as input parameters to our model. Our more fundamental goal here is to determine whether simple, parameterized queueing network models can be used to accurately predict various aspects of network- and user-level performance.

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 input are the mean arrival rate, mean residency time at each AP, and the inter-AP transition probabilities, which are determined from empirical traces. The question we address is the following: can such a simple open queueing network model, with its many simplifying independence assumptions, accurately predict various measures of network-level performance (e.g., user population distribution at APs) and user-level performance (e.g., mean sojourn time) in the wireless network?

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, abstract model can indeed be used to accurately predict a number of performance measures of interest. However, in order to accurately predict certain performance measures, we find that it is necessary to distinguish 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 II describes the traces we use, and how we pre-process the trace. Section III presents our proposed queueing network model, which is validated in Section IV. Related work is discussed in Section V and Section VI concludes this paper.

II. THE TRACES

There are several publicly available traces of long term user activity in a wireless LAN (WLAN) [2] [3] [4]. 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. The trace we will use to parameterize our model, and against which we will validate model predictions, is the Dartmouth trace [3], which records wireless user activity for a 17-week period, from 11/2/2003 to 2/28/2004.

A. Trace Description

To better understand the trace, let us begin by defining several key terms in wireless networking. When a mobile user first enters the system, searching for an available AP to connect to, a validation of the user's identity (user name and password) is performed. If a user is validated, then we say the user is *authenticated*, and is then *associated* with the connected AP. If a user moves from one place to another, he/she will need to find a new AP to connect to; the user will first need to *disassociate* from the old AP, and *roam* to (and associate with) a new AP. Last, if a user decides to go offline, and manually disconnects the wireless device from the network, we say the user first *disassociates* from the current AP, and then will be *deauthenticated* while manually logging-off the system or gracefully shutting down the device.

The Dartmouth trace consists of syslog events and Simple Network Management Protocol (SNMP) polls. The syslog contains records sent from APs to a central server whenever mobile users (clients) authenticate, associate, roam, disassociate, or deauthenticate. We found, however, that the syslog is an unreliable source for observing users' disassociations from an AP - users rarely disassociate their devices from an AP manually, and rarely shut down their laptops gracefully (which results in explicit deauthentication). Therefore, the exact timing of a user's departure from the network is not known on the basis of the syslog alone. The SNMP trace, on the other hand, passively records useful related information. The wireless LAN's mobility controller (i.e., a central server that coordinates all APs on campus) polls each AP every five minutes. At each such SNMP poll, an AP reports to the controller those clients that are currently associated with that AP. Although this information still 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.

B. 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 removed all weekend, holiday, and inter-session periods as well[1].

We 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 becomes disassociated from all network APs. Although we extract only the 7 am to 7 pm non-holiday-non-weekend periods from the 24/7 traces, there are users who had joined the network prior to 7 am who appear in the traces, and users in the 7-to-7 extracted trace who leave the network after 7 pm. We handle these boundary cases by artificially having all users who are active at 7 am arrive to the network at 7 am, and have all users who are in the network at 7 pm leave the network at that time. This pre-processing step results in a bulk-arrival at 7 am, and a bulk-departure at 7 pm. We are currently investigating a mixed (open and closed) queueing network model where these bulk-arrivals/departures are replaced by a closed population of users who always remain in the network. Details about this mixed queueing network model of mobility can be found in [5].

The processed trace contains 547 APs across 6 main different types of buildings (as listed in Table I), with 6,452 distinct users.

As discussed in Section II-A, 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)

[1]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.

are considered to be two separate sessions for that user.

Figure1 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.

Besides the issue of identifying departures, we also observe that in the 5-minute SNMP reports collected by the controller, a specific user is sometimes concurrently associated with multiple APs. This occurs when a user is associated with one AP for part of the five-minute interval, and then a different AP for another part of the same five minute interval. When such conflicts occur, we assign this user to the AP that most recently reported the user as being associated with that AP and remove the user from other APs for this time interval. We process the trace from start-to-finish, sequentially applying this rule as needed.

After resolving the issues above, the exact departure time of the user from one AP (and its departure from the network or its associating with a new AP) is still unknown - we only know the five-minute interval in which this departure occurred. In this case, we randomly choose the departure time from a uniform distribution across this five-minute interval.

The last issue that arises when processing the trace is the presence of a "ping-pong effect" - the phenomenon where 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 [6]. Since it is difficult to identify precisely when a user starts to exhibit the ping-pong effect, and how many APs are involved in this effect, we do not consider this phenomenon in this paper. Instead, we treat each movement as a regular transition from one AP to another.

III. 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. When a user associates with an AP, he/she will be served immediately, regardless the bandwidth each user is allocated. Each AP behaves as if there

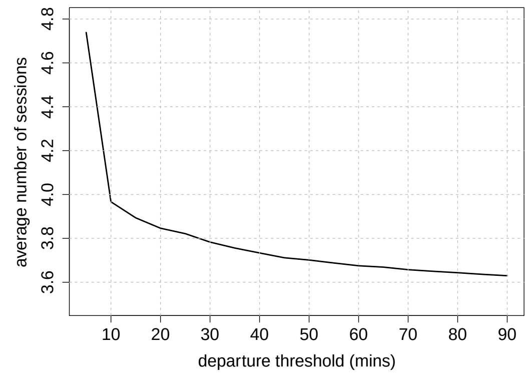

Fig. 1. Average number of sessions for various departure length threshold.

are infinite number of servers for each queue, and the AP can thus serve an infinite number of users.[2]. Before discussing the details of our model, let us first introduce the key parameters in our model:

- 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 a user moving from AP_i to AP_j. P denotes the corresponding $N \times N$ transition matrix.

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:

$$P(U_i = u_i) = e^{-\rho_i} \frac{\rho_i^{u_i}}{u_i!} \tag{2}$$

Hence, the joint steady state population probability distribution of those APs on campus is of the following

[2]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.

product form:

$$
\begin{aligned}
\pi(\vec{u}) &= P(U_1 = u_1, \ldots, U_N = u_N) \\
&= \prod_{i=1}^{N} \frac{\rho_i^{u_i} e^{-\rho_i}}{u_i!}, u_i \geq 0; 1 \leq i \leq N
\end{aligned} \tag{3}
$$

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

IV. 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. An analytical study of user mean sojourn time and average path length (number of visited APs) is in [5].

Algorithm 1 Simulation

Input: γ_i, $\frac{1}{\mu_i}$, P
Output: sojourn time, list of traversed APs per visit

for $visit := 0$ to ∞ **do**
  empty trajectory list
  initial arrival at AP_i with probability: $\frac{\gamma_i}{\sum_{i=1}^{N} \gamma_i}$
  add AP_i into trajectory list
  sojourn time:= $\frac{1}{\mu_i}$
  while sojourn time $\leq$ 12 hours **do**
    make a transition to AP_j according to P_{ij}
    if $j! = 0$ **then**
      increment sojourn time by $\frac{1}{\mu_j}$
      add AP_j into trajectory list
    else
      exit while;
    end if
  end while
  if sojourn time $>$ 12 hours **then**
    sojourn time := 12 hours
  end if
  record sojourn time, trajectory list for this $visit$
end for

A. User Population Distribution

At first glance, we would like to know how well the model matches the empirical results. We found that the most heavily loaded APs are those in residential buildings, followed by academic buildings. Figure 2 shows the user population distributions of the five most heavily loaded APs in residential buildings. In each plot, the red dashed line is the result predicted by the model (with the load, ρ, 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.

AP Type	# passed K-S test	total No. APs	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%
Total	506	547	92.5 %

TABLE I

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 (i.e., predictions from our $M/G/\infty$ queueing model) matches the empirical distribution, and is not sensitive to the binning of our data (in our case, the number of users), as it is in Chi-square test [7].

The Kolmogorov-Smirnov statistic is defined by:

$$D_n = \sup_x [|F_n(x) - F_0(x)|] \tag{4}$$

where $F_n(x)$ and $F_0(x)$ are the cumulative distribution functions (CDFs) of the empirical data and predictions from our model. D_n, the K-S statistic, is the least upper bound of all point differences $|F_n(x) - F_0(x)|$. In other words, the smaller D_n is, the closer the two distributions are. The exact distribution of the statistic D_n and its corresponding significance level can be found in [7]. In this paper, we set the significance level of K-S tests to 0.05 (i.e., with 95% confidence level). Table I 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.

B. Sojourn Time Distribution

In this subsection, 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 as shown in Alg.1, with exogenous arrival rate γ_i to AP_i determined from the trace. When a user associates to AP_i, he or 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., $j = 0$).

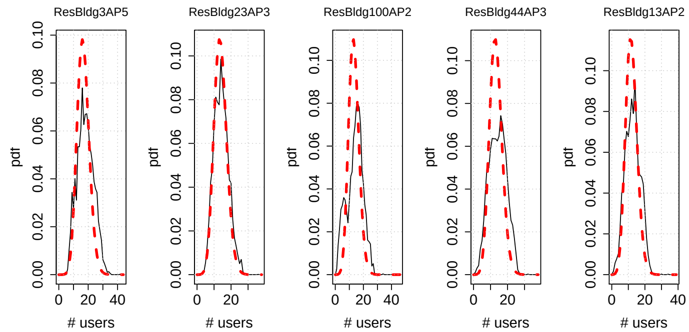

Fig. 2. Population distribution of the most heavily loaded 5 APs in residential buildings.

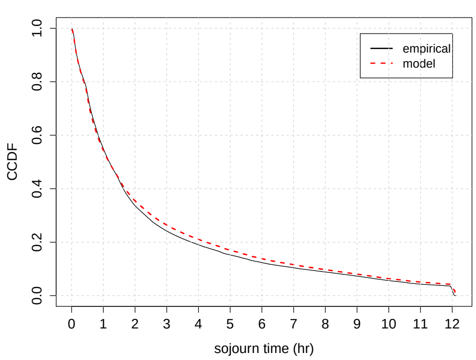

Fig. 3. Sojourn time distribution

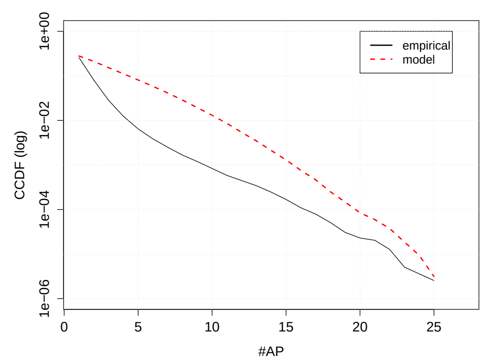

Fig. 4. Distribution of number of distinct visited APs

Figure 3 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.

C. 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 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. For this behavior, we say that the user has a *single stay* in the network. 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 $P_s(i)$ (which is 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 when first entering the network. With this additional distinction, Figure 4 plots the distribution of number of distinct visited APs from the trace and the model. The agreement here between the model predictions and the empirical observations is not as close as for our previous metrics. We are currently investigating these differences.

V. RELATED WORK

As wireless networks have become ubiquitous, researchers have been trying to understand fundamental properties of these networks by analyzing traces from WLAN users [2] [3] [8]. However, most of these efforts focus on obtaining specific statistics of users, or certain user behavior patterns [6] [8] [9]. To our knowledge, this paper presents the first analytical model with a simple open $M/G/\infty$ queueing network for a campus wireless network with empirical validation. Previous works on network modeling with $M/G/\infty$ queues are mostly focused on wired networks or modeling network applications such as peer-to-peer networks [10] [12] [13].

The most closely related works are theoretical models from cellular networks [14] [15]. Kim et al. [14] developed a multi-cell mobility model for cellular network under the assumption that cell dwell time has hyper-Erlang distribution. They started with a $M/G/c/c$ model to compute the loss probabilities of calls, but due to the computational complexity, they down-model the service time from a general distribution to an exponential distribution (with the property of Erlang insensitivity of service time). Hence they, in the end, use $M/M/c/c$ queues to model cellular network mobile users. On the other hand, Ashtiani et al. [15] model the spatial traffic of cellular networks and focus on *active users* only (users always in connected mode). Their work leads to a closed queueing network with fixed population.

All these models were designed for different assumptions, to exhibit mathematical properties, with no empirical validation. Our proposed model, on the contrary, starts with empirical trace data, and can accurately predict various network and user-level measures in a simple yet efficient manner.

VI. 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, with input of mean arrival rate, mean residency time at each AP, and inter-AP transition probabilities determined from the empirical trace, 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).

Acknowledgment

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

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

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

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

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

[5] Yung-Chih Chen, Jim Kurose, and Don Towsley. A mixed queueing network of model of mobility in a campus wireless network. Technical report, Department of Computer Science, University of Massachusetts at Amherst http://gaia.cs.umass.edu/networks/papers/Chen11_mixedQueue.pdf.

[6] Minkyong Kim and David Kotz. Periodic properties of user mobility and access-point popularity. *Personal Ubiquitous Comput.*, 11:465–479, August 2007.

[7] Robert Hogg and Elliot Tanis. *Probability and Statistical Inference*. Prentice Hall, 6th edition, 2001.

[8] Wei-jen Hsu, Debojyoti Dutta, and Ahmed Helmy. Mining behavioral groups in large wireless lans. In *Proceedings of the 13th annual ACM international conference on Mobile computing and networking*, MobiCom '07, pages 338–341. ACM, 2007.

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

[10] F. P. Kelly. Networks of queues with customers of different types. *Journal of Applied Probability*, 12:542–554, 1975.

[11] James R. Jackson. Jobshop-like queueing systems. *Management Science*, 10:131–142, November 1963.

[12] Daniel S. Menasche, Antonio A.A. Rocha, Bin Li, Don Towsley, and Arun Venkataramani. Content availability and bundling in swarming systems. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, CoNEXT '09, pages 121–132. ACM, 2009.

[13] Di Wu, Yong Liu, and Keith W. Ross. Queuing network models for multi-channel p2p live streaming systems. In *INFOCOM*, pages 73–81, 2009.

[14] Kyungsup Kim and Hoon Choi. A mobility model and performance analysis in wireless cellular network with general distribution and multi-cell model. *Wirel. Pers. Commun.*, 53:179–198, April 2010.

[15] Farid Ashtiani, Jawad A. Salehi, and Mohammad R. Aref. Mobility modeling and analytical solution for spatial traffic distribution in wireless multimedia networks. *IEEE Journal on Selected Areas in Communications*, 21:1699–1709, 2003.